UNIVERSITY OF SOUTHAMPTON

The Public Health Analogy in Web Security

by

Huw Fryer

A thesis submitted in partial fulfillment for the degree of Doctor of Philosophy

in the

Faculty of Physical Sciences and Engineering

School of Electronics and Computer Science

March 2016

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING, SCIENCE AND MATHEMATICS
SCHOOL OF ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

by Huw Fryer

Traditional law enforcement methods have proven inadequate against the current levels of cybercrime we are experiencing. This is due to the ease of automating attacks, and also that even a single jurisdiction prepared to ignore or unable to prosecute cybercriminals mean that they are usually beyond the reach of local law enforcement. This has led to different analogies to attempt to describe the phenomenon, and one of these is that of public health. In the past, this was used to describe the propagation methods of computer "viruses", which exhibited similar characteristics to biological viruses. Whilst other malware also had a similar propagation pattern, these no longer apply given the popularity of drive-by downloads, where Web pages attack users who visit them. A consequence of this new method of propagation is that "infected" machines do not have any contagion, so one infected machine on a network does not mean that an other machine on the network will become infected as well.

This thesis proposes a novel interpretation of the public health analogy, which focuses on the notions of efficacy and rights, so that these guidelines can continue to be used. This is considered in the context of the major stakeholders who could intervene in the drive-by download process, where it is concluded that hosting providers are best placed to intervene to make a difference. It is proposed that they should proactively search for vulnerable websites they host, and warn the operator, implementing blocking procedures if the operator does not respond. An agent based model is then used to assess the efficacy of such an intervention.

Contents

List of Figures

List of Tables

Listings

Acknowledgements

Doing a PhD can often be a lonely endeavour, but nonetheless I have had a great deal of support over the last few years as I have made this journey. There are far too many people to mention everyone individually, but I wish to take this opportunity to mention some of the people I'm particularly grateful to.

First of all, to Les Carr and everyone at the DTC, for giving me the opportunity to do the PhD and have a great time in Southampton. I have really enjoyed the experience, and have met some fantastic people along the way.

Thank you to my supervisors, Tim Chown and Sophie Stalla-Bourdillon who have been incredibly helpful in giving me ideas, and politely helping me get rid of the overly stupid ideas I had myself. Roksana Moore as well, who supervised me for the first two years and put me in a good position to finish it.

Everyone in the WAIS group has been great to be around, a constant source of ideas and inspiration. A particular thanks to Rikki Prince and Yvonne Howard who were prepared to give me a chance to work as a teaching assistant, and Jon Hare who helped me greatly with the simulation chapter.

Outside of work, there have been people who have helped on a personal level. Everyone who was prepared to listen to me complain, in particular Areeb, Dalal and Charlie. My housemates over the last few years: Sheena Au-Yeung, Declan Kiernan, Jenny Bannister and Sophie Carr, who have taken the quirks of a PhD student in their stride and been incredibly patient.

Most importantly, Dixie Thamrin, who met me six months before the end. You kept me going at the most difficult time, prevented me from suffering a complete breakdown, and have made, and continue to make me really happy. Thank you so much.

Finally, what acknowledgement section would be complete without a special thanks to my parents? Thank you for supporting me for so long, letting me hide out at your house to finish writing, and most importantly believing in me when I didn't believe in myself. Looks like you were right!

Chapter 1

Introduction

I am the thief
And you may run
Yes, and you may hide
But I will get you yet
Yes I will
Because of you, 'cause it's just your pride
Just your pride! (Can, 'The Thief' (1981))

1.1 Background

Like any other technology, computers have turned out to have a significant amount of use by criminals as well as legitimate use. The problem has been more severe than with previous technology, due to the combination of two factors. Firstly, computers have greatly increased the speed at which a task can be automated. Secondly, the Web has got rid of the majority of the geographic limitations towards finding more victims so this automation can be put to good (or rather malicious) use. The combination of these two factors make it difficult to deal with crime in the traditional manner.

An example of this automation in action comes from the volume of spam, which despite having reduced considerably from a high of 92.6%, still represents 75.2% of all emails(Trustwave, 2013). The main way that criminal groups are able to maintain infrastructure which can send this volume of spam, or perform other undesirable actions is through the use of malicious software (malware). Malware takes over a victim's computer, and having done that can target either the users, or recruit them into a botnet, i.e. a distributed network of computers which is of great to value to an attacker.

A compromise may result in an infected system that is used in multiple criminal activities, and the cumulative effect of these activities can be devastating. This research

explains how the phenomenon of drive-by downloads has evolved to become a significant threat to both Internet users and third party systems. Targeting the users might include something as simple as altering search results to gain advertising revenue, or spying on the browsing habits to target adverts. More seriously, it can steal credentials to online banking; or render a user's computer unusable (e.g. through encrypting all their files) unless they pay a ransom. The ability for distributed computing which exists with botnets offers the opportunity to conduct activities such as distributed denial of service attacks; sending spam; and more recently mining bitcoins.

To effect a compromise via a drive-by, a criminal will create a malicious Web page which, when visited, attempts to exploit vulnerabilities on the users' computer automatically. This will typically be done by compromising otherwise legitimate websites, which means that traditionally 'suspect' sites are not the only ones which will include this malicious software (Yen et al., 2014). Websites such as those which include pornography are actually responsible for a reasonable amount of malware (Provos et al., 2007; Wondracek et al., 2010), but are far from the only ones that do. In addition, in contrast to "worm" based malware, such attacks are invisible to the operator of the victim's network, and Web browsing offers a way through the firewall to increase the probability of infection. These can be used in phishing attacks, or by taking advantage of trending topics, which increase the risk of exposure significantly (Moore et al., 2011a).

1.2 Traditional Approaches to Crime

Traditionally, after a crime has been committed, the police would expend resources in performing investigative work and catch the criminal. A sufficiently high conviction rate would be sufficient to ensure that the criminal justice system provided a deterrent to make sure that the criminal element in society remained at a manageable level. Although international crime has been possible for centuries, for the most part geographical limitations meant that any crime committed would generally concern people in the same area.

As technology develops, occasionally new legislation is required, as has been the case to cope with the increased use of computers in society. The two main statutes in relation to Web based crime in the UK are the Computer Misuse Act 1990, and the Fraud Act 2006.

The Computer Misuse Act (CMA) was a private member's bill which was introduced in 1990, to criminalise hacking following the failure to convict in *R v Gold and Schifreen* under the Forgery and Counterfeiting Act. Section 1 introduces the offence of unauthorised access to computer material, if a person "*causes a computer to perform any function with intent to secure access or causes such access to be secured*". Section 2 consists of unauthorised access, but with the intention of committing further offences.

Section 3 governs dissemination of viruses and denial of service attacks (modified by s36 Police and Justice Act 2006), and Section 3A governs the creation or supply of malware and other similar items for use in computer misuse(introduced by s37 Police and Justice Act 2006).

Where "phishing" is concerned[1]. This simplified the law related to fraud, creating a general "fraud" offence, with ways of committing it. The focus has shifted considerably from a results based crime, to one of representation. Section 2 "Fraud by false representation" is the provision most likely to be the one used for phishing. The key element is intention, requiring an intent to make a profit or cause another to make a loss. Possession of malicious software for use in fraud(s6), or making/supplying is also an offence(s7).

With crime being limited by scale and geographic area, acquisitive crimes would require a series of identifiable victims within an area the attacker could reach. Increased reliance on the Web and Internet means that this assumption no longer applies, and provides advantages for criminals in that they can attack someone in a completely different country with the same ease as they could attack someone in their own country[2]. The scalability enabled by computers has also increased opportunities for criminals, that the same attack can be used against millions of victims, meaning that even attacks with a low yield and low success rate can still generate a significant profit.

Since criminals can scale attacks, and target people in different countries, two significant issues get in the way of law enforcement being able to pursue these attackers. Firstly, either the victim or the attacker is outside the jurisdiction any law enforcement operation is based. If the attacker is not based within the jurisdiction, they will not necessarily have the power to pursue the attacker. For example, if they are based in a different country the perpetrator may not have committed a crime, such as in the case of the the "ILOVEYOU" virus Chien (2001), or law enforcement might be faced with uncooperative local law enforcement. Even with co-operative nations, they will be forced to justify the expenditure of attempting to solve a crime of a victim who does not pay tax in their jurisdiction. Secondly, the scalability means that losses are often distributed over many victims with a small amount of loss to each individual victim or company. This means that if the crime is reported at all, there is still the issue of attempting to justify expenditure on a case with such small losses.

In addition, the process of gathering evidence for the purpose of a prosecution is also beset by procedural difficulties, in particular where there is no mutually binding assistance provisions between the two jurisdictions (Weber, 2003). These limitations have been

[1]Phishing is discussed in Section 2.1.4, the main Act to be considered is the Fraud Act 2006. For present purposes, phishing can be described as an imitation of a trusted institution for the purposes of obtaining user credentials

[2]That said, despite the ease of attempting, it is likely to be less convincing if the attacker lacks knowledge of local cultural conventions or language.

recognised by the Council of Europe, and their response was to create the Convention on Cybercrime in an attempt to harmonise legislation between countries and improve co-operation. Alongside the laws to be harmonised, consisting of computer integrity, fraud and content style offences (Art 2 – 13), procedures related to investigation and prosecution were introduced, as well as requirements for international co-operation.

Unfortunately, whilst this may improve co-operation between the states involved – and maybe even put political pressure on other states to participate, there remains a weakest link problem. In a weakest link game, each player has a "veto" over the possible overall collective success(Hirshleifer, 1983). In this case, as long as there are states who do not deal with cybercrime (whether through incompetence or corruption), then it is possible for criminals to continue their activities without fearing law enforcement. If there are a lot of weak links then they can quickly shift operations or infrastructure to the second weakest link and continue as before. Increasing the states ratifying the convention will gradually reduce the problem, but as long as there is even one state who turns a blind eye then the overall effectiveness will be minimal. The current total of signatories of the Cybercrime Convention is 53 of which 45 have ratified (Council of Europe Treaty Office, 2015).

Anecdotal evidence suggests that a considerable amount of attacks originate in Eastern Europe, or China. For example, Conficker A searched for a Ukrainian keyboard layout on victim computers and would die if it found it Porras et al. (2009b), and McCombie et al. (2009) found evidence of a connection to organised crime in Russia. It is possible that any connections such as this make it more difficult for prosecutions. In 2011, Microsoft's investigation into the ringleaders of the Rustok botnet led them to Russia and offered $250,000 for information leading to their arrest (Meisner, 2011), but this was not enough for anyone to offer information about it. More recently, accusations have been made by various parties in the USA recently about China sponsoring cybercrime. This has included issuing warrants for the arrest of members of the Chinese military for stealing corporate secrets (Schmidt and Sanger, 2014). Both Russia and China have declined to be signatories to the Cybercrime Convention, and they represent a significant portion of the world's population and (arguably) are the source of a considerable amount of computer related crime.

Whether the Cybercrime Convention can remain applicable given the general difficulties in attribution, and the availability of obfuscating criminal tools is another potential issue (Maurushat, 2010). Technologies such as Tor[3], and the ability to use zombie computers as proxies[4] make detection very difficult – unless the criminal in question makes a mistake. This has been known to happen – it was reported that the operator

[3] https://www.torproject.org/

[4] This involves hijacking another user's computer, so that the criminal activity appears to be coming from somewhere else. See Section 2.3.1 for more details

of the "Silk Road" marketplace made some basic security errors which allowed him to be caught(Goodin, 2013)[5].

Despite the difficulties with pursuing criminals like this, there are periodic examples of attackers getting caught. In 2013, the suspected author of the Blackhole exploit kit was arrested in Russia (BBC, 2013), which led to a 15% reduction in its market share the following year (Trustwave, 2014). Other examples include those suspected of using the ZeuS crimeware kit (BBC, 2010), and ransomware attacks (BBC, 2013). These successfully managed to obtain international co-operation, and countries such as Russia and China have been included in these sorts of attacks (e.g. (BBC, 2013, 2012a)).

Anderson et al. (2012) suggest that economically speaking it is far more efficient to expend resources on actually catching criminals rather than trying simply to prevent the crimes from occurring. They point out the disparity in costs between the amount of money spent on prevention, compared to the profits made by attackers and how much money is lost by society. Whilst not disagreeing with them, this research is working on the assumption that the weakest link problem is currently insurmountable, and that supplemental approaches are required to make the problem more manageable.

1.3 Alternative and Supplemental Approaches

The economic elements of security have been the subject of a considerable amount of study, in particular following research by Anderson (2001), and Varian (2000), and is now referred to as "security economics". The discipline analyses the fact that security decisions are often not necessarily made with the aim of improving security, but could be used as a means of dumping liability onto another party, or could even counter-intuitively lead to worse security. One example is that of trust certificates, where Edelman et al. (2006) found that due to a lack of diligence by the providers of such certificates, the uptake led to a state of adverse selection where the majority of websites which did use these trust seals were in fact malicious.

Some of the concepts from security economics will play a role, particularly some of the models discussed by Florencio & Herley in relation to decisions by attackers Herley (2010); Herley and Florêncio (2009). However, the economic problems are widely known and studied (Moore et al., 2009), so this thesis will make use of some of the literature and models but will not attempt to make a contribution to the field.

Instead, this thesis will focus on the strategies which a government could apply to attempt to solve the problem. Rather than attempting to solve the problem piecemeal, by prosecuting one criminal at a time, the more general idea of an analogy to guide strategy will be applied. An analogy provides a conceptual understanding of the problem, and

[5]Although there is some dispute about the case put forward by the FBI, see e.g. (Krebs, 2014).

by extension offers guidance as to how it can be solved (Betz and Stevens, 2013). There are two main analogies which are commonly used and will be considered in this thesis in relation to cybercrime: those of war, and of public health.

The space, or war analogy is based around the idea of strong outer defences in order to prevent attackers from coming in. Similar rhetoric is also used in regards to the phrase "attack", and arms race, in an "us vs them" description (Betz and Stevens, 2013). This approach has its limitations for several reasons. To begin with, attribution is very difficult when it comes to any attacks, so a military style response offers little deterrence due to the difficulty of responding to the parties involved. Given that it is possible for individuals, not affiliated to any country's military to also conduct these attacks (with or without the authorisation or knowledge of the country), then this adds to the complexity. Whilst it can make for great rhetoric and/or scaremongering (e.g. (Gross, 2011)) in terms of concrete policies or strategies it is of limited value.

The alternative analogy, that of public health has far more flexibility (Charney, 2012), and enables government to make decisions which will affect the overall "health" of the nation, possibly at the expense of individual rights. Traditionally, this strategy was centred around the use of epidemiological concepts to limit the spread of viruses and worms, owing to the similarity of the spread of the two e.g. (Kephart and White, 1993). Also inherent in the analogy was the need for recognition that there was no way which the spread of malware could be completely contained, and that it should be used as a means of simply managing the problem to a controllable level (Murray, 1988). The worms and viruses which caused problems in the last two decades no longer represent the issue which they once did. Modern malware will frequently either induce victims to download it voluntarily, or by attacking the browser once they visit a malicious website. This pattern means that the traditional epidemic models used to study the spread of malware no longer apply to the same extent, since the mere fact that one computer on a network is infected does not necessarily mean that another computer on the network will also be infected.

However, that is not to say that the public health analogy has lost its utility. Like with war, the preservation of health is regarded with suspicion from those concerned about state encroachment into individual rights (Epstein, 2003b), and there has been a considerable amount of debate as to exactly how public health should be defined and what power should the government be granted in order to ensure the health of the majority of the population (Coggon, 2012). Modelling such as this can also be used as a tool to inform policy, and to decide the appropriate balance to be made between security and rights (Edwards et al., 2012). This thesis analyses the similarity between creating public health policy, and how it could be used to create appropriate law in relation to cyber security. In particular, the focus will be on efficacy and rights as an appropriate method to decide on policy.

It is something which could have general application, but for this research it is limited to drive-by downloads, and related Web based malware. This is one of the major problems currently afflicting the Web, and represents a good example of a propagation method which does not follow the traditional models for viruses and worms.

In order to get the greatest amount of effectiveness from any particular "public health" intervention for the issue of drive-by downloads, Internet intermediaries will play an important part. These have been popular targets for liability in the past, as a result of their visibility, deep pockets, and their ability to have an impact (see e.g. *SABAM v Scarlet* (2011), and other related court cases relating to Article 15 and Article 12 of the E-Commerce Directive), and targets of legislators in order to prevent issues such as copyright infringement (e.g. Digital Economy Act (2010) in the UK) and prevent access to pornography (B.B.C, 2013). These issues will be discussed in more detail, in terms of which intermediary would be the most effective to require action from; and how this could be done in order to ensure that they remain competitive internationally.

1.4 Motivation and Challenges

This thesis is different to many others, in that it's interdisciplinary. The main disciplines are law and computer science, nominally 50% each, but of these there is inspiration taken from economics as well. Cybercrime is a complex problem, and to attempt to do anything to solve it requires understanding from several areas. It is not only law which can solve technological problems, but it is possible to influence behaviour quite considerably.

The idea is not new that the market can influence behaviour, or even that behaviour can be manipulated by placing technological barriers in the way (Lessig, 2006), but it is one which can have significant implications for cybersecurity. Consider the Great Firewall of China, which enforces the wishes of the government that information not be available for general consumption. Rattray et al. (2010) also cites a story comparing the effect of the Witty worm in U Cal Berkeley and Lawrence Berkeley labs, where the latter only allowed machines onto the network who had installed the patch for this worm. They only got one infection, compared to the former who got 800.

Whilst the idea of doing interdisciplinary research is important, the practicalities of doing so presented some significant challenges. Firstly, the distance between the disciplines is considerable, and as such it is difficult to transfer the details of any particular concept over to the other discipline very well. A key element of the thesis is the simulation I conducted in Chapter 6, as a means of validating the efficacy of the proposal. Unfortunately, the background required to be able to understand the methodology and analysis of the results is very different to what a legal academic would be expected to know. Even background material in any particular discipline which assumes any knowledge cannot be easily understood by someone in another discipline. For example, most law students

will have heard of *Donoghue v Stevenson* where the modern law of negligence starts. Most computer scientists will have a reasonable idea about how HTML and JavaScript works. However, that knowledge cannot be assumed.

Funding for this research was done on the premise that disciplines be split approximately equally across a particular subject. This has also not proved easy to implement in this case, for several reasons. Whilst there are both technological and legal issues, it is only in a few small areas where they intersect and knowledge of both is important. Given time constraints inherent in doing a PhD, this can mean that the increased breadth occasionally leads to the appearance of rigour being sacrificed in some areas. Whilst interdisciplinary research definitely needs to play an increased role as we seek to solve problems, the equal splitting of disciplines is something which should be reconsidered in the future.

1.5 Hypothesis and Research Questions

The hypothesis for this thesis is:

> The public health analogy remains a viable framework for combating drive-by downloads, even though malware no longer spreads in a manner similar to a virus

There are many different types of cybercrime, so consideration must be given as to why the hypothesis should be restricted to a specific example of a type of cybercrime rather than something more general. As Section 1.3 discussed, the public health analogy as traditionally thought no longer has the same utility given the change in methodology by criminals, so what's to stop the same from happening to drive-by downloads? Although it is conceded that strategies could change again, the nature of the Web is such that drive-by downloads are a problem which will are unlikely to go away.

Consider the following categories commonly cited computer threats:

1. Insider threat e.g. (Greitzer et al., 2008)
2. Offline crimes supported by computers
3. Denial of service e.g. (Ferguson, 2000)
4. Advanced Persistent Threats (APTs) e.g. (Tankard, 2011)
5. Infection (including drive-by downloads) e.g. (Caballero et al., 2011)

Offline crimes supported by computers, although beset by practical issues of their own, are crimes which are best dealt with by the police in the traditional manner. There is debate about some of the difficulties which they are faced with in terms of obtaining evidence e.g. relating to encrypted communication, yet the crimes themselves have changed very little. There is an established method of going about solving these crimes, and being local many of the challenges relating to the use of ICT and the Web do not apply.

The insider threat is a similar case. Here, inside a company or other entity, a person with a privileged position uses their position to conduct malicious behaviour. This could be done for either personal gain, or with political motives (e.g. whistleblowers such as Edward Snowden). Although serious, this again is an issue best handled by the police, since it is not possible to automate and requires proximity to the target so the pool of possible suspects will be minimal. It is not an issue which the general population has any connection to – except possibly having to pay extra for a product as a company attempts to recover losses.

A denial of service (DoS) attack seeks to cause the services provided by the victims to be made unavailable to their legitimate customers. Although associated with computers and the Internet, it could also be achieved in certain circumstances by continually calling a person's number to prevent them from answering their phone (this happened to Mr Zeran in the infamous *Zeran v AOL* case discussed in Section 5.6). Its application in regards to the Internet makes it far more suitable for analysis than the previous two points, since many attacks rely upon weaknesses in network configuration and are a negative externality of Internet use. However, a significant element of denial of service attacks is the victim computers enlisted in order to make the attack a Distributed DoS (DDoS) attack (See Section 2.3.2.4). The manner in which these computers are enlisted requires that they be infected, and hence participate in these attacks without their knowledge (unless they are voluntarily participating due to political causes.)

An APT is an attack which will often make use of infections to achieve their aim, for example to host the spear phishing email or to simply report upon the content of the victim computer in the company network. Similar to the insider threat, these attacks will often target companies or wealthy individuals with a greater level of effort but with a far greater reward anticipated. It is something which is investigated as situations occur, owing to the less frequent nature of the attacks.

Aside from the first two items in the list, the remainder of the items are all centred around the requirement for a criminal being able to infect computers. Having done this, there are ways that these machines can be monetised, from attempting to grab trade secrets in an APT to the array of techniques discussed for indiscriminate attacks discussed in Section 2.3.2. There are other ways of infecting a computer, but the wide

applicability of being able to infect using drive-by downloads is significant (Caballero et al., 2011) and as such is chosen as the hypothesis for this thesis.

Having chosen this hypothesis, the following research questions will be answered in order to validate the hypothesis:

1. How can the public health analogy be re-imagined given the current threat landscape?
2. Which stakeholder is the best to target in order to minimise the damage caused by drive-by downloads?
3. Are intermediary obligations to combat drive-by downloads appropriate?
4. Can actions by a single country, or group of countries, have a statistically significant effect on the worldwide prevalence of infections from drive-by downloads?

1.6 Structure

Chapter 2 introduces the general issues raised by malware: what it is, and how it is useful to criminals. The different propagation methods are discussed, and the ways in which they differ from each other, prior to introducing drive-by downloads which is the main focus for the rest of the thesis. The drive-by download issue is divided into two broad problems: exploitation of the server (website), and exploitation of the client (browser). Methods to exploit the website are explained, such as SQL injection cross site scripting (XSS), and how these can be used to then exploit the client once the user visits the compromised website. The state of the art relating to the defence against these attacks are also presented, and it is explained why this is not enough. The chapter finishes by introducing the specialisation and "cybercrime as a service" which exists in the criminal economy. Methods in which the criminal can make use of infected machines, such as recruitment to a botnet, ransomware and theft of banking credentials are introduced.

Chapter 3 introduces some of the economic inefficiencies associated with security, and makes the claim that market forces alone are inadequate to ensure that adequate security standards are maintained. The various stakeholders who play a part in drive-by downloads are introduced in more detail, and the steps which they could take to mitigate the problem are discussed. Criminals are also included as stakeholders, albeit one who opposes the efforts of the others. The criminal economy is considered in more detail, in particular potential weaknesses which defenders could exploit. Private law is presented as a possible solution to the inability of market forces to solve the problems, and some literature is presented which has attempted to do this. It is concluded however, that private law is inadequate on its own, and offers a case study of Article 4A U.C.C in

the USA, where there have been difficulties in ensuring the optimal level of security investment primarily owing to the high costs of litigation. A class action solution is considered as a means of mitigating this problem although it is considered the benefits may be limited. It is argued that this should be supplemented with state regulation, for which the public health analogy is a good place to begin assessing regulatory choices.

Chapter 4 considers two possible analogies, that of war and that of public health, which could aid government decision-making in relation to security. The war analogy is dismissed, and public health is argued as the best approach to guide security regulation. Some background to "real world" public health is introduced, and the trade-offs which are required in order to best preserve rights but also be able to adequately respond in the event of an emergency. Some of the epidemiological modelling techniques are introduced, and it is discussed how they have been used in relation to Internet security over the past twenty years. More conceptual ideas relating to the role of the state in public health are considered, and how this can use techniques such as epidemiological modelling to inform the decisions they make within a public health style framework.

A novel interpretation of the analogy is introduced here, where the criminals are regarded as the pathogen as opposed to the malicious software. The malicious software is considered to be merely a symptom of the fact that the real pathogen, i.e. the criminal exists. In order to eradicate this, it is proposed that the environment needs to be made more hostile to the criminal, to reduce the viability of cybercrime. The best way to do that, it is proposed, is to minimise the amount of vulnerable websites. However, the importance of emphasising efficacy and rights in deciding on an intervention to allow this to happen is the focus.

In chapter 5, interventions by the stakeholders introduced in Chapter 3 are considered for the potential efficacy, and rights issues. It is concluded that hosting providers are the best target for regulation, both in terms of efficacy and in relation to fundamental rights. It is proposed that legislation should be introduced to compel hosting providers to proactively scan the websites that they host for customers for vulnerabilities and unpatched software; and that they take some sort of action. The exact nature of this required action is left open, although Chapter 6 assumes that they are charged with blocking malicious websites.

This chapter also acknowledges that there are potential dangers in failing to comply with Article 15 of the E-Commerce Directive, and analyses the current case law in the ECJ and the domestic courts and concludes that such an obligation would be consistent with the Directive. The system in the USA under §230 of the Communications Decency Act 1996 is also briefly discussed, in particular §230(c)(2), which provides protection for wrongly taking down content in "good faith". It is concluded that a provision such as this would be necessary as well, given the inevitability of occasional false positives in detecting malicious software.

Given the international nature of the cybercrime problem, it is necessary to consider the issue of whether actions by a single country such as the UK could actually have any global impact on the level of drive-by downloads. In Chapter 6, an agent based model is used to determine the efficacy of an intervention by hosting providers in various countries according to various parameters, such as levels of compliance, effectiveness, and which countries the intervention was adopted in. It is concluded that there is a statistically significant drop in the level of infected machines where the USA or the EU intervene, although this is not the case were the UK to intervene unilaterally. Other parameters were modified, and one other observation was that raising the vigilance – of users or websites – could have a similar result.

1.7 Contributions

Parts of this work have appeared in a paper for the Workshop on the Economics of Information Security (WEIS) 2013 (Fryer et al., 2013), and have been accepted for publication in the Computer Law and Security Review (Fryer et al., 2015). The main contributions of the thesis itself are as follows:

1. Analysed the public health analogy using combined expertise from both law and computer science. This enabled the analogy to be considered on a national and international scale, as opposed to being restricted to individual networks.

2. A novel interpretation of the analogy between public health and Web security is introduced, where the criminal as opposed to malware, is regarded as the pathogen. This is a more flexible framework to work from, and allows concepts concerning the public health analogy to continue to be used, despite the different propagation methods of current Web based malware. In addition, it allows the focus of the analogy to move more towards the impact of the trade-off between efficacy and rights rather than specifically public health style strategies such as quarantine.

3. Empirical data was used to create an agent based model of the efficacy of intervention by hosting providers. Simulations were used to assess the likely effect under different conditions, and demonstrated that it is possible to have a significant effect on levels of malware with only a small amount of countries conducting interventions.

Chapter 2

Background

> *What we've got here is a failure to communicate. Some men, you just can't reach. So you get what we had here last week, which is the way he wants it. Well, he gets it. I don't like it any more than you, men*" Cool Hand Luke (1967)

This chapter will begin by introducing the issues of malicious software (malware) in general, and some of the methods it can use in order to successfully get installed on victims' computers. This will lead to the notion of the drive by download method, which is the focus of this thesis. Drive by downloads will be pointed out as a major threat, due to their ability to attack computers without the victim necessarily engaging in risky behaviour such as viewing pornography, or participating in peer to peer (P2P) downloading.

The role of both the client and the server in Web based attacks will be considered, how the attacks work and how they can potentially be mitigated. Whilst drive by downloads are the focus of the thesis, the problem is considered in the context of a slightly more general issue. When a server is controlled, it is only on some occasions that the attacker will choose to conduct drive-by download attacks (Canali et al., 2013b). Other attacks are possible, such as phishing, and enabling other forms of attack e.g. distributed denial of service (DDoS) attacks, or black hat search engine optimisation (SEO) (John et al., 2011). The cumulative effect of these attacks is significant, and this is offered as a justification for the focus on drive-by downloads and website compromise.

Having established the level of the threat, the criminal economy will be discussed. The increase in specialisation over the past ten years has mirrored the "X as a Service" model, with an attacker able to subdivide elements of the operation to many different operators (Grier et al., 2012). Ultimately, however, a large amount of the economy relies upon the availability of infected machines for various uses (Caballero et al., 2011). These

uses will be discussed, and some of the methods which exist for a criminal to monetise them.

Many types of software could be considered malware, depending on the definition one wishes to use. This could include merely unwanted software, installed alongside another legitimate program (e.g. "spyware" which collects Web browsing data, or "adware" which spams the user with a lot of adverts, possibly based on data collected by spyware). Frequently, this software will be installed without the user's consent, but not necessarily. It may be that in exchange for an offer, or the use of a free piece of software the user chose to give away their data, or be bombarded with adverts in exchange. For this report, any references to malware will include only software which can be used to take over control of a computer without consent[1]. It is malware like this which is used for the creation of botnets, defined as "networks of 'bots', compromised hosts, that are remotely controlled by a [criminal] master host via one or more controller hosts" (Karasaridis et al., 2007).

In order for an attacker to successfully infect a computer, it is necessary for the malicious code to run on the computer, and there are two main ways of making this happen. For an attacker, it is ideal if the victim voluntarily installs the malware on their computer. This is often achieved by "social engineering", which is an attack against the user rather than against the underlying technology. This could be as simple as including malware as part of another application (e.g. peer-to-peer downloading), or some form of phishing or spear phishing attack (see section 2.1.4).

If this is not possible, then it is necessary for the attacker to exploit a vulnerability on the computer, causing the computer to execute instructions of the attacker's choosing.

A vulnerability is a flaw, or "bug" in a piece of software which amounts to a security weakness. Vulnerabilities will be classified as having various levels of severity, e.g. under the Common Vulnerabilities and Exposures (CVE)[2] list. The most serious are those which allow Remote Code Execution (RCE), which will typically be rated 9 or higher, since these vulnerabilities allow an attacker to run their own code rather than the code intended by the application. This is usually done by confusing the program into accepting user input as commands to be executed rather than as data to be manipulated.

An exploit is a piece of code which takes advantage of the vulnerability, in order to run the desired code. In traditional computer based applications, this will often be done by corrupting the memory (One, 1996), but in Web applications there are many other methods with which this can be achieved, and other less serious vulnerabilities which exist.

[1] Because it is possible to "enrol" in a botnet voluntarily, for purposes such as participating in a political DDoS attack

[2] https://cve.mitre.org/

2.1 Propagation

The means by which the malware executes on the victim computer is one thing which needs to be considered, but so too is the means by which it spreads. The following sections will examine some of the different types of malware, and their methods of propagation.

2.1.1 Viruses

Early malware was mostly computer viruses, a phrase now incorrectly used as a layperson's term for all forms of malware. Early viruses would spread from file to file on a computer, and then transfer to different computers through the physical transfer of floppy disks between users (Kephart and White, 1991). This was the most logical way for it to spread, since use of networks with computers was in its infancy. Cohen (1987) introduced a proof of concept virus (see Listing 2.1) which had the ability to replicate itself through any network in which files were shared, and demonstrated that there was no practical way in which the spread could be completely prevented. In the same publication, he defined a virus as "a program that can 'infect' other programs by modifying them to include a possibly evolved copy of itself" (Cohen, 1987).

```
program virus :=
{1234567;

subroutine infect executable :=

        {loop: file - random-executable :=;
        if first-line of file = 1234567
                then goto loop;
        prepend virus to file;

        }
subroutine do-damage :=
        {whatever damage is desired}
subroutine trigger pulled :-
        {return true on desired conditions}
main-program :=
        {infect-executable;
        if trigger-pulled then do-damage;
        goto next;
        }
next;
}
```

LISTING 2.1: Pseudo code of a virus, from (Cohen, 1987)

As the use of email increased, email based viruses became increasingly effective and would spread via contacts in a victim's address book. A notable virus of this ilk was the ILOVEYOU virus which spread to millions of computers in 2002 and caused considerable

disruption, by overwriting a variety of files with the source code for the attack[3] (Chien, 2001). Later attacks developed to "macro" based viruses, which was code which could be included from within Microsoft Office documents (e.g. spreadsheets making complicated calculations).

2.1.2 Worms

A "worm" is another form of self-replicating code. The difference between the two, is that whilst a virus was constrained in having to attach itself to individual files, a worm could install itself only once on a computer and then scan other computers to push itself onto. Rather than relying on physical devices to propagate, it automates the process through using network or Internet connections. No user intervention is required for worm propagation, it simply scans for vulnerable machines and exploits the same vulnerability on each one. As usage of the Web began to grow significantly in the early 2000s, worms were incredibly common and incredibly effective, see e.g. Staniford et al. (2002).

Few of the computers connecting to the Internet had adequate security to deal with these attacks. Firewalls, a piece of security software designed to regulate the flow of network traffic, were not commonly installed, meaning that it was possible to attack computers without any defences. Another issue, was that many computers were directly visible on the Internet to an attacker, since they had globally routable IP addresses, so the targets were easier to find.

Even diligent users were unable to get the appropriate security patches installed in time given the rate of infection (Granneman, 2004). Operating system vendors also took time to adjust to the nature of the threat, in that their products were under such relentless attack. For example it was not until 2003 that Microsoft introduced a regular patching cycle, and even then the update mechanism required users to opt-in rather than be done automatically, which meant that a lot of the time updates never happened

The spread of worms was often measured by the speed at which they successfully managed to infect all vulnerable hosts, e.g. (Staniford et al., 2002). The first major worm was known as the Morris worm, which was written in 1988 by a student in an attempt to see how many machines were connected to the Internet (Orman, 2003). The time most associated with worms, however, came in the early 2000s. Notable worms of this period included Nimda, Blaster, Sasser, Code Red (v 1 and 2) and Slammer. Although slightly later than the others mentioned, in 2008 – 2009, Conficker was an example of one of the more sophisticated worms to emerge during the last decade(Porras et al., 2009a). Despite its rapid spread it did not make many changes to the victim machine, and the bandwidth it expended was very low to be almost undetectable.

[3] It is here where the different definitions of malware become difficult to separate. It could plausibly be described as either a virus or a worm.

At its peak it had somewhere between 5 and 13 million infected hosts, and continues to have between 1 and 1.5 million unique IP addresses attempting to contact the C & C servers daily(Conficker Working Group, 2011). Its spread was as a result of a vulnerability in the Microsoft Windows operating system, known as MS08-067[4]. This vulnerability was to do with Windows handling of remote procedure call (RPC) requests, and failed to handle a specially crafted request, which the advisory claimed, could:

> [A]llow remote code execution if an affected system received a specially crafted RPC request. On Microsoft Windows 2000, Windows XP, and Windows Server 2003 systems, an attacker could exploit this vulnerability without authentication to run arbitrary code. It is possible that this vulnerability could be used in the crafting of a wormable exploit(Microsoft, 2008).

Despite the general success of these worms, their use as a propagation method has since fallen out of favour, and there are several possible reasons for this. The first, is that operating system vendors have caught up with the threats, and have introduced additional security and updates into their products. As such, worm exploits which attacked operating systems - are not so easy to find. Modern operating systems also have firewalls installed by default (Davies, Microsoft), which largely solves the problem of malware pushing onto a machine. Similarly, the depletion of IP addresses also led to the adoption of Network Address Translation (NAT) hardware, which enabled multiple computers on a local network to share the same IP address on the Internet (Srisuresh and Holdrege, 1999). A side effect of this is that the NAT hardware will not accept unsolicited communications from the Internet, and the private IP addresses used cannot be seen from outside the Local Area Network (LAN) which made worm based attacks more difficult.

2.1.3 "Trojan" Malware

Trojan malware is a means of social engineering, i.e. where the victim is tricked into running malware voluntarily. This is a reference to the Trojan horse of Greek legend, where the Greeks besieging the city of Troy persuaded the defenders to open the gates to let in a wooden horse, supposedly as a tribute to their valiant foes. Within the horse were said to be a group of soldiers hiding, who were then able to open the gates which let the rest of the army in.

In the context of malicious software, this might be from an application which performs one function, whilst at the same time executing the malicious code to take over the computer. This might be done through applications for sharing files through peer to peer

[4]Conficker came in several different variants, and did not only use MS08-067 but also used other techniques such as brute-forcing passwords and USB infection (Conficker Working Group, 2011).

networks, which many users have an incentive to load, since it removes the requirement for paying for the content they wish to consume. Similarly, a video file may decline to play unless the victim chooses to download a "codec" to install on the computer. Alternatively, a Trojan could be executed through impersonating a well-known application which might be downloaded, see e.g. (Symantec, 2012).

It is not only through downloading software, USB drives can also be used as a means of executing code since many versions of operating systems will automatically open certain files on them. A method that has therefore been used to attack systems which are not connected to the Internet has been to leave USB sticks around in somewhere which would cause curiosity to the person finding it. By seeking to read the sensitive data purported to be on the stick, the unwitting user sticks it into their computer which then runs the code. It was reportedly a technique such as this which enabled the Stuxnet malware to infiltrate the defences of an Iranian nuclear program (Falliere et al., 2011). Conficker also used this technique, once again demonstrating how malware could have more than one propagation method (Conficker Working Group, 2011).

2.1.4 Phishing

The original Simple Mail Transfer Protocol (SMTP) dates from 1982 (Postel, 1982), and as such was from a time when it was not necessary to worry about security to any great extent. One feature missing from the specification, was a requirement to authenticate any email being sent, and as such, an attacker can claim to be a trusted entity the individual may have had contact with. Figure 2.2 demonstrates how this can be achieved, observe that at no stage is the user prompted for identification. Originally, "phishing" was an attack aimed merely at user credentials for accounts, which had value in early hacking forums. It was also possible to use this technique to merely get passwords in order to extract money from a victim's bank account (see Figure 2.1), although since security has now improved it is generally necessary to use malware to successfully extract money.

A commonly used method for this is known as a man in the browser (MITB) attack. This takes the idea from a traditional man in the middle (MITM) attack, but rather than intercepting the communication between the user and the bank (which is encrypted) it intercepts the communication between the user and the Web browser they are using. By manipulating the content being displayed, they can require the victim to add extra details than would otherwise be required, in order to perform their own actions.

In order to induce the victim to download malware, whether MITB or something else, some form of social engineering is required. This could be combined with some form of exploit, where opening a document with commonly used software could trigger some sort of exploit affording control of the system to the attacker. One type of attack such

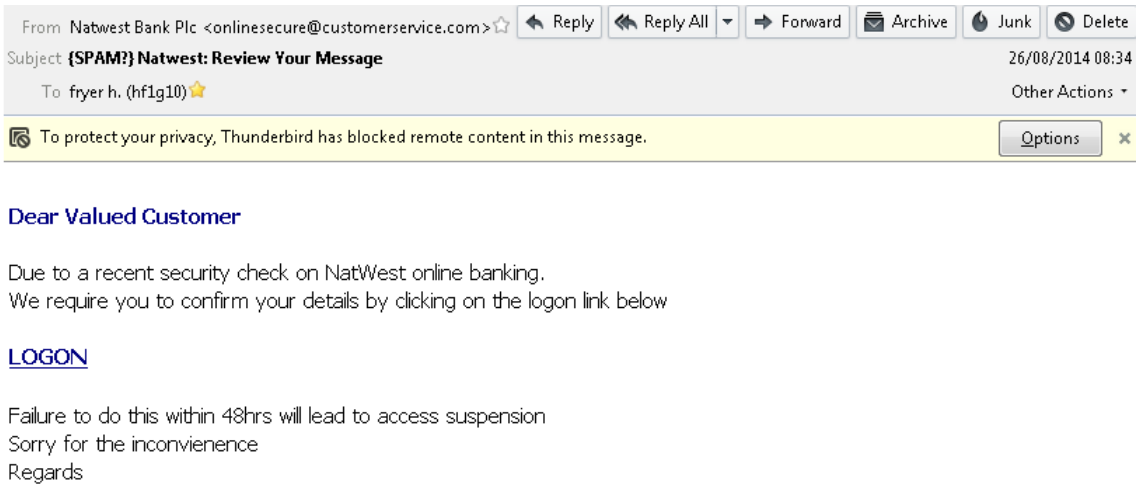

FIGURE 2.1: Example of a simple email, aimed at getting banking credentials

```
huw@huw-VirtualBox:~$ telnet smtp.ecs.soton.ac.uk 25
Trying 152.78.68.136...
Connected to gander.ecs.soton.ac.uk.
Escape character is '^]'.
220 gander.ecs.soton.ac.uk ESMTP Sendmail 8.13.8/8.13.8; Thu, 8 Mar 2012 13:47:41 GMT
HELO smtp.ecs.soton.ac.uk
250 gander.ecs.soton.ac.uk Hello huw-pc.ecs.soton.ac.uk [152.78.64.12], pleased to meet you
MAIL FROM:"YouTube Service"<info@youtube.com>
250 2.1.0 "YouTube Service"<info@youtube.com>... Sender ok
RCPT TO:< xxxxxxxxx @googlemail.com>
250 2.1.5 < xxxxxxxxx @googlemail.com>... Recipient ok
DATA
354 Enter mail, end with "." on a line by itself
TO:  xxxxxxxxx @googlemail.com>
Subject:YouTube Account Upgrade Reminder
Content-Type: text/html; charset=ISO-8859-1;
<html>
<body>
<p><img border=3D"0" alt=3D"YouTube" width="175" height="33" src="http://s.ytimg.com/yt/img/email_logo_no_tagl
ine.png" /></p>
<p>As a reminder, you created a YouTube account for  xxxxxxxxx @googlemail.com but haven't completed the signup
 process.
Complete your registration and upgrade your account by clicking the link below:</p>
<p><a href="#">Click this evil link!</a></p>
<p>Thank you,<br/>The YouTube Team</p>
</body>
</html>
.
250 2.0.0 q28DlfgM029630 Message accepted for delivery
```

FIGURE 2.2: Example using telnet to send a phishing email

as this is a malicious PDF document claiming to have details about the location of a parcel previously delivered to their address. A real world example is the attack which targeted RSA, which used a zero-day exploit within a Microsoft Excel document titled "2011 Recruitment Plan.xls" (Rivner, 2011).

2.1.5 Drive-by Downloads

A drive-by download takes the opposite approach of worm-based malware which "pushes" malware, and waits for the victim to come to them in a "pull" based approach. The process requires that a user visits a website which is under the control of an attacker. Once the website has been visited, malicious code on the website will attempt to subvert

the user's browser, and run its own set of instructions to take over the computer. This might be done through using JavaScript to corrupt the browser itself, or will use one of the plugins which the browser is running to similar effect. Plugins are additional features added to the browser, often to play multimedia content, such as Adobe Flash, Adobe Reader, Microsoft Silverlight, and Java. Like with any application, these will occasionally contain vulnerabilities which can be exploited in the same way by an attacker to get their malicious code to run.

This offers some significant advantages to an attacker over a worm-based attack. The logs of a user visiting a malicious website will be virtually indistinguishable from normal Web browsing, so the compromise has less chance of being discovered. Related to this is that, although a firewall can completely block attacks from the Internet, it has to let some traffic through in order to make Web browsing possible. Using a website can therefore offer the attacker a way through the user's firewall, and therefore increase the potential victim base (Provos et al., 2007).

The phenomenon of drive-by downloads is not a new one, but it has remained a significant threat. Provos et al performed a detailed analysis of drive-by attacks as early as March 2006-07(Provos et al., 2007), and also January - October 2007 (Provos et al., 2008). They found 1.3% of results in Google search results were malicious; and 0.6% of the most popular 1 million URLs had, at some point, been used as malicious hosting. The slightly later papers of Caballero et al. (2011) and Grier et al. (2012) both describe the uses that drive-by downloads have within the criminal economy (see the discussion in section 2.3), demonstrating that this mechanism continues to play a key role.

A typical attack will use a previously benign website which is compromised by the attacker to display malicious content. This is a separate part of the attack, before the victim browses to the website, and will use some weakness in the website or the server it is hosted on, commonly including out of date software; malicious advertising; or exploits using unchecked user data. Following the exploitation of the website, the content will then be changed to include malicious code, usually to redirect the victim to an attack website, which contains the code performing the exploit (Caballero et al., 2011).

There are benefits to an attacker in compromising a legitimate website as opposed to buying cheap throwaway domain names[5]. The existing reputation contained by a legitimate website means that it is harder to shut down than a malicious site, and also it becomes more likely that a potential victim will visit it. A legitimate website will already have a certain amount of traffic, and this can be enhanced through taking advantage of trending topics. Major news events will lead to a lot of people searching for information about it, so controlling a website about major news events is an advantage. Moore et al. (2011a) found evidence of this practice with advert filled and malicious sites effectively

[5] Although this strategy remains common in China (Aaron and Rasmussen, 2013)

exploiting trending terms on both Google and Twitter generating considerable profits for criminals.

2.1.6 Discussion

In many instances, the distinctions between various types of malware can be considered to add no benefit and may even be harmful. In this instance, however, the distinctions are necessary in order to distinguish the current attack vector of drive-by downloads from older methods. With older propagation methods, simple steps can be taken. Similarly, it is well known that Trojan style malware can be run by opening untrusted executable files, so one can simply avoid running applications unless they are sure it's trustworthy and appropriate usage of firewalls can prevent the spread of worms.

With drive-by downloads things are slightly more complicated, since there is constantly a risk with going to any website that the computer might become infected, so even a careful user's computer could become infected. Both users and websites can mitigate the threats to themselves through patching[6], but the separation of concerns in regards to the effect are different. At a high level, there are three steps which need to happen for a successful drive-by download to occur:

1. A previously benign website needs to be taken over and redirect the user to an attack site, or embed the attack site within the victim website;
2. A user with a vulnerable Web browser (or browser plugin) has to visit it;
3. The vulnerability within the browser gets successfully exploited by the malicious code.

2.2 The Problem of Drive-by Downloads

This section will describe the ways in which websites are exploited, and what can be done to prevent this from occurring. A malicious advertisement can be another way of making a drive-by attack possible (Li et al., 2012). This is also an important problem, but is regarded as being out of the scope of this research. Instead, the decision was made to focus on attacks which compromise the Web server in order to perform their attacks.

Once the server is compromised, it can not only be used for drive-by downloads, but also for phishing, or hosting or promoting questionable content. Indeed, only a small amount of compromised websites are used immediately for a drive-by download (Canali et al.,

[6]Or through the use of security products, which can detect changes made to certain files or identify known variants of malware.

2013b). Web borne attacks are serious in general, in that they can offer a foothold into an otherwise secure network, or more generally affect people browsing the Web. Compromised websites are a more popular choice for criminals rather than buying domains, the APWG reported that somewhere close to 90% of phishing websites were legitimate, compromised websites (Aaron and Rasmussen, 2010), whilst the Google Transparency Report demonstrates a similarly large percentage (Google, 2015).

There are few figures available as to the overall prevalence of drive-by downloads in comparison to other methods of attack, but there are several reports indicating their utility for criminals e.g. (Grier et al., 2012; Caballero et al., 2011). Research for by Osterman also found that Web based malware was the most common form of infection faced by corporate networks (Osterman ReGoogle, 2014).

2.2.1 Web based Exploitation

Websites are made possible by the integration of a range of applications, and as such are prone to attacks in the same way that conventional desktop applications are. The Web server is responsible for accepting the requests coming from the client, and responding with the content. Beyond the very simplest websites, there will almost always be interaction with a database. This allows the display of dynamically created pages based on specific queries, and functions such as login to take place.

The display of the website itself is done using HyperText Markup Language (HTML) and Cascading Style Sheets (CSS), which the user's Web browser understands. Creating content using HTML is also frequently automated so that novice operators can also add content. This is done by a content management system (CMS), which presents the operator with a WYSIWYG interface, and stores the content in a database which can be retrieved and displayed when it is requested by a user. Another component of a Web page is JavaScript, which is an event based scripting language. It waits for things to happen, such as a page load or a user's click, and then manipulates the document structure of the HTML.

This complicated interaction between different software makes website particularly prone to attacks. There are traditional memory corruption attacks which can be used against the software running infrastructure, such as the Web server or the database. Like with home computers, a website also offers a way through a firewall, since it accepts traffic from outside and interacts with the underlying infrastructure in order to respond to requests.

These attacks will often rely on mixing data with application instructions, so websites which do not correctly validate their user input to prevent this can be attacked in this way. Injection based attacks are an example of this, and would often target the database by finding out what values the application uses to query the database and change them

to be commands. For example, a Web page might query the database with the condition: everything with a type value of 'product'. An attacker could simply modify this value to be `' UNION SELECT UserId, CreditCardNo, CCV From CreditCards--`. This appends a quotation mark which signifies the search should be for an empty string, and `UNION SELECT` is the command to fetch more information from the database in addition to the original request – in this case all the credit card details stored. The `--` is a comment, which tells the computer to ignore anything after it, so can exclude the rest of the intended command.

This would potentially make the complete query: `SELECT ProductId, Name, Price FROM Products WHERE ProductName = '' UNION SELECT UserId, CreditCardNo, CCV FROM CreditCards--'`. Note the trailing `'`, this is now ignored by the query.

Another attack would use JavaScript in a similar way. As was discussed previously, JavaScript is used as a means of corrupting the user's Web browser when they come and visit the website, to install malware on it. This would not place the website under the attacker's control, but instead would directly target the users who visited it. The effect is the same – it turns the benign page into a malicious one. Alternatively, JavaScript can be used to disguise the content on the page, or to redirect the user to a malicious page, e.g. for phishing. Attacks which embed malicious JavaScript in the page are known as Cross Site Scripting (XSS), and take advantage of the fact that the way the W3 specification works, means that it is very difficult to filter all forms of attack, whilst simultaneously allowing desired content to be displayed.

HTML itself has at its core the concept of retrieving documents and resources from other locations. Whilst this is essential for the Web to function, there are also attacks which use this as well. Requesting an image on a restricted page can assist in deanonymisation, by seeing if the user is currently logged on to a service (if a 404 error is returned they are not, otherwise they are). Similarly, given that it is not possible to reliably guess whether a file is an image by the extension, any request for an image could be a script on another website, for example `<img src="http://evil.com/evil.php />"`, which would load a PHP page. This is used as a means of tracking for email campaigns. Possibly most serious of all, is that it is possible to embed the entire HTML of another page by using an `<iframe>` tag. These features are all necessary for the Web to continue to function, yet offer opportunities for attackers as well (Grossman, 2013).

The code in all the different layers of the Web server and application will, like any other software, have bugs which will occasionally turn out to be vulnerabilities. Well supported software will be updated, where the vulnerabilities will be fixed, but then the weakness in the server remains if that updated software is not installed. This leaves the website in a vulnerable state which could be exploited. A website being vulnerable is more serious than a normal user being vulnerable, because the user can hide to some degree behind the protections offered by their ISP, or the fact that the probability is

generally pretty low that they will visit an attack website. A website by contrast, exists entirely so it can be found and so has none of this protection. It is also more serious than an individual machine being infected, because of the damage it can cause by having many users visiting it.

The sort of vulnerabilities described above are a small subset of those reported by OWASP[7], a not for profit organisation, as being the most prevalent in Web applications. In particular, they periodically publish a "Top 10", detailing the most common, and rank them according to their impact (OWASP, 2013). Somewhat concerning, is the apparent prevalence of vulnerabilities from the OWASP top 10 on the Web.

WhiteHat security's 2013 Global Security report found that 86% of websites had a "serious" vulnerability on their website, defined as the ability to compromise at least part of their site (WhiteHat Security, 2013). Trustwave's Global Security Report 2013, found that password practices are weaker than might otherwise be expected, as well as many other vulnerabilities (Trustwave, 2013). Checkmarx analysed the source code of the most popular WordPress plugins, and found that 20% of them, and 70% of the most popular e-commerce plugins, contained similar, serious, vulnerabilities (Checkmarx, 2014). Although these are all security vendors, possibly with their own interests in presenting the data in certain ways, the amount of data breaches which have occurred do lend credence to their analysis as to the seriousness of the situation.

Automated scans are made by criminals to detect vulnerable websites and to seek in exploiting them. One strategy is the use of search terms indicating the presence of vulnerable components, or a website which is already compromised. This was demonstrated by Moore and Clayton (2009), who showed that there was a correlation between these search terms and the compromise of websites for phishing attacks. A search for `phpizabi 0.848b c1 hgp1` would return websites powered by an old version of phpizabi, which contained a vulnerability allowing the upload of files to the server (vulnerability CVE-2008-0805)[8]. Websites which were already compromised might have an uploaded "shell", which is a piece of functionality designed so that the attacker can maintain control of the server. The phrase `inurl:c99.php` would locate websites which had any URL containing c99.php, which is a popular shell used by attackers, and would demonstrate that the site was already compromised[9].

Vulnerable websites are at constant risk of compromise, and it is simply a matter of time before the compromise happens. Like with home zombie computers, there is a wide range of uses a website can be put to, demonstrated by Canali & Balzarotti who deployed web based honeypots to analyse exactly what an attacker would do following a

[7] https://www.owasp.org/index.php/Main_Page

[8] Successfully uploading a file to a server could have the effect of RCE on the server, since by browsing to that location the code will be executed as the page is loaded. This could then allow more permanent control to be obtained.

[9] This search no longer returns any results.

successful exploitation[10]. They used 500 websites with different characteristics, mostly based around vulnerable CMS software. In the event that the attacker could upload a shell, then they would on 46% of occasions, and then use that to log on average after 3 and a half hours. Only 1.1% of attackers specifically sought to add a drive-by download onto the website, but nearly 49.4% attempted to gain more permanent control of the machine, and 27.7% tried to get the machine into an IRC botnet, which would also enable drive-by downloads alongside other more general malicious activity such as phishing, sending spam, or hosting illegal content (Canali et al., 2013b). This demonstrates the variety of ways an attacker can use a compromised website as a means of facilitating drive-by downloads.

2.2.2 Existing Approaches to Drive-by Download Mitigation

The majority of the literature has focused on mitigating the effect of drive-by downloads to the clients, after the initial compromise (to the website) has already happened. This can be roughly split into two categories: pre-emptive approaches which search through websites in advance of the victim visiting the site so as to warn them, and real-time approaches are methods of detecting whether a page is malicious at the time that the user visits the page and attempts to prevent any damage from occurring, although both use similar methods for detection.

The identification of a malicious page enables a sign to be placed to warn users from visiting it, or to prevent the execution of the malicious code if they visit it anyway.

To identify pages in advance, a client honeypot would typically be used, e.g. (Nazario and Holz, 2008). This is an application which mimics vulnerable browsers and visits websites in order to induce them to attack. Depending on the level of information which the researcher wants, the site can either be classified as malicious (or not); or further details could be discovered such as in what ways the attack site attempts to interact with the browser. For example, using a virtual machine with a known good state, changes to the state system can be identified as malicious software.

Provos et al. used high interaction honeypots for their investigation into the level of malicious pages described earlier, which looked for any changes to the state of the machine; suspicious redirects, or suspicious downloads. Pages identified by Google (either through the use of client honeypots, or as part of their website scanning process) are presented with a warning when in the search results through the safe-browsing API, which is also used by browsers such as Google Chrome and Mozilla Firefox.

In the event that a user tries to browse to a malicious page, they are given further prompts to attempt to prevent them from going onto the page. Despite general criticism

[10] A honeypot is a server deliberately designed to look vulnerable, in order to induce attacks. These can be used to gather intelligence about current attack trends and strategies

about the effectiveness of browser warnings (i.e. that users ignore them), a recent study by Akhawe and Felt (2013) suggested that these warnings are actually effective, with only between 9% and 23% of users going through malware or phishing warnings.

Detection methods for malicious websites also contain trade-offs between accuracy and processing time required for classification. Simple analysis can classify pages based on characteristics of known malware. These are simple and quick to do, but can easily be circumvented with minor changes in strategy by the attacker. Applications exist for this exact purpose in the criminal market, which simply change a few lines of code, or make some changes to the structure in order that they cannot be identified, known as packers.

More sophisticated methods analyse the characteristics of a page, and search for characteristics which are out of the ordinary and use those to assign a probability that a certain page is malicious. This is known as anomaly detection, and an example includes Cova et al. whose analysis viewed certain programming techniques or a large amount of redirects as suspicious. These techniques can be enhanced by using machine learning techniques, where a program learns to recognise these malicious characteristics and adapt to new variants as it gets trained. Examples of applications using these techniques include Cujo (Rieck et al., 2010), ZOZZLE (Curtsinger et al., 2011), and SurfGuard (Sachin and Chiplunkar, 2012).

These techniques are not confined to analysis of the page content. John et al. identified a drive-by campaign where pages on compromised websites suddenly increased in popularity in the search rankings, making use of black hat search engine optimisation (SEO) techniques (John et al., 2011). They found that this was a good way of identifying some malicious attacks, and whilst blocking based on these criteria would be easy to get around, it would require the popularity of the page to be reduced making the pages less likely to be viewed. Zhang et al. sought to identify compromised Web pages through their links to attack servers hosting the malicious content, which as attackers currently work, is usually hosted on a different server (Zhang et al., 2011). Through combining knowledge of IP addresses and domains related to those particular servers, then a network of compromised pages could be identified.

Rather than attempting to identify the malware from the characteristics of the server, it is also possible to identify through actions in the user's computer. This has the advantage that there is no need for any knowledge about how the malicious code is constructed, but simply relies upon the observed effects of the browser or operating system after visiting the page. One way this can work is through relying on the fact that RCE usually requires memory corruption to occur, and then for the attacker's code (known as shellcode) to be executed. Egele et al. sought to identify shellcode in output from the Web page, whereas other techniques have been to examine the download of files after a page has been visited (Egele et al., 2009). If it appears that they have been downloaded without consent, then that would suggest that they were unwanted at best

and likely malware. In these cases then the code to execute these programs could simply be ignored (Hsu et al., 2011).

That said, the use of client honeypots or pre-emptive detection in this way does have certain limitations. Firstly, the browser which is being simulated might not be the target of the malware, such as if a honeypot used a version of Internet Explorer when the malware targeted Mozilla Firefox it would probably not attempt to execute. Similarly, "IP Centric" malware would only appear to users from certain IP addresses, and potentially avoid detection longer. An IP address can identify which network a request is coming from, and so would know if it is a client honeypot, or a security company. These companies have known IP address ranges, and possibly known user agents (e.g. in the case of a search engine such as Google), which could conceivably make it harder to identify. Given that there are a finite amount of companies with the expertise or resources to check for drive-by downloads on a regular basis, this is something which could limit discovery. On the other hand, it might require more control over the website than required by some of the typical attacks, because it would need to add the IP Centric code on the server side, e.g. Listing 2.2.

```
$ip_addresses = array();
$ip_addresses[] = '192.168.0.1';//etc.
//Add IP ranges to array...
if(in_array($ip_addresses,$_SERVER['REMOTE_ADDR'])){
        //do nothing
}
else{//only want to display this if it's safe...
?>
<iframe src="evil.com/evil.php"></iframe>
<?
}
```

LISTING 2.2: Sample IP centric code an attacker could use to selectively attack

Another approach which will be discussed in more detail in chapter 4 is that of detecting potentially malicious websites whilst they are vulnerable. Two main contributions have been made to this by Vasek and Moore (2013) and Soska and Christin (2014). Vasek and Moore (2013) identified characteristics of websites which were likely to become compromised, defining them in terms of the public health terminology of "risk factors", finding amongst other things that using a popular CMS is one of the factors likely to be an influence. Soska and Christin (2014) performed a machine learning operation on a series of websites which later became vulnerable, and were able to guess with a reasonable degree of accuracy which would. The use of risk factors allows the allocation of resources in order to best apply the protection required. These are discussed in more detail in Chapter 4.

2.3 The Criminal Economy

Having analysed the way in which malware can be used, this section will make the connection between malware and the criminal economy.

Many types of attackers exist, each with different skills, resources, and motivation. Early "hackers" or "crackers" were largely interested in the technical challenge of hacking, and the kudos gained from being part of the community. There are still attackers such as this, although now typically they can make a reasonable living by marketing their skills to defend firms from attacks. Occasional groups such as Anonymous, or Lulz Sec appear occasionally with the intention of causing havoc[11]. Attribution can often be difficult, there have been allegations of state backed malware or hacking, most recently in the attack on Sony Pictures in December 2014 (B.B.C, 2015). The event which originally brought this concept to the public consciousness was the Stuxnet worm targeted at Iranian nuclear facilities (Falliere et al., 2011).

The group which this thesis focuses on is perhaps the most common – the professional criminals whose intention is simply to make money, and who are frequently connected to organised crime. Much has been made of the sophistication of the market behind this class of attacker, and the specialisation and division of labour which makes the enterprise efficient. For example, it is now possible to engage in cybercrime without any technical skills whatsoever, due to both the software having been constructed and then sold on as a product. The ZeuS banking Trojan was an example of this. This was one of the MITB types of malware discussed earlier, except that it was sold as a piece of software (with support) to other criminals, allowing the purchaser to interfere in the online banking sessions of its victims. Like with a lot of legitimate software, its source code was leaked online and could be used for free. More examples of the products and services available to criminals can be seen in Goncharov (2012).

2.3.1 PPI and Botnets

A large part of the economy revolves around botnets, and the infrastructure required to operate them. This section will look at the concept of botnets, and the Pay Per Install (PPI) market which exists to support it. Botnets were mentioned briefly in section 2 as a network of compromised computers, this section will explore the concept a bit more deeply.

Having been able to execute malicious code on the victim's computer, the attacker needs to retain control. This control is achieved by requiring the victim machine, known as a bot or a zombie, to periodically send messages to a command and control (C & C)

[11] It transpired that one of the main Lulz Sec leaders was working with federal police in the USA for several months (B.B.C, 2014)

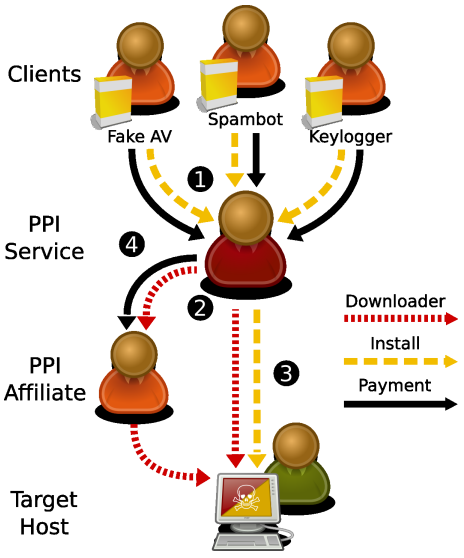

FIGURE 2.3: How the PPI system works, taken from Caballero et al. (2011)

server to ask for instructions. An attacker can make use of these zombies in a number of ways, for example to hide their tracks or to make use of the distributed power of a botnet "army". The distributed power could be used, for example, to brute-force cryptographic keys or passwords, or to overwhelm a server in a distributed denial of service attach (DDoS).

Their versatility means that they are an essential part of many criminal endeavours. Consequently, the obtaining of victim machines for an attacker is a central part of the economy, and this has been outsourced. As Caballero et al. (2011) observes:

> "At the heart of this ecosystem lies the *infection* of victim computers. Virtually every enterprise in this market ultimately hinges on access to compromised systems. To meet the demands for a wholesale infection of Internet systems, a service called *pay-per-install* (PPI) has risen to predominance." (Caballero et al., 2011).

A PPI operator works through obtaining purchases from clients, who wish to have their software installed on various hosts. They will have a downloader program, which will install itself on the victim host, and once it has permission will download and execute the program provided by the client. They might do this themselves, or they may outsource this to affiliates and act only as a middleman. If they use affiliates, then the affiliate will have a unique downloader which will do the same, but identify them as the ones who caused the software to be downloaded. Following successful installation, the PPI will pay the affiliates (if required) and download the client's application onto the host (Caballero et al., 2011). An illustration of how this system works can be seen in figure 2.3.

On occasion, it may be more specialised than this, as the affiliates can also outsource the exploitation of the machine in the first place in order to get their downloader onto the host. Where drive-by downloads are used, exploit kits such as Blackhole can be rented in order to find a means of successfully exploiting the machine in order to run their software in the first place (Grier et al., 2012). An exploit kit contains a library of exploits, such that the customer can create custom versions of malware (in this case, a downloader). Once rented out, the "customer" can embed the content from the exploit kit into the Web page of the compromised website they are attacking as part of a drive-by download.

The price that might be paid for installations of a particular branch of malware being installed on a victim machine varies depending on their "quality", which is to say, the likelihood of them being a high value target (Goncharov, 2012). Machines in affluent, western countries or within a high value organisation are more likely to be worth more, and have a greater amount of bandwidth available (e.g. for sending spam). On the other hand, machines in poorer regions are likely to be worth far less, since they are less likely to have a lot of extractable money, and connections are likely to be less good as well Goncharov (2012). There is similarly a market for traffic with which to infect with drive-by downloads; and the exploit packs themselves, with a significant amount of malware families using that model to propagate Grier et al. (2012). These bots will then be monetised in various ways, depending on their value, discussed in detail in 2.3.2.

In addition to monetisation, the botmaster needs to ensure that they can retain control of their botnet, both against "white hats" as well as other criminals. It is the C & C server which is often the part of the botnet which is the most vulnerable, and so additional levels of redundancy are often required to ensure that it remains resilient to takedown attempts.

Early botnets used Internet Relay Chat (IRC), which is a protocol which facilitates chat in the form of plain text. By requiring the bot to connect to a specific IRC channel, and to recognise various lines of "chat" as commands, it was possible to use this method to control a botnet. This is no longer used all that often, because it is easy to detect, and for a network administrator to block access to well-known ports which run IRC should they choose. Instead, HTTP Is commonly used as a means of communication which is more difficult to simply block, but requires additional effort to ensure that it is possible to recover in the event that these are taken down. If IP addresses are simply hard coded into the communication commands, then those can simply be blocked and lack flexibility. If DNS is used, then the registrar can block access to the DNS address being used.

One way of attempting to get around this issue was for attackers to use a technique known as fast flux DNS, in order attempt to increase the time it took for malicious domains to be located. It works by cycling through a series of IP addresses with a short

Time To Live(TTL) on the domain name, which act as proxies to the true location of the content. This makes it more difficult to locate and investigate, and also enables a single server (in a bulletproof hosting location) to provide content for all the flux servers requesting information from it (Nazario and Holz, 2008). This is one way where having a large amount of "disposable" bots can assist an attacker in retaining their infrastructure. Any machine detected as hosting malicious content, or acting as a controlling node, is simply a low value victim computer, and it can be switched to another.

Related to this is "domain flux", where a botmaster makes use of potentially thousands of different domain names changing on a regular basis, and requiring the bots to attempt to connect to each one in turn until they receive a response. This is designed to make it difficult or expensive for an adversary to retain control of the botnet, since they would have to keep control of all the domains whereas the botmaster would need only one. In the past, security companies have lost control of botnets this way. With this strategy, it is also necessary to protect against anyone who does manage to obtain a domain name from simply taking control. Researchers at Santa Barbara University succeeded in briefly taking over the Torpig botnet, by reverse engineering the communications protocol and purchasing one of the domain names (Stone-Gross et al., 2009). Since Conficker in particular, it is now standard to defend against this through the use of public key encryption (Conficker Working Group, 2011).

2.3.2 Monetisation

2.3.2.1 Spam

One of the more popular methods of monetising a botnet is through sending "spam" emails. Spam can take many forms, but in general it can be considered as unsolicited email, one taking advantage of the easy way in which the process of sending email can be automated. Phishing could be regarded as a subset of spam, as could emails with malicious attachments or "419 scams"[12]. The sort of email most commonly associated form of spam, however, is the sale of goods which are not genuine, or not strictly legal, such as Viagra, fake Rolex watches, or casino membership.

In the past, "open relays" were more common, which are mail servers which would allow anyone to send email on them (recall that SMTP did not initially require authentication), but these are now less common, and blacklists such as the one maintained by Spamhaus ensure that if a server is used too much in this way it effectively cannot send or receive email at all. Instead, this is a use for botnets, since the variety of IP addresses mean that large amounts of spam can be sent without falling foul of an ISP's blocking policy.

[12] 419 scams are emails which induce the victim to engage in a transaction with the attacker which requires some sort of administrative fee to be paid. Such examples could include money laundering, winning a lottery, or an inheritance from a distant relative

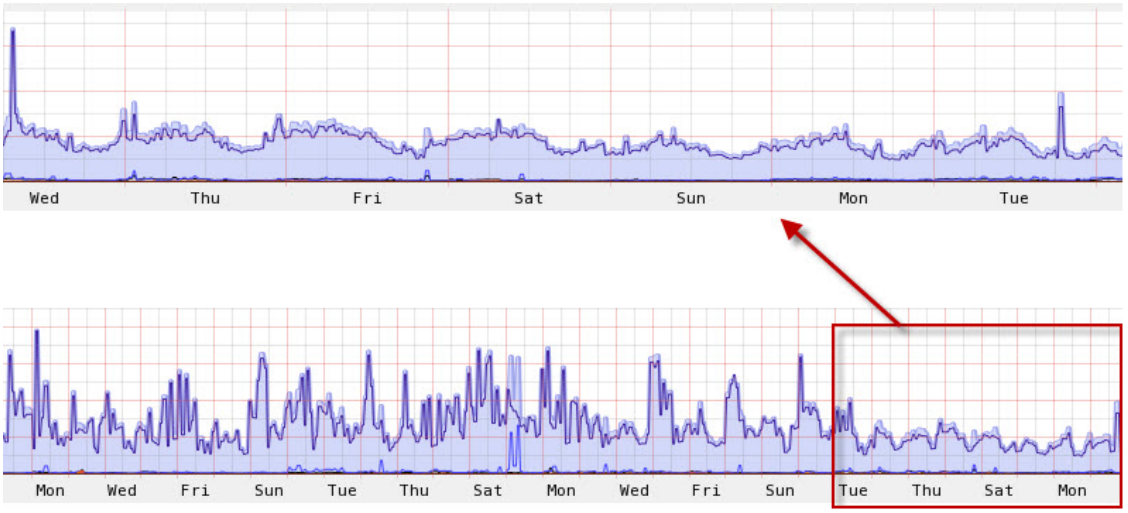

FIGURE 2.4: Effect of the Rustock botnet takedown, from (CYREN Security Blog, 2011)

The vast majority of all email is spam, that is over 75% (Trustwave, 2013), and the vast majority of it never reaches the inbox of potential victims. Figure 2.4 shows the effect that a botnet takedown can have on global spam levels, it illustrates global spam levels right before and after the takedown of the Rustock botnet.

2.3.2.2 Direct Monetisation

There are a few ways in which individual victims can be used. As the name might suggest, fake anti-virus (fake AV) mimics the appearance of a genuine anti-virus product, and provides a list of all the pieces of malware the "scan" has picked up on the victim's machine. In order to remove these "threats" the victim will have to pay a fee. Another popular version of the attack, is to display a message on the victim's computer claiming to be from the local law enforcement office. This will claim that illegal activity of some sort has been detected from the victim's computer, and that they must pay a fine or face court proceedings or more severe sanctions. These crimes might be as simple as downloading copyrighted material (which a lot of people have done), or even child pornography which would undoubtedly scare the victim since such accusations (even if false) can ruin careers.

Ransomware uses a similar strategy, but goes further in that it physically prevents the user from being able to access files or perform functions on their computer. For example, one type simply locks the screen until a payment is made. Recently, Cryptolocker and other similar strains of malware encrypt all files that the user has write access to. Whilst this was not a new idea, the increased sophistication is evident as RSA 2048 keys were used, and in a way without any obvious implementation errors. This meant that there

was no obvious way to get the files back without paying the ransom, particularly with some users having write access to company servers causing commercial damage[13].

Another form of direct monetisation is to use the malware to obtain banking credentials, and to require the machine to download MITB malware or similar in order to get money from the victim's account (see Section 2.1.4). Depending on the value of the target this can yield substantial profits, although often money mules will be required to physically remove the money from the victim's account which can make it difficult to do at scale. A notable example of this attack in practice was "Operation High Roller", which was observed by McAfee (2012). This operation targeted individuals with high net worth, and made use of ZeuS and SpyEye to interfere with their banking sessions. According to the report, a considerable increase in sophistication was evident, and the attackers made $78 million. This is a different type of attack to the low effort/low yield attacks described earlier, since the attackers targeted only specific accounts, but is a good illustration of the monetisation of this sort of attack.

The ways these work with a botnet is by instructing the victim computer to download and install a program which will allow the computer to be affected in this way. The ZeuS botnet, for example, was doing this with the Cryptolocker ransomware. This is a direct form of monetisation, and whilst reasonably effective, it might fail, since it alerts the victim to the fact that there is malware on their computer, and prevent long term exploitation. On the other hand, many victims will continue to fall for the same attacks more than once, and so it can continue to be profitable. It is also possible to hide banking fraud from the victims for a considerable period of time as well, for example by changing the display of transactions on the screen so that they are what the user might expect them to be.

2.3.2.3 Bitcoin Mining

Bitcoin[14] is a pseudo-anonymous, distributed currency, which does not rely on a central agency such as a bank to verify transactions. Instead, the transactions are visible in a public ledger known as a block chain where each transaction is chained together based on a hash of the previous transaction. The process of confirming these transactions is known as mining, which is designed as a computationally expensive process to prevent double spending (in order to change the records, every block in the chain would have to be changed as well). The mining process requires that a SHA-256 hash be generated of the block's head, and that this be less than or equal to the target value (Nakamoto, 2008). If this fails, then the nonce in the block's head increments, generating an entirely

[13] Recently it has emerged that several of the victims of Cryptolocker are actually able to get their files back for free as a result of a law enforcement operation which took over part of the botnet where the private keys were stored.

[14] https://bitcoin.org/en

different hash. Those who create a block are rewarded with 25 bitcoins (around £4,466, 5 January 2014), and can claim transaction fees from that chain.

Although the 25 bitcoins reward can generate a reasonable amount of income, mining is a very competitive process meaning that considerable investment has to be made in specialised hardware in order to have a chance. In addition to hardware, the amount of power needed to run these operations is also considerable. Some botnets have begun to take advantage of the distributed power offered by their bots in order to solve these problems. Whether this is more profitable than other botnet based endeavours is not clear, McAfee's report suggests that it is not (McAfee, 2014), although this was in relative terms compared to other possible sources of income. By the same logic, other sorts of attacks, such as password cracking could be solved in the same way.

2.3.2.4 Distributed Denial of Service Attacks

A denial of service (DoS) attack involves an attempt to overwhelm a server such that it is rendered unusable, or suffers a serious reduction in its functionality. Such attacks do not necessarily require a botnet, as long as the attacker has got enough computing power or bandwidth to throw at the target. However this is usually relatively easy to stop, since a single IP address could be blocked and then the attack fails. By instructing bots to attack, then it becomes a Distributed DoS attack (DDoS). This is far harder to stop as far more computational power is available and can be difficult to distinguish attacks from genuine traffic.

A notable example of this from 2013, was when Spamhaus were attacked with bandwidth totalling 300 GBPs (Leyden, 2013) with a DNS amplification attack, whereby machines in a botnet would make DNS requests but "spoof" their source IP address, so that it appeared that the requests were coming from the victim. What makes the attack successful is that it takes 64 bytes to make a DNS request, but since RFC 2671 up to 4096 bytes can be returned. The potentially huge amplification factor means that it does not take much to generate enough traffic which could potentially trouble even well protected servers.

The use of a Web application rather than a network based attack to enable a DoS attack is enough to classify it as a vulnerability. Whilst there are probably more serious problems if there is an SQL injection vulnerability, this could be used to this purpose by repeatedly performing queries against big tables. For example, consistently joining a large table against itself, e.g. `'; SELECT * FROM ActivityLog a1, ActivityLog a2, ActivityLog a3...--`. If all database resources are tied up with queries such as those, then other operations such as logging in could be impacted. Alternatively, simply repeatedly loading content heavy pages from many sources could affect the bandwidth available for the site, resulting in reduced performance.

DDoS attacks are often political in nature rather than necessarily immediately profitable (BBC, 2011). On the other hand, the ability to rent access to a botnet to someone engaging in a protest is a means of monetisation. Alternatively, extorting websites could be another method, as could selling to a business who wishes to push their competition offline. On a related note, criminals are in competition with each other, so the ability to DDoS a competitor could impact their ability to do business.

2.3.3 How the Criminal Economy Works

This section will first consider how the way that the criminal economy manages to function, in order that they can make money. It is suggested that there are some weaknesses in the current arrangement, and suggests how this could possibly be exploited using the PPI market as an example.

A plausible model has been proposed by Herley (2010), that there are only two types of attackers: those who engage in highly scalable attacks; and those who use targeted attacks. A scalable attack is one like spam, phishing attacks or botnets as described in Section 2.3 where following an initial investment the cost per victim does not rise linearly with the revenue to be gained from them. This is a very inefficient attack, since there is no scope for any additional per-victim effort to increase the chances of success or adaption in the event of failure. This is the case for success as well, having emptied someone's account this method would continue to try and target them again – even though they didn't have any money. An example of a scalable attack like this was presented by Kanich et al., who discovered that sending 350 million spam emails sold only $2,800 worth of "pharmaceutical product" (Kanich et al., 2008).

For a targeted attack to be profitable, then the attacker must ignore most potential victims where the value of an attack cannot be ascertained, otherwise they are likely wasting their effort. Given the distribution of wealth, the vast majority of Web users will never need to worry about this sort of attack since they have below average income. An attack of this nature might include spear phishing, a "watering hole" attack[15] or researching the target to determine the answers to security questions or make better guesses at the password.

Given their indiscriminate nature, the vast majority of attacks which we experience are the highly scalable attacks. Given the ability of computers to automate tasks means that the cost per user barely increases at all and so the amount of attacks which can be launched are huge, yet with such inherent inefficiency the figures reported by Kanich et al suggest that in that case the attackers might only just be making a profit. This raises an interesting possibility for a slight increase in costs to the attacker being enough to

[15]Guessing which website a potential target might visit, and compromise that, possibly using a zero-day exploit.

reduce incentives for participation. This is an idea which has been explored in the past in relation to honeypots adding additional uncertainty to the chance of success of an attack and therefore increasing the cost per user (Li et al., 2009).

On the face of it, the criminal economy appears to be a sophisticated arrangement, and it has been widely reported as such. Reports such as the one by Fossi et al. (2008) demonstrated the availability of a wide variety of goods on IRC in 2007; more recently Goncharov (2012) presented a report classifying the goods and services available from the Russian underground. However, it has been hypothesised that it is not as efficient as has been widely reported. This presents the possibility that interventions at certain elements of the market could have a significant effect.

Work by Herley and Florêncio (2010) hypothesised that the criminal market was a lemons market, and that the initial discoveries of a sophisticated market could simply have been full of suckers. Separately, they argued that the profits described from cybercrime were largely exaggerated based on faulty statistics and methodologies (Florêncio and Herley, 2013), and that this was a factor in encouraging criminals into an otherwise profitless industry causing losses for everyone else (Herley and Florêncio, 2009). In relation to the PPI business model, even whilst merely discussing the scope of it and how it increased the specialisation, Caballero et al. (2011) pointed out the fundamental conflict of interest which exists within it: that affiliates get paid for every installation, yet each additional installation they get can degrade the quality of the product (e.g. by competing with other affiliates for the same machine, and increasing the chance of discovery).

2.3.3.1 Example using the PPI Market

McCoy et al. (2012) and Stone-Gross et al. (2013) both point out that these enterprises rely on affiliates for generating leads. Using PPI, an affiliate can generate leads such as this through making available the victim's computer for installation of the particular malware binary, as was discussed in Section 2.3.1 (Caballero et al., 2011). For both affiliates, it is important to have infected machines: for spammers it is necessary to send a large amount of emails to get purchases (Kanich et al., 2008), whereas for fake AV, access to the machine is needed to install the software which will persuade the victim to pay for their "product".

In order to ensure a continuing revenue stream, the supply of bots needs to be repopulated as machines are lost through, for example through remediation, or having alerted the user to their presence of malware through direct monetisation. The amount at which infected machines are sold at depends upon their "quality" – their ability to generate revenue (Goncharov, 2012). Whilst it is less important for a spamming botnet to worry about direct income from these machines, it is of more concern to a Fake AV operator. A victim in a developed country is more likely to have the resources available to pay, with

Country	**Price per 1,000 bots**
Australia	$300 – $350
UK	$220 – $300
Italy	$200 – $350
New Zealand	$200 – $250
Spain, Germany, or France	$170 – $250
USA	$100 – $150
Global mix	$12 – $15
European mix	$80
Russia	$100

TABLE 2.1: Cost per 1,000 bots, reproduced from (Goncharov, 2012). All dollar amounts are US$

someone in a developing country less so and so less worth targeting. This was shown by Stone-Gross et al. (2013), with 76% of the victims being from the United States, despite representing a considerably smaller proportion of world Internet users.

To analyse the costs a PPI supplier can expect to gain, figures from the Goncharov (2012) report of the Russian Underground will be used to provide an illustration. For simplicity, this model will assume that the only means of infection is that of drive-by downloads; and that there are no transaction costs. This is not realistic, Caballero et al. (2011) described the PPI market as having a "middleman", who have their own affiliates and then sell on the installs. No effort has been made to verify the accuracy of the figures from Goncharov (2012), but since the figures are compiled by the same authors it can be assumed to be at least consistent for the purpose of an example.

In order to generate a sufficient amount of traffic, someone selling infections would need to be able to achieve the following:

- Source an exploit for a website
- Locate a website which is vulnerable to the exploit
- Source an exploit for a client/exploit kit
- Get traffic to arrive at the website
- Set up an attack server to host malicious content
- A means of laundering the money they obtain

This is not necessarily all going to be done by the same individual, and this is largely the point of the specialisation in the market – another example of cybercrime as a service. Whilst performing a task like sourcing an exploit, or packaging it into a binary may not be a direct cost if someone chooses to do it themselves, it remains a cost in time – the time will be assessed in cost, again according to the Goncharov (2012) report.

Traffic to the website would be approximately \$7 – \$15 per 1,000 unique users (from the USA). Assuming a success rate of 10%[16] and the median cost \$11, then the attacker has to spend \$110. In addition they would have to spend \$25/day for an exploit bundle and a cheap hosting platform of \$1.

Therefore, in order to infect 1,000 American hosts they have spent \$135 for revenue of between \$110 and \$150 (Goncharov, 2012). Subsequent infections would be cheaper, because they would not have to invest in the exploit pack, but recall that malicious websites have a very short average lifespan (e.g. phishing websites have a mean of 32:32 hours uptime, and a median of 8:42 hours (Aaron et al., 2014)) so additional infrastructure would have to periodically be re-purchased. Finally, to avoid undue attention from law enforcement, any revenue they obtain will have to be laundered where they can likely expect to lose a considerable amount (assuming half for the purposes of this example) so the money launderers can make a profit as well.

In this analysis, it appears that the market for infections only just breaks even, so even a small amount of additional effort might make it unprofitable. For example, if the already finite population of vulnerable machines or websites were reduced, then the scarcity would naturally require prices to go up owing to the effort of obtaining the required level of infections. Assuming the PPI vendors as being on the supply side, and spammers or fake AV vendors being on the demand side, how might the market react to a small increase in price?

Research by McCoy et al. (2012) suggests that the majority of spammers would likely fare badly, with the majority earning a very small amount and only a few big players making the majority of the money. If the smaller players were forced out of business, the pharmacy programs might struggle as well, since affiliates generate a large amount of their income, and the heavy tailed distribution shows the big players make only a small amount of that income. Indeed, they suggest "such organisations are fragile to economic disruption of even a modest scale". Similarly, as described, Kanich et al. (2008) believed that the spam botnet could only be viable if fully vertically integrated.

The fake AV market might react a bit differently, but could again be successfully disrupted in the event that vulnerable machines in developed countries were reduced in availability. The high level of reliance upon "customers" from the USA demonstrated by Stone-Gross et al. (2013) indicate that they have to be more discriminating in the traffic they wish to continue generating revenue. This means that there are fewer options for replacing "high quality" victims, which could quite conceivably cause them problems.

[16] Grier et al. (2012) mentioned an attack on the `mysql.com` website , which had a success rate of 9–14%

2.4 Conclusions

This chapter provided a general background of the issues of malicious software, and provided terminology for framing the discussion in future chapters. Drive-by downloads were distinguished from other sorts of propagation methods, and the state of technological solutions which have been deployed against them. It was suggested that a more appropriate method would be to prevent websites from being compromised in the first place, thereby restricting the ability of the attacker to serve malicious software. Some ways in which a website could become compromised were considered, and how the attacker might be able to take advantage of the compromised site that they had control over.

A brief look at the criminal economy revealed that they are likely operating on thin margins and, assuming they are rational, significant disruption could be caused by a small increase in price. The next chapter will look at the other stakeholders: those who would be charged with causing this disruption. The phenomena making this disruption a difficult proposition will be considered, and then an analysis of private law methods which could possibly be used to induce these stakeholders to play a more active part in Internet security.

Chapter 3

Stakeholders and Economic Background

> *Riot shields, voodoo economics. It's just business, cattle prods and the I.M.F*
> (Radiohead, Electioneering (1997))

The previous chapter established some context and terminology relating to criminals and malware in general, and drive-by downloads in particular. This chapter will build on that foundation by conducting an analysis of the stakeholders who are involved, or could intervene in the drive-by download process. To begin with, some background and terminology will be introduced relating to some of the economic phenomena which exist in relation to Web security. This is fairly uncontroversial, but will provide some background to the terms and literature which exists in this area. This will be expanded upon specifically in relation to the individual stakeholders.

The process of the drive-by download is represented in Figure 3.1. Anarchania represents a jurisdiction with a limited capability to enforce cybercrime laws. Jurisdictions such as these are common places for attackers to host their C & C infrastructure. The purpose of including Anarchania in the process is to reaffirm the point that getting rid of the malicious content itself is not easy. Recall that the nature of the Web is a series of linked documents and content means that the location of the malicious content is separate from the website itself. Embedding content in an <`iframe`> tag enables the content to be stored in such a jurisdiction, and there is little local law enforcement can do about it.

Alice represents the end-user, who has not patched her computer. She connects to Bob's website using her ISP for Internet access and browser to use the Web following a recommendation from a search engine[1]. Because Bob's website has already been

[1] Also possible would be to click on a link in a malicious email, or social media content. These have been excluded from analysis, the social media content falls under the "website" stakeholder, and usual function of Web browsing involves the use of a search engine to locate content

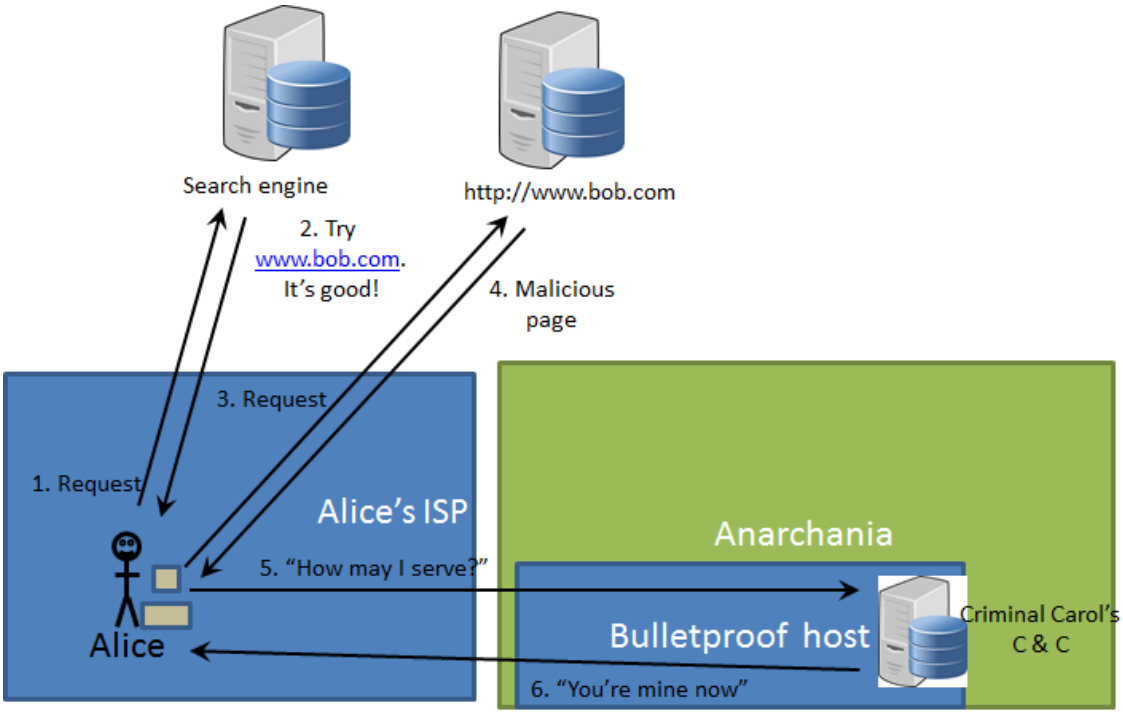

FIGURE 3.1: The process of a drive-by download

compromised and is serving malware, Alice's computer then also becomes infected and communicates with the C & C infrastructure which is hosted in Anarchania. Carol herself is also based in Anarchania, and as such is beyond the reach of our law enforcement agencies.

Each of the stakeholders involved in this process will be considered in terms of their capabilities and motivations, bearing in mind the economic phenomena discussed. Carol, the criminal, will not be considered, since by definition she is outside of the legal system, and therefore not worth pursuing from a law enforcement point of view. A discussion of Carol's motivations, and the criminal economy in general can be seen in Section 2.3. Using these models, measures which can be implemented will be proposed which might severely limit their capabilities.

General economic theory contains the assumption that, in general, the market is the best way of ensuring efficient behaviour (Gans et al., 2011). As a result, there is a considerable amount of literature relating to various private law based deterrence methods for Web security (See discussion in Section 3.2). As a means of attempting to avoid government intervention, this literature will be analysed. However, it will be seen that there are both legal and practical difficulties behind this. The legal difficulties will be discussed in the context of UK case law, and the practical difficulties will be illustrated through a case study from USA law – that of Article 4A U.C.C which relates to commercial funds transfers.

It will be concluded that given the stakeholders and their motivations, alongside the difficulties presented in terms of market forces and private law, some form of government intervention is required. This will lead on to the next chapter, where a framework will be discussed in which regulation to this end can be formulated.

3.1 Market Forces

The economic difficulties associated with security are well known, and the resulting "security economics" discipline has numerous publications associated with it (e.g. Anderson (2001); Edelman et al. (2006); Moore et al. (2009)), and an annual workshop dedicated to discussing these issues[2]. This section does not attempt to add to the literature on the subject, but rather use it in order to add some definitions to the terms which will be used in describing preventing the stakeholders from pursuing security solutions efficiently.

This is not to say that the whole of the issue of Web security and drive-by downloads can be reduced to rational economic analysis. Individuals and companies can have altruistic motivations, and difficulties in security may frequently be down merely to ignorance rather than a calculated economic decision. Similarly, in certain circumstances there is a market for security – for example, in order to handle credit and debit cards an organisation has to be compliant with the PCI-DSS standard (PCI Security Standards Council, 2013), which is something which can lead to a demand for a higher level of security. Nevertheless, the economic analysis provides a plausible reason for some of the problems, and, as has been shown, enough insecurity exists that a greater effort is required to improve the situation.

3.1.1 Public Good and Freeriding

Goods are classified in two ways: in terms of whether they are *excludable* and whether they are *rivalrous*. With an excludable good, one can exclude its use by others (Gans et al., 2011). A rivalrous good means that its use by one diminishes its availability for others, and so people are rivals for it. Security is often regarded as a public good, which is to say it is *non-rivalrous* and *non-excludable*. This distinguishes it from resources like the Internet, or the Web which are frequently regarded as common goods (they are *non-excludable* but *rivalrous*).

Both create potentially inefficient situations unless managed properly. Hardin (1968) coined the phrase "tragedy of the commons", where he argued that individuals acting rationally come to a suboptimal outcome – the resource inevitably becomes depleted. The illustration is of a medieval common, where it is permitted for farmers to graze their cattle. In the interests of the population it is appropriate to limit grazing on the

[2] `http://weis2015.econifosec.org` is the next one

common in order to preserve it, however it is not possible to reach this situation through rational play. By grazing an extra cow on the common, the individual gains the whole benefit of this, whereas the detriment is split through the whole population.

Public goods have a different issue – that of freeriding. Since it is not possible to exclude others from the benefit, there is little incentive for an investment in the good. Rather, the optimal strategy is to freeride off the investments of others. Commonly cited goods which are regarded as public goods are health and security. This is an example of market failure, and one for which the state will often take control and provide. Consider national security – it is not possible to exclude from others the benefits of paying for a standing army (less danger of being invaded, and possibly a lower crime rate), so it makes more sense for national security to be paid for out of general taxation rather than leaving defence of the city to the market. Health as a public good is considered to be analogous with Web security Mulligan and Schneider (2011), and is discussed in more detail in Chapter 4.

Security retains its public good characteristics when applied to the Web as well. In particular, this relates to bargaining with those whose operations have negative externalities (discussed below, see 3.1.2). As it relates to the issue of issues like botnets and the consequences, as is the focus of this research, it is a particular problem. The problem is a global one, so any intervention by a particular stakeholder – whether that's one of those discussed below in Section 3.3, or actions by an individual country the same issue remains. An example of this can be seen with click fraud, which affects many different companies classified as advertisers, or advertising distribution networks. Any one of those who chose to invest in improvements would be unable to prevent everyone else from benefiting from their investment. As such, each of them has the incentive to freeride rather than invest in preventing the issue from occurring.

3.1.2 Negative Externalities

A negative externality is defined as follows: "One party's action will be said ... to create an *externality* – if it influences, or may influence with a probability, the well-being of another person in comparison to some standard of reference" (Shavell, 2004). In general, this relates to some form of nuisance or risk creating behaviour adopted by one party who is not subject to the negative effects of the externality. A common example might be that of pollution from a factory affecting a neighbour's property, or the sound or smell which are side-effects of industrial activities.

Externalities are a common issue in relation to security on the Web and the Internet. For example, vulnerabilities in a piece of software are problems for the users of the software as opposed to the vendor selling the software. Whilst this could adversely affect their reputation, in general the costs are borne by the users (Anderson, 2001). Similarly, those

running websites which get compromised do not suffer from having a drive-by download on their website. They are largely invisible, and the costs for clean-up will go to the users, or to different organisations that have to deal with the ongoing issues like botnets and click fraud. Whilst they will eventually have to clean up the website, it is likely a cost they will only have to bear if their site is blocked or it is noticeable[3]

Similarly, UDP based DDoS attacks leverage vulnerabilities from operators who generally do not feel the full effects of the attack. An attack such as this works by "spoofing" the source IP address of a request to be that of the intended victim. A request will be made for a service such as DNS, which gives responses far bigger than the original request. These can be further amplified by using zombie computers (although these are not strictly needed).

This means that there are three separate negative externalities in relation to this attack. It has been regarded as bad practice to allow source IP address spoofing at least since 2000 (Ferguson, 2000), yet there remain enough networks which have not been appropriately configured for the attack to be effective (Beverly et al., 2009). Similarly, incorrectly configured DNS resolvers can also be possible to fix yet this has not always been done, see (Open Resolver Project, 2013). An example of an attack using these was the attack against Spamhaus in 2013 (Leyden, 2013).

Where possible, it is believed that negative externalities should be solved through bargaining. In his influential paper, Coase (1960) argued that many cases which place the blame solely on one person should actually be thought of as a reciprocal problem between two people. Restricting one from carrying out their activity in order to prevent harm to the "victim", harms them by depriving them of enjoyment of their property. In a world without transaction costs, the parties would always agree a mutually beneficial bargain in order to enable both to enjoy their property rights. Indeed, in many cases the transaction costs are sufficiently low that the parties can bargain to a beneficial arrangement. Yet, the situation relating to the Web in general, and security on the Web in particular, is not one of them.

The transaction costs would be particularly high, since the amount of people and companies adversely affected by a negative on the Web is incredibly high. This naturally increases the transaction costs substantially, so as to be particularly difficult to bargain, particularly since they can be geographically distributed and not readily identifiable. In addition, even were it possible to work around these high transaction costs, there is a lack of information to the cost of the externality, making it more difficult to reach a bargain. A final issue returns to the public good phenomenon discussed above, in that it is not possible for one party to make a bargain and exclude other parties from the benefits of it. It is in their interests to wait for this to occur and freeride the benefits.

[3]Defacement by organisations such as Anonymous or other political factions are a different matter, and might hopefully begin to internalise these externalities associated with websites. However, for the purposes of this research they are out of scope.

In situations like this Coase (1960) argued that courts are necessary, or even occasionally governments, but they should only decide based on a principle of wealth maximisation. Private law solutions are discussed in Section 3.2.1, and regulation in Chapter 4 and 5, although it should also be noted that this remains a public good issue for the state as well. Encouraging the problem to be solved costs money, and there is no way of preventing the benefits from affecting other states – thereby preventing the need for them to take action.

3.1.3 Information Asymmetries

An information asymmetry is a situation where the buyer of a product or service has a lower level of information than the seller. This leads to a situation described by Akerlof (1970) as a "market for lemons", whereby the only products remaining on the market end up as lemons[4]. The used car market was used as an illustration of the phenomenon by Akerlof as an explanation of the difference in quality between a new car and an old car. Whilst the seller of the car has had time to appreciate the value of the car, the prospective purchaser has no means of telling whether or not the car will turn out to be a "lemon", which will stop working at some point after purchase. As such, the prices converge on the average price of a lemon since that is all customers are prepared to pay for them, which drives the good quality cars off the market since they are unprepared to sell them for that price, until all that is left is "lemons".

With cars, the difficulty lies in the hidden mechanical parts and general lack of knowledge about the mechanical parts by users. In relation to the stakeholders discussed in this chapter, the issue arises in two obvious locations: the security of intermediaries, and the quality of software. Both require specialist skill and knowledge in order to be able to appreciate whether they are more secure than their competitors. This has led ISPs to compete on price with thin margins, and consequently unprepared to invest more in security than is necessary (see Section 3.3.4 for more detail).

In relation to software, an example will be provided to illustrate the scale of the difficulty. A lot of software is closed source, which is to say they provide a working application out of the box, but the source code used to generate that application is regarded as a commercial secret. For example, the popular Microsoft Windows operating system is closed source, as are all Microsoft products. Where this is the case, it is not possible at all to ascertain the quality of the source, and the likelihood it will contain vulnerabilities. Particularly on the Web, a large proportion of the software is open source which provides an opportunity to illustrate the difficulties which might be faced should a customer wish to examine it[5].

[4] In the USA, "lemon" is slang for a bad second hand car

[5] It is possible to examine the source code of a closed source product to a certain extent, through a process called reverse engineering. This is also a skilled task, and requires niche knowledge even amongst technical customers, that of Assembly programming language

Product	Files	Blank	Comments	Code
WordPress 4.0	904	51055	91K	233715
Joomla! 3.3.3	3925	121247	240651	623720
Drupal 7.31	1413	63155	157510	331272
Django 1.7	2641	95966	139132	417604
Apache server 2.4.10	768	38085	54754	211524
Apache server 2.2.29	1215	63120	88570	358016
PHP 5.6.0	2562	163746	174927	1127715
PHP 5.5.17	2484	159085	171949	1093547
PHP 5.4.32	2364	151862	162321	1045462

TABLE 3.1: LOC for popular open source software products

Given that commonly used software is maintained by teams of developers, intuitively this would suggest that even a skilled operator would have to spend a considerable amount of time in analysing it to be able to get an idea of the extent of the vulnerabilities in it. To demonstrate the scale of the difficulty, a test was conducted using the free CLOC tool[6] to obtain values of the lines of code of open source software commonly used on the Web. Table 3.1 contains the outcome of this, and demonstrates that even software with "simple" functionality such as WordPress has 233k lines of code. The source of the PHP language clocks over a million, and a whole operating system would comprise tens or hundreds of millions of lines.

On its own, LOC means very little, but it provides an idea of the scale. Given the amount of different paths which would exist to get through the program flow, the difficulty in ascertaining its quality or security is quite significant.

3.2 Private Law

Even in a system dominated entirely by market forces, it is necessary to have some form of legal system so individuals have a remedy in the event that someone causes them a wrong. Contract law, for example, forms a basis through which bargains between individuals (or companies) can be enforced in a court of law. Tort law is used where one party causes a loss to another party, whether through carelessness or through some other common law defined wrong. The main form of tort to be discussed will be negligence, which applies where one party has caused a wrong to another due to a lack of care. Many of the same principles and difficulties apply to strict liability as well, including difficulties in proving that one particular party *caused* the damage.

Arguably, tort law serves the function of reducing transaction costs (Shavell, 1993). It would simply be impractical to require everyone driving a car to have a contract with everyone else driving a car (and all pedestrians), for example, so instead negligence

[6] http://cloc.sourceforge.net/

law provides a framework where someone causing an accident is liable to pay damages. Whilst originally the system was probably merely a method through which corrective justice could be served, there is a body of legal theory about its function within wider society due to the deterrent effect of paying damages (Goldberg, 2002). Through precedent, this allows a recognised standard of behaviour to be understood and the optimal level of precaution to take. This is a balance which must be struck, because preventing accidents entirely could simply prevent an activity providing social benefit from being conducted if the costs of prevention are too high (Landes and Posner, 1980).

To use tort law as a means of correcting these conditions is something which has generated a considerable amount of interest, particularly in the USA, although there have been a limited amount of actual cases. This means that a lot of the ideas discussed are theoretical, and based on analogies e.g. a server being like land. In this section, some of the theories about the application to tort law for this function the few cases which have reached the courts will also be considered.

3.2.1 Literature on Tort Law for the Internet

Lichtman & Posner counter two common arguments against the imposition of liability that 1) ISPs will overreact and stop accepting risky users; and 2) ISP liability removes the incentive from users to take appropriate care[7]. To the first, they argue that some form of tax relief or support may be appropriate to assist smaller operators (Lichtman and Posner, 2006) who may otherwise struggle to cope with the additional costs for litigation. To the second they argue that it is necessary to tailor the liability such that users retain the incentive to be secure. Given that transaction costs are low ISPs can use contract law to enforce conditions on the users (Lichtman and Posner, 2006).

Johnson performed a detailed analysis on existing legislation and judicial opinion to examine the possibility of a tort for the victims of identity theft from cyber security failures by database operators (Johnson, 2005). He concluded that general tort principles did support the idea that there was a relationship between the two parties (the victim and the database operator) so there was no reason in principle why liability could not be applied. Public policy would also support the notion of holding database operators liable, because losses would be minimised as a result of greater investment in database security. The rule prohibiting recovery for economic loss, a frequent stumbling block towards successful claims, should not apply to losses in regards to cyber security, but that claims should be limited so as to encourage operators to investigate for when a breach occurred so victims could take action. Since the paper was written however, there has been a considerable increase in the amount of US states with breach notification laws so this possibly does not apply to the same extent.

[7] They seem to largely be speaking about access providers, but the analysis is more general.

Citron's proposal also argued that the way we view property needs to change. However, the argument advanced was that negligence for database operators would be insufficient to cope with the problem for three main reasons (Citron, 2006). Firstly, the rapidly changing rate of technology means that it is difficult to know what the current optimal level of precaution to take would be, therefore leading to a potentially inefficient outcome. Secondly, there are no clear norms to guide for future behaviour because of the constantly evolving tactics of criminals making the idea of negligence "*shaky*" on its own (citing (Abraham, 1997)). Finally, since there is a level of residual risk of data leakage in any event, strict liability should be used to discourage marginal operators from entering the market, and efficiently allocating the risks(Citron, 2006). As such, the courts should adopt a *Rylands v Fletcher* (1868) strict liability model. The databases used to store personal information were argued to be an analogy to the reservoirs of the 19th century, and that the natural consequence of the "escape" of personal data is identity theft, or harm for the victims[8].

The providers of software have also been the target on several occasions, Microsoft in particular due to its huge market share and the global impact vulnerabilities in its software has, e.g. (Kuwahara, 2007). Holding developers liable was also central to Rustad and Koenig's proposal of a tort for "negligent enablement of cybercrime" based on product liability (Rustad and Koenig, 2005). Their proposal was that providers of services and products used on the Internet be liable for knowingly marketing defective products and services. This they argued is much like the producers of cars, whose products became noticeably safer and more secure once they were the subject of product liability. Security practitioners have also argued this, for example Bruce Schneier argues

> "*Software vendors are in the best position to improve software security; they have the capability. But, unfortunately, they don't have much interest*(Schneier, 2007)".

The report by Anderson et al. for the European Network and Information Security Agency (ENISA) considers several of these issues in the liability section of their report. They rejected the notion of liability for ISPs in favour of a scale of fixed statutory damages for damages of any malicious users. In relation to liability for defective software, it was suggested that a simple approach might have worked in the past, there are too many products which have use software for an approach like this to be viable as an overall solution (Anderson et al., 2009). Instead, they recommend that vendors be liable for vulnerabilities in order to encourage an improved rate of patching.

[8]The notion of a computer/land analogy has been discussed as well in the context of cybertrespass(Epstein, 2003a). This was largely in response to the decisions in *eBay v. Bidders Edge* and *Hamidi v. Intel*. This revolves around whether a computer is sufficiently analogous to land (and hence not requiring actual damage) as opposed to chattels (where actual damage would be required). A discussion of trespass is beyond the scope of this research.

That users might be subject to liability is also an issue which has been considered. Henderson suggested that a zombie computer which was "*knowingly insecure in the face of a well-known threat*"(Henderson and Yarbrough, 2002) could be considered negligent, describing it as akin to driving a car on the road with a known defect. It was suggested that, with the threat of lawsuits, the ordinary level of care taken by users would evolve into a level of care sufficient to make DDoS attacks less practical. The notion of duty was considered in response to the costs incurred by the actions of "Mafiaboy" in February 2000, as a means for the victims to obtain damages from someone solvent in a manner similar to the manufacturers of firearms for victims of their use(Henderson and Yarbrough, 2002). De Guzman by contrast considered liability a means of forcing users to internalise the costs of their own insecurity. It was suggested that finding liable someone who left the keys in a car door for damage when the car got stolen could be extended to leaving a home computer unsecured(De Guzman, 2009). The analogy was further extended, that whilst a stolen car didn't have the owner present, a hijacked computer did and as such was similar to a car on the road with a consequent duty to other users(De Guzman, 2009).

3.2.2 Discussion of Tort Law Application

Whilst section 3.2.1 analysed some of the literature in this area which suggests tort as a possibility, this section will consider some of the chief practical difficulties. A major issue is that of causation, and also the difficulty in persuading claimants to go through the courts to recover damages. To conclude, a case study will demonstrate these difficulties in connection with Article 4A U.C.C in the USA.

There is authority which suggests that it is possible, in principle, to hold an operator of a website, or intermediary, liable in tort law although the circumstances which this has been tested have been limited.

Patchett v Allied Trades Association Ltd (2009) is one example demonstrating that, in principle a website can be liable in negligence. However, given that the loss is purely economic a higher standard is required to successfully claim. *Patchett* was a case where the claimant relied on information on a website as to the quality and liquidity of a company which manufactured outdoor swimming pools. It turned out that the company was experiencing financial difficulties and left prior to completing the job leaving the claimant with a loss of £44,000. The Hedley Byrne principle (*Hedley Byrne & Co Ltd v Heller & Partners Ltd* (1963) was applied in that provided there is a sufficient "*assumption of responsibility*" it is possible to claim for pure economic loss. Although, as in *Hedley Byrne*, the court held that because there was a disclaimer then there was no liability.

Similarly, in *L'Oréal v eBay (UK)*, Arnold J confirmed that it is would be possible to hold an intermediary liable for the torts of another according to the existing authorities. In relation to the particular case against eBay about trade mark infringement, it was not possible to hold them liable as joint infringers. The two main reasons for this were that, eBay did not offer preference to illegal content (at [377]), and that there was no common law obligation for them to intervene (at [373–375]).

Although neither of these authorities directly supports the proposition that a website could be liable for damage as a result of infection, it does leave open the possibility that they could be liable for a security breach causing actual loss. The primary point to emerge from *Patchett* is that a disclaimer would exempt them from liability. Although the requirements for pure economic loss are more stringent, there are no opportunities to agree (or decline) to the terms and conditions prior to an attack by a drive-by download, so may not be enough to protect websites or other stakeholders from liability.

Despite it being possible, in principle for a website (or some associated secondary liability) to be liable, challenges do remain. A significant challenge is that of proving causation, in that the negligence of one particular website caused losses to a malware victim. This would first require identification of the website in question – which is no easy task, currently Netcraft has the amount of active websites at 178,164,265[9]. In addition, it would have to be shown that this particular website was negligent by becoming compromised, and that through that negligence the victim's computer became infected. Finally, proving that particular malware itself caused losses is frequently difficult because there are multiple malware instances on the same machine. It is also likely that the victim would be liable in contributory negligence since it would usually only be possible through a lack of care on their part as well.

The difficulties and costs associated with merely proving causation leads to a state where the costs of obtaining relief are likely to substantially outweigh the losses from any one particular instance. The incentive to litigate is contingent on the damages being greater than the cost of the lawsuit. As Ogus observes, "The chief problem afflicting private law is that of transaction costs. Rationally, individuals and firms will only seek to enforce rights where the expected benefits exceed the expected costs, which include not only legal expenses but also time and trouble" (Ogus, 1994). The combination of lawyers and court fees are expensive, so there must be a considerable amount of loss in order for it to be worthwhile. This applies in particular where the outcome of the case is uncertain, for example in the event that there are few existing precedents. Where the losses are spread amongst a large amount of victims, then a serious misallocation of resources is apparent.

[9] http://news.netcraft.com/archives/2015/03/19/march-2015-web-server-survey.html, last accessed 22 March 2015

Consider seeking damages of £5,000, with lawyers fees of £2,000, where there is no precedent suggesting a 50% chance of winning. This reduces the expected outcome to £2,500, and then even the small chance of an appeal is likely to increase the fees to more than the expected outcome (a 25% chance of appeal increases the expected outlay to £2,500, with the probability of losing remaining the same because there is no precedent.) This is particularly the case where the potential plaintiff has lower income, and hence places a higher marginal utility on their money than a defendant who (frequently) will have more money available to spend.

The precautions which could be taken in order to avoid suit would also ideally be according to a social optimum. Learned Hand J articulated a rule which is broadly in accordance with this: that the cost of precautions should equal the product of the probability of damages and the amount of harm to result (*United States v. Carroll Towing Co.*). So to spend £100 to prevent harm of £10,000 would be an adequate precaution, if the harm had a 1% chance of occurring.

Whilst a lot has been written about the social, or economic theories of tort law (Goldberg, 2002), and how they can be used to encourage economic efficiency or socially desirable behaviour, ultimately the mechanism through which these are decided is in a court where there are two parties who have to pay the fees required by the court. Given that the cases involve individuals, then there is a considerable possibility of a divergence between what is personally optimal and what is socially optimal. One of these situations is where costs are spread such that, although there is a considerable social cost to a lack of precaution there is no incentive for the injurer to take the precaution since there is no risk of suit. The other is where there is little that can be done to reduce the probability of harm, yet the damage caused on the occasions that it does happen is worth litigating over. Whilst it is possibly in the interests of society for parties to settle rather than waste resources on litigation, this means that there remains no precedent where the courts have set out what they deem to be acceptable behaviour. For more detailed discussion, see Shavell (1982); Kaplow (1986); Menell (1983); Rose-Ackerman and Geistfeld (1987); Shavell (1997, 1999).

This is a problem which could be solved, at least partially, through a form of collective redress or class action lawsuit. In such a system, a group of claimants pursue a case together such that the costs are shared between them enabling the pursuit of a claim which would otherwise not be financially viable. Most commonly thought of as a system in the USA, it is also something which has received attention in the EU as a means of combatting anti-competitive behaviour (European Commission, 2005, 2008). Although the system in the USA has improved access to justice for individuals, and arguably remove the bias in favour of the defendant (Rosenberg and Spier, 2014) it has also been regarded with suspicion in Europe from business as being symptomatic of the worst excesses of compensation culture (Hodges, 2010). The use of such a system would have little to do with compensation in relation to drive-by downloads, and would be almost

exclusively to do with deterrence – even with a class action, the amount to be recovered is still far too small to be worth it in most cases (Wagner, 2011). This is good for the purposes of this thesis: although the possibility for victim to obtain redress is important, the primary purpose is for an improvement in incentive for more secure Internet systems.

Some practical issues remain prior to the implementation of such a system. Firstly, obtaining enough claimants into the class for the pursuit of an action. Aside from a few big websites, the amount of visitors to a website reduces very quickly[10]. The procedure for identifying the potential members of a class is also not straightforward, since the majority of Web browsing is done pseudononymously. Assuming a website is compromised and serving malware, access logs would be able to identify the IP address and browser user agent of any visitors. Further information would be required to identify the victims by contacting ISPs, who would then need to contact individual subscriber accounts to ascertain the identity of the victim in particular. The class has still yet to be identified, since it would also be necessary to show that the malware was successfully installed onto the victim's machine.

The problems described here are thought of primarily in terms of negligence, but mention should also be made of strict liability torts. It is possible that the use of a strict liability tort could lead to a more efficient outcome. In particular, the difficulty of determining what amounts to a reasonable level of precaution could be reduced (Citron, 2006) and that it might reduce transaction costs (Shavell, 1980). However, this remains problematic. It is still necessary to get past the causation issue, and the distribution of losses remains a problem. This means that there remain considerable difficulties in making use of tort law, even if there is no need to prove blame.

3.2.3 Case Study: Article 4A-U.C.C

One area which illustrates clearly the difficulties faced by private law in ensuring an efficient system has been the area of commercial bank transfers in the USA. Whilst American consumers are limited in the liability they owe for fraudulent transactions, the same protection does not apply to commercial transactions, where it is governed by Article 4A-U.C.C. According to this statute, under Article 4A-202(b), if the bank is adopts a "security procedure" to identify the customer, then the customer is liable for the loss if:

i) the security procedure is a *commercially reasonable* method of providing security against unauthorized payment orders [emphasis mine], and

ii) the bank proves that it accepted the payment order in good faith and in compliance with the security procedure and any written agreement or instruction of the customer restricting acceptance of payment orders issued in the name of the customer.

[10] See discussion about the website rankings and distribution in Chapter 6

This section will illustrate how the three major cases have not provided clarity in terms of what constitutes a "commercially reasonable" level of security; and the suboptimal level of litigation inherent in the system.

The creation of Article 4A in 1989 was specifically designed anew for these sorts of transactions, due to the complex balancing of interests and the requirements of assessing risk inherent in dealing with large sums of money like this without the need for concern (French, 1990). This led to some interesting provisions, like the fact that the banks were not required to have the "best" security procedure in place but rather one which was "commercially reasonable". The "commercially reasonable" requirement is decided by law (per §4A-202(c)), allowing the bank to decide the amount to invest in security considering the preferences of the customers, the size, type and frequency of payments, and what other similar banks are doing. It is specifically stated that "[t]he standard is not whether the security procedure is the best available. Rather it is whether the procedure is reasonable for the particular customer and the particular bank...." (per §4A-203, comment 4)

Offering the banks immunity from liability if a procedure was followed was intended as a means of improving security, by providing an incentive to the banks to introduce procedures where previously they had not. The intention was that the customer would discuss their requirements in detail with the bank based on the factors described above, and a procedure would be agreed for each one. As Internet banking became increasingly common, there have been some difficulties in regard to whether the customer authorised the transfer if their computer had malware or they were the victim of a phishing attack.

Two cases at around the same time sought to answer these questions, but provided little practical guidance as to what exactly constituted an acceptable level of behaviour as regards to banks' security systems.

3.2.3.1 Patco Construction v Ocean Bank

In *Patco Construction v Ocean Bank*(2011 and 2012, hereafter *Patco Construction First Instance*, and *Patco Construction Appeal*) the plaintiff suffered a net loss of $355k following a string of withdrawals which totalled $588k, of which $243k was blocked or recovered. This was as a result of the ZeuS Trojan being installed on the machine which they did their banking transactions[11]. At issue was the degree to which the bank's security procedures were "commercially reasonable". The parties agreed that the FFIEC's document "Authentication in an Internet Banking Environment" would be regarded as an adequate determination of commercial reasonableness. At the time of the loss, the latest version of the guidance was from 2005.

[11]This was disputed by the bank, due to the forensic information being inadequately preserved. It was a question of fact, and not ruled upon in this case.

Key amongst the guidance was the notion of two factor authentication, with the document advising that single factor authentication alone (e.g. ID/Password) was inadequate for "high-risk transactions involving access to customer information or the movement of funds to other parties", and reminding that many such compromises had happened as a result of this level of security. The document recommended that the security system have "layers", and that authentication constitutes at least two of the following three kinds of authentication:

1. Something the user ***knows*** (e.g., password, PIN);
2. Something the user ***has*** (e.g., ATM card, smart card); and
3. Something the user ***is***(e.g. biometric characteristic, such as a fingerprint) (**?**) – emphasis original.

In *Patco Construction*, the bank used the Jack Henry NetTeller system. They elected to make use of six features: device cookies, risk profiling, challenge questions, user ID/PW, dollar rule, and subscription to the eFraud network. However, they declined the opportunity to implement out of band authentication; user picture; tokens; and (most importantly) monitoring of risk reports. A decision taken by the bank in relation to the "dollar rule" was to lower the threshold to $1, beyond which all transactions would add the "challenge questions" to the procedure. The challenge questions were also the mechanism by which suspicious transactions were challenged. These red-flags were logged, but no action was taken to prevent them.

The plaintiff claimed that as a result of certain decisions made by the bank, the security procedure was not commercially reasonable. These decisions included lowering the dollar rule to $1 and asking the challenge questions increased the likelihood of keylogger malware (which was known about at the time) successfully obtaining this information before it was discovered. This problem was compounded by the fact that the extra "factors" and layers fell back to the challenge question as the sole additional challenge, thereby essentially limiting the authentication to a single factor.

The bank argued that they implemented the latest security system, which complied with the recommendation by including layers and multiple factors of authentication. They observed that the decision to lower the dollar rule was in response to a series of low dollar thefts, and that was a reasonable response. Similarly, they argued that other banks in their circumstances were operating similar systems.

At first instance, the court held in favour of the bank (*Patco Construction, First Instance*). They emphasised the nature of the security product used by the bank, and how it was designed specifically in order to comply with the FFIEC guidance. The judge also noted with approval the additional measures offered by the system. Whist agreeing

that the lowering of the dollar rule to $1 was in hindsight "suboptimal", it was held that it was "highly significant that Jack Henry permitted its bank customers to adjust the Dollar Amount Rule threshold to any level, including as low as $1", and that its use in that manner in response to another type of fraud was essentially an act of balancing the risks as required in the banking environment. Finally, the court was persuaded that even had the changes suggested by the plaintiffs have been implemented, they would not have made any difference (*Patco Construction, First Instance*).

This was overturned on appeal, as a more balanced approach was taken by the judge in viewing security as far more like a *process* as opposed to a *product*. On appeal, the court held that the decisions made by the bank, such as lowering the dollar rule to $1, as well as failing to do anything to chase suspicious transactions despite the ability to do so, meant that the system was not commercially reasonable. They regarded the security as an overall process which the bank failed at, holding that "it was these collective failures taken as a whole, rather than any single failure, which rendered Ocean Bank's security system commercially unreasonable" (*Patco Construction Appeal*, at 211). Given the background, with knowledge of keylogging malware in the industry, and the reduction of the dollar rule which "essentially deprived the complex Jack Henry risk-scoring system of its core functionality" should have led to extra security measures from being taken. They observed that at the time, hardware tokens were being used within the industry, and where they were not, manual verification was taking place.

3.2.3.2 Experi-Metal v Comerica

In *Experi-Metal*, the plaintiffs were victims of a phishing attack which caused a loss of $1.9 million, of which $1.3 million was successfully recovered by the bank. They sought to hold the bank liable for the losses under Article 4A. On the date the attack occurred, at 7:15 the victim gave information relating to his secure token and then "[b]etween 7:30 a.m. and 2:02 p.m., ninety-three fraudulent payment orders totalling $1,901,269.00 were executed using Mr. Maslowski's user information". In order to facilitate this, the attacker moved money around, totalling approximately $5.6 million in twenty transfers, including the creation of an overdraft. Only three of these transactions were rejected, as a result of "funds not available"

Having detected the fraud at 11:30, and after confirmation from the plaintiffs that the wire transfers were not authorised, steps were taken to withdraw all wire transfers from the queue and to recall as many as possible. Actions were quickly taken to disable the accounts, but they failed to disable the account holder already logged in for another hour and a half until at 2:05 that session was completely killed. Between 12:30 and 2:05 a further set of transactions were made, all were recovered bar one which was worth $49,300.

Unlike *Patco Construction*, the level of "commercial reasonableness" of the security procedure was not called into question. Instead, the two questions at issue for the court were: 1) Was the person whose information was stolen authorised to make transactions on behalf of Experi-Metal, and therefore did Comerica adhere to their security procedure?; and 2) did Comerica accept the payments in "good faith"?

As to question 1), the court held that despite the absence of a direct document authorising him to conduct ACH transactions, Experi-Metal acknowledged that he was. Despite the evidence of another employee claiming she believed that he was no longer entitled to do so on the date of the phishing attack, the court was unconvinced. As such, it was held that Comerica had adhered to their security procedure.

Question 2) had a far more detailed discussion, and is likely why the court chose to go into such detail as to the exact response of the bank following the initial discovery. According to the U.C.C., "good faith" is defined as "honesty in fact and the observance of reasonable commercial standards of fair dealing.". The court held that there was no suggestion of dishonesty in the Comerica staff (the subjective "honest heart, empty head" test), however they also had to prove that the staff acted according to "observance of reasonable commercial standards of fair dealing.".

In this, Comerica appeared to make an error in the presentation of their case. Both parties agreed that the burden of demonstrating good faith lay with the bank, yet Comerica failed to present any evidence to this effect. The only suggestion that there was of the commercial standards of the staff was by an expert witness for another area – one who was not qualified to make such an assessment. The court held that in order to prove this fact, they would be required to:

> "Comerica was required to present evidence from which this Court could determine what the "reasonable commercial standards of fair dealing" are for a bank responding to a phishing incident such as the one at issue and thus whether Comerica acted in observance of those standards."

Having failed to do this, they failed to rebut the presumption, and as such were deemed not to have acted in good faith.

3.2.3.3 Choice Escrow & Land Title v BancorpSouth Bank

In *Choice Escrow*, $440,000 was stolen from an account that the plaintiff held with the defendants. They sought to recover the funds on the grounds that the security procedures instituted by the bank were not commercially reasonable. The bank counter-claimed for the attorney's fees in relation to defending the court case, based on an indemnification agreement they had signed with the plaintiffs.

The banking platform InView, offered by the defendants included four security features: UID/PW; PassView, to ensure that the computer was identified; placing dollar limits on the wire transfers; and "dual control", where two individuals were required to authorise themselves for a payment order before the bank would send it. The plaintiff declined the opportunity to implement the last control, and signed a waiver acknowledging this. A further opportunity was also declined by the plaintiff to implement this, when they enquired about limiting offshore wire transfers to limit the impact from phishing, they were told it was not possible and the bank recommended that they implement the "dual control" procedure.

In March 2010 the employee's account was compromised by malware, which caused $440,000 to be transferred from their account to an account in Cyprus. The plaintiff argued that although the bank complied with the security procedures, they were not commercially reasonable, arguing instead that the transactions should be manually inspected. The court rejected this argument, and concluded that the dual control was commercially reasonable. The distinction made by the court between the objective element of the "good faith" test and the commercial reasonableness test was as follows:

> "While the commercial reasonableness inquiry concerns the adequacy of a bank's security procedures, the objective good faith inquiry concerns a bank's acceptance of payment orders in accordance with those security procedures."

This is a distinction which makes sense, and the court held that the bank had processed the transactions in good faith, and that the loss had to be accepted by the plaintiff.

The interesting element in this case was the indemnification that the bank agreed with the plaintiff, that:

> As long as BancorpSouth has performed as provided in Section 8 above, the Customer shall indemnify and hold BancorpSouth harmless from any and all claims, damages, losses, liabilities, and costs and expenses, including reasonable attorney's fees, which relate in any manner to the Services performed under this Agreement. (Cited by *Choice Escrow v BancorpSouth*, at 625).

The court overturned the first instance decision, and held that since this was not inconsistent with Article 4A the statute permitted additional indemnifications such as this – and as such the plaintiff was liable for the bank's attorney fees.

3.2.3.4 Analysis

Cases relating to Article 4A have been sparse, particularly since Internet banking became more common, and so has the academic literature on the subject. The lack of academic

literature is likely due to the fact that we have yet to see a ruling from a court higher than a circuit appeals court.

Burrow (2013) argued that the *Patco Construction* was incorrectly decided, since it ignored precedent that the FFIEC guidance was adequate to protect a bank under the "commercially reasonable" standard, citing several divisional decisions. Similarly, he argued that the decision failed on policy grounds, by leaving the required actions an unspecified level beyond the FFIEC document, and argued that the financial burden on small banks could be considerable. Hene (2010) presented a review of cases concerning the determination of "commercially reasonable" security prior to the three cases discussed. It does appear to show that the obligations on banks were indeed limited in the offline world, but also discusses some of the issues regarding the increased move to online banking. Somewhat of a concern was that in 2010 the "widely used" password systems were regarded as inadequate – but more as a result of password guessing, or the MySpace phishing attack, as opposed to commonly used malware (Hene, 2010).

French (1990) discussed the provisions of the new (as it then was) Article 4A. Interesting to the discussion here, are some of the ideas relating to the security procedures. Firstly, he suggests that the customer may actively negotiate with the bank for different security procedures. Whilst in 1990 this was likely the case, when a large amount of banking transactions are done online, the expertise resides with the banks placing the customer at a significant disadvantage. The disadvantage is likely magnified by the following opinion:

> The sender can be expected under Article 4A to exercise greater scrutiny regarding the security procedure than may be the case under current law. One reason for this is that under Article 4A the sender of a payment order bears the risk of loss for certain unauthorized payment orders verified pursuant to the security procedure. Therefore, it is in the best interest of the sender to ensure the security procedure is adequate. (French, 1990)

Security relating to Internet banking in general, and financial transactions in particular, is a specialise area. The probability of a client being in a position to adequately assess the risks, and therefore negotiate, is small.

Another issue, is the amount of time it can take to get certainty in the law. So far there have been three major cases, and a few minor ones (see Burrow (2013)). Yet, there is still no certainty as to exactly what constitutes a "commercially reasonable" security procedure. By 2012 we knew that the implementation of Ocean Bank's procedure in 2009 was not reasonable according to a document written in 2005. *BancorpSouth* had a procedure that would probably have been reasonable (having rejected the security procedure, the plaintiff sought to claim an unrealistic requirement which was rejected).

It is also worth noting, that when faced with the same set of facts, the two judgments in *Patco Construction* came to completely opposite conclusions.

Similarly, there are also no guidelines as to exactly what the customer is supposed to do (Burrow, 2013). In *Patco Construction* there were some disputes as to the facts of the case relating to the communications between the customer and the bank, and so they could not rule on that issue. Similarly, it was not tested in *Experi-Metal*, and in *BancorpSouth*, the bank's procedures were found to be commercially reasonable, so there was no need to try it. In this setting, it takes two to tango in order to conduct a funds transfer, and as Compton observed "Banks literally spend hundreds of thousands of dollars on their security systems, and as between [the customer] and the bank, it is almost always the case that it was the customers system which was hacked...." (Cited by Hene (2010)).

A large part of the problem, is the cost of bringing suit being prohibitively expensive. In *Patco Construction*, with $355k riding on the outcome, the judge concluded by suggesting that "the parties may wish to consider whether it would be wiser to invest their resources in resolving this matter by agreement". In the end, Patco reached a settlement with the bank, whereby the bank paid the money lost, but they both paid their own legal fees (Valigra, 2013).

The opportunity for customers to bring suit has also likely been reduced as well, with the finding that they can be liable for the bank's legal fees (*BancorpSouth Bank v Choice Escrow*). Whilst this is perhaps an inevitable consequence of Article 4A, where the customers can only claim the money lost (with interest), it is an illustration of exactly the sorts of issues which we can expect in other cases relating to security. The amounts of money involved here were large, $355k, $560k and $440k, which is a situation exceptional to this sort of transaction (which is why Article 4A was required as a separate body of law). This is simply an illustration of the difficulties discussed in the previous section about bringing suit to ensure the optimal level of precaution be taken.

3.3 Stakeholders

This section considers the major stakeholders involved in general Web use, and considers what is their role and what it is that they might be able to do in order to improve security overall. It is concluded in Section 3.3.7 that hosting providers should take action in order to mitigate drive-by downloads. A small proposal of what a reasonable hosting provider might be expected to achieve is proposed.

3.3.1 Software Vendors

Software enables complex functionality to be performed on computers, by abstracting away a series of electrical gates into something useful. The instructions required to write software have become more and more abstract, such that now commands that look similar to natural language can be used. It remains a very complicated and time-consuming process, so an industry has emerged where customers will pay for specific functionality provided by the software. An inevitable consequence of the complexity inherent in large pieces of software, is that the software will contain bugs, which will occasionally be vulnerabilities which allow an attacker to exploit the system. The motivations for the generic "software vendor" are difficult to define, because there are so many types of vendor offering different types of software, so a few broad categories will be discussed.

A major distinction between software is whether it is an open or closed source product. A characteristic of an open source piece of software, is that the source code is made available, and often will have contributions by developers in different locations. By contrast, a proprietary piece of software is developed by a single company, and the source code is kept as a trade secret. The difference between the two has been referred to as the cathedral (for proprietary) and the bazaar (for open source) (Raymond, 1999). It is not completely as simple as this, because whilst an open source product can generally not be charged for in itself, it can still be a for-profit endeavour with licensing restrictions or support costs for other uses. In the context of this group as stakeholders, a difference worth thinking about might be that the developers of an open source product may be less likely to be identifiable or have deep pockets. For example, the Heartbleed bug connected to SSL was created by developers who were largely working voluntarily on an incredibly small budget (Marquess).

In terms of the relative security, it is probable that there is not in fact any difference between the two types of software. Ross Anderson argued that (in an ideal world) these arguments on either side cancel each other out and that the openness (or otherwise) of the source does not matter Anderson (2005). This is supported by the empirical analysis provided by Schryen has performed empirical analysis of vulnerabilities and patching of open and closed source products. 17 widely adopted packages were compared including operating systems, web browsers, email clients and database packages Schryen (2011). Using data from the National Vulnerability Database[12], the mean time between vulnerability and disclosure, severity of the vulnerability and the patching (or not) behaviour of the vendor. It was found that there were no significant differences between open and closed source.

Although software has different purposes, these largely align with the general motives based on the business models described above. For those who are commercially minded, being first to market offers a substantial advantage, because once it reaches a critical

[12] https://nvd.nist.gov/

mass then it becomes very difficult for the end users to switch to other software. For example, because Microsoft had the dominant operating system, then more people would write software to work on that operating system meaning that there was a greater level of functionality, and as such more people would buy it. Similarly with an application type, switching becomes difficult when it comes to transferring files between format (the dominant format is likely to win), or training costs to move to the rival piece of software. As such, vulnerabilities within their software are an afterthought to the immediate requirements of releasing the software. Similarly, the end users are unlikely to notice vulnerability being fixed and better appreciate additional features being added, further lowering the priority of vulnerability fixing (Anderson, 2001).

However, the major software vendors already appear to spend considerable resources in fixing vulnerabilities in a timely manner, and releasing the updates to their users, suggesting that there are competing incentives at play. This is likely to do with limiting reputational damage, which is something which has affected certain companies in the past. For example, Microsoft previously had a reputation for not caring about the security of their products as they struggled to cope with the transition to an environment where their operating system went from one designed for a single user to one connected to a network where it was under constant attack. Similarly, Oracle's Java has come under fire due to the amount of security flaws in it, leading to advice to abandon or severely restrict its use[13].

3.3.2 End Users

The end users here are the people who use their Web browsers to visit websites. Although the population of users will have a wide range of ability, technical competence and education, there is a general low opinion as to the "weakest link" in the security chain. Many dismiss users' ability to make correct security choices, and agree with McGraw et al. (1999) quip that "given a choice between dancing pigs and security, users will pick dancing pigs every time". By acting in this way, then they place themselves at risk yet the costs they create are usually borne by someone else. Is this general assessment of users fair? It is certainly true that there are some users like this, who will download everything, not patch their computers and decline to run security software but what about the population in general?

There is some evidence that Web users place a low economic value on the security of their machine. Christin et al. (2012) and Kanich et al. (2011) both ran studies on the crowdsourcing tool "Mechanical Turk" https://www.mturk.com/mturk/welcome, which allows companies to offer to pay users a small amount of money to complete tasks easily accomplished by humans but difficult for computers to do. Both studies asked

[13] Journalist Brian Krebs, who reports on Banking fraud and other security related issues is one such advocate of this, see e.g. http://krebsonsecurity.com/2014/07/java-update-patch-it-or-pitch-it/

Exploit	Platform or technology	1Q12	2Q12	3Q12	4Q12
Win32/Pdfjsc*	Documents	1,430,448	1,217,348	1,187,265	2,757,703
Blacole	HTML/JavaScript	3,154,826	2,793,451	2,464,172	2,381,275
CVE-2012-1723*	Java	–	–	110,529	1,430,501
Malicious Iframe	HTML/JavaScript	950,347	812,470	567,014	1,017,351
CVE-2010-2568 (MS10-046)	Operating system	726,797	783,013	791,520	1,001,053
CVE-2012-0507*	Java	205,613	1,494,074	270,894	220,780
CVE-2011-3402 (MS11-087)	Operating System	42	24	66	199,648
CVE-2011-3544*	Java	1,358,266	803,053	149,487	116,441
ShellCode*	Shell code	105,479	145,352	120,862	73,615
JS/Phoex	Java	274,811	232,773	201,423	25,546

TABLE 3.2: Reproduced from Microsoft Security Report July - December 2012. *Vulnerability also used by the Blacole kit, the totals for this vulnerability exclude Blacole detections

users to install untrusted programs on their computer for amounts as little as $0.01 and found a significant percentage were prepared to do so. The popularity of file-sharing tools similarly demonstrates a lack of care for a machine's security.

It would also not be unfair to suggest that there is (or at least was) a significant percentage of people who do not adequately patch their machines. Conficker, for example, exploited the MS08-067 flaw (amongst others) which had already been patched by Microsoft (Conficker Working Group, 2011). This did not stop it successfully infecting millions of computers, and many computers appear to still be infected by it as they attempt to communicate with the sinkholed C & C servers. Similarly, Table 3.2 shows the most common exploits that were being attempted against clients in 2012. Most of them had already been patched for a considerable period of time, indicating that there are many unpatched computers since there would be no point in attempting the exploit otherwise.

A few points should be clarified about the data presented in this table. CVE numbers are unique identifiers for a particular, reported vulnerability, prefixed by CVE, the year it occurred, and then a unique number from that year. Blacole is an exploit kit, which identifies the versions of software being used by the users visiting a website, and then chooses an appropriate exploit based on that information. The persistence of some of the exploits over the course of the year demonstrates that they are still successful long after they are known about, otherwise they would not continue to be used. This suggests that any solutions which require installation of additional software to prevent the execution of attacks are unlikely to have success on their own, because the only people likely to install them are people who are already largely safe.

This is a problem which has got worse in 2014, since Microsoft stopped supporting Windows XP which continued to retain a significant market share (netmarketshare.com, 2014). Whilst businesses can afford the high additional costs to keep security patches, this is unlikely to be possible for most home users and as such there is a significant proportion of computers which are likely to be subject to additional security flaws. There

are also practical difficulties in distributing patches, and enabling users to understand security issues on mobile devices (Android in particular), which could lead to a big increase in mobile malware over the coming years (Jeter and Mishra, 2013).

On the other hand, this is not representative of the whole problem. For example, Herley argued that many of the security warnings presented to them are protecting against theoretical rather than actual problems, and that adopting them would lead to a considerable loss in terms of time outstripping the potential loss from an attack. He also pointed out, that the "dancing pigs" comment is unfair, in that users not actually offered "security" but rather are offered a set of complicated guidelines for managing risk (Herley, 2009). Consider the choice of password strength for a website account. A stronger password is harder to remember but the only time that this is likely to offer an additional advantage to the user is in the event that the website database is breached (with the password safely stored) and an offline attack can be mounted against it. In most other circumstances there is no advantage: if they are phished; a keylogger is installed on their computer; or if the website stores the password in plaintext. Herley also observed that by spending more than 0.36 seconds a day checking links to see if an email is a phishing email then an individual user will lose more in lost time over a year than they would if they fell victim to a phishing attack (Herley, 2009)

The advice offered to users is frequently out of date, and they are victim to badly designed interfaces which could offer them additional security. Adams & Sasse argue that it is these policies, which are incompatible with working practices, which is what causes bad security decisions rather than the users themselves. When they can see the rationale behind security requirements with well-designed software, their security practices are good (Adams and Sasse, 1999). The introduction of automatic updates being enabled by default has led to a reduction in the amount of machines which are too out of date. Firefox and Chrome will automatically update, and Windows as early as XP Service Pack 2 turned on the firewall by default and made Windows Update an opt-out rather than an opt-in service.

3.3.3 Website Operators

As was established in the previous chapter (Section 2.2), a compromised website is a key element for a successful drive-by download attack. Since it is the operator who is responsible for the content on the website then this places them in a position where they should take care to avoid their website becoming compromised. However, this position is complicated due to the fact certain things are taken care of by their hosting provider. The responsibilities of the two when dealing with drive-by downloads are inevitably linked, given the variety of different arrangements and allocations of responsibilities. Nevertheless, they are to be analysed as distinct stakeholders, since their motivations and business models are [usually] fundamentally different.

When considering how a website is to be defined, the definition is limited to a series of Web pages with their own domain name, intended to be viewed by the public. Hosting models such as Dropbox which offers the ability to store documents is not included within this definition, because the usual intention is not to publish Web pages, even if it is possible to do so, and use in social engineering attacks (e.g. by placing a link in a phishing email). There remain many types of websites, ranging in complexity from a tiny amount of static pages, to specify contact details and opening hours of an offline business, up to Web applications where the whole business is built around providing an online service. The technical expertise that the operator might possess can vary from a complete novice to a team of developers working full time to maintain it.

A team of developers would often be working on a website more critical to the business, yet this is no guarantee that they are in a position to make a website entirely secure. Consider the discussion in the previous chapter (Section 2.2.1) relating to XSS, where difficulties with the specifications the support different browsers need to provide to malformed code, and the huge amount of cases to be checked for make it incredibly difficult. Where a website relies on displaying user generated content (UGC), such as Facebook[14], YouTube[15] or Twitter[16] this becomes increasingly difficult.

Depending on the purpose of a website, an option is to use a CMS, which enables novice users to add content to a website, and can provide additional functionality, such as validating user input as just described. A CMS with a large community will be scrutinised and regularly updated, and are often regarded as being generally secure. However, recall that this also means that attackers can scrutinise it for weaknesses, and having a large user base has been hypothesised as being a "risk factor" for compromise due to the increased reward from successfully finding a vulnerability (Vasek and Moore, 2013). Some 35% of websites do follow this model, with nearly 70% of them using WordPress (W3techs.com, 2015), so a successful exploit would affect millions of websites, for the same effort that an exploit against a custom made website might require.

The functionality provided by different CMSes also varies. WordPress is generally basic, having been initially designed as a blogging platform. Drupal[17] is a more complicated platform, which is designed for more complex websites. Both offer the ability to add "plugins" in, which offer additional pieces of functionality. Frequently it is these rather than the CMS itself which are the source of vulnerabilities, see e.g. (Checkmarx, 2014). As security goes, operators might be reluctant to update these due to fears it might break website functionality, if they are even aware of the requirement to upgrade.

[14] https://www.facebook.com
[15] https://www.youtube.com
[16] https://www.twitter.com
[17] https://www.drupal.org/about

3.3.4 ISPs

Depending on the context, the phrase ISP (Internet Service Provider) can have many different definitions. For example, it could be used to include providers of Internet infrastructure (including hosts), or even more generally providers of any service on the Internet such as a search engine. In this case, however, an ISP will be regarded as the operator who physically provides access to the Internet (the access provider).

The intervention of an ISP is generally regarded as an effective strategy in minimising the effect of users participating in botnets. Unfortunately, an ISP has got limited incentives to perform any security actions on behalf of its users because the market is based on thin margins and is characterised by information asymmetries. This makes interventions costly, and as it is generally believed that users are not prepared to pay more for security there is little reason to improve it more than necessary.

They could identify compromised machines in their space, and have a remediation policy to prevent future attacks. This (usually) doesn't happen, because relative to the profits that they might make per subscriber, the cost of a remediation policy is quite high in terms of infrastructure and support costs[18]. Because users cannot tell which ISP is more secure, competition is fought on price, which makes it even more difficult to spare the resources (Van Eeten et al., 2010). There are also fears about what would happen if they made a mistake taking action against a certain subscriber. Not only would this cost the ISP money from the intervention, but inconvenience the subscriber and possibly even lead to legal action since, unlike in the USA, there is no protection for "Good Samaritan" (See more general discussion in Chapter **??**).

There are two broad approaches an ISP could take to mitigate the problem of drive-by downloads: that of content filtering, and that of remediation.

The first option is content filtering, where the ISP blocks objectionable websites to prevent access by the user. There is precedent for ISPs blocking websites in the UK, Cleanfeed and similar products reportedly used to prevent access to child pornography using the list maintained by the Internet Watch Foundation[19]. A recent agreement was also reached between the government, and major ISPs that they will require an "opt-in" from their subscribers before they will display (legal) pornographic websites (B.B.C, 2013). Although responsible ISPs would have blacklists for overtly malicious websites, doing so for compromised websites is not something which appears to happen to any great extent. As was alluded to in Section 3.3.2, this is something which could prove more effective than merely advising the users of a danger.

[18] See analysis in the paper by Clayton (2011) which suggests that the claim of one support call being a year's worth of revenue is an exaggeration, but not all that much (page 5 footnote 2)

[19] https://www.iwf.org.uk/about-iwf/remit-vision-and-mission. See again discussion in Chapter **??**

The other way in which an ISP could intervene would be by identifying infected customers as their machines start to talk to botnet command and control (C & C) server, and then by notifying them. This is possible because a customer's Internet traffic is all routed through the ISP, enabling them to see which servers they are communicating with. That a device is infected could also be noticed by other entities, for example those who observed the IP address as part of a denial of service attack against them.

This could go further by seeking to prevent the compromise from happening in the first place. This could be done by preventing user access to the network in the first place, or by filtering content. The idea of preventing access to the network is something which can be done based on the "posture" of the machine, and is known as Network Access Control (NAC), or Network Endpoint Assessment (NEA) (Sangster, 2008). This is something generally reserved for corporate networks. Scaling this for the Internet, is regarded as being specifically outside the scope of the IETF NEA Working Group (Sangster, 2008) and includes its own set of problems and considerations outside the scope of this research (such as privacy concerns, and lying endpoints, see e.g. (Anderson; Schneier, 2010)).

Given the lack of incentive for an ISP to intervene in their users' activities, there are a few occasions where legal obligations are placed on them, although often not for security. In relation to content filtering, rights holders have noticed the advantage of targeting an ISP to prevent access to websites which, they claim, participate mostly in copyrightable materials. ISPs generally no longer fight these injunctions, there are costs required in order to implement blocks against these websites – in *Cartier v BSkyB* (2014), at [61 – 65] Arnold J set out the claims by the ISPs of the costs of implementation of such a system, for example at [61] BT estimated that 60 days of employee time were spent enforcing the injunctions to date (discussed in Chapter **??**).

Occasional other systems have either been proposed or implemented in relation to remediation. Clayton (2011) is one example of an academic proposal, where public subsidies were suggested to assist ISPs in participating in remediation. In other countries, there have been agreements between ISPs, or other initiatives which require action. In the UK, The Digital Economy Act 2010 introduced the requirement for the courts to order ISPs to undertake technical measures against repeat copyright infringers (ss. 17 – 18), but these were never implemented and are due to be repealed (Department for Culture, Media & Sport and Vaizey, 2012). In Australia, ISPs are voluntarily signed up to a standard which requires them to notify customers (up to and including quarantine) (IIA, 2010), and the Dutch anti-botnet treaty covers 98% of the market (Buenaventura, 2009). Both have worked to some extent, though neither has been a stand out success. Microsoft's security report rates them as being better than the worldwide average, though nowhere near the best. Also, van Eeten et al. (2011) demonstrated the difficulty faced by ISPs, finding that they succeeded in contacting approximately only 10% of their infected customers.

3.3.5 Search Engines

Search engines provide listings of websites in order of their relevance to the search terms provided by the user. The most popular search engine is Google[20], with close to 90% of the search engine market in the UK. Their PageRank algorithm gives the rank of a website by working out the amount of other Web pages linking to it; but also adds to the calculation by a determination of the quality of the links to it (Page et al., 1999). In order to do this, the search engine will "crawl" every website it has listed and traverse through the links on the Web page. A website will usually specifically register to be indexed in this way, or if there are enough links to a page then that can be enough for it to be included in the search engine results.

In addition to use as drive-by downloads, a frequent use of compromised websites is to manipulate search engine rankings. This could be to advertise pharmaceutical products on a website which will be seen by more people than a spam email; or alternatively to raise the prominence of advert filled websites or attack websites to gain more victims (Moore et al., 2011a; John et al., 2011). This demonstrates the importance of a search engine in limiting the effect of drive-by downloads. Luckily, a search engine does have the incentive to eliminate drive-by downloads, because they reduce the quality of the search results, and consequently damage the product they are attempting to sell (Edwards et al., 2012). In addition, one of the consequences of a lot of infected users is click fraud. This is another thing which degrades the quality of the product for the search engine's product, since they sell the opportunity to advertise. This requires them to spend considerable resources on detection, and in addition, court cases have been brought. In 2005, companies like Google and Yahoo were forced to settle class action lawsuits (*Lanes and Others v Yahoo! and Others*(2005)).

A lot of work is already done by these companies, for example, the results discussed earlier, presented by Provos et al. (2007) are from Google. Google also provide a "Safe browsing" service, which allows developers to check pages they are about to visit to see if that page is listed as being malicious. An inevitable consequence of indexing almost all websites that exist on the Web is that they are in a position to see pages which are vulnerable, or which have already been compromised. As such, they are in an ideal position to minimise the effect of drive-by downloads.

A search engine could include elements of the security posture as part of their ranking algorithms, to make it more difficult for users to find allegedly malicious websites in their search results. In February 2011, Google rolled out a significant change to their ranking algorithm designed to remove low quality websites Google (2011), so it is not inconceivable that it could apply to security as well. An approach such as this was examined by Edwards et al. (2012), who simulated the effect of lost traffic and amount of malicious websites in response to this consideration in the rankings. Search engines

[20] http://www.google.com

already do this, the key point of this study was to examine the consequences of an increased level of false positives. An interesting idea they presented was the concept of "depreferencing", where the ranking of a website would be lowered by a function of the certainty of the search engine that the website was malicious. This paper is examined in more detail in Chapter 6.

In addition to the work already done by search engines to aid in the quality of their search results, there is precedent for a search engine performing an intervention. The DNSChanger malware caused its victims' DNS settings to point to a malicious DNS resolver under the control of the attackers. When this was taken offline, DNS would cease to work for many people, so the ownership was transferred to the control of the FBI for some time. This would run out eventually as well, so Google decided to warn people it suspected of being victims that they might be infected (Menscher, 2012).

3.3.6 Hosting Providers

As indicated in Section 3.3.3, website operators will frequently outsource a lot of the responsibility to the hosting provider. As the website grows in complexity, the more likely it is that there will be some requirement for greater control by the website operator over the infrastructure provided by the host. Some of the popular options will be enumerated here, followed by an analysis of the incentives and existing solutions. On occasion, the website operator will be regarded as the hosting providers as well, where they are in charge of the administrative functions. Otherwise, if they are only in charge of content they will be regarded as distinct.

The least control a website operator might retain over their website content is the use of CMS as a Service option, where the CMS comes ready installed. In this case, the hosting provider would be in charge of running administrative tasks on the server, and the CMS. `http://wordpress.com` is an example of such a service, the flexibility for the user is limited in that they cannot install custom themes or plugins on their website.

One of the more common hosting option would be shared hosting (Canali et al., 2013a), where the provider offers FTP access to a part of a Web server, and allows the user to make use of pre-installed Server and database software. This would be shared with other customers, with associated limits on the space, bandwidth and control offered. In the event that they were to want to install a CMS, then it would be possible to do so manually, or "one click" installations through software such as Softaculous[21] The administrative responsibilities would generally be with the hosting provider in this case.

Finally, a hosting provider could offer full access to a physical server, or virtual machine (VPS), where the customer gets root access and installs the applications they wish in

[21] `http://www.softaculous.com/`

order to serve content. This is something which would offer higher performance than having to share resources with other websites (like with shared hosting), and offers full flexibility for the operator. To serve content, they can install the Web server as they wish and customise it to their own specifications, but would have to perform administrative and security tasks themselves.

As can be seen, the trade-off for the website operator is flexibility vs additional effort. Where additional effort is required, then they begin to take responsibility for the security of the server. In those cases (VPS and dedicated server models), the website operator would be regarded as equivalent to the hosting provider.

Research has shown that the operators of websites frequently appear to be caught unawares when their website becomes compromised. StopBadware and CommTouch (2012) conducted a survey of website operators whose sites were compromised, finding that 63% of their sample didn't know how their website got compromised, and only 6% were able to identify the fact that their website had been compromised. Almost half were only notified when they were faced with a browser warning screen. Regarding the recovery, some 26% of the websites remained compromised, and only 46% were able to solve the problem themselves.

This would suggest that the hosts are better placed than the operators to respond to security threats. Intuitively this makes sense, since they are likely to have greater permissions on the server and better technical knowledge. Indeed, it is regarded as a responsibility of a good provider to be able to respond in the event that a website is compromised.

Whether hosting providers are currently doing a good enough job in this, however, is open to question. From the responses received by StopBadware in their survey, there was general dissatisfaction at the lack of support from the hosting providers once websites become compromised. Some believed that the host was directly at fault, in that by becoming compromised themselves that led to all the customers on the server being compromised as well.

This was a motivation behind a study by Canali et al. (2013a), who examined the ability of popular Web hosts to deal with attacks and compromises, and appeared to show that the security protection on offer is inadequate. They created websites at a series of different providers, and simulated different types of attacks on them in order to see how the hosting providers responded. Whilst some providers were able to block or mitigate some of the attacks, none of them provided complete protection. Of those who failed to mitigate these attacks, the researchers made an abuse report to the provider, which is a complaint to the hosting provider about inappropriate content on the websites. Of those, only 50% of hosting providers replied to the report, and only one chose to notify the website operator.

The factors behind a hosting provider not taking more proactive action, or possibly greater care are likely similar to those of ISPs. There is no way for the user to perceive whether the hosting provider is taking additional care, making them unlikely to pay extra unless they are under some specific requirements, and even where the customer might choose to pay for additional security services, they frequently cannot detect them (Canali et al., 2013a).

That is not to say that there are not occasional companies who do extra. For example, in 2011, Dutch hosting provider Antagonist announced a service for finding and fixing vulnerabilities in their customers' websites (de Vries, 2012). Upon enquiry, Antagonist declined to elaborate further on the technology or effectiveness, but this is a demonstration of the position the hosting provider is in to make a difference where the website operator lacks the requisite knowledge or desire to do so themselves.

This might be as simple as upgrading the CMS, to completely changing the code used to communicate with the database. In the event that the operator chooses not to respond, the most extreme sanction that could be applied would be to completely block the website. But whatever the method employed, whether using tools to automatically update, or blocking entirely, the hosting provider is in a great position to provide assistance.

3.3.7 Conclusion: Hosting Providers Should Take Responsibility

The preferred option is that hosting providers be considered responsible for drive-by downloads. Conventional legal theory would agree with this – they are the "deep pockets" who actually serve the content. Chapters 5 and 6 will demonstrate that they are the most effective in terms of being able to mitigate it.

This section offers a proposal for actions which should be taken by hosting providers in order to mitigate drive-by downloads. The discussion is presented around "reasonableness" – i.e. what a reasonable hosting provider should be doing. This is something which could be the form of regulation, whereby the state could fine operators who fail to adhere to the requirements, or an indication of negligence which could be pursued by victims. The discussion about Article 4A in the USA suggests that it could be difficult to adequately incentivise people to bring suit.

Some possible requirements which hosting providers could be required to adhere to are as follows:

1. Outdated/vulnerable software be identified and patched as soon as possible after a fix is available (whether by informing users or conducting the update themselves).

1	Injection
2	Broken Authentication and Session Management
3	Cross-Site Scripting
4	Insecure Direct Object References
5	Security Misconfiguration
6	Sensitive Data Exposure
7	Missing Function Level Access Control
8	Cross-Site Request Forgery (CSRF)
9	Using Components with Known Vulnerabilities
10	Unvalidated Redirects and Forwards

TABLE 3.3: OWASP Top 10 Vulnerabilities 2013 (OWASP, 2013)

2. Vulnerabilities in the OWASP Top 10 be identified, and the customers informed and be required to fix as soon as possible. The current Top 10 can be seen in Table 3.3.7[22].

The first requirement, of ensuring that software is kept up to date does appear in the Top 10 list itself, under both item 5 (security misconfiguration) and item 9 (using components with known vulnerabilities). Nevertheless, this is regarded as more important for the current proposal. The first reason for this is that, for an attacker, a vulnerability within a piece of software has considerably greater impact than identifying a vulnerability on an individual website. The websites which would be affected by this proposal are generally websites which make use of CMSes – and amount to nearly 40% of all websites (W3techs.com, 2015). As such, this can be regarded as the easier "win" for reducing the overall level of compromises.

For the second requirement, OWASP is an appropriate choice since it is technologically neutral, and the high amount of members they boast ensures that it can be regarded as a widely applicable pseudo-standard. It is also referenced in the current PCI Data Security Standard under Requirement 6 "Develop and maintain secure systems and applications" (PCI Security Standards Council, 2013). Given its popularity, tools already exist which can audit websites for compliance with the OWASP Top 10, any of which can be used for this purpose[23]. That Antagonist are contemplating offering a service such as this to their customers indicates that it is feasible[24].

In Chapter 6, the efficacy of the proposal of hosting providers checking websites for outdated software is tested through simulation. The exact requirement of exactly what a hosting provider should be expected to do is left open, but essentially there are two options. Whatever the outcome, to begin with the hosting provider has to let their

[22] There is more detail about each vulnerability on the OWASP website https://www.owasp.org/index.php/Top_10_2013-Top_10.

[23] See, for example, the list compiled by OWASP themselves at https://www.owasp.org/index.php/Category:Vulnerability_Scanning_Tools.

[24] Antagonist declined to offer any details about implementation or costs relating to the service following my request, or whether it had led to any change in their customer base

customer know that they have discovered a vulnerability in their website. This would likely have to be done through an email, which would minimise costs and also ensure that it happened with reasonable speed. Having to access the Internet already in order to make any changes to the website, it is far more reasonable to expect a website operator to expect to receive emails relating to the service online (as opposed to the issues discussed for ISPs (Livingood, Jason and Mody, 2012)).

This could be enough on its own – if a website operator were to get informed about a vulnerability on their website with a proof of concept method of exploitation, then it in many occasions that may be sufficient to incentivise them to solve the problem on its own. Alternatively, if this is not enough, then it might be necessary for the website to be either blocked or fixed by the hosting provider. Since version 3.7, WordPress automatically update blogs which are installed, unless the website operator specifically opts out (WordPress, 2013). This suggests that it is technically possible to do so in many occasions without breaking functionality.

Blocking is an alternative option, the equivalent of a quarantine procedure as far as the Web is concerned. Ideally requiring a website to be blocked has an impact on freedom of expression, and following public health principles, ideally restrictions of rights such as this should be limited. On the other hand, if other options fail then this is a viable solution. Another suggestion that might be feasible, would be to limit the amount of bandwidth available to the website, so that it takes a long time for the website to load. This was a possibility in the DEA (Section 124G(3)(a)) that ISPs might have had to implement to implement to repeat copyright infringers. This could serve as part of an escalating response to a non-response from a customer prior to blocking the website completely.

The other issue which needs to be decided is that of cost. From the progress of the Digital Economy Act (DEA), it can be seen that one of the major issues for the ISPs was that of cost (British Telecom, 2012; everything everywhere, 2012; TalkTalk Group, 2012). In particular, it was suggested that one of the costs which Ofcom had underestimated was that of customer support following the sending of letters (British Telecom, 2012), and this is likely to be an issue here as well, especially if the hosting provider needs to block the vulnerable or malicious websites. In relation to the DEA, the issue was between the ISPs and media rights holders over who should bear the loss for copyright infringement, in this instance it is an issue for the government. As part of their aim to protect the population online (see Section 4.3), then it is not unreasonable to expect that the government would be able to offer subsidies to providers such that they can remain competitive compared to other countries.

3.4 Discussion

This chapter has discussed the roles that various stakeholders could play in improving security online, and discussed some of the market conditions which exist online meaning that it is unlikely that they are going to do much. Private law, particularly negligence, was discussed as a possibility of mitigating these unfortunate conditions using theories of liability which have previously been applied in the offline world. However, it was concluded that the distribution of losses mean that it is unlikely that enough people will bring tort claims meaning that the level of care required to be taken is not socially optimal. A class action solution is a something which might reduce the issue of distributed losses, although the benefits it might offer appear to be minimal.

The difficulties presented in the application of tort, or other private law, lead to the notion that alongside a private law solution, it would be necessary to have a form of public regulation alongside it.. Shavell (1983) argued that the determination of whether to use liability or regulation came down to four determinants:

- Difference in knowledge in relation to the risky activity;
- Administrative costs incurred;
- Inability to pay for the harm done;
- Possibility of escaping suit for harm done.

The decision as to the optimal choice will be based on the relative advantages of each one. Ordinarily, the first two would be favourable to tort liability. The administrative costs happen only in the event of harm rather than all the time; and for most activities the individual will know better about the specific merits of the activities. On the other hand, there are circumstances where they are less advantageous – the incentive to take care to avoid liability is reduced if the maximum cost is greater than the assets the tortfeasor possesses. In addition, in circumstances where it is not individually optimal to bring suit or difficult to prove causation, then regulation may be the more favourable option.

In relation to the stakeholders as described in this section, public regulation does appear to be the better option. Section 3.2.3 demonstrated the incentives to bring suit, and the causation issues were touched upon in Section 3.2.2 reducing the risk of harm successfully being tied to the tortfeasor. In addition, the difference in knowledge assumption does not apply to the same extent, because there are information asymmetries (see Section 3.1.3). Shavell (1983) suggested as much himself that circumstances where expert knowledge is required does in fact place the regulator in a better position in this case owing to the better resources they possess to gather the required information.

Specifically, three difficulties with the use of private law were discussed: causation; reasonableness; and distribution of losses. These were demonstrated with a case study, which examined the major cases connected to U.C.C. Article 4A in the USA. The *Patco Construction* cases demonstrated the issues which exist with *ex post* cases for determining an adequate standard of behaviour when the cost of litigation is so high. It also appears to demonstrate that providing an incentive by exempting certain actors from liability does not work properly, where the prospect of a claimant bringing suit is low. The lack of cases which have been through the courts in relation to this means that there is still no certainty as to which cases are likely to succeed, and the limited incentive to bring suit means that there is no incentive for adequate security standards to be maintained.

Strict liability was considered as a possible alternative to negligence law and, whilst it may be slightly more efficient, the same sorts of problems remain. This demonstrates that, for Internet security in general, an *ex ante* system of regulation is required to supplement the private law and market based solutions we have discussed. As a means of guiding government regulation, the next chapter introduces the idea of a public health analogy, which is a similar situation where government intervention is required in order to correct market failure. Traditionally, this analogy has focused more on specific techniques for implementing strategies on individual networks, but recently there has been a wider interpretation of this analogy which offers guidelines based on the fact that cybersecurity is a public good.

Chapter 4

The Public Health Analogy

Ring-a-ring o' roses,
A pocket full of posies,
A-tishoo! A-tishoo!
We all fall down. (Children's nursery rhyme, Origin Unknown)

Having analysed the stakeholders involved in Web and Internet security, and some economic background (Chapter 3), it was concluded that no satisfactory security standard has emerged with market forces or private law alone. As a means of guiding policy and setting out a framework for public regulation, it was decided that analogy was a useful tool. The discussion presented by Betz and Stevens (2013) provided a summary of the two popular analogies for cybersecurity: that of cyberspace as a space; and biological analogies. These are the two analogies which will be considered.

The cyber "space" analogy is frequently used alongside "cyber war", where it is viewed as a "space" to be defended against an adversary, and these will be considered together. The idea is that by setting a tight perimeter, one can keep out the "enemy" whether that is worms; viruses, or other malicious software. This strategy is evident in the common installation of firewalls on personal computers; and at a national level to the "great firewall of China". The "great firewall", however, is more about preventing its citizens from accessing content outside, so the enemy it's preventing from coming in is external influences.

Naturally enough, military rhetoric pervades discussion in terms of cyberspace or cyber war. There is an "arms race" between attackers and defenders as they seek to generate technologies or strategies which will defeat the other one. A cyber "attack" can occur, possibly having the effect of "digital armageddon", or "digital Pearl Harbour". The Stuxnet attack has even been described as "the Hiroshima of cyber war" (Gross, 2011).

Whilst this is a commonly used analogy, there is some scepticism as to the war analogy. Rid (2012) argues that "an act of war has to have the potential to be lethal; it has to be

instrumental; and it has to be political", and that none of the attacks so far encountered have enough to meet that criteria. Even the Stuxnet malware (Falliere et al., 2011) was simply an act of sabotage. DDoS attacks such as those against Estonia in 2007 (Traynor, 2007), whilst inconvenient, certainly do not amount to an act of force – even when (if it could ever be proved) they were authorised by a state actor, see e.g. general comments by Bruce Schneier (Schneier, 2013), most recently in relation to the recent Sony attack (Schneier, 2015) Finally, Libicki (2007) made a point worth noting : "there is no forced entry in cyberspace. If a destructive message gets into a system, it must be entirely across pathways that permit such a message to get through" (Cited by Betz and Stevens (2013)).

Another issue relating to the war analogy is that it is unlikely to be effective. Spear phishing campaigns, such as the one against RSA (Rivner, 2011) demonstrated that access to a person inside the company is possible, and drive-by downloads are let through the firewalls because Web browsing is necessary for most networks. It has also been shown to be possible to undermine the "Great Firewall of China" as well (Clayton et al., 2006). Defence in this manner is ineffective, but attack equally so. The "enemy" referred to within this research is generally non-state actors (and even if their attacks are state sanctioned, it is difficult if not impossible to attribute) and as such cannot be targeted in a conventional warlike manner. Even identification of these attackers is very difficult. Without the ability to adequately attack or defend, the war analogy does little to help us in this instance.

The other analogy presented by Betz and Stevens (2013) is a biological analogy. The element of the biological approach to be taken for discussion in this chapter is that of public health. The idea has a long history in computer science, starting with von Neumann's work on self-replicating autonoma, and then the concept of a computer "virus" in the 1980s (Bilar and Filiol, 2009; Murray, 1988). In this chapter, however, more influence will be taken from the more general regulatory approach to public health interventions rather than the "virus" analogy.

The reason public health is particularly appropriate, is that due to the balances required between the good of the individual and the good of society, there are rigorous procedures which are followed before the formation of any intervention (e.g. Haynes (1999)). Although restrictions of rights are less serious in the online equivalent, there are still trade-offs and consequences inherent in creating legislation to combat security. For example, it is frequently opined that a reduction in individual privacy is required in order to adequately defend. This is a decision which also has to be made in relation to public health – whilst consultations between a patient and their doctor are confidential, the state needs to know if someone has a communicable disease so as to plan national responses and reduce the probability of an epidemic.

The chapter here assumes that the state at least has a right, and probably a duty to enact laws and enforce restrictions upon rights in order to best preserve the population's health. This is not a universally adopted approach, but in this country the formation of the NHS in 1948 out of general taxation demonstrated some precedent for government intervention in health. Even Epstein, suspicious of state intervention recognised this, for matters such as sanitation and prevention of epidemics, agreeing that neither the market or tort law were appropriate to deal with these issues (Epstein, 2003b).

This chapter will begin by examining the role of the state in relation to public health. This will begin by considering the evolution of this role from the industrial revolution where sanitation became a key issue, and an early instance of central government taking responsibility for a public health issue. Current policy will be considered, and some of the issues in relation to the good of the individual against the good of the overall population will be considered – and how many rights the government could justifiably take in order to promote public health.

Having established this background for health, early ideas of the analogy will be discussed. In relation to network security, the majority of the focus has been on the application of epidemiological models in order to predict or minimise the spread of computer viruses and worms. These models will be discussed as an essential element of a public health strategy, but then an analysis of more recent research which considers some of the wider policy issues such as the notion of computer "hygiene" (Carlinet et al., 2008), and a consideration of public health as a public good.

Finally, a reimagined version of the analogy will be presented in the context of efficacy and rights, where the attacker is regarded as the pathogen, and the malware merely a symptom of it. An example is presented of a website being akin to an infected Cholera pump in the 19th century.

4.1 The Role of the State in Public Health

Public health has not always been regarded as a priority for the state. Previously, this role would have been confined to using coercive powers to limit the spread of contagious disease through quarantine or other similar methods. The modern role of the state in relation to public health can be traced back to the middle of the 19th century, where a series of legislative initiatives were introduced with the aim of improving sanitation. This would previously have been the responsibility of the local authorities, but with the Public Health Act of 1848 a centralised body was set up with the power to take measures to improve sanitation should the local population desire it – or if the mortality rate exceeded 23 per 1,000 (Wohl et al., 1983). Further reforms continued, with the Sanitation Act of 1866 being the first empowering the central government to take action (Wohl et al., 1983).

As of the mid twentieth century, the role of the state has gradually expanded, such that now health is a key part of government policy[1]. In 1945, the NHS was formed providing free healthcare paid for out of general taxation and has survived to this day taking a considerable proportion of taxation revenue. The government has also regularly taken issue with elements of society which negatively affect the health of the population such as cigarettes, alcohol and, more recently, obesity. This can be argued as being excessively paternalistic, whilst others argue that the government has a responsibility as a "steward" to promote good health (Nuffield Council on Bioethics, 2010).

Why might a government care? We could argue that in a similar manner to crime, an unhealthy population is in principle a **Bad Thing**. Like Ashworth and Horder (2013) comments:

> "the chief concern of criminal law is to prohibit behaviour that represents a serious wrong against an individual or against some fundamental social value or institution"

And then goes on to say...

> "the decision to make conduct into a crime implies that there is a public interest in ensuring that such conduct does not happen and that, when it does, there is the possibility of state punishment"

Early proponents of state interference in relation to public health, however, did so on utilitarian grounds. The argument was that the amount spent on rates towards the Poor Law could be significantly reduced if the poor were less likely to succumb to disease. Chadwick, one of the architects of the Poor Law, investigated the effect of ill health on the economy. In 1842 he published his *Report on the Sanitary Condition of the Labouring Population of Great Britain*, generating a collection of "lurid details and evocative descriptions, damning statistics and damaging examples into a masterpiece of protest literature" (Wohl et al., 1983). These findings drew attention to the inadequacy of current sewage and sanitary systems – and their connection to epidemic disease[2]

Since 1945, given that the government has committed to spending a not inconsiderable proportion of the national budget in maintaining the NHS, it is also plausible to argue that by restricting the rights of its subjects it is making an economic decision. Consider cigarettes and alcohol, which are heavily taxed because they have a high social cost. The government does not wish to have to foot the bill when the inevitable health costs

[1] In the UK at least. Other countries like the USA have a far more libertarian attitude towards healthcare provision

[2] And the tendency for immorality such conditions created, which was probably a lot worse than the death and suffering.

return to bite those indulging in destructive behaviours. On the other hand, research does suggest that the amount of revenue gained from taxing cigarettes is a considerable amount more than the direct cost to the NHS (compare HM Revenue & Customs (2007) and of the Royal College of Physicians et al. (2005), cited by Nuffield Council on Bioethics (2010). When considered like this, then it could just as easily be seen as a cynical ploy to gain more revenue from products with an inelastic demand curve.

Alongside the general principle, it does appear that public health decisions are likely to have some economic motive one way or the other. Aside from direct costs, a workforce which has a significant proportion of its population being unhealthy could have negative effects on the economy if they are unable to work. There is also the [expensive] danger of public disorder in the event of an epidemic, as a scared population seeks to target those suspected of carrying the infection (Brazier and Harris, 1996). These decisions also need to acknowledge that government does not act in a vacuum, they need to account for exactly what is a politically viable solution.

This comes into conflict with the prevailing attitude in the UK courts, which is that the liberty of the individual should be protected. This has been seen in regards to the development of medical case law, in cases of refusing treatment by patients. Whereas in *Leigh v Gladstone* it was ruled that forced feeding of prisoners was acceptable, *Re. T* held that a competent adult may not be treated against their will, even if that is to "protect them from themselves" (*Re T(Adult: Refusal of Medical Treatment)*, at 661), even if it will result in their death[3]. Whilst this principle might work fairly well in individual cases, when applied to the population there are inevitably cases where the autonomy ordinarily afforded to an individual come into conflict with the rights of the population as a whole.

A common starting point for considering the circumstances where infringing an individual's rights is from Mill's On Liberty: (Mill, 1999)[4]

> "As soon as any part of a person's conduct affects prejudicially the interests of others, society has a jurisdiction over it, and the question whether the general welfare will or will not be promoted by interfering with it, becomes open to discussion"

An example of when this might apply is in the case of an epidemic, (actual or preventative) where the state might feel that they need to act in a certain way to protect the population. Under the Public Health (Control of Diseases) Act 1984 and associate

[3] It is different in the cases of "incompetent" adults and children where strong degrees of paternalism still exist. See *Re. R* (1991) and *Re. W* (1992)

[4] Whilst Mill is commonly cited as a general assumption for liberal democracies, that does not necessarily mean that every intervention needs to be justified. Coggon (2012) argues that libertarians should also be forced to justify their position, given that the default is anarchism so any form of a state constitutes some form of intervention.

regulations Public Health (Infectious Diseases) Regulations 1988, the government has got a right to override almost all of an individual's rights, in the event that they have a "notifiable disease". The right to confidentiality is removed (ss. 10–11, 1984 Act), and those known to be suffering from such a disease commits a criminal offence if they expose others to that disease in a public place (s. 17, 1984 Act). Similarly, if they are suspected of having one of these diseases[5], they may be forcibly tested (ss. 35 – 36, 1984 Act) and a magistrate may order the detention of an individual to a designated hospital (ss. 37 – 38, 1984 Act).

An interesting point, is that HIV does not appear on this list, despite the high level of interest in its spread. The removal of confidentiality might have reduced the willingness for people to come forward and speak to doctors, and therefore allows the state to keep better control on the spread (Brazier and Harris, 1996). This demonstrates the discretion that the state can apply in order to best protect the population. An example where this is applied is in preventative vaccinations for certain diseases. In the United States, there is a set of laws indicating that a child cannot enrol in a public school without having done this (Gostin, 2010), although this could lead parents to postponing the vaccinations until the child is close to school age – leaving them exposed whilst they are most vulnerable (Moran et al., 2006), cited by (Nuffield Council on Bioethics, 2010). On the other hand, in the UK there is no compulsion for parents to vaccinate their children. Whilst this caused some issues with the media paranoia about side-effects of the MMR vaccine, there is not necessarily a difference in uptake of vaccinations within Europe despite different regimes[6].

The level of intervention justified by government is something which has been debated to a great degree in the bioethics literature (see e.g. references in Coggon (2012)), so little attention needs to be paid to it here, merely to describe that there are different ways which public health can be considered. The approach taken here is that of Gostin (2010), who argues that public health policy operates on several levels with the goal of a healthy population:[7]

1. *Interventions* designed to influence behaviour through education, incentives/disincentives, or deterrents;

[5]Including cholera, plague, relapsing fever and typhus; and the powers for the Secretary of State for Health to extend these have led to polio, meningitis, diphtheria, dysentery, most common childhood diseases, rabies and tuberculosis also being included (Brazier and Harris, 1996).

[6]See review in Nuffield Council on Bioethics (2010), observing that although in individual countries, e.g. France there is a difference in uptake between quasi mandatory and voluntary vaccinations, Sweden has the highest rate of uptake where vaccinations are voluntary. They also note the effectiveness of small incentives, but observe that too high of a payment could lead to people taking risks they are not comfortable with.

[7]For the sake of completeness, it should be noted that this view of public health is not universal. A competing idea is that expressed by academics like Epstein (2003b), who argue that the state should keep restrict its intervention for the "traditional" public health – i.e. using coercive powers only for containing epidemics.

TABLE 4.1: Intervention Ladder, taken from Nuffield Council on Bioethics (2010)

8	Eliminate choice
7	Restrict choice
6	Guide choices through disincentives
5	Guide choices through incentives
4	Guide choices through changing default policy
3	Enable choice, e.g. offer services
2	Provide information
1	Do nothing, or monitor

2. *Regulation* to change behaviour, e.g. safety standards;

3. *Altering environments* relating to town planning, licensing, or advertising.

It is inevitable that the majority of regulation will infringe upon an individual's liberty to some description, since it is backed by coercive state power. Without such restrictions – or at least the ability to restrict rights, then the situation would likely be anarchy rather than a state (Coggon, 2012). The Nuffield Council report on bioethics relating to public health described the role of the state as a "steward", with a duty to not only to prevent harm to others but also with a duty to promote health. In their report, they presented a useful conceptual model illustrating the sort of interventions a state could make, and their relative degrees of severity, reproduced in Table 4.1 (Nuffield Council on Bioethics, 2010).

In general, according to Mill's argument (which Nuffield Council on Bioethics (2010) agrees with, the state should aim for the least possible restriction on rights wherever possible. On the other hand, this can be necessary in certain circumstances (e.g. see discussion above). In addition, it has also been observed that doing nothing about a public health issue also implies a value judgment by the state (Nuffield Council on Bioethics, 2010), implying the position is that they do not view the issue as worthy of an intervention.

Deciding appropriate restrictions on rights is one thing, another issue is to decide whether it is worth performing an intervention in the first place. Pressure on resources means that decisions have to be made as to interventions most likely to make an impact. The NHS has consistently had real terms increases in budget since its creation in 1948, yet it is not operating on limitless funds, and so decisions have to be made as to where this budget goes. This is far from a trivial task, see e.g. Asthana et al. (2004), and in the event that a decision based on resource allocation is challenged in the courts, the courts are likely to come on the side of the NHS trust. The authority is *R v Cambridge Health Authority, ex. parte B (A Minor)*, which provides:

> "Difficult and agonising judgments have to be made as to how a limited budget is best allocated to the maximum advantage of the maximum number of patients. That is not a judgment which the court can make. In my judgment, it is not something that a health authority such as this Authority can be fairly criticised for not advancing before the court", per Bingham M.R.

In general, three steps could be taken to decide whether to choose one form of intervention over another:

1. Is the proposed intervention effective on an individual level? I.e. does it actually work at all?
2. Given the realistic state of affairs, how much of an effect would the intervention have?
3. Given the cost, and the level of efficacy (and possibly restrictions of freedom), is it worth running the intervention at all? (Haynes, 1999)

Whilst this is a good starting point for rational decision-making, it must be recognised that inevitably there will be other influences on decisions which are made. In particular, the media plays a significant role in informing/distorting public opinions on policy. For example, in relation to resource allocation, decisions to prevent access to "life-saving" cancer drugs can cause a considerable amount of reporting in the media, see e.g. (Cooper, 2015; Donnely and Walton, 2015; Hope, 2015). Similarly, the reporting of scares about the MMR vaccine led to a significant reduction in uptake (Begg et al., 1998), which has caused occasional measles outbreaks across the country (Godlee et al., 2011). Even without media bias, politicians have their own ideas and ideologies, and do not necessarily have a detailed understanding of the processes of healthcare (Alaszewski and Brown, 2011).

This section has provided a background of how the state will likely need to occasionally restrict rights in order to manage health emergencies, decide on the efficiency of appropriate policy, and also be pragmatic about other influences such as the media. Government has played an increasing role in relation to the nation's health since the 19th century, in particular following the formation of the NHS. These limitations and trade-offs are ones which can also be considered in the context of Web security. The next two sections will demonstrate the way that the public health analogy has been used in the past, followed by how it has expanded, and then what lessons we can take from health policy into creating sensible cybersecurity policy given all the inherent limitations.

4.2 The Benefits of Public Health Techniques

One of the major benefits to the use of public health as a strategy to decide on Web security policy is the scientific nature of the regulation. This means that decisions as to what regulation to introduce are are decided based (in theory) on whether there is a scientific case that such an intervention would make a significant difference[8]. For health, this is necessary because of the potentially significant adverse effects on either individuals or the population if it goes wrong. In assessing whether or not to engage in some form of intervention, the efficacy of the intervention needs to be assessed in both ideal conditions and more realistic real-world conditions. Having established the efficacy, it is also necessary to analyse whether it's worth doing from either the point of view of cost, or the restrictions upon rights which it might cause. One example of an implementation of public health policy is the choice of vaccinations for the NHS to allocate. Similarly, the ban on smoking in public places was also introduced only after there was evidence of secondary smoke causing cancer in non smokers.

One of the main ideas which allows scientifically based legislation to occur is the use of compartmental models of epidemiology. By abstracting away some of the details, the effect of an intervention on a whole population. For theses, a homogenous population is placed on a graph, each with the same likelihood of interacting with each other. Each of the individuals will have a particular state, usually restricted to Susceptible, Infected, Exposed or Recovered in various combinations. For example, for an illness like measles, the model would be described as SIR, since upon recovery one gains an immunity to the illness and can be described as recovered. The common cold by contrast, might be described as $SIRS$, since recovery could confer a temporary immunity to the individual cold, but the individual would then likely turn to a susceptible state for the next variant to come along. By using these concepts, the flow of the disease can be modelled across the population. One of the central contributions to this field was Kermack and McKendrick (1927).

4.2.1 Application of Compartmental Epidemiological Models

The advantages of epidemiological approaches like this were quickly noticed by security researchers as computer viruses began to spread across individual networks, and worms spread across the Internet. This section describes some of the research in this area.

Kraus, in a masters thesis considered self-replicating code as a form of life in the biological sense. He concluded that whilst it could reproduce and mutate it lacked the elements of metabolism to be classified as alive. A virus on the other hand represents a far simpler form of life, which needed to rely on the metabolism of other organisms to survive, and

[8] Whether this happens in practice, given the realities of politics is, unfortunately, far more difficult to demonstrate.

hence a similarity could be shown between these programs and a virus. Early malicious code was overwhelmingly of this "virus" type, for which Cohen's definition "a program that can 'infect' other programs by modifying them to include a possibly evolved copy of itself" will be used (Cohen, 1987). Cohen performed a series of experiments demonstrating the potential threat posed by viruses and their ability to spread through any network. He also noted the difficulty in stopping them, observing that many solutions were impractical (e.g. complete isolation of machines (i.e. quarantine)), or were imprecise and resource intensive to implement. These are problems which have still not been solved today!

The analogy was expanded to make use of epidemiological concepts to contain the growing threat of viruses. An early example was Murray (1988), who noted the similarities between the spread of viruses in the real world and computer viruses, and concluded that the field of epidemiology could offer some solutions to contain the spread. He observed:

> "community, population, carrier, portal of entry, vector, symptom, modes of transmission, extra-host survival, immunity, susceptibility, sub-clinical, indicator, effective transfer rate, quarantine, isolation, infection, medium and culture are all terms from epidemiology that are useful in understanding and fighting computer viruses".

Murray (1988) then analysed some of the defences to prevent epidemics such as isolation, quarantine, and immunisation, also suggesting that the system manager should be viewed as an epidemiologist and keep on top of these threats. Kephart and White (1991) were also influenced considerably by the macroscopic approaches of choosing individuals to cure from the point of view of the population rather than the individual, and the mathematical approaches to back that up. They used a standard SIS model on a directed graph, and identified the likelihood of the eventual extinction of the virus as $t \to \infty$.

Having been observed that there was a certain threshold for virus spread to become an epidemic, one of the possible solutions suggested by Kephart and White (1993) in their future work was a "kill switch", where each machine which recovered from the virus would inform its neighbours who would in turn scan themselves and inform their neighbours if they found trace of a virus. Targeted immunisation was also a concept which was considered, contrasting with the essentially random immunisation of anti-virus strategies (relying on the vigilance of the user). Wang et al. (2000) observed a significant improvement on both hierarchical and clustered topologies, and Newman et al. (2002) analysed address books and email propagation finding that targeting 10% of the vertices for immunisation would make any epidemic negligible, as opposed to random targeting having hardly any effect.

As worms started to become more prevalent, there was research analysing the spread across the Internet. Wang and Wang (2003) consider the effects of adding timing parameters (infection delay and user vigilance) to propagation models. They argue that the classic SIS and SIR models are too simplistic. Their parameter for infection delay means that nodes can be both infected and infectious – but the two are different, modelled universally as a constant. The vigilance is modelled with two parameters: the vigilance coefficient [0-1], and the vigilance period – the length of time for which this increased vigilance applies. The period is also universal and constant. They demonstrated the accuracy of their model through running a simulator built on top of Network Simulator. The vigilance property is something which is also considered by Kelley and Camp (2012) in a later paper. Their model had two sub-populations of vigilant and non-vigilant nodes, with nodes able to move between the two, for example in response to recovery from a compromise.

Staniford et al. (2002) analysed the propagation of Code Red I, Code Red II, and Nimda, and explored the prospects of improved scanning of worms so that they could successfully infect each susceptible host before defenders could respond. They considered hit-list scanning, where a prior list of potentially vulnerable machines is obtained to enable the spread from there; permutation scanning, or combinations of the two. They hypothesised a "flash" worm, which could infect all vulnerable servers within tens of seconds. They proposed the idea of a "Cyber CDC" which would enable co-operation and co-ordination in response to a worm outbreak, although more as a concept than as any firm ideas for implementation.

The properties of scale free networks and propagation are also something which has been considered. A heavy tailed distribution, with a distribution Pastor-Satorras and Vespignani (2001) $P(k)\ k^{-2-\gamma}$, with various values of γ implies that "each node has a statistically significant probability of having a very large number of connections compared to the average connectivity of the network (k)". Where $\gamma <= 1$ there is no threshold and any infection can eventually reach the whole network. Madar et al. (2004) also considered the difficulties in applying targeted immunisation due to the requirement of knowledge. Instead, they propose a solution which immunises random acquaintances of a node on a network. Therefore, the strategy requires only local knowledge, and appears to be reasonably effective.

Other papers have sought to review these ideas and offer suggestions to either make more generally applicable or realistic. Serazzi and Zanero (2004) performed a review of various propagation models, and presented their own which was validated on data from the Slammer worm. They dismissed the other models without expending much time in explaining why the assumptions were invalid. Their model extended the random constant spread (RCS) model, whilst adding for the difficulties in increasing spread between different ASes. Within a single AS, the propagation of the worm proceeds unhindered, however outside the propagation slows down.

Yang and Yang (2012) presented a generic model $SLBS$ to counter the flaws they perceive with previous work in modelling epidemics. The main change here is to replace the single I with two separate states: latent and break out to model the fact that a virus may infect a computer but not cause any adverse effects to other computers, which will at a later date "break out" and begin to attack others. Given that the model is designed to be generic, they choose not to include the R element from the model (another flaw they perceive). This makes sense in terms of a generic model, to simply model the level of infected computers but unfairly criticise other models which consider only the susceptibility to a single vulnerability e.g. Slammer. They do not present any validation of their model.

A lot of this work no longer applies, however. Scanning worms are no longer as effective as they previously were, due to improved firewall adoption, and possibly as a side effect of the use of NAT. The last major worm to spread around the Internet was Conficker, which was over five years ago. The current model for malware propagation (drive-by downloads) does not rely on the physical spread of malware between computers like the old worms and viruses did. This means that the compartmental epidemiological models have less to offer us than they previously did, because the fact that one machine on a network is infected does not mean that other machines will necessarily be infected[9]

Whilst machine to machine propagation is no longer the same threat as it was previously[10], there are still models being produced in order to test the efficacy of various Internet scale interventions. Edwards et al. analysed the effect of an intervention by a search engine, specifically to combat drive-by downloads. An interesting solution they considered was to "depreference" some results based on a function of the probability that the page was infected (Edwards et al., 2012). As ever, a trade-off was required between false positives and false negatives, but essentially it retained one of the major principles of epidemiology: minimising exposure. Although servers are different to clients, they are still nodes on the same graph, and hence epidemiological modelling can still be useful.

Similarly, many of these models are of limited application when considered in a global context. Whatever the effectiveness on individual networks, Conficker succeeded in spreading to millions of computers (after the MS08-067 vulnerability it exploited had already been patched). That said, the response to Conficker did demonstrate global co-operation in successfully rendering the C & C inoperative by successfully persuading over 100 registrars to pre-emptively shut down the domains that the bots would be contacting in the future (Conficker Working Group, 2011).

[9] Although note research by Shin et al., who found a similar distribution of infections of the Conficker botnet (a worm) and the MegaD and Srizbi botnets (predominantly propagating by drive-by) (Shin et al., 2011) ; and Moura et al. who conceived the concept of "bad neighbourhoods", blacklisting traffic from a network with a high amount of infected machines making it more likely that the traffic from other machines would be malicious (Moura et al., 2011)

[10] There are occasions where this can happen, for example DNSChanger was able to take over the router (FBI, 2011) and in effect compromise all victims on the network. In addition, there were a few reports about exploitation of the "Shellshock" vulnerability, (Goodin, 2015)

Since Conficker, however, the problem of lack of a highly susceptible population has already been solved to a large extent with increased use of patch Tuesday and automatically updating software. Microsoft led this, and other major vendors have followed suit such that the majority of vulnerabilities are already fixed before they are known about by attackers. As such, the majority of devices are patched, so the idea of an epidemic is not the concern, but rather the fact that there is continually a series of machines which are compromised. Epidemics might emerge as a problem again, however, since there has become an increased amount of devices on the Internet which cannot be adequately patched.

On the other hand, it is increasingly the case that devices publicly accessible on the Internet are constrained in their ability to be updated. When Microsoft stopped support for Windows XP, it still accounted for over 26% of the desktop market (netmarketshare.com, 2014). There are still vulnerabilities being found in XP, and Microsoft software running on it, so this presents difficulties. Similarly, a large amount of Android phones, and increasing adoption of the "Internet of Things" means that devices without any user interface are also online, and being used in attacks (Krebs, 2015a). Limited capacity of these machines is also problematic, since it complicates any updating procedures. Whilst an important issue and concern for the future, it is currently outside the scope of this research.

4.3 A Wider Interpretation of the Analogy

Whilst compartmental epidemiological models are one area which has seen considerable research, there are other elements of the analogy which have been comparatively neglected. These mostly relate to concepts which aid in the understanding of a particular part of Web security, and can be adapted to the varying types of threat.

Epidemiological approaches such as this have had mixed levels of success. On the one hand, their use across a single network has been successful in being able to control the spread of malware epidemics. On the other hand, many of the strategies have been designed from the point of view of a network administrator who might wish to limit the impact of a particular malware outbreak on their network. The difficulties of successfully implementing an international solution to prevent this propagation have been modelled by Moore et al.. That said, the use of posture checks, quarantine, isolation, and the other concepts discussed by Murray (1988) have been used to some degree of success on corporate networks. An example might be a NAC approach of requiring the installation of anti-virus software on a computer and the latest signatures in the anti-virus software, and refusing access to a network without them.

At a higher level than the mathematical models, it is possible to use the analogy to analyse concepts and approaches. One of these is the notion of risk factors, which can

be used to identify which devices may be at risk of infection. Having established these, resources can be expended accordingly to appropriate parts of the population.

Maier et al. (2011) analysed the effect of client side infection, by conducting analysis on a university network. They compared the difference between the browsing habits of the users with their inclinations to install the latest version of the software. They found that it was more likely to be "risky" browsing behaviour rather than lack of updates to software or anti-virus. The way that they defined risky was by visiting Web pages after the browser had already warned them. Yen et al. (2014) analyses the characteristics of malware encounters in a corporate network found that encounters were most common off-site, and that nearly half of Web based malware fell into the "business" or "travel" category, further supporting the fact that drive-by downloads are not confined to suspect websites. Other studies include Carlinet et al. (2008), who identified peer to peer activity and having a Windows operating system as risk factors; Canali et al. (2014) identified late night browsing, and longer time spent online.

On the server side, Vasek and Moore (2013) analysed the likelihood that a Web server would become infected, and found that contrary to popular belief websites which had the latest versions of CMS software (WordPress was the one they analysed) were more likely to be infected than older versions (Vasek and Moore, 2013). Although they were unable to determine cause and effect, they suggested that this was likely because attackers sought the most popular version of the software to attack hence pushing the numbers in this way. A possible reason that they were successful, despite the absence of known security flaws, was because of a seatbelt effect encouraging more risky behaviour. Alternatively, it could have been due to the CMS being generally secure but the plugins to the CMS not, and attackers seeking exploits which are likely to be effective against the widest range of victims.

Another initiative with public health characteristics is that of remediating bots once they have been infected. This will be discussed in more detail in Chapter 5 as a solution which ISPs can achieve, but the concept will briefly be discussed here. The idea of remediation is for an ISP to notify the end-user when their machine is compromised. Maurushat (2010) argued that rather than attempt legal actions against or takedowns against botnet operators, that a better solution would be remediation. Whilst this often doesn't happen (see Section 3.3.4), there are a few instances where it does.

COMCAST is one notable ISP who does offer a service where they inform customers when they detect malicious activity from their machine (Livingood, Jason and Mody, 2012). Clayton described a system which could enable remediation to be funded by government for a comparatively small fee. This involved using bulk purchasing of security software as a means of saving money, and encouraging users to sort it out themselves (Clayton, 2011). He argued that it would likely be a fairly small cost per person, analogous to something along the lines of putting fluoride in the water.

In Chapter 3, the difficulties of improving security due to its public good characteristics were discussed, which is to say it is both *non-rivalrous* and *non-excludable*. Another good which could be considered a public good is that of health, for which a couple of examples will now be presented. Firstly is sanitation, like in 4.1. Another example is that of vaccinations, which is a freerider issue, based on the public good of population health. The optimal strategy of the individual is to freeride off the herd immunity of the rest of the population rather than take the small risk of an adverse reaction.

Mulligan and Schneider (2011) came up with a concept of "public cybersecurity", which is an analogy of public health due to their shared public good characteristics. They argue that there is a need for action to improve cybersecurity, which they believe to be an issue for collective action based on politically agreed solutions. Having discussed the notion of "public cybersecurity", they go on to consider a series of examples of how the two things are similar. Like public health, they argue, any application for cybersecurity will have trade-offs, and limitations on an individual's ability to do something, and that the level of interference should be based upon the benefit to the population as a whole. It is suggested that a doctrine of public cybersecurity is any policy with the goals of"(i) producing cybersecurity and (ii) managing insecurity" (Mulligan and Schneider, 2011), and provide several examples of where they believe that public health strategies map onto similar cybersecurity policies.

Amongst these arguments is that the heterogeneous nature of the population allows humans to be well defended against biological attack, whereas computer networks are usually homogeneous in order to facilitate interoperability. Approaches of artificially introducing variety into applications is one possible solution, because it decreases the success rate of exploit attempts[11]

Similarly, a government needs to be able to collect information about health in order to best manage or prevent outbreaks of disease. This is done on various levels, from providing information to the public to make decisions which prevent the spread of illness, to requiring doctors to report incidences of certain illnesses. So too could monitoring network activity, or user device posture provide information to contain the threat of malicious software – yet the problem exists in monitoring network traffic. It is something which could be easily abused by government; but in some jurisdictions privacy or competition laws prevent ISPs from sharing data for a potentially more efficient response (citing Van Eeten and Bauer (2008)).

Patching is also regarded as a problem whose solution would depend on the reason that patches are insufficiently deployed. For example, if the failure to patch is due to bandwidth issues, possibly government could subsidise bandwidth costs for patching sites; or require legislation to prevent anti-piracy measures being bundled into patches with

[11] A homogeneous Internet ecosystem remains a problem, but it is not as bad as it previously was when home (Microsoft) PCs devices were in excess of 90% of the devices connected to the Internet. Now, there is a wide variety of devices connected from mobile phones up to and including light bulbs and toilets

security fixes (Mulligan and Schneider, 2011). Related to this, is the issue of software requiring an excess of patches in the first place. It is suggested that the feasibility of encouraging greater education of software developers, and incentives for adopting best practices. Whilst patching is the main solution advocated in this research, the effectiveness of improving the quality of code is doubtful. Section 5.2 makes the case that liability in such cases is unlikely to be the most effective, whilst the practicalities of developing software suggest that even with the best of intentions for developing quality software a greater level of education is unlikely to have a great effect – particularly considering the large amount of legacy software currently in production.

The concept of isolation or air gapping networks is considered in terms of blocking malicious packets at an ISP or national level, a result of which could be the separation of the Internet into several different states. Whatever the desirability of such a method (and it is written with a warning about potential concerns and note of the Great Firewall of China) it is a method which is unlikely to be effective. Such filters may be able to block worms, but are less viable of a solution in the current climate[12]. Even for denial of service attacks, the possibilities of false positives are significant, given that it is possible to use otherwise legitimate traffic to conduct a denial of service attack, so measures such as this would require serious consideration.

Finally, ISPs are considered as potential intermediaries, which could be used as a means of preventing malware spread or informing customers. These are issues which have been discussed in Chapter 3, and will be further discussed in Chapter 5, so the details will be omitted from this section. In summary however, the paper raises a lot of interesting issues, and offers a framework from which this research has developed. Similar contributions are analysed for the remainder of this section.

Governance online takes into account the fact that there are similar applications for a population based approach online as there are in the real world. This might include using information gathered from the modelling approaches described in the previous section, or merely following the general principles relevant to the economic phenomena which make security so hard.

Sales (2013) is another example of considering the public health model as a potential regulatory framework. Rather, a part of a regulatory framework alongside other potential elements such as environmental law. The difficulties in establishing Coasean bargaining due to transaction costs is noted, whilst also considering the difference between public health as an ex ante [preventative] solution as compared to medial practice as an ex post solution. The primary issue considered here is that of surveillance and monitoring, as a means of providing information as to the level of infections such that action can be taken. Whilst acknowledging the potential issues in regard of public health practices in terms of (US) constitutional law, it is suggested that a reduced focus on sensitive

[12] See discussion in Chapter 2

medical data could allow surveillance in such a strategy as a means of combatting cyber security.

At a level of few restrictions in relation to public health, there has been some thought that the model for public health "norms" could be used as a model for cyber security norms at either an individual or a national level (Hunker, 2010). He observes that the Cybercrime Convention is an adequate starting point for states, and that socially desirable norms such as hand washing prior to eating, or getting vaccinations could be carried over. It was also suggested that the surveillance policies adopted by the CDC could be transferred to cyber security. Charney (2012), at Microsoft, also considered these smaller ideas from an educational point of view, moving up to more systematic strategies up to and including forced isolation or quarantine.

His recommendation was to require every machine connecting to the Internet to be able to assert its "health" before being allowed to connect. In the event that it failed to pass the appropriate tests, then it should be quarantined (Charney, 2012). This is essentially the NAC solution discussed above scaled to the whole Internet, although the rationale is different. Rather than doing this as a means of preventing epidemics, the argument was that these machines were what enabled attackers to send spam, and continue successfully to attack other machines.

This faces a few significant problems, not least of which is the issue of deciding what the criteria are for allowing a machine to connect. Whilst this is possible to do in a corporate environment, where the owner of the network will have some level of control over the machine to connect, this is not the case when scaled to the general public. Another issue which has to be considered is the fact that it is not possible to trust a machine which says it is clean. A rootkit could fool an operating system into thinking that it did have an updated anti-virus running, for example, or the latest version of different software. Consequently, the machine would report its state as being clean when it wasn't. Although it is possible to get around this using cryptographic techniques, the "trusted computing" concept is widely mistrusted and might be difficult to gain acceptance (Anderson). This is something which may be more suitable for corporate networks where they can adequately lock down their employees' machines.

Rattray et al. (2010) also analysed the Internet through a public health model, regarding cyberspace as a commons much like the sea and space and with no overall government control. They suggest that "security in the cyber commons is often self-provided by users rather than by a central authority", and that current models of governance are "a messy amalgamation of ... agreements – all of which are sufficient for cyberspace to flourish but insufficient to make it safe". They analysed some of the similarities between various techniques which were used in public health, which also had application in cybersecurity.

This was done in the context of maintaining the power of the United States (Rattray et al., 2010). Whilst this is not the goal of this thesis, they did present a good argument

as to why an individual country (the United States) should take the lead on cybersecurity issues despite the freerider issues. By cleaning their own space, they can hold countries accountable for insecure systems in theirs, and begin to achieve norms which the international community can act upon. A cleaner cyberspace also affords fewer opportunities for malicious actors to take advantage of the insecure systems and attack them. The UK government has also expressed a desire to make the UK "one of the most secure places in the world to do business in cyberspace" (Cabinet Office, 2011)

4.4 Reimagining the Analogy

So far within this chapter, two different elements of the analogy have been emphasised. Firstly, the tools used to track disease being used to great effect with prediction and response in relation to virus and worm outbreaks were discussed. This was followed by ways in which these could be used to formulate policy. But the use of analogy can bring with it a certain amount of flexibility, there are different elements of it which can be focused on for different purpose. The common element which the majority of the literature has focused on has been the malware – the virus, worm, Trojan propagated in certain ways. This has been viewed as the disease, and that adopting various approaches can confer an immunity (temporary or permanent); or a recovery from the disease. Even approaches viewing the Internet as a commons view a scenario whereby "herd immunity" could be achieved from enough clean machines.

An analogy is a linguistic device which enables us to understand concepts which we might otherwise not be able to understand (Betz and Stevens, 2013). In this case, the design was to solve a problem: that of malware spreading through a network – and how best to stop it. In a strictly computational sense, this is undoubtedly the most useful metaphor which can be applied, and the models which are used can be effective in limiting the damage caused by it. However, when applied to the problem such as the public good as we have in dealing with the effects of botnets, another analogy is more appropriate. Lawyers deal in analogies regularly, particularly in "hard" cases, where authorities are not conclusive in favour of either party (Weinreb, 2005; Vandevelde, 1996a).

It is proposed that it is the criminals themselves that should be regarded as the disease, and the malware merely as a symptom of a larger issue. As such, we should view the availability of unpatched computers as a characteristic of the environment which enables the parasite to thrive. Malware propagation, and infected computers are merely symptoms of their existence. By doing this, the methods of attempting to control malware on the Web takes on a different approach – that of making the environment hostile for a criminal, much like something which could be attempted for a biological parasite. This represents a different approach to the majority of the literature (see Section 4.2), so the

remainder of the chapter will be devoted to explaining what improvements viewing the analogy in this light can offer.

It should be made clear that these analogies are not competing with each other – it is intended that they supplement each other. They are designed for a different purpose, the compartmental epidemiology discussed in so much detail in the literature (Section 4.2 is designed to solve a particular problem – that of virus spread through a network. The way the analogy is presented here is designed to be used by government in order to guide legislation in a manner consistent with the public health principles described in Section 4.1, namely a focus on efficacy and rights.

Analogies by their nature share characteristics with the thing they are analogised from, but the important element is the relevance of the characteristic to the underlying policy requirement, and why characteristics of the alternative analogy are less like it (Vandevelde, 1996b). To make the case for an environmental version of the public health analogy, an example of an environmental problem will be considered – the unsanitary conditions of the 19th century which allowed the Cholera bacteria to thrive.

Here the similarities between a website and a water pump are considered as a means of demonstrating this analogy. In particular, we will consider how weaknesses in them can be used to cause injury to others. To begin, both can be compromised by a flaw in the underlying infrastructure. A water pump is an object whose content is also vulnerable to compromise as a result of infrastructure – the overall sanitation system of the town it is located in. In the 19th century, the issue was that sewage would overflow, and end up mixed with the drinking water. Through this manner, bacteria entered the drinking water and caused illness to those who drank the water. Section 2.2.1 discussed in great detail some of the common attacks against websites. Those weaknesses can be used to exploit a website, such that it can be used to host drive-by downloads see (Provos et al., 2007, 2008; Grier et al., 2012; Caballero et al., 2011). Having done this, then the users are attacked.

Another similarity is in the nature of their use: both require their users to come to them in order to provide something. A water pump relies upon a person to come to it with a bucket, pull the pump handle such that the mechanism for releasing water works and it goes into their bucket. They then go back to their home to use the water they have collected as they wish. Similarly, a website requires a connection from a client, which then sends a request along the network which will then communicate with the server. The server will then send content back along the network, and the client will be able to make use of it. In both cases, if they are compromised, then alongside the content they return to the user/client, then something else is returned alongside what they wanted – the Cholera bacteria or exploit code to attack the machine.

These two examples demonstrates a similarity between the nature of the two different "things". Whilst not directly relevant in this instance, the context also provides an

example of how the similarities between two different times and legislative contexts could also inform us (Vandevelde, 1996b).

In the nineteenth century the industrial revolution was taking place, as Britain changed from being a largely agrarian society to one dominated by industrial cities. This had a considerable effect on the population, such as the public health issues described in this chapter, and the new technology which was on offer. The courts arguably effected a grand compromise in relation to the law of private nuisance, (e.g. in *St Helens Smelting* (1865)) such that industry could carry on even whilst causing discomfort to their neighbours (Brenner, 1974). The increased production and opportunities for employment meant that the population were largely happy with it (McLaren, 1983), and hence the side effects such as pollution were not considered in any great detail.

This is a similar situation to the one we are in now. A recent disruptive technology (around twenty years old) has improved the ability to communicate, changed employment practices, and had a big effect on how people live their lives. It has also taken time for regulation of this new technology to take effect, Rattray et al. (2010) described the Web as being akin to the wild west, and observed that the there was regulation "enough to thrive but not enough to be safe". Given the improvements to the state of sanitation, and the general state of peoples' health over the last century, it is hopeful that the Web can also gain considerable improvements in security.

The literature discussed in this chapter has illustrated, there is already an established analogy between public health and security. Despite the complementary nature of this analogy, why is it more useful to government than the established one? To begin with, considering the criminal as a pathogen as opposed to the malware offers a far greater level of flexibility. Considering infections as the main enemy is likely to prove a losing battle, as the amount of "strains" of malware have increased steadily over the last 10 years. It fails to take account of the different motivations and capabilities of the adversary, and the different strategies which might be required to overcome them. For example, if the adversary is not motivated by money (and so falls outside the scope of this research), then a small increase in difficulty is not necessarily going to dissuade them from attacking. This was something alluded to by Charney (2012), although he kept the focus on the malware rather than the attackers.

It is also worth noting that whilst compartmental epidemiological models can still be used, they have lost a lot of their usefulness since the viral like transmission now generally doesn't apply. They can still be used[13] by considering them as interacting subpopulations where members of the same population do not interact with each other.

For a government, it is more appropriate to take a step back and consider the wider picture, with as much context as possible. A main difference, is that there is a lot less

[13] Indeed, see Chapter 6 where these techniques are applied

control over "immunisation" policies for individual machines compared to a corporate environment where it is possible to use technologies such as NAC. If they are to make an intervention, they need to account for more a greater amount of variables, notably the economic impact a decision might have on the stakeholder they choose to target for it. In addition, governments are (or at least, they should be) far more reluctant to interfere in the privacy or liberty of an individual. These are not necessarily unjustified but, once again, the government must be held accountable to the population and consider these matters in an open and scientific way.

Unlike analogical reasoning in regard to case law, or statutory interpretation where it might often occur, there is no individual statute to draw upon to determine the policy objectives of the government and to see whether this analogy could be useful in supporting the policy. Therefore, other sources are necessary to ascertain what it is.

The UK government takes security seriously, recently claiming "The security of our nation is the first duty of government. It is the foundation of our freedom and our prosperity." (HM Government, 2010). This is not unexpected, as described above, security being a public good means that it is best that government provide it. Similarly, in the same document, they emphasise the importance of the Internet as being inextricably bound to the "real" world: "Cyberspace is already woven in to the fabric of our society. It is integral to our economy and our security and access to the internet, the largest component of cyberspace, is already viewed by many as the 'fourth utility', a right rather than a privilege." (HM Government, 2010). Although government protection within cyberspace has been limited, it is now recognised as a "tier one threat", on a par with terrorism; natural disasters; and an international military crisis (HM Government, 2010).

Following this observation, the government released another document detailing their strategy which responds to this "tier one threat". Within it there are four primary objectives:

- The UK to tackle cybercrime and be one of the most secure places in the world to do business in cyberspace;
- The UK to be more resilient to cyber-attacks and better able to protect our interests in cyberspace;
- The UK to have helped shape an open, stable and vibrant cyberspace which the UK public can use safely and that supports open societies;
- The UK to have cross-cutting knowledge, skills and capability it needs to underpin all our cyber security objectives (Cabinet Office, 2011).

These four objectives can be considered in terms of the reimagined public health analogy. The first objective and the third objective can immediately be regarded as consistent

with adopting the reimagined analogy. By making the UK one of the most secure places to do business, this directly impacts the ability of criminals to make money owing to increased effort required. As a result, the environment is less conducive to cybercrime, and better for the public. The UK is a developed country, meaning that people living here are a greater target to cybercriminals owing to greater potential rewards. By the UK having improved its environment, it has the potential to impact criminals globally. Getting over the freerider issue related to the public good of cybersecurity, the UK will have led the way in shaping this form of cybercrime. Objective two can also be supported by considering public health, a bad environment is good for criminal attackers as well as state attackers. Improving the environment will serve to increase the resilience. The previous three objectives inevitably lead on to the fourth objective, which is to have the knowledge, skills and capability to achieve all this.

High level objectives are good, but recent government initiatives demonstrate similar principles. The "Cyber Essentials"[14] is one such initiative, where the government has introduced a series of simple measures to protect data for which organisations can be certified. This is now compulsory to be able to bid for government contracts. Education into cyber "hygiene", and personalised action plans offered to organisations to help them improve their security posture.

4.5 Limitations of the Analogy

Although the analogy has many points which provides us with an excellent framework for regulatory strategy, there are some elements which either do not quite fit, or may present slightly different policy issues.

The most obvious issue is that, for the most part (although not exclusively), malicious software does not threaten anyone's life. Notable exceptions to when infection or other malicious activity might have a life or death impact could be the reported example of Conficker infecting computers in hospitals in Sheffield (Williams, 2009), or the (unconfirmed) suicides associated with the Ashley Madison leak (Baraniuk, 2015). By contrast, the notifiable illnesses enabling mandatory treatment mentioned in Section 4.1, and those where there are vaccination programs are ones with the potential for high morbidity or mortality (Department of Health, 2006).

Considering the comparatively small effect, it could be argued that there is less of a case for the state to intervene, since there is no direct harm to others. The effect of not having access to the Internet could cause significant difficulties to many, given society's current reliance on them. The cost savings from moving government services to being online only have led to the government considering such a migration (Asthana and McVeigh, 2010), which could make the matter worse.

[14] https://www.cyberstreetwise.com/cyberessentials/

Another limitation is that in public health, a government has got control over all entities within its territory. It can introduce regulation, or perform other interventions which could conceivably have a significant impact on any particular public health problem. As previously described, central government began compelling local authorities to improve sanitation in cities, and as a result the infrastructure gradually improved. This becomes more difficult in relation to Internet infrastructure, because many of the offending servers are outside the jurisdiction of the individual state[15]. The international dimension of the Internet also introduces a potential freerider situation for government, as the incentive is to simply freeride off the benefits of other states doing work to improve the overall level of online security[16].

While monitoring and surveillance are an important element of public health, the potential effects of using these techniques should be seriously considered, as there is the potential for abuse by government. Although quarantine measures have been subject to abuse (Gensini et al., 2004), knowledge of an individual's status is not something that a government can readily abuse. By contrast, increased surveillance and monitoring of an individual's Internet communications could equally be used to ascertain their political opinions (and persecute them), or to prevent access to certain information. Whilst a case is made for being able to access this information from a law and order point of view, this is something which should be done transparently rather than under the ambit of a "public health" strategy. That said, this opinion is not universally held, it is suggested that the less sensitive nature of information about computer networks compared to public health information would make it more acceptable from a civil liberties point of view since the restrictions on rights would be less severe (Mulligan and Schneider, 2011; Sales, 2013). As Mulligan and Schneider (2011) comment, there is likely a level where society can determine what is an acceptable trade-off.

Another consideration which needs to be made, is that the institutions which we would rely on to perform public health style interventions are not government controlled. Sales (2013) observes that it is possible for a country to use "individual nodes within a far flung network" in order to assist with this – a role which is not so easy for cyber security. When considering the previous point, this is a good thing, however assuming that there were adequate safeguards to prevent abuse, additional complications would arise. Part of the role of government is to ensure a healthy population, whereas this is not a part of the role of Internet intermediaries. Their role is to provide a service, and to make a profit from doing so. By contrast, (in the UK at least), there are institutions whose primary purpose is to improve health, such as the NHS or the CDC in the USA. Where other public institutions such as schools are required to perform interventions, it is something which a government can provide additional resources for, or legitimately include as part of a wider policy. In the USA, for example, there are laws preventing enrolment of

[15] In the next chapter, this effect will be analysed through simulation

[16] For more details about public goods and the freerider problem, see Section 3.1.1

children into public schools until they can demonstrate that they have had vaccinations (Gostin, 2010).

4.6 Conclusions

This chapter considered the use of analogy as a potential method to guide the development of regulation to combat the drive-by downloads, given that market forces and private law are inadequate due to the inefficiencies discussed in chapter 3. Although the analogy of war is popular, where the defenders seek to protect a "space" in cyberspace, the public health analogy was regarded as being the preferred option due to its flexibility and scientific methods. The application of compartmental models was discussed, and regarded as a useful method in both the real world and with malware – as evidenced by the large amount of literature.

It was suggested, however, that the slightly sparser literature about the application of concepts from public health policy was something which should be looked at alongside the practical application. An extension to the existing literature was proposed, where the attacker was regarded as a pathogen rather than the malware itself. A different focus – that of efficacy and rights was proposed, within this context. No analogy is perfect, and the limitations discussed here reflect this. That does not, however, stop the idea of using public health as a general framework for making decisions about Web security, and the similarities between the two are more similar than they are not. The reliance on the Web and the Internet means that "Cyber-security is too important, and too intricate, to leave to the criminal law and the law of armed conflict" (Sales, 2013), and the public health analogy as described offers a useful alternative narrative.

In the next chapter, the stakeholders introduced in Chapter 3 will be considered in terms of efficacy and rights, in terms of how they can reduce drive-by downloads. Chapter 6 will make use of an agent based simulation of drive-by downloads to assess the efficacy of the intervention.

Chapter 5

The Role of Intermediaries

Previous chapters have discussed the stakeholders involved, and their motivations for any action relating to security, and concluded that reliance on market forces alone was insufficient. Chapter 4 then considered the public health analogy as a means of guiding the formation of regulation which could be used to supplement a private law and market based solution.

This chapter will analyse an application of the public health analogy for drive-by downloads. The previous chapter described the focus of the analogy for this research in terms of efficacy and rights, as opposed to specific public health style interventions such as quarantine or forced treatment. That does not preclude their use – indeed, the proposal advocates a blocking procedure – but merely considers that supplemental to the general approach. Whilst efficacy is mostly considered in terms of "could" it work, when considering possible scenarios to apply the efficacy to, consideration must be made as to whether it would, and also to the cost. The analogy as described refers to the criminal as a pathogen in an environment which they can thrive – that with a large amount of unpatched computers. As such, it follows that the best solution is to implement an appropriate patching regime, and this is the idea to be explored in this chapter.

On the client side, modern software has increasingly started to automatically upgrade in order to fix vulnerabilities. Microsoft's "patch Tuesday" provides monthly updates to Windows and other software on the second Tuesday of every month (Microsoft), which Adobe aims to sync their updates with (Dignan, 2012), and Oracle have quarterly update cycles for their products (Oracle, 2015). Web browsers also update automatically, with Microsoft Internet Explorer being updated as part of patch Tuesday, while Chrome and Firefox also automatically update by default. This means that vulnerabilities do not exist for very long, suggesting that people who are falling victim to drive-by-downloads are using very old versions of software before the automatic updates were introduced, or

have consciously decided to prevent the updates from running (e.g. if they are running a pirated version of the software)[1].

Mention needs to be made again to Vasek and Moore (2013), who found that the conventional opinion that older web servers are more likely to be compromised is not supported by evidence from their study. They found that instances of the latest versions of CMSes were over-represented in the lists of compromised websites compared to older versions. This is a surprising finding, which suggests that something outside of patching is causing a problem. It could be that engaging in risky behaviour was more noticeable, or that they used a service such as `wordpress.com`, and then forgot all about their website. Either way, this is an issue outside of the model. The attacker will therefore presumably be running the same sorts of attacks against all the websites they can find (following the model described by Herley (2010)), and so the principle remains the same.

Having considered the benefits that the public health analogy can do for assessing effectiveness of interventions, or for deciding on regulatory strategies, specific public health techniques will be considered in the discussion related to different intermediaries. The three main stakeholders we wish to consider are hosting providers; access providers; and search engines. They can all work to achieve the same end, but do so by virtue of subtly different public health strategies. Hosting providers are looking largely at a vaccination strategy, whereby the software on the Web servers is patched prior to infection. The action an access provider would be able to do would consist of quarantine or isolation. This is largely in the case of their end users, although it would also be possible to isolate certain websites from their customers would also be a sort of quarantine[2] Finally, search engines do not have the ability to restrict access either to websites or end users. As such, they are offering a more informational element, such as in warning, or by making it more difficult to locate the websites.

The remainder of the chapter will be structured as follows: Firstly, following from Section 3.3.7 the case for hosting provider liability will be made. Following this, other stakeholders will be analysed, and it will be shown that they fail offer an alternative which is as good as this. Firstly, the stakeholders who should be excluded from analysis will be discussed, including website operators; end users; and software vendors. The remaining stakeholders will be considered: the ISPs (access providers); hosting providers; and search engines as viable options to implement interventions. These will be considered in terms of fundamental rights, and efficacy of the intervention.

[1] Occasionally, after an easily exploitable zero-day has been discovered and made publicly known, there will be a period of time where everyone is at risk. Mitigations are usually published in these instances, and, whilst it is rare that many will follow these mitigations, the situation is also comparatively rare (Bilge and Dumitras, 2012)

[2] However, this would not be strictly correct, because they do not have the ability to isolate infected or potentially infected websites from the general population, only their own population. The one to be discussed in greatest detail will be the proposal to quarantine users.

European legislation relating to intermediary immunities and will be analysed to determine the potential obstacles to the introduction of legislation. A provision in the USA, §230 of the Communications Decency Act will also be discussed, because its Good Samaritan protection offers a potential safety net for operators acting in good faith – and such a provision will likely be necessary should any obligations for security be introduced.

5.1 The Case for Hosting Provider Intervention

Section 3.3.7 discussed the steps a "reasonable" hosting provider might take, in order to ensure that the website remains free of infection. They were defined as follows:

1. Outdated/vulnerable software should be identified and patched as soon as possible after a fix is available (whether by informing users or conducting the update themselves).
2. Vulnerabilities in the OWASP Top 10 be identified, and the customers informed and be required to fix as soon as possible. The current Top 10 can be seen in Table 3.3.7[3].

This section will outline how this proposal for acceptable behaviour could be set out in the form of regulation. The regulation would have two primary features:

Firstly, it would establish that hosting providers would have a duty of care for any visitors to websites they host, and a responsibility for any losses obtained as a result of drive-by downloads on websites they host. However, performing the steps above and being able to demonstrate that they had performed these steps, would be regarded as evidence of acting as befits a "reasonable" hosting provider. This would lead them to be immune from any tortious claims relating to the security of the websites they host.

As the discussion in Chapter 3 has indicated, it is likely to be rare that a victim would lose enough money to make a lawsuit worthwhile. This could potentially have the effect of preventing precedent from being established, as well as arguably removing the incentive for an operator to take due care. Consequently, the second element of this regulation would allow for the regulator to inspect the procedures adopted by the hosting provider, and to issue fines should they be failing to perform to an adequate standard. This is something which could be unenforced if norms lead to a general increase in security.

Holding hosting providers responsible is supported by conventional liability theory, where tort law is viewed as being a means of achieving a certain end. There are many competing

[3]There is more detail about each vulnerability on the OWASP website https://www.owasp.org/index.php/Top_10_2013-Top_10.

theories as to exactly what this end is, and indeed the best way to go about it (see Goldberg (2002) for a detailed discussion). Despite these different theories as to the function of tort law, one requirement is that it be effective. In order for that to happen, there needs to be someone with deep pockets who can internalise the loss, possibly through negligible price increases for their customers. Similarly, if they are not in a position to make a difference, then there is little point – no matter how blameworthy (Varian, 2000)). In addition, it needs to be possible to identify them with the minimum amount of transaction costs.

These are arguments used in favour of vicarious liability, despite judicial acknowledgement that it essentially conforms to "social convenience and rough justice" (*Shatwell* (1965), per Lord Pearce at 685). Atiyah argued that it "it is simply the most convenient and efficient way of ensuring that persons injured in the course of business enterprises do not go uncompensated" (Atiyah, 1967). The exact same thing applies in relation to security decisions (see Lichtman and Posner (2006)).

In addition to being the optimally placed in terms of liability, hosting providers are also the best choice from the point of view of the public health analogy as well.

From the perspective of a public health technique, the role of a hosting provider would be akin to a vaccination program, where susceptible servers are found prior to becoming infected so as to prevent an epidemic from emerging in the first place. This has been shown to have been effective, in particular the eradication of smallpox as well as a significant reduction in instances of measles in the 20th century. The concept of state involvement in vaccination programs was discussed in more detail in Section 4.1

5.1.1 Efficacy

The first point to make about the efficacy of scanning websites for vulnerabilities, is that it will probably not catch all of the vulnerabilities which exist on that site. Some vulnerabilities will require very specific input to exploit it, but reduction in vulnerabilities will naturally make the job of an attacker more difficult and therefore reduce the requirement for ISP or search engine monitoring. Consider in economic terms the plight of an attacker in choosing which websites to attack. He will seek to maximise his profit, and will hence seek to gain the maximum amount of traffic for his effort. At present, given that vulnerabilities are known, and it is possible to fingerprint websites to see which software they are running, he can run his own automated scan and any vulnerabilities will be found. The homogeneity of the Web and software industry mean that existing vulnerabilities will work for potentially thousands of websites, therefore minimising the work the attacker has to do.

In the event that all known vulnerabilities are patched, the attacker is left with two choices. Firstly, he can visit websites manually and attempt to find the vulnerabilities

which automatic scanners miss and lose out on the advantages of automation which computers offer. Secondly, he can attempt to find new vulnerabilities (zero-days) on his own and use an automated scanner to find sites vulnerable to these vulnerabilities. These vulnerabilities take considerable resources to find, and their value decreases rapidly, as soon as they are found out about. The model discussed earlier from Herley (2010) suggests that this would eliminate the viability of drive-by downloads to all but a subset of attackers who choose to use targeted attacks.

Hosting providers are in a similar position to search engines, except that they are at an advantage due to having more information about the websites in question. They can see more files, and details about them, and hence make better assessments. In addition, being able to observe traffic coming in and out which might characterise an attack can also help with this determination. As discussed in Section 3.3.6, Canali et al. (2013a) show that they are not necessarily doing as much as they could to prevent these attacks. As such, an extra obligation could also have a greater efficacy than search engines that already have an incentive to minimise compromised sites in their listings.

5.1.2 Rights

The monitoring of hosted websites to detect malicious webpages should be easily justified for data protection purposes: in such a case the hosting provider would pursue a legitimate interest as per Article 7 of the Data Protection Directive. What could be slightly more problematic would be the decision to block access to an allegedly malicious website if the website operator does not implement the necessary security measures. Automated individual decisions under Article 15 of the Data Protection Directive would need to be explained to users who should be given an opportunity to oppose the decision. What is more the blocking of a website would also have repercussions for the exercise of the right to freedom of expression and the right to conduct one's business. This being said, because hosting providers are private actors and are not providers of publicly accessible electronic communications services they should enjoy more freedom in terms of private ordering through the means of their terms of use. Besides, from the perspective of freedom of expression, restricting the activity of content providers might be more acceptable than restricting the activity of content consumers, those willing to access content.

From a privacy point of view, actions taken by a hosting provider has an advantage over access providers, because it is possible to completely separate the code, or malicious content from any personally identifiable data. The identifiable data is stored in a database, and it is the code which grants access to that data. However, there is no requirement to view any of this data in order to identify a vulnerability. Even obtaining evidence of

an SQL injection vulnerability (i.e., connected to the database) can be done through an arbitrary piece of text rather than the data itself[4]

Similarly, "fingerprinting" an outdated version of software can be done by examining a small amount of files which, again, have no connection to the user data. By identifying these files, one can perform a cryptographic hash on their content, and use the combinations of these files to find the version. A cryptographic hash, such as md5 or sha1, transforms one piece of content into another with the additional properties that:

1. Different pieces of content will never have the same hash;
2. It is not possible to guess the original content from the hashed content[5].

The first property is the one which can be exploited here, since it provides a quick method of identifying the file. For example, on the majority of WordPress installations, the `/readme.html` will remain on the server, which changes regularly as new versions are released. Rather than having to parse through the text for identifying features, by performing a hash on that file it will be unique every time. The current version (4.1.1 at time of writing) has a SHA-1 hash of `8eeaf68a9e75e1dad72ee1eaa42ebee970851602`. Even if this file is not on the server, there are many files which will change from version to version allowing a version to be "fingerprinted". Tools such as Blind Elephant[6] already exist to do this, and a similar tool was also developed by the author during his second year, details of which can be seen in Appendix A.

There are some drawbacks which would need to be carefully considered, however. Firstly, there is a considerable danger that rushing to patch websites with vulnerabilities could lead to changes which are difficult to implement as they might break existing functionality. For example, a WordPress website could make use of a plugin which gets abandoned by its developers, for which a vulnerability is subsequently discovered. There is no obvious way to upgrade it, and site functionality could be dependent on the particular plugin. In addition, the updated versions could themselves break functionality on the website. This is arguably more critical than a client side upgrade going wrong[7], because businesses can rely upon their online presence. On the other hand, this is a problem which a business would have to fix in any case, so is a risk which they should already have to manage.

[4] For example, using `' UNION SELECT 'foo', 'bar' FROM Users--` could demonstrate that the vulnerability existed by placing `foo` and `bar` at certain parts of the document, and not touch any user data.

[5] A common use of the second property makes hashing a useful means of storing passwords, since it adds an extra layer of security in the case of a breach. For example, performing a sha1 transformation on the phrases `password1` and `Password1` yields the completely different hashes `e38ad214943daad1d64c102faec29de4afe9da3d` and `70ccd9007338d6d81dd3b6271621b9cf9a97ea00`.

[6] `http://blindelephant.sourceforge.net` The software also supports doing this for the plugins of websites which can also be used for exploitation.

[7] Such as happened in March 2015 for the regular Microsoft update, (Krebs, 2015b).

There is also a danger that a system like this could cause complacency with the website operators and actually increasing the danger of compromise. A similar phenomenon has been observed on the client side, with Christin et al finding a correlation between anti-virus installation and malware infections (Christin et al., 2012). Similarly, WhiteHat Security hypothesised that complacency was a reason for a higher level of vulnerability in websites which used static source code analysis tools compared to those who did not (WhiteHat Security, 2013). This could also be a possible explanation for the phenomenon observed by Vasek and Moore (2013) where latest versions are more likely to be compromised. Given that there are no known vulnerabilities, then there must be an additional explanation for what enabled the compromise to occur, and complacency having performed some security task could conceivably be a reason. Similar effects have reportedly been found in relation to seatbelt legislation (Peltzman, 1975), known as risk compensation theory. It is contended, however, that the risks of this are exaggerated (Cohen and Einav, 2003).

5.2 Stakeholders Excluded from Analysis

The three stakeholders to be excluded from analysis are: end-users; website operators; and software vendors. At first glance, it might appear that these are well placed to make an impact, since it is the software vendors who create the software and the end-users and website operators who install it. However, when considered on a global scale, these fail largely on efficacy grounds, and, given the lack of efficacy the potential restrictions on rights cannot be justified.

The arguments against both end users and website operators are similar in the main respect: there are too many of them for an individual to make an impact[8] Netcraft (2015) believes there are 178 million active websites at time of writing, and whilst exact figures are difficult to come by, Web users likely number in the billions. The majority of users are vigilant with respect to patches by default – unless they actively choose not to be – due to the majority of popular software being set to automatically update. The percentage of websites which are vigilant is smaller, but website operators likely want to avoid having their site compromised as long as it does not require too much effort.

The remaining population are non-vigilant in both users and websites. Whilst a requirement on them to improve vigilance might have a small effect, this is the part of the population who are less likely to comply in the first place! These factors added together make the chance of an intervention like this having an impact very small indeed. This is the other side of the liability theory ideas discussed in the previous section: they

[8]This is not strictly true for some websites, where those are amongst the most popular in the world. However, as Chapter 6 will demonstrate, the heavy tailed distribution of the Web makes this number very small indeed.

are not the cheapest cost avoider, and it would be more efficient to place responsibility somewhere else.

The other side of this is that the developers of software are generally doing a good job, and are likely to have to rely on the non-vigilant users to update any fixes which might be required. As Table 3.2, and associated discussion in Section 3.3.2 demonstrated, this is not being carried out, and older exploits were still successfully being used to attack users. Actively maintained CMSs like WordPress or Drupal are periodically updated to fix bugs or vulnerabilities, but the high volume of websites which use them (WordPress is said to run on around 23% of the websites on the Web (W3techs.com, 2015)) means that an attack against one website will work against many websites, making them an attractive target for attackers. It does appear to be a significant problem, since many websites continue to use out of date CMS software. A recent example of this is the large amount of Drupal websites which were successfully hacked, following the disclosure of a significant SQL injection flaw in the core of the CMS (Hess and Rudge, 2014). Although Vasek and Moore (2013) appeared to show that current versions of server side software were more likely to be attacked, that older software can be exploited remains an issue.

The huge complexity inherent in popular software products does suggest that making bug free software is not possible. Developers are human, and cannot be expected to consider every possible set of outcomes in a large codebase. Even if we were prepared to accept significantly increased costs in future software to eliminate vulnerabilities, this does not solve the problem of software which already exists. The developers of this software would then be faced with the issue of fixing known vulnerabilities in a timely manner, or considerably longer for potential future bugs or vulnerabilities. Software development is inherently to the advantage of an attacker, since they need to find only one vulnerability whereas the defender needs to find all of them (Anderson, 2005). Currently, major software vendors are largely doing a good job in fixing vulnerabilities quickly, so there is little need to impose an additional obligation.

Whilst this section has shown that an intervention by any of these three stakeholders would likely fail on efficacy grounds, it is also likely that enforcing any obligation against users would result in an unacceptable compromise of their privacy – particularly given the likely lack of efficacy. This is discussed in more detail in the section on ISPs (Section 5.3), since it is something which would need to be enforced by them.

Having excluded some of the stakeholders from analysis, this section will now consider the issues concerning the remaining stakeholders: ISPs, search engines, and hosting providers. Each stakeholder will be considered in terms of efficacy, and potential legal issues – particularly those connected with fundamental rights.

5.3 ISPs

There are many ways an ISP could choose to notify their customers that their machines have been compromised. These are listed in RFC 6561, written mostly by people affiliated to COMCAST an ISP based in the USA who has attempted to introduce such a system (Livingood, Jason and Mody, 2012). They found that no method was guaranteed to be 100% successful and each has its own set of limitations. For example, an email might be quick and possible to automate, but there is no guarantee that it would be ever be read, whether due to spam filters, users not using that account, or simply ignoring it[9]. On the other hand, blocking Internet access would alert the user to the problem, but there may not be anything they could do about it, such as if there was more than one machine on the network, or the infected device didn't have any interface with which to solve the problem (e.g. Internet of things type devices).

Of these potential approaches, within in the context of public health it is quarantine that holds the most interest. Recall from Section 4.3 that this is an idea which has been applied to local networks, and has been proposed to expand to the whole of the Internet (Charney, 2012). In history, quarantine was a process which separated those infected – or suspected of being infected – with a contagious disease[10]. The conventional historical wisdom was that any disease would have made itself apparent within 30 or 40 days, and that absent disease by this time the individual[11]. As understanding improved, a more nuanced strategy was able to emerge, with diseases known to have different incubation periods. For example, the severe acute respiratory syndrome (SARS) outbreak in 2003 had an incubation period of 2 – 9 days (Gensini et al., 2004).

The costs and effectiveness of quarantine is something which can vary wildly depending on the circumstances. In relation to SARS, there is little doubt that it was expensive. Singapore, for example, is estimated to have spent $5.2 million in implementing its quarantine procedures (Mubayi et al., 2010). Similarly, the amount of people quarantined in Canada for every instance of SARS was around 100, whereas in Hong Kong it was only approximately 12 (Mubayi et al., 2010). Sattenspiel and Herring (2003) analysed the effectiveness of quarantine controls using the influenza epidemic of 1918 in Canada as data. They concluded that the effectiveness varies depending on how long it lasts, when it is introduced, and argues that it has to be highly effective in small communities to have an appreciable effect on the pattern of the disease – although it was found to be effective in some circumstances.

[9] It is also a bad idea to begin getting users into the habit of reading an email, and following a link to "fix" the problem with their security, which is already a vector for attacks and making profit known as scareware.

[10] This is distinct from the concept of isolation, where only those *known* to be infected are separated (Gensini et al., 2004)

[11] Whether person or livestock. A considerable amount of the literature relating to quarantine has been to do with preventing the spread of illness in cattle. This is not something which generalises well to humans, since it is possible to implement strict movement controls, and even cull the population, to control the spread, see Sattenspiel and Herring (2003) and references within

Although legislation exists in the UK, and other countries, to control the spread of disease, these are around specific circumstances where the potential consequences are deemed grave enough to restrict the rights of individuals. However, there are several legal principles which suggest that this may be difficult to implement in application to the Internet. In this case, Article 15 of the data protection Directive is of relevance as it in principle attempts to protect data subjects from automated individual decisions without adequate safeguards. In addition, even if ISPs are private actors, because they act as necessary gateways between their subscribers and the whole Internet they are regulated differently from other Internet actors to make sure users basic rights can be exercised. A right to access information (as well as to receive information) including information online has also been recognised by judges (or through statute) at the national level [12] and by the European Court of Human Right (ECtHR). It could thus be argued that if the quarantining amounts to a suspension of Internet access it needs the setting up of a judicial proceeding to be implemented.

There are also some practical problems with the implementation of a quarantine for the Internet, to control the overall number of infected devices rather than to merely prevent the spread. Firstly, a characteristic of outbreaks such as SARS was that it was a temporary situation where individuals were deprived of liberty for only a couple of weeks. To do so for Internet access, this is something which would be done constantly, and would likely be regarded as an intolerable restriction of rights, given the difficulties which can arise for an individual who does not have access to the Internet. Unlike visiting a doctor, it is not easy or free to fix a computer which is infected, which poses additional problems. Schneier (2010) also points out, that there would need to be a means of appealing against a decision for a device incorrectly regarded as being infected – particularly important, given the figures for quarantine (Mubayi et al., 2010).

Finally, the effectiveness that a policy such as this could have is doubtful. One characteristic of modern computer use is that it is mobile, so the decision to block access to a device for one network is not something which would necessarily work for another network. Detecting symptoms from an infected device is also potentially more difficult than in an illness, because malware is designed by people who wish it to evade detection. Whilst the incubation periods of infectious illnesses are known, it would be a simple matter for an infected device to change the period of time it takes to become active. Any quarantine approach would also likely lead to an arms race, as attackers sought to react to any developments by the access providers (Schneier, 2010).

Less extreme measures by ISPs may have some effect, and are less onerous from a legal standpoint.

[12] See e.g. Conseil Constitutionnel Decision n 2009-580 of June 10th 2009, at http://www.conseil-constitutionnel.fr/decision/2009/2009-580-dc/decision-n-2009-580-dc-du-10-juin-2009.42666.html. In Estonia a law was adopted in 2000 stating that Internet access was a human right

From a right to private life and data protection perspective, the processing including the collection and retention of the customer's traffic data including IP addresses could probably be justified. Under Article 6 of the E-Privacy Directive, providers of public communications networks and publicly available electronic communications providers such as ISPs can process traffic data for traffic management purposes. This should include the safeguarding of network security and fraud detection[13], and since these can be deemed as an ISP's legitimate interests, no consent from their users should be required.

The justification of the processing of traffic data might be problematic if the data are transferred to third parties, even if it is for the very same purpose: safeguarding network security and fraud detection since Article 6 only targets data controllers acting under the authority of ISPs. For example, there might be lists of IP addresses of known infected devices such as those displayed on the website of the Honeynet Project or the lists maintained by Spamhaus or other similar companies about addresses which are known to have sent spam[14], or known C & C ZeuS servers which could in many cases be ordinary hacked computers[15].

Organisations like Team Cymru also notify ISPs when they detect compromised machines coming from the ISP's network. Here an argument can be made that in the case of the sharing of personal data to detect infections informed consent on the parts of the ISPs' subscribers is needed. In any case it is important that the principles of data minimisation and limited duration of the Data Protection Directive are complied with by the recipient of the personal data. And information relating to the recipient of the data should be provided to the subscribers, who should be able to exercise their rights (e.g. right to access, or rectification – see Article 12 of the Data Protection Directive).

Blocking a website is comparatively straightforward, and is already done in certain circumstances, notably copyright infringement and child pornography. The efficacy of such a move was argued in the recent *Cartier* case in relation to the blocking of The Pirate Bay following a previous injunction in 2012 (at [218–237], per Arnold J). Particularly persuasive for the judge, was the drop in the amount of users outside of the UK who continued to access the website compared to a striking drop from within the UK (at [223-224]). The two graphs presented in the opinion can be seen in Figure 5.1, which appears to demonstrate that there is at least a small amount of efficacy in blocking websites. It is also plausible that those who are prepared to download illegal content are prepared to try and work around any filtering efforts, whereas here the primary aim is

[13] See Recital 39 of the proposed general data protection Regulation. Proposal for regulation of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data (General Data Protection Regulation) COM(2012) 11 final.

[14] http://www.spamhaus.org/sbl

[15] https://zeustracker.abuse.ch/blocklist.php

merely to prevent victims from visiting pages which will attack them, so it might even be more effective.

That said, there are subtle differences in blocking compromised websites for security purposes compared to blocking websites for objectionable or illegal content. Websites who are usually the subjects of blocks – websites which infringe intellectual property rights; and websites which contain [legal or illegal] pornography. By contrast, the websites who are the subject of these blocks are in a variety of different categories, whose status potentially shifts from malicious to benign on a regular basis. This would require frequent assessments by the ISPs in order to ascertain the level of maliciousness, potentially creating a large amount of useless traffic to do so. A more appropriate stakeholder to do this on such a scale would be a search engine, who is actively invited to scrape most websites in any case.

Some of the potential legal issues an ISP might come up against would relate particularly to end users' fundamental rights and liberties. These include the right to freedom of expression including the right to access information (Article 10 of the ECHR and Article 11 of the European Charter); the right to private life and data protection (Article 8 of the ECHR and Articles 7 and 8 of the European Charter); and the right to conduct one's business (Article 16 of the European Charter).

Content filtering is particularly problematic in relation to these fundamental rights.

There is the usual issue of data protection and right to private life involved with this as well, since this requires the ISP to monitor the customers' traffic in order to determine which websites they are visiting to prevent access to malicious sites. On a technical level, it is true that preventing access to malicious domains is relatively trivial. A domain name is required to identify a website, and in general web use would not be possible without this. An ISP could use DNS based blocking, but such a technique is likely to amount to over-blocking and to become problematic from the perspective of freedom of expression. At this stage, the freedom of websites operators to conduct their business is at stake as well as the right to freedom of expression of both the users and producers of content. The ability to conduct one's own business, in particular if one is victim of a false positive in terms of their compromised or vulnerable status is likely to be jeopardized in a significant number of situations.

5.4 Search Engines

A search engine displays results in order of relevance to the query input by the user, so the method discussed in Section 3.3.5 which they could use to limit the effectiveness of drive-by downloads is to either warn the user, or not display it in the results, or

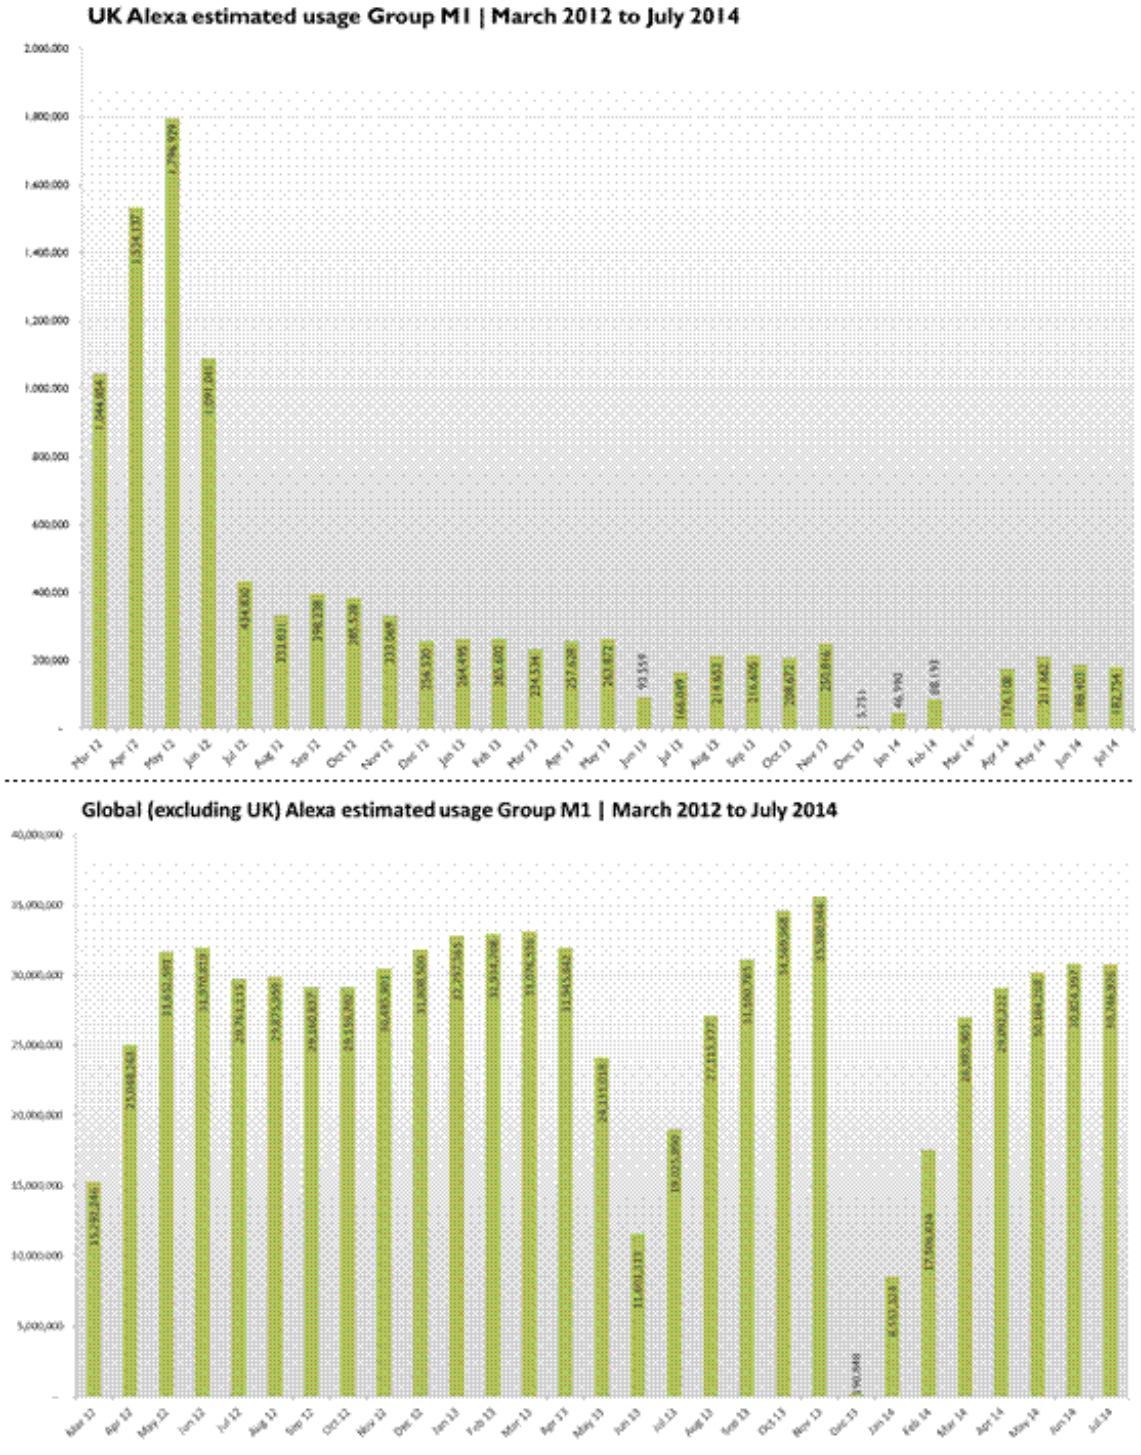

FIGURE 5.1: Comparison of UK and global visits to The Pirate bay website, taken from the judgment of *Cartier v BSkyB* (2014)

to "depreference" the results of malicious (or potentially malicious) websites (Edwards et al., 2012).

Due to the incentive of having a relevant set of results, search engines already perform work on doing their best to ensure that the Web pages in their listings are not malicious. However, there are limits as to what exactly they can do. Like any analysis of malware, a choice needs to be made between false positive and false negative errors. The consequences of falsely declaring a website to be malicious could potentially be serious to the site in question, given the essential nature of search engines for browsing the Web. Search engine providers need to be conservative about which pages they choose to classify as malicious to avoid mistakes, because the erroneous classification of a website could be seen as defamatory.

Search engines are hampered in their ability to detect malicious pages because of the nature of the service they offer. For example, it is possible for malware to hide from search engines, such as IP centric malware which was discussed earlier (Section 2.2.2). It is also the case the search engine cannot see parts of the website. There is a protocol, which, though not officially recognised, is commonly followed, and consist of keeping a page called robots.txt in the main folder of a website (as in `http://www.google.com/robots.txt`) indicating which automated scanners are permitted to visit the website; and which pages of the website they are allowed to visit. A scanner would identify itself through its user agent, for example Googlebot or Bingbot identify Google or Bing, two of the major search engines. On most occasions, the website operators would be keen to have these search engines visiting their website, but in the event that they do not then the search engines would follow the protocol and decline to visit certain pages.

Finally, although search engines are essential for the functioning of the Web, they do not have the ability to withdraw access to a website. As mentioned above, Akhawe and Felt (2013) found that between 9 and 23% of people ignored browser warnings about websites. Even at the lower bound of 9%, this is a not insignificant amount of traffic continuing to visit these websites. Pages which the search engines do not see would not necessarily be high in the rankings, but could still be used to host malware linked to from emails or social media. As discussed earlier, phishing emails use compromised websites in some 90% of cases, and social media websites are becoming increasingly popular with attackers because, in part, of the higher conversion rate. Research has shown that some 8% of all URLs on Twitter are spam of some description with a reasonably successful click-through rate (Grier et al., 2010). Facebook is also targeted through clickjacking, where a "like" button or similar is hidden on a (spam) page a person visits which persuades the victim's friends to visit the page as well (Faghani et al., 2012).

The approach of using a search engine do not so easily map onto a specific public health technique, and the restrictions on rights which might occur as a result of search engine interventions are limited. On the one hand, it could be argued that they restrict the

right to carry on one's business, and potentially freedom of expression. On the other hand, one could argue that as regards the issue of the freedom of expression of Internet users including content providers and the freedom to conduct one's business, search engines should be treated differently than ISPs since their activities are not essential to get access to the Internet and the delisting is not tantamount to a blocking.

Nevertheless, by providing an informational service such as this, there are certain risks which would need to be considered by the search engine. The first of that of defamation, where a search engine classifies a website as malicious. It would be possible to avoid defamation, by merely blocking the website rather than issuing a warning about it in the search results, or pushing it down the rankings without offering an explanation, although there are often warnings in the listings, see e.g. (Akhawe and Felt, 2013). In addition, a current service offered by Google is the "Safe Browsing API"[16], where it is possible to query to see whether a website is known to be infected. An increase in false positives there could potentially lead to further claims of defamatory content.

This risk can be seen, by analogy, in the case of *e360 Insight v The Spamhaus Project* (2007) Spamhaus maintain listings of known spamming addresses, and send them to ISPs so that they can be blocked. After doing this to a "marketing company", (i.e. classifying them as spammers) they were sued for defamation and had to fight a five year legal battle (For more details see Section 5.6). Although the situation is different, the effect on the plaintiff was very similar. Since they were blacklisted, they were unable to send any email which therefore deprived them of revenue. A company declined access to search engine results could well have similar issues in attracting customers. On the other hand, if any sort of immunity is too generous, then companies could be left without recourse in the event that their website is damaged, see e.g. (Ezor, 2010).

Whilst not blocking, having a good search engine ranking is essential if one wishes to compete online. It is this essential nature that leads to the next issue, where a search engine is in a dominant position on the market, then anything it does will be scrutinised for potential abuses of this position. This is particularly an issue for Google, who have a market share of around 90% in some countries, e.g. in the UK, (BBC, 2012b). Given their dominance, by excluding a website from their results, they are severely impacting upon that websites ability to participate in whichever market they aim to compete in. Even by using depreferencing, they risk accusations of bias and indeed, Google has been investigated in several jurisdictions for unfairly placing its own sites higher than opposition.

In the USA the FTC investigated "whether Google manipulated its search algorithms and search results page in order to impede a competitive threat posed by vertical search engines" (FTC, 2013). The investigation noted that enhancements made to Google's algorithms were done to improve the product and even though they hurt rivals there was

[16] https://developers.google.com/safe-browsing/?hl=en

no anticompetitive action by them. Google were also investigated by the EU Commission, and agreed to make changes to the display of their results rather than continuing with an adversarial procedure (Almunia, 2014).

In reality though, the circumstances for the scenario of blocking or depreferencing a malicious Web page are different. The cases brought by the Commission and the FTC related to abuses against Google's competitors rather, than the mere effects of pursuing some action which affects markets they are not participants in. Search engine intervention will only be effective and thereby reduce the exposure to malicious/vulnerable websites for a significant amount of people if all search engines with a high market share implement it, so a dominant provider such as Google is an advantage from a practical point of view.

5.5 EU Legislation

In this section, EU legislation and case law which could potentially impact the ability to enforce obligations for hosting providers. In the EU, different levels of protection are available, depending on the services provided by the intermediary. These are set out in the E-Commerce Directive Articles 12 – 15, under Section 4: Liability of intermediary service providers. Articles 12 – 14 detail the immunity from liability that intermediaries can claim, and Article 15 prohibits the imposition of a general obligation to monitor. The issue of fundamental rights will also be considered, which is something the CJEU has also been keen to emphasise in their judgments on intermediary liability.

The immunities are dependent on prima facie liability already having been established, so it is worth noting that there are situations where these are not relevant. Part of the proposed regulation would impose liability to hosting providers in the event of loss were they to fail to adhere to certain standards. The issue discussed below, whereby the intermediary were to lose their immunity as a result of obtaining actual knowledge through performing their obligations from this legislation. As such, it should briefly be considered whether current case law would permit hosting providers to be found liable in other situations.

As Section 3.2.2 showed, it is possible, in principle, for a website to be liable in negligence, or for the torts of another. This is a principle which could quite easily extend to hosting providers since it is them who are actually serving the content rather than the websites themselves. In *L'Oréal v eBay* (2009), it was held that eBay could in principle be jointly liable for facilitating tortious acts, and cases such as *Bunt v Tilley* and *Godfrey v Demon Internet* demonstrated that in diverse areas it might be possible to consider a hosting provider liable.

As with all of the immunities (Art. 12 – 14), they do not prevent injunctions being imposed by states. In this case Article 12(3) does provide that domestic legislation can be created to terminate or prevent infringements, so on condition these obligations do not amount to a "general obligation to monitor"[17], then there is little argument which can be made against injunctions. Recital 45 also provides:

> (45) The limitations of the liability of intermediary service providers established in this Directive do not affect the possibility of injunctions of different kinds; such injunctions can in particular consist of orders by courts or administrative authorities requiring the termination or prevention of any infringement, including the removal of illegal information or the disabling of access to it.

The main provisions will be enumerated here, and then the major cases will be discussed.

5.5.1 E-Commerce Directive

5.5.1.1 Articles 12 – 13

Articles 12 and 13 are designed to protect access providers from liability. There are no major cases in connection with Article 13, so the emphasis here will be on Article 12. Article 12 provides as follows:

> Article 12
>
> "Mere conduit"
>
> 1. Where an information society service is provided that consists of the transmission in a communication network of information provided by a recipient of the service, or the provision of access to a communication network, Member States shall ensure that the service provider is not liable for the information transmitted...
>   (a) does not initiate the transmission;
>   (b) does not select the receiver of the transmission; and
>   (c) does not select or modify the information contained in the transmission.
> 2. The acts of transmission and of provision of access referred to in paragraph 1 include the automatic, intermediate and transient storage of the information transmitted in so far as this takes place for the sole purpose of carrying out the transmission in the communication network,

[17] See discussion about Article 15 in 5.5.1.3

and provided that the information is not stored for any period longer than is reasonably necessary for the transmission.

3. This Article shall not affect the possibility for a court or administrative authority, in accordance with Member States' legal systems, of requiring the service provider to terminate or prevent an infringement.

The liability exemption for Article 12 does not allow the intermediary to do anything with the data, even cache it for improved performance. Whilst adequate for access providers, who are "mere conduits", certain services are not practical unless they do employ some sort of caching. Consider a DNS service, where once a the location of a particular domain is found, the server will hold onto it until the TTL (time to live) on the record expires. This time can be, for example, 86,400 seconds (24 hours), or even longer, and therefore falling outside the protection of Article 12. Given the volume of DNS queries, it would be impractical to run the service without caching, so for services such as this, Article 13 provides similar protection.

Article 13 provides additional protections for "automatic, intermediate and temporary storage of that information, performed for the sole purpose of making more efficient the information's onward transmission to other recipients of the service upon their request". This is also alongside conditions such as expeditious removal of certain types of content once the intermediary has actual knowledge of it. Using the DNS example, a cache poisoning attack is something which would likely come under this provision, that it should be fixed once they have been alerted to it. Whilst there have been tests as to the limits of Article 12, intermediaries relying on Article 13 in response to litigation are rare (Eecke, 2011).

5.5.1.2 Article 14

Article 14 offers protection for intermediaries who host information, and provides as follows:

Article 14

Hosting

1. Where an information society service is provided that consists of the storage of information provided by a recipient of the service, Member States shall ensure that the service provider is not liable for the information stored at the request of a recipient of the service, on condition that:

   (a) the provider does not have actual knowledge of illegal activity or information and, as regards claims for damages, is not aware of facts

or circumstances from which the illegal activity or information is apparent; or

(b) the provider, upon obtaining such knowledge or awareness, acts expeditiously to remove or to disable access to the information.

2. Paragraph 1 shall not apply when the recipient of the service is acting under the authority or the control of the provider.

3. This Article shall not affect the possibility for a court or administrative authority, in accordance with Member States' legal systems, of requiring the service provider to terminate or prevent an infringement, nor does it affect the possibility for Member States of establishing procedures governing the removal or disabling of access to information.

Like Article 13, this also provides for a notice and takedown procedure once the provider in question has actual knowledge of the illegal activity. It is in connection with this provision that the greatest questions about intermediary liability have arisen in the European Court. This is unsurprising perhaps, given the wide variety of possible actors who could be involved. On the one hand, a provider who merely provides access to a server so their customers may run websites has very little control over the information stored within it. On the other hand, the livelihood of many Web 2.0 companies depends on the ability to manipulate data provided to them by customers and other sources.

5.5.1.3 Article 15

Article 15 of the E-Commerce Directive is also concerned with the immunities, though more in regards to the limits to which the state can require obligations of them. Naturally the introduction of any form of legislation to impose a duty on hosting providers, or any other intermediary, would be more likely to fall within Article 15 rather than Articles 12 – 14. Article 15 provides:

Article 15

No general obligation to monitor

1. Member States shall not impose a general obligation on providers, when providing the services covered by Articles 12, 13 and 14, to monitor the information which they transmit or store, nor a general obligation actively to seek facts or circumstances indicating illegal activity.

2. Member States may establish obligations for information society service providers promptly to inform the competent public authorities of alleged illegal activities undertaken or information provided by recipients of their service or obligations to communicate to the competent

> authorities, at their request, information enabling the identification of recipients of their service with whom they have storage agreements.

5.5.2 Balancing Fundamental Rights

Another issue which needs to be considered, is that of fundamental rights enshrined in the Charter of Fundamental Rights of the European Union. The issue becomes what should be done when rights conflict is one which emerges in several cases. Given that the majority of current case law is related to intellectual property, this involves the application of the right to property against other rights. Article 17 reads:

Right to property

1. Everyone has the right to own, use, dispose of and bequeath his or her lawfully acquired possessions. No one may be deprived of his or her possessions, except in the public interest and in the cases and under the conditions provided for by law, subject to fair compensation being paid in good time for their loss. The use of property may be regulated by law in so far as is necessary for the general interest.
2. Intellectual property shall be protected.

In *Promusicae*, the CJEU held that it was a matter for national courts to decide on which rights take priority, and did not offer any guidance as to which should prevail. The court held:

> That being so, the Member States must, when transposing the directives mentioned above, take care to rely on an interpretation of the directives which allows a fair balance to be struck between the various fundamental rights protected by the Community legal order. Further, when implementing the measures transposing those directives, the authorities and courts of the Member States must not only interpret their national law in a manner consistent with those directives but also make sure that they do not rely on an interpretation of them which would be in conflict with those fundamental rights or with the other general principles of Community law, such as the principle of proportionality[18].

Many of these rights simply restate European Convention rights, and analysis in the relevant decisions is based on that jurisprudence[19]

[18] *Promusicae* (2008), at [68]
[19] See e.g. *Scarlet Extended v SABAM* (2011)

5.5.3 SABAM Cases

Two cases on what constitutes a general monitoring obligation involve SABAM – , and their disputes against an ISP, and a social network platform in order to protect the copyright of their members. These cases are *Scarlet Extended v SABAM* and *SABAM v Netlog*. In *Scarlet Extended v SABAM*, the ISP, Scarlet was ordered by the Belgian court to comply with an injunction sought by SABAM, representing rights holders. The injunction required them to take action to make future file sharing impossible, by installing a system deemed by an expert witness to be feasible, which would detect and block peer to peer traffic. Scarlet appealed, on the grounds that its effectiveness had not been proved; that it was contrary to Article 15 E-Commerce Directive, and that it was contrary to EU legislation relating to protection of personal data and the secrecy of communications. The appeal court referred the case to the CJEU.

The court considered the terms of the injunction, and ascertained the requirements to be that Scarlet had to:

- Identify all peer to peer traffic;
- Find files which were subject to intellectual property rights of the rights holders;
- Determine whether they were being shared unlawfully;
- Block the sharing which was to be regarded as unlawful[20]

Given the terms, the CJEU held that: "Preventive monitoring of this kind would thus require active observation of all electronic communications conducted on the network of the ISP concerned and, consequently, would encompass all information to be transmitted and all customers using that network"[21], and that this was contrary to Article 15(1) of the E-Commerce Directive[22]

The court also considered the *Promusicae* judgment, and the requirement that there be an appropriate balance between the rights of intellectual property rightsholders and fundamental rights[23]. Given the unlimited time that this injunction would be required for, and thecomplex and complicated nature of it, there would be restrictions on the ISP's right to conduct its business[24], and potentially restrictions on Articles 8 and 11 of the charter in regards to the ISP's customers.

Another case involving SABAM concerned an injunction against a social network "Netlog", where SABAM alleged that Netlog's customers were making copyrighted material

[20] *Scarlet Extended*, at [38]
[21] ibid. at [39]
[22] ibid. at [40]
[23] ibid. at [45]
[24] Article 16 of the European Charter, and in addition contrary to Article 3(1) of Directive 2004/48 requiring the measures to be introduced not be unneccessarily complicated or costly.

available without permission. SABAM gave notice that it should cease and desist making this material available, and demanded a fine for every day without compliance. The court queried whether such an injunction would be contrary to Article 15, or fundamental rights of the European Charter or ECHR.

The court considered the filtering system to be implemented would require the following:

- first, that the hosting service provider identify, within all of the files stored on its servers by all its service users, the files which are likely to contain works in respect of which holders of intellectual-property rights claim to hold rights;
- next, that it determine which of those files are being stored and made available to the public unlawfully; and
- lastly, that it prevent files that it considers to be unlawful from being made available[25].

The court determined the extent of the monitoring requirements which were "active observation of files stored by users with the hosting service provider and would involve almost all of the information thus stored and all of the service users of that provider (see, by analogy, Scarlet Extended, paragraph 39)."[26]. The reasoning from *Scarlet Extended*[27] in relation to the obligation being contrary to Article 15 and fundamental rights. The court also considered that this could potentially undermine freedom of information owing the difficulty in distinguishing between lawful and unlawful content[28].

5.5.4 20th Century Fox v BT

In *20th Century Fox v BT* (2011), the plaintiffs sought to obtain a court order under s97A Copyright Designs and Patents Act, requiring the respondents to take action to block access to the Newzbin2 website, which offered the ability to illegally download copyrighted material. The order requested that the following technology be adopted:

1.

   i IP address blocking in respect of each and every IP address from which the said website operates or is available and which is notified in writing to the Respondent by the Applicants or their agents.

   ii DPI based blocking utilising at least summary analysis in respect of each and every URL available at the said website and its domains and

[25] *SABAM v Netlog* at [36]
[26] ibid. at [37]
[27] ibid. at [38, 46]
[28] ibid. at [51]

> sub domains and which is notified in writing to the Respondent by the Applicants or their agents.

2. ...[P]aragraph 1(i) and (ii) is complied with if the Respondent uses the system known as Cleanfeed...

Essentially, the order sought "that BT should implement the same measures with regard to the Newzbin2 website as it already operates with regard to URLs reported to it by the IWF"[29]. BT objected to these on a number of grounds, most pertinant for the case here under fundamental rights, Article 12 E-Commerce Directive; and Article 15 E-Commerce Directive.

In relation to fundamental rights, BT argued that Section 97A was not "prescribed by law", and hence contrary to Article 10 ECHR[30]. This argument was based on the reasoning of Advocate General Villalón in *Scarlet Extended v SABAM* (see Section 5.5.3), this being prior to the judgment of the CJEU on the matter. This claim was rejected, where it was held that there was a distinction between what Advocate General Villalón deemed to be excessive, Arnold J. held:

> On the contrary, the order sought by the Studios is clear and precise; it merely requires BT to implement an existing technical solution which BT already employs for a different purpose; implementing that solution is accepted by BT to be technically feasible; the cost is not suggested by BT to be excessive; and provision has been made to enable the order to be varied or discharged in the event of a future change in circumstances. In my view, the order falls well within the range of orders which was foreseeable by ISPs on the basis of section 97A, and still more Article 8(3) of the Information Society Directive.[31]

It was also submitted that the order would be contrary to Article 12(1) of the E-Commerce Directive, with it being common ground that BT were a mere conduit[32]. Although not liable for the copyright infringement themselves, it was argued that the text of Article 12(3) allowing the possibility for a court to issue an injunction should be read restrictively given that Article 14(3) offered more latitude to the state to "procedures governing the removal or disabling of access to information"[33].

This argument was also dismissed, and it was held that Article 12(3) should be read with Recital 45 as well as Article 18(1) concerning the power of member states to make

[29] ibid. at [70]. The Internet Watch Foundation (IWF) is an organisation which aims to combat child sexual abuse content online. See *20th Century Fox v BT* (2011) at [65 – 69] for more details
[30] *20th Century Fox v BT* (2011) at [163]
[31] ibid. at [177]
[32] ibid. at [159]
[33] ibid.

use of courts and injunctions[34] and observing that there was no limit on the type of injunction a state could require under Article 12(1). Article 14 was regarded as distinct from Article 12 given the different context to do with hosting providers, and so did not call for a restrictive reading of Article 12(3)[35]. *Dramatico Entertainment v BSkyB* (2012) followed a similar line of reasoning, and was not contested by the ISPs.

Article 15 was also argued alongside Article 12. The court briefly considered both the notion of "monitoring" and the notion of "general". It was held that blocking access to a website did not require active monitoring, but merely block through automated means without inspecting any data about subscribers. In addition, it was held that even were this to be regarded as "monitoring", it would be specific monitoring rather than general and as such not contrary to Article 15[36].

5.5.5 Google France v Louis Vuitton Malletier SA

The first major test of the Article 14 immunity in the ECJ came in *Google France v Louis Vuitton Malletier SA* (2010). In this case, trade mark holders claimed that Google was liable for infringement through its AdWords service by allowing advertisers to use words associated with these trademarks. Google AdWords[37] is a facility provided to advertisers who can have their link placed at the top of the search results, provided they pay Google for each click received. Given that there was no limit to the amount of advertisers who could choose to have their advert displayed, when a user searched for these trademarks they were given links to websites selling counterfeit products. The option also existed within the AdWords to advertise "imitation" or "copy" or other similar terms [38].

In order for Google to fall under the protection of Article 14 of the E-Commerce Directive, the court considered whether Google constituted an "information society service", requiring: "any service normally provided for remuneration, at a distance, by electronic means and at the individual request of a recipient of services"[39]. By reference to Google's AdWords service[40] rather than Google itself, the court held that Google did meet the

[34] Recital 45 reads: "The limitations of the liability of intermediary service providers established in this Directive do not affect the possibility of injunctions of different kinds; such injunctions can in particular consist of orders by courts or administrative authorities requiring the termination or prevention of any infringement, including the removal of illegal information or the disabling of access to it.". Article 18(1): "Member States shall ensure that court actions available under national law concerning information society services' activities allow for the rapid adoption of measures, including interim measures, designed to terminate any alleged infringement and to prevent any further impairment of the interests involved."

[35] *Twentieth Century Fox v BT* at [160]

[36] ibid. at [162]

[37] https://www.google.co.uk/adwords/

[38] *Google France* at [29]

[39] Article 2(a) of the E-Commerce Directive refers to Directive 98/34/EC as amended by Directive 98/48/EC for the definition

[40] ibid at [23]

requirements and could be considered for immunity from liability according to Article 14 of the E-Commerce Directive[41].

Both the opinion of the Advocate General and the court held that there were further requirements for an information service provider in order to benefit from Article 14 immunity. Advocate General Maduro read into the goals of the E-Commerce Directive as being "to create a free and open public domain on the Internet.... by limiting the liability of those which transmit or store information" [42] in general, whilst considering about Article 15, "the very expression of the principle that service providers which seek to benefit from a liability exemption should remain neutral as regards the information they carry or host" [43]. As such, Google's "natural" search engine rankings would fall under the immunity but their AdWords service would not [44].

However, the court also held that the limitation of the "intermediary service provider" needed consideration according to the goals of the E-Commerce Directive. The court chose to interpret this in terms of recital 42 (E-Commerce Directive), which reads as follows:

> "The exemptions from liability established in this Directive cover only cases where the activity of the information society service provider is limited to the technical process of operating and giving access to a communication network over which information made available by third parties is transmitted or temporarily stored, for the sole purpose of making the transmission more efficient; this activity is of a mere technical, automatic and passive nature, which implies that the information society service provider has neither knowledge of nor control over the information which is transmitted or stored."

The inevitable consequence of the Court's decision to do this, meant that for an intermediary to qualify for Article 14 protection "it is necessary to examine whether the role played by that service provider is *neutral*, in the sense that its conduct is merely technical..." [45]. They declined to rule on what constituted neutrality, holding that it would be a matter for the national courts to establish on a case by case basis [46].

5.5.6 L'Oréal v eBay

The other major European case concerning Article 14 is that of *L'Oréal v eBay* (2011). In this case, L'Oréal claimed that eBay, an online auction service, was jointly liable

[41] ibid at [110]
[42] *Google France*, opinion of Advocate General at [142]
[43] ibid. at [143]
[44] ibid. at [144-145]
[45] *Google France* at [114], emphasis mine
[46] ibid. at [119]

for infringements of its trade marks with the individuals who were selling them using the service [47]. eBay already took measures to remove items contravening its policies, one of which was rights violations, and offered help to rights holders in removing these items under a program referred to as VeRO (Verified Rights Owner), although L'Oréal declined to take part in the program claiming it was inadequate [48].

The question arose as to whether eBay had immunity from liability under Article 14 of the E-Commerce Directive. Following *Google France*, the court held that that eBay was an information society service:

> "It is apparent from the definition of 'information society service', cited at paragraphs 8 and 9 of this judgment, that that concept encompasses services provided at a distance by means of electronic equipment for the processing and storage of data, at the individual request of a recipient of services and, normally, for remuneration. It is clear that the operation of an online marketplace can bring all those elements into play."[49]

They additionally followed that an intermediary requires still to act "neutrally by a merely technical and automatic processing of the data provided by its customers..."[50]. Optimising the presentation, or offering other assistance to the users of the service would be deemed to not be neutral, though they left it to the national courts to decide on individual cases.

The other element of the decision relating to Article 14, was the point at which the intermediary, if entitled to immunity, would be deemed to have obtained "awareness" of the illegal content or activity and as such be required to remove it. This requirement under Article 14(1)(b) was interpreted as being "aware of facts or circumstances on the basis of which a diligent economic operator should have identified the illegality in question"[51], however the information is obtained[52], e.g through an internal investigation or through a notification, though noting that merely being notified is not automatically constituted as awareness, e.g. if it is insufficiently detailed or substantiated [53].

The court also discussed whether the intermediary should prevent future infringements, and if so, what they should be. The court held that, Article 15(1) prohibited "active monitoring of all the data of each of its customers in order to prevent any future infringement of intellectual property rights"[54]. However, "injunctions which are both effective

[47] ibid at [32 – 34]
[48] ibid at [46]
[49] ibid. at [109]
[50] ibid. at [113]
[51] ibid. at [120]
[52] ibid. at [121]
[53] ibid. at [122]
[54] *L'Oréal v eBay* at [139]

and proportionate may be issued against providers such as operators of online marketplaces", citing the analysis in *Promusicae*[55]. Therefore, it was held that national courts must be able to ensure an online marketplace "take measures which contribute, not only to bringing to an end infringements of those rights by users of that marketplace, but also to preventing further infringements of that kind"[56]

5.5.7 Digital Economy Act (UK) Cases

Following the passage of the Digital Economy Act through parliament, BT and TalkTalk sought a judicial review in relation to the provision for obligations placed upon ISPs, in *R(on the application of British Telecommunications and another) v The Secretary of State for Business, Innovation and Skills, and others* (hereafter *DEA First Instance*, and the associated appeal *R(on the application of British Communications and Another v Secretary of State for Culture, Olympics, Media and Sport* (hereafter *DEA Appeal*). Amongst other grounds, the claimants argued that such obligations would constitute an infringement of Articles 12 and 15 of the E-Commerce Directive.

The claim relating to Article 12 rested on the legislative history, and purpose of the directive, and invited the court to take a wide interpretation of Article 12 in order to achieve the assumed purpose from the legislative history.

The history in question was the decision of the Commission to reject the European Parliament's proposal to decline to include a section concerning "notice and takedown" in relation to Article 12 [57]. Further, it was argued that Article 12(3) would otherwise prevent a mere conduit from being required to terminate or prevent an infringement committed by another, supporting their interpretation of 12(1) extending beyond mere liability to the responsibility of the ISP [58]. As well as arguing that the fines imposed for non-compliance under the DEA would constitute liability, the construction of Article 12(1) sought by BT was:

> "an economic or other burden falling on an ISP that would not have so fallen if there had not been a transmission of information in breach of copyright, etc." [59]

The judges were unconvinced by these arguments. At first instance, Parker J held that the legislative history in fact illustrated the careful balancing act done by the Commission meant that one should be more reluctant to change the words beyond their natural meaning rather than interpret it in a different way. As a result "liability *for the*

[55] *Promusicae* at [65 – 68], cited ibid at [141 – 143]
[56] ibid. at [144]
[57] *DEA Appeal* at [56]
[58] ibid. at [57]
[59] *DEA First Instance*, at [108]

information transmitted was a carefully delineated and limited concept", [60]. Choosing to construct Article 12(1) in the manner suggested by BT would not be possible "without doing violence to that language and thereby upsetting the careful balance represented by the text" [61].

The construction of Article 12(3) was also regarded as limited, merely enabling member states to implement legislation as long as the ISP is not made liable "*in respect of to the infringement itself*" [62]. The penalties which could be imposed on the ISPs under the DEA were regarded as completely dependent on non-compliance of the DEA itself and hence were not to do with the infringement of the customers themselves [63]. The Court of Appeal agreed with this assessment[64]

In relation to Article 15, the court defined a general obligation as follows:

> "A "general" obligation refers to a systematic arrangement whereby the putative "monitor" is inspecting or examining information randomly, or by reference to particular classes... and is not focusing on a specific instance that has... been brought to its attention" [65]

Given "monitor" its natural meaning, it was held that there was no obligation for the ISP to do so, and as such it was not going contrary to Article 15, on account of the fact that it was the rights holders who were doing the monitoring, and then passing reports onto the ISPs.

The court held that the obligation required by the DEA was not monitoring in that sense, and that the ISPs had an essentially "passive" role, whereas it was the rights holders who were performing the monitoring [66]. It was also argued that the legislation imposed "a general obligation to seek facts or circumstances indicating illegal activity" [67]. This again was held to be consistent with Article 15, because it is merely acting on reports it recieves rather than actively seeking facts for evidence of infringement.

5.5.8 Cartier v BSkyB

Previous cases requiring injunctions have required the copyright holders to use s97A Copyright, Designs and Patents Act 1988 (CDPA) which provides that "(1) The High Court (in Scotland, the Court of Session) shall have power to grant an injunction against

[60] *DEA First Instance*, at [102], emphasis original
[61] ibid. at [108]
[62] ibid. at [103], emphasis mine
[63] ibid. at [107]
[64] *DEA Appeal*, at [50 – 60]
[65] *DEA First Instance*, at [114]
[66] ibid. at [115-116]
[67] ibid. at [118]

a service provider, where that service provider has actual knowledge of another person using their service to infringe copyright". This was the implementation of Article 8(3) of the Information Society Directive into domestic legislation which related only to copyright. There being no equivalent provision relating directly to trademarks in UK law, the court in *Cartier* (2014) examined a more general set of principles for obtaining an injunction for blocking websites. Instead, the court relied on Section 37(1) Senior Courts Act which provides: "The High Court may by order (whether interlocutory or final) grant an injunction . . . in all cases in which it appears to be just and convenient to do so." [68]

It was held that the court did have the authority to grant an injunction in favour of the rights holders, confirming that a more general power to do so did exist. The requirements for granting an injunction amounted to a balancing approach between existing rights, and proportionality of the injunction sought (at 192 – 261). Specific measures proposed by the ISPs for consideration were regarded as important in the calculation of whether an injunction was proportional such as alternative measures [69], efficacy [70], and costs [71]. Arnold J concluded that "In my view the key question on proportionality is whether the likely costs burden on the ISPs is justified by the likely efficacy of the blocking measures and the consequent benefit to Richemont having regard to the alternative measures which are available to Richemont and to the substitutability of the Target Websites", and that in this instance the injunction was justified (at [261]).

Given that the *Cartier* case was for an injunction in relation to access providers, it is unnecessary to go into detail about the costs, benefits, and alternatives. A couple of points worth noting, however, was the proposition that "It is common ground that, in principle, the most effective means of removing offending websites from the internet is takedown by the host. " (at [199]), although in this case the ability of the fraudulent website to switch provider was regarded as problematic. The drawbacks relating to search engines intervention were also mentioned, relating to malware the relevant one was that the content was still online even if it was not visible in search results (at [214]).

5.6 USA Intermediary Law

The applicable law in this country is EU law, but the provisions in the USA also deserve some attention, §230(c) of the Communications Decency Act offer safe harbours for intermediaries. This is not an attempt to draw any sort of comparison between the two regulatory regimes, but it is interesting for two primary reasons. Firstly, the statute predates the E-Commerce Directive, and influenced its development (Eecke, 2011), and

[68] Cited by *Cartier*, per Arnold J. at [74]
[69] ibid. at [217]
[70] ibid. at [219]
[71] ibid. at [253]

the volume of litigation could offer additional insight. Secondly, §230(c)(2) which offers protection to intermediaries who incorrectly take down or block questionable material in good faith. In the event that legislation compelling action were to be introduced, it is worth considering becuase false positives are likely to occasionally occur making such a provision potentially valuable.

§230(c)(1) provides as follows:

> 1) Treatment of publisher or speaker
>
> No provider or user of an interactive computer service shall be treated as the publisher or speaker of any information provided by another information content provider.

The first major test for this legislation came in the case of *Zeran v America Online* (1997). In this case, an anonymous user posted about bad taste t-shirts relating to the Oklahoma bombing, and provided the claimant's phone number for people to call. The claimant, not connected with this enterprise, began to receive abusive messages and death threats and asked AOL to remove the posts, and specify that they were false. He alleged that AOL was liable for defamation, and AOL cited §230 as a defence (*Zeran v America Online*, at 329).

The court held that §230 did protect AOL from liability as a publisher of the information, but went further and also claimed that the liability of a distributor was also exempt (*Zeran v America Online*, at 332). This meant that even in the face of actual knowledge, §230 continued to protect the defendant. The rationale for this was twofold. Firstly, that following any report of defamation a service provider would likely be overly restrictive on free speech in order to protect themselves. Secondly, the aim of the statute included the intention to self-regulate, and the court held that liability with actual knowledge would weaken the incentives for intermediaries to self-regulate or monitor since they would then be liable (*Zeran v America Online*, at 333).

There has been little appetite on the part of the judiciary to reduce its scope. Similar results followed in *Johnson v. Arden* with the court holding a general immunity from content created by a third party rather than the joint liability afforded by Missouri law for a wrong "done by concert of action and common intent and purpose" (at 790). The court declined to follow this reasoning, and held that §230 allowed immunity from content created by a third party. *Barnes v Yahoo!* (2009) also viewed other claims which relied upon the role of publishers (such as the creating or deleting of content), and confirmed that despite the intention of the act being to regulate in defamation law, it also applied in other contexts (*Barnes v Yahoo! Inc.*, at 1101-02), and *Goddard v Google* (2009) where a user fell victim to a fraudulent advertisement.

Nevertheless, it has been noted that the *intention* of the Act, and the overall effect it has had on subsequent actions has not been consistent. Amongst the reasons for enacting §230 was in order to overturn the decision from *Stratton Oakmont, Inc. v. Prodigy Servs. Co*), (See e.g. *Zeran v America Online* at 331, *Barnes v Yahoo!* at 1101, *Roommates* at 1163), in which a message board provider was found to be liable for defamation on account of the fact that they undertook editorial decisions by removing obscene content.

In *Doe v GTE Corp* (2003) the court observed the difference between the supposed intention of the statute (to protect those that seek to remove questionable content from liability as a publisher), but that its application has been that these providers are afforded protection from doing nothing. They offered a reading of the seeming inconsistencies between the caption (good Samaritan blocking) and the text, which would allow a state to add intermediary obligations onto the provider whilst still protecting them from defamation claims. *Barnes v Yahoo!* explained this contradiction by suggesting that the protection is offered if they do not make modifications; but offers additional protection in the event that they do (*Barnes v Yahoo!*, at 1105).

That is not to say that the immunity is infinite, there are occasions where it does not apply. *Fair Housing Council of San Fernando Valley v Roommates.com* (2008) held that by **requiring** the third parties to enter certain information then they lose the immunity for that content. In *Fair Housing*, this was asking the users their sex and sexual orientation and whether they had any children, and then requiring users to express their preferences about potential roommates, which was contrary to California law (*Fair Housing*, at 1161). Whilst Roommates.com were liable for those parts of the sites, the comments which were entered by the users on their profiles fell within the protection. Similarly, actively encouraging illegal content from users also had the effect of depriving the intermediary of immunity *Jones v Dirty World Entertainment Recordings* (2012).

§230(c)(2) CDA offers protection for blocking acts, done in good faith. It is this provision which is particularly worthwhile considering for the purposes of legislation which requires hosting providers to block websites which might have malicious content owing to the potential for false positives. It provides:

> No provider or user of an interactive computer service shall be held liable on account of –
>
> (A) any action voluntarily taken in good faith to restrict access to or availability of material that the provider or user considers to be obscene, lewd, lascivious, filthy, excessively violent, harassing, or otherwise objectionable, whether or not such material is constitutionally protected"

This provision has not received the same amount of attention as has §230 (c)(1), yet some important cases have come through. *Zango v Kaspersky* held that anti-virus software was protected under §230, and that spyware was objectionable content which could be included in the protection as being blocked. Consequently, the classification of the claimant's software as spyware by the defendant was not actionable. This is an important decision, because otherwise security companies could be inundated by claims of defamation.

There have also been cases dealing specifically with spam email, where operators have either blocked legitimate mail as well or if the sender denied that they were sending spam email. It is possible that the wiser course for Spamhaus in the case of *e360 v Spamhaus* (2007) might have been to have simply use §230 immunity rather than deny the jurisdiction of the US courts. This resulted in court cases lasting five years. In *Holomaxx v Yahoo!* (2011), the court held that all doubts should be in favour of immunity, and also suggested that Holomaxx pointed out no industry standards that Yahoo! had breached.

There have been other judgments which recognise the reality of the situation in the fight against spam email. *Holomaxx Technologies v Microsoft* (2011) claimed Microsoft filtering their emails was done in bad faith due to Microsoft wishing to cut costs by filtering both good and bad emails. However, it offered no support in defence of this claim and so Microsoft could not be liable due to §230. In *e360 Insight v Comcast* (2008), the court observed:

> "Under the law, a mistaken choice to block, if made in good faith, cannot be the basis for liability under federal or state law. To force a provider like Comcast to litigate the question of whether what it blocked was or was not spam would render §230(c)(2) nearly meaningless."

That is not to say that the intermediaries have had it all their own way, however. For example, *National Numismatic Certification v eBay Inc.* (2008) held that no general protection applied where specific examples existed in the statute, and in *Smith v. Trusted Universal Standards* (2010), Microsoft and Cisco both failed in whether they should be classified as software access providers for protection under §230.

5.7 Discussion

Although interventions at the level of an access provider have got the potential to be effective, technologies which could be implemented are likely to amount to an excessive restriction of the user's fundamental rights. A more general issue though, is whether we should be considering ISPs as policemen or merely enablers of access to the Internet.

Whilst some content such as child pornography is universally regarded as unacceptable, should ISPs really be performing additional steps to prevent access to it when their function is merely to control access. An appropriate analogy might be the postal system, which allows objectionable or harmful materials to be sent, or a seller of cars which can be used to facilitate crimes or cause serious accidents. Slippery slope arguments also apply, given the extension of this filtering from the IWF list, to copyright infringement, to (possibly) "extremist" content in the future (B.B.C, 2014).

The scope of this research does not allow space for an appropriate contribution to the debate relating to content filtering or censorship in general. However, the potential for abuse of such a system, and rights restrictions which would be inherent in its implementation mean that other options should be pursued before making requirements of access providers.

Whilst the restrictions on rights from a search engine are considerably less than that of an ISP, there does not appear to be a likelihood that the effectiveness of an intervention would improve significantly if an obligation were imposed. The incentives are in favour of them taking a small amount of action (See Section 3.3.5) to making their results as good as possible, so it is concluded that there should be no obligation on search engines beyond what they already do.

Hosting providers by contrast offer the best of both worlds when compared to search engines and ISPs. The effectiveness is higher than that of a search engine, in that they can effectively block the malicious website or page, and have a greater amount of information to make the assessment. Compared to an ISP, they offer very few restrictions on individual rights, whilst likely offering the same efficacy[72]. As the following analysis will show, it is also likely that such an obligation could very well be consistent with Article 15 of the E-Commerce Directive, and a proportionate response to the threat from malware.

On an initial reading of *Scarlet Extended* and in particular *SABAM v Netlog*, it might appear that an obligation on hosting providers to scan for vulnerable or infected websites. However, a careful reading of the judgment indicates that there is considerable scope for less onerous obligations. As described in Section 5.5.3, the obligations required montitoring of: "*all* of the information thus stored and *all* of the service users of that provider"[73] (emphasis mine). This implies that it was merely the broad nature of the requirement, rather than the requirement itself, which was contrary to Article 15, and that there is still scope for a similar obligation with a smaller scope (Kulk and Zuiderveen Borgesius, 2013).

This is supported by the other cases discussed, demonstrating that there are many situations where an intermediary could have obligations imposed upon them. In *L'Oreéal*

[72] A comparison between the two is not considered in this thesis, and remains a topic for future work
[73] *SABAM v Netlog* at [37]

v eBay, the requirement on an online marketplace to take measures to prevent further infringements is particularly stringent[74]. The discussion in *Twentieth Century Fox v BT*[75] also demonstrated the limits of Article 15 protection. Whilst blocking individual websites may not a "general" obligation, it is still necessary to inspect every single request made by a user, in order to determine whether to block the website.

Intellectual property rights holders have had a considerable amount of success in the UK, in obtaining injunctions to require access providers to block websites which infringe copyright, and more recently trademark (See summary in *Cartier* (2014), per Arnold J. at [52 – 57]). Following the *Twentieth Century Fox v BT* decision, we are left with a situation where an access provider can arguably be required to "both systematically monitor users' activities and eventually modify users' behaviour", leading to Article 15 protection being essentially redundant and that instead, Article 8 and Article 10 ECHR should be used (Stalla-Bourdillon, 2013).

On the other hand, the judgment of *DEA First Instance*, whilst not prohibiting an obligation on ISPs, did so only because the monitoring which was occurring was done by the ISPs, and that there was no requirement for them to do so. Recall that the definition of a general obligation was a "systematic arrangement whereby the putative "monitor" is inspecting or examining information randomly, or by reference to particular classes...". It is exactly this sort of obligation which we aim to impose on ISPs. However, this is not a decision which binds courts as to the nature of Article 15, and prevailing judicial opinion appears to be far less inclined to consider a general obligation to be something akin to the *SABAM* cases rather than anything else.

If, as it appears, there is a distinction being made between the content of the monitoring/filtering and the metadata, then the proposal still may have difficulty since it needs to analyse the content of some files to ascertain what version is running. However, it is submitted that the obligation is small relative to the "general" scanning of metadata by systems such as Cleanfeed. There are a very small amount of files which would need to be considered at all, so it is submitted it would be consistent with Article 15.

Whilst it may be possible to introduce legislation compatible with Article 15, there are further considerations which should be made about intermediary liability in general. Amongst the probable purposes of the immunities in the Directive was to shield intermediaries from the increasing amount of cases in the 1990s against them (Eecke, 2011). Whether a duty like the one proposed would be consistent with the purposes of the Directive remains to be seen, and the consequence of an obligation like this, whilst consistent with Article 15 could have knock on effects on Article 14 which might be unfortunate. As Eecke (2011) observed, Article 15 is required in order for Article 14 to be effective, since a general obligation to monitor would imply actual knowledge of the

[74] at [144]
[75] at [162]

content, therefore causing the intermediary to fall outside of the protection of Article 14.

This is particularly true, as a result of the decisions in *Google France* and *L'Oréal v eBay* regarding the "neutrality" with which intermediaries are required to act in order to benefit from the immunity from liability. The opinion of Advocate General Jääskinen in *L'Oréal v eBay* drew attention to the flaws in the current reasoning. He suggested that the interpretation of recital 42 by the ECJ in *Google France* was in error. Instead, he argued, the recital was probably referring to Article 13 following the previous recital's mention of caching considerations and that it would be more sensible.

Regrettably, as described, the main judgment of the ECJ chose not to follow his opinion, and we are now faced with the implied consequence that a provider which voluntarily (or otherwise) undertakes measures to stop an infringement of some description – or to improve their security – would lose their immunity from liability. As Advocate General Jääskinen himself commented

> "I would find it surreal that if eBay intervenes and guides the contents of listings in its system with various technical means, it would by that fact be deprived of the protection of Article 14 regarding storage of information uploaded by the users." (*L'Oréal v eBay* at [146]).

By imposing on hosting providers a duty to check for security, and remove/block as appropriate, then this places them in an editorial role. Given current judicial thinking, that "neutrality" is required in order to fall within Article 14 protection, this is likely to fall outside. The following section considers legislation in the USA which offers a greater liability shield to intermediaries who take action in good faith to remove obscene content. A similar provision might be required along side the current proposal, for the following two reasons:

1. Any form of security requires trade-offs between false positives and false negatives. In the event that a website is falsely blocked as a result of a false positive, then there could be considerable consequences for the operator.

2. An obligation to detect these sorts of issues could lead to an otherwise general obligation for other infringements, since the level of monitoring has increased. This is not the intention of the proposed legislation – should [potentially] infringing content be accidentally discovered, it should not be for the intermediary to decide upon its status but rather should be an issue for the courts.

As discussed, §230 and the associated decisions hold that even with actual knowledge the intermediary has got immunity from liability.

In the EU, however, this level of protection does not apply, and the situation would place intermediaries between a rock and a hard place in the event of a duty as described. In order to prevent action, an intermediary would have to move more aggressively to remove content. On the other hand, this raises the inevitable prospect of false positives, for which there is no protection. A possible reason for the lack of this provision could have been an emphasis on free speech, where an intermediary is likely to err on the side of caution prior to removing content (Stalla-Bourdillon, 2012), the opposite approach taken by the American legislators. In the event of an obligation, this is a situation which would need to be changed.

However, as the US experience has arguably taught us, this goes against another possible purpose of the Directive, which is to promote self-regulation (Eecke, 2011). There, it is argued, since the intermediaries are protected whether they do intervene or whether they choose not to intervene, then they are still immune from liability. As such, given that it is cheaper *not* to intervene, then this is the better strategy for them, owing to the phenomena described in Chapter 3. This would therefore have to be a very narrow protection, exclusively for security purposes. Even this would have to have some form of limitation, as the potential to be blocked for no reason could have a severe impact on someone's business and livelihood.

5.8 Conclusions

This section has analysed the role that intermediaries would play in implementing a public health regulatory strategy. It was argued that hosting providers are best placed to perform this, on account of their being the cheapest cost avoider; and offering an effective solution which had a minimum impact on rights. The other stakeholders from Chapter 3 were considered, but although they might be in a good position to make an impact it was argued that they were not in such a good place as hosting providers. ISPs might have been able to make a significant difference if they were able to identify vulnerable or infected machines on their networks, but the restrictions on rights were significant in applying a quarantine style solution. A search engine has less in common with a public health strategy, but might be in a position to offer information to Web users.

Having shown that hosting providers were best placed, the regulatory framework in the EU was considered, in particular Articles 12 – 15 of the E-Commerce Directive and the Charter of Fundamental Rights. Case law in this area was discussed, and it was concluded that the E-Commerce Directive would be consistent with an obligation on hosting providers. However, the danger of losing immunity from liability under Articles 12 – 14 motivated a discussion of §230 of the Communications Decency Act 1996 in the USA. In particular with the current thinking about "neutral" or "passive", combined

with the danger of false positives, it was also concluded that any legislation should also offer protection to the hosting provider from these issues.

The next chapter will conduct a simulation to determine whether an obligation for hosting providers would be effective in the real world. Naturally, prior to implementing any new regulation, it is preferable to have some indication that it would have a reasonable prospect of making a significant difference to the overall level of malware.

Chapter 6

Simulation

> *Have you ever had a dream, Neo, that you were so sure was real? What if you were unable to wake from that dream? How would you know the difference between the dream world and the real world?* The Matrix, (1999)

In Chapter 5, it was established that hosting providers are best placed to solve the drive-by downloads issue. It was proposed that they offer vulnerability scanning services to their customers, and that they ensure that their customers have got the latest version of whatever software they are using. In this chapter, a simulation is performed to assess the efficacy of such an intervention by hosting providers in relation to minimising the global prevalence of drive-by downloads.

Despite being based on the ideas presented in the models described in 4.2.1, this is different in that it is an Agent Based Model (ABM), or more specifically, the simulation of an Agent Based Model. The compartmental, or Equation Based Models (EBM), have historically been used for epidemiological research, but increased computational power has led to a greater amount of ABM such as this (Rahmandad and Sterman, 2008). Chief amongst the arguments in favour of such a model is its ability to capture heterogeneous behaviour. EBM will assume that the population is homogenous and well mixed (Parunak et al., 1998), but that doesn't adequately capture position of the Web when considering the effect of more general malware as opposed to a single type which would target all systems. The distribution of the Web may approximate a scale free network (Barabási et al., 2000) also benefits from this, since it enables the agents (end-users) to select which of the websites they would choose to visit. Naturally, the trade-off is in efficiency (Rahmandad and Sterman, 2008; Cheng et al., 2011). It is not unheard of in the malware propagation field either, for example Hofmeyr et al. (2011) use ABM as opposed to EBM.

According to the three requirements for adopting an intervention – the efficacy (can it work?); the effectiveness (will it work?); and the efficiency (is it worth it?) (Haynes,

1999) – this is an assessment of the efficacy of the intervention. This will answer research question 4: "Can actions by a single country (or group of countries) have a statistically significant effect on the worldwide prevalence of infections from drive-by downloads?" The simulation will represent an intervention by a hosting provider of blocking a vulnerable or infected website a fixed period of time after the hosting provider notices.

6.1 Background

The primary advantage of modelling an intervention is that it is cheap (Hofmeyr et al., 2011), and consequently allows novel ideas to be investigated (Edwards et al., 2012). This is important, since it would take a considerable amount of both financial and political capital to address the issues of drive-by downloads without some evidence as to the likelihood of effectiveness, and the considerable cost which could be imposed on various stakeholders.

Another significant advantage, is that the heavy tailed or power law distribution of the Web (Barabási et al., 2000) means that a considerable degree of variance can be expected from any set of results, which can obscure the effects of experiments conducted in the real world (Edwards et al., 2012). As an example, consider in September 2014, a high profile flaw was discovered in the popular eBay website [1] which could have been used to hijack accounts. This is an atypical event, but with a significant user base this – or other events with this magnitude – could hide the effects of any intervention. By repeating simulations based on models, it is possible to get a far more accurate idea of the potential effectiveness of the intervention.

Finally, in this instance in particular, the aim is to analyse the global effect on the Web. On a local area network, it might be possible to segment the population to examine the effect, but with the Web examining a small subset of it will not make a difference to the overall criminal economy. In particular, if the high ranking websites are not included then the effect will likely be negligible.

As discussed in Section 4.2, there has been a considerable amount of literature using differential equations to model infection rates. Of these, however, very few have considered the issue of drive-by downloads or the more general effect of botnet interventions. Notable exceptions are Edwards et al. (2012), who looked at the effect of search engine intervention; and Hofmeyr et al. (2011) who looked at the effect of AS level intervention of traffic filtering. It is these two papers, as well as Kelley and Camp (2012) which the model presented in this chapter is based on.

Hofmeyr et al. (2011) modified the existing agent based model ASIM (Holme et al., 2008) to model the effect of Autonomous System (AS) level interventions in terms of

[1] https://www.ebay.co.uk. In November 2014, eBay were 22nd on the Alexa Web Information Services rankings

reducing the flow of "wicked" traffic. This model represented the Internet at the AS level, with the ASes acting as primitive economic agents who control traffic over a geographic network profiting from traffic flow through their network. The overall network grows, as agents expand where there is sufficient population, and the profit from the additional traffic is invested in expanding the network until the whole of the population is covered. To generate their network, it was left to grow on one occasion up until there were 10,000 ASes, and this was used on each occasion for each of their experiments. Of the customers, they set a *wickedness* level of 0.1, to signify that 10% of the population were participating in botnet activity, deliberately being slightly higher than the results obtained by Van Eeten et al. (2010).

The *wicked* traffic flows from AS A to AS B, their model provided five options for what each agent could do acting either unilaterally, or in co-ordination with other ASes – whether by country or by size:

- Do nothing;
- Reduce egress wickedness;
- Reduce ingress wickedness;
- Reduce egress and ingress wickedness;
- Blacklist wicked traffic sources. (Hofmeyr et al., 2011)

The difference between the work by Hofmeyr et al. (2011) and a lot of the other work, is that it models the whole of the Internet as opposed to a single network. This enabled them to examine the effect of the skewed distribution and view the effect of different ASes acting in concert. They found that only 0.2% of the biggest ASes were more effective than 30% of the ASes acting at random, as well as the diminishing returns demonstrated by increasing the amount of ASes acting in concert; and more aggressive filtering (with consequently increased traffic loss).

Looking directly at drive-by downloads, Edwards et al. (2012) considered the effect of an intervention by a search engine. Two scenarios were modelled: where the search engine blacklists malicious websites; or where the malicious website was "depreferenced" in the search engine rankings. The distribution of website popularity was obtained through calculating an exponent of a heavy tailed distribution from a random sample of websites obtained through the Alexa Web Information Services API, and compared it with a uniform distribution. This was modelled by reducing the reducing the traffic to the website to either 0 (for blacklisting), or by reducing the traffic according to a depreferencing parameter σ at time $t + \beta$ where β was the delay between the infection of the website and the detection by the search engine.

Kelley and Camp (2012), developed the work of Kephart and White (1991) developing a model with vigilance parameters. Kephart and White (1991) used the notion of a kill switch, where, upon recovery, an infected node would notify all its immediate neighbours to check themselves. Kelley and Camp (2012) extended this by having two interacting sub-populations, where a node could transition to the vigilant population depending on the global infection level η. These same nodes transition back to the non-vigilant population at constant rate δ. The probability of recovery is governed by social cues to recover, for example a browser reminder to update, so nodes responding to these parameters will not actively scan their systems unless they receive the reminder.

6.2 Description of the Simulation

6.2.1 Overall

This section presents a brief description of the simulation. Further detail concerning the roles of each different will be discussed below. In public health terms, there are two interacting populations (clients and servers[2]) each of which contains two sub-populations (vigilant and non-vigilant). Both populations follow an $SIRS$ pattern, which stands for ($Susceptible \rightarrow Infected \rightarrow Recovered \rightarrow Susceptible$). The website population is never really *Recovered*. A residual probability remains that they will become infected, although for ease of distinction between the other *Infected*, it shall be described as *Recovered*. The population of websites was set at 1,000 and the population of clients 100,000.

At the beginning of the simulation there are no infected nodes. The transition into an *Infected* state occurs with a fixed probability on each turn, representing the probability of a vulnerability being found in either the browser or the CMS depending on the population. The node may also recover with a fixed probability, which represents them patching their software and confers on them a temporary immunity. This probability depends upon whether the node is vigilant or non-vigilant, the probability of a vigilant node recovering is far higher than a non-vigilant node. Although described as *SIRS*, websites retain a residual probability of infection to represent extraneous factors such as losing credentials, although this is very low.

The simulation lasts for 150 turns, with the following events occurring on each turn:

1. Host scans website, and blocks them if appropriate
2. CMS vulnerabilities are discovered, and this changes the state of all websites using that CMS.

[2]The phrase clients is used interchangeably with users, and the phrase servers with websites.

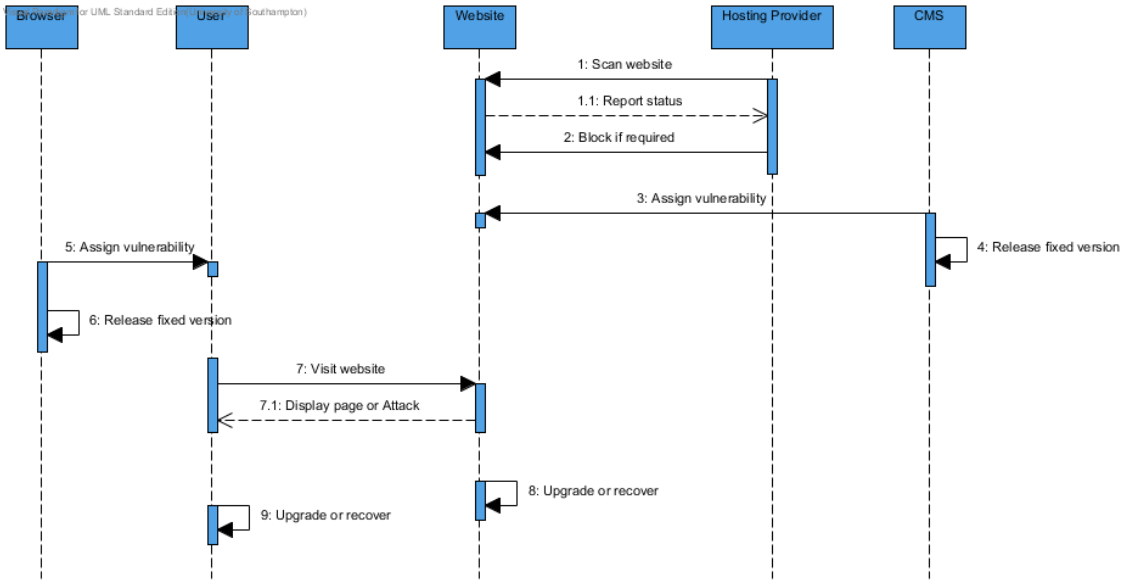

FIGURE 6.1: Sequence Diagram demonstrating the interactions for each turn.

3. Browser vulnerabilities are discovers, and this changes the state of all users with that browser.

4. Users visit a selection of websites, determined at random prior to the start of the simulation.

5. Users state changes if appropriate.

6. Websites state changes if appropriate.

A representation of this process using UML can be seen in Figure 6.1.

6.2.2 Hosting Providers

Since the goal of the simulation is to assess the effect of hosting providers in blocking vulnerable or infected websites on the overall population of infected users, the hosting providers are the most important agent in the model. It is accepted that there reasons why a hosting provider might choose not to implement such a procedure. For example, if the cost of implementing the infrastructure for this intervention is too expensive then it is less likely that they will be compliant. Similarly, if they think they avoid detection by the authorities, or the fine is inadequate deterrent, i.e. if $P(D) * F < C$ where D is detection, F is fine, and C is the cost of implementation (Ogus, 1994).

Prior to the beginning of each simulation, hosting providers are randomly assigned countries. In the event that their country is one of the countries which is implementing legislation to require that they block vulnerable or infected websites, they will be assigned a *Compliance* and *Effectiveness* value. *Compliance* is a Boolean value, which determines

whether or not the host will comply with the legislation in the host country. The *Effectiveness* value is a probability assigned uniformly across all hosts, which determines the probability on successfully noticing the vulnerable or infected state of the website.

Where the hosting provider has a country which is not participating, then compliance is always false (by definition, effectiveness is 0). As such, they play no role in the model, other than that websites which are assigned to them are not hosted by compliant hosts.

Where the host is compliant, they will assess each website on every turn. First, they will see whether the website is blocked. If it is, and the website has a status of *Recovered*, then they will unblock the website with 100% effectiveness. If not, then they will assess whether the website is vulnerable, with the effectiveness parameter. If it is, then the website will be regarded as "notified". Once the turns notified exceeds the specified model value, then if the website is still vulnerable, then the website will be blocked. If the website is infected, then it will be blocked on the same turn. This increases the chances of the website recovering to a clean state, regardless of whether they are vigilant or not.

When the host implements the blocking procedure, then that is regarded as completely successful – it is not possible for any users to visit the website.

Figure 6.2 represents this procedure in the form of an activity diagram.

A final point to note, is that as a simplifying assumption, there will not be any switching between hosting providers. Whilst it might be the case that website operators choose to switch to a cheaper provider, there is no evidence either way. For example, consider a hosting provider wishing to move to a more favourable jurisdiction. In the event that they had to move outside the developed world, then it would be necessary for them to invest in a considerable amount of infrastructure which would cost a considerable amount. Given that page loading speed is a part of Google's ranking algorithms (Google, 2010) as well as being important for usability, it is argued that this is a reasonable assumption to make.

6.2.3 Browsers and CMS

The focus of the simulation is the effect of patching. Both clients and websites have some form of software – either a browser (for users) or a CMS (for websites). Each piece of software has two versions: a *Current* version and an *Old* version. An old version is regarded as vulnerable, and places the agent in question into a *Susceptible* state. There is a fixed probability of a new vulnerability being discovered on each turn, and every agent with that product will then enter a *Susceptible* state, i.e. their browser/CMS will be regarded as *Old*. This does not apply to websites who do not have a CMS, who will be

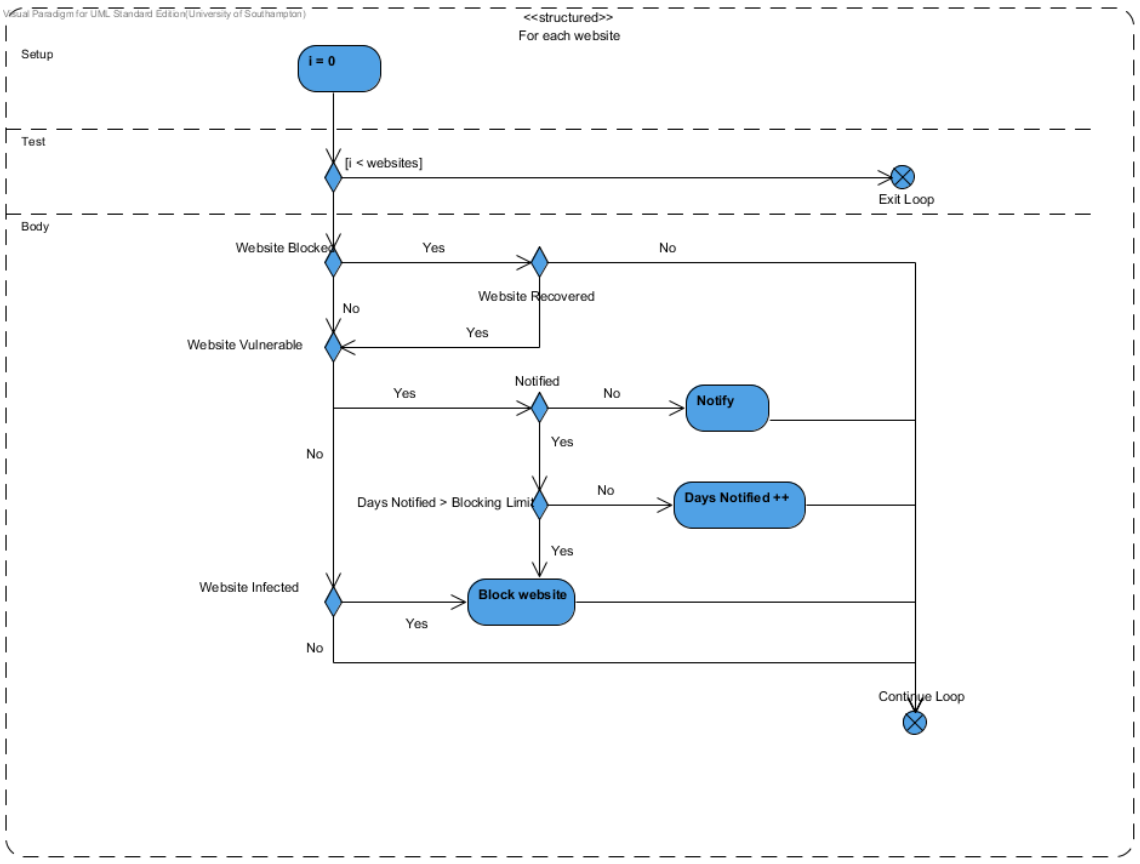

FIGURE 6.2: Activity diagram for hosting providers, with in the simulation

regarded as *Infected* until they separately become vulnerable at a far lower probability[3]. This is assuming the premise that an attacker is far less likely to attack a website with custom software, because the effort expended is the same as attacking a CMS for a smaller reward (i.e. only a single vulnerable website as opposed to thousands).

The environment is not homogenous, i.e. both browsers and CMSes are divided into separate groups to represent the variety of different products which exist. A vulnerability discovered against one browser or CMS will only affect agents with that particular software, any others will remain in a *Recovered* state until their software becomes vulnerable.

Software only remains in a *Vulnerable* state for one turn. At the beginning of the next turn, the vulnerability is considered to be fixed, and any agents who then choose to upgrade return to a *Recovered* state. A new *Current* version is only released in response to a new vulnerability being discovered, so any agent on an *Old* installation is *Susceptible*.

[3]In this instance, a CMS is representative of all software on the server, including e.g. the database

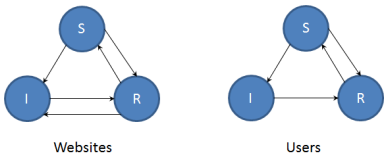

FIGURE 6.3: State transitions for websites and users. The only difference is that a user cannot transition from $R \rightarrow I$

6.2.4 Users and Websites

The probability of either a user or a website updating their browser/CMS is dependent on the sub-population that they are in. Both users and websites are sub-divided into *Vigilant* and *Non-vigilant* populations. Both populations are constant, and there is no transition between these populations. Each member within the population is uniformly likely to recover from either a *Susceptible* or *Infected* state. Prior to the beginning of the model, this will also impact the software that they will have – *Non-vigilant* members will start with an *Old* version, and *Vigilant* members will start with the current version. Prior to the simulation, each user will select websites which they will visit every turn. The selection is not uniform, but rather is based on the website's popularity parameter. Having selected these, the user will visit the same websites on each turn (unless they are blocked).

Both user and website become susceptible from their software becoming *Old*. If the website is *Infected*, and the user *Susceptible*, then the user has a fixed probability of entering an *Infected* state. *Vigilant* users and websites have a high probability of upgrading immediately after their software becomes vulnerable. For users, this represents an automatic update such as with Patch Tuesday; for websites this represents a notification from their hosting provider or their CMS configuration which they then [usually] manually install[4]. *Non-vigilant* users will generally only upgrade in response to an infection, likely something which directly impacts them such as a Cryptolocker or ZeuS infection. A *Non-vigilant* website will usually upgrade only in response to entering an *Infected* state, or being blocked by a hosting provider.

Users and websites both being a subclass of Agent, the transition between states is similar (see Figure 6.3). The only difference, is that there is a residual possibility that a website could become infected even though they are not susceptible. This was not modelled for the user, because the simulation is only concerned about users falling victim to drive-by downloads. The manner in which the website got infected is not important for these purposes. The infection process requires a susceptible user to visit an infected user, where they will become infected according to a certain probability. This process is described in more detail in Figure 6.4.

[4]The difference between the automatic updates, such as those done by WordPress, and other manual updates is not considered.

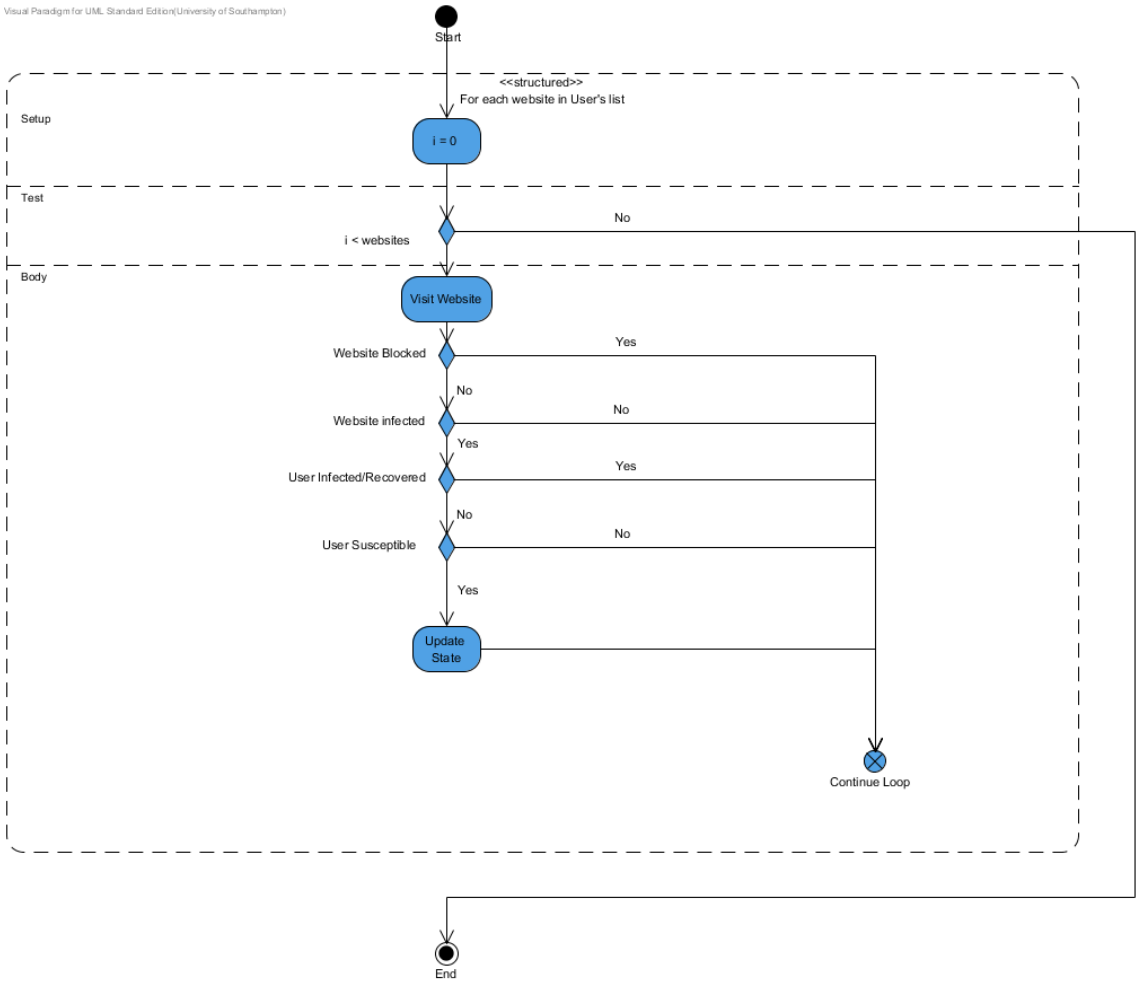

FIGURE 6.4: Activity diagram representing the User infection process.

6.2.5 Attackers

Although attacks on websites, and the consequent attacks on users were a part of the model, the attackers themselves were not modelled as agents. This was because, given their nature, data about attackers was difficult to come by, so the role of the attacker in the model was exclusively limited to the fixed probabilities for new vulnerabilities for pieces of software; and infection of websites.

6.3 Selection of Parameters and Implementation

6.3.1 Randomness

The model is probabilistic in nature, and relies upon a series of Bernoulli decisions to determine the outcome. The Bernoulli function takes the parameters discussed in this section and compared against the generation of a pseudo-random number. Where the

number generated is less than the parameter, then the outcome is true. The Bernoulli function can be seen in the code listing:

```
public static Boolean SetBoolByProbability(Double probability)
{
       //Randy is a single global instance of a Random object
       return Globals.Randy.NextDouble() <= probability;
}
```

Listing 6.1: Bernoulli function as it appears in the C# code

The Microsoft .NET `System.Random` class was used to generate numbers [5]. This pseudo-random number generator is based on Donald Knuth's subtractive random number generator algorithm, which, although it contains a bug which impacts the level of randomness (Arazi, 2011), (that has been acknowledged by Microsoft (Ytosa, 2011)) it likely still contains sufficient randomness for the purposes of this simulation.

The random class is seeded by the value of ticks since the system was started by default, or it is possible to manually seed the class. Two instances of the random number were required. To generate the network, the same random number was used every time. This number chosen was 1233077111, which was the tick count from the first time the simulation was run. Following the generation of the network, a new random instance was generated for the decisions to be made within the simulation. For this, the seed was unique on each occasion. Because there were many simulations being run in parallel using different processes, it was necessary to check for duplicate values. As such, a uniqueness constraint was placed on the Seed column in the Simulation table in the database. Any simulation which attempted to insert a value into the database would recover and sleep for a random amount of time from a new `Random` instance.

On each occasion that a distribution needed to be calculated, where possible this was done using empirical data. The distributions were generated at random at the beginning of the simulation, by using the cumulative distribution values as keys for a dictionary. Where the generated random number appeared between the two figures then this key was chosen. The code for this can be seen in Listing 6.2:

```
public static double GenerateObjectIndex(double randy, List<double> probabilities)
{
    //Returns the key to choose a certain type of object in a dictionary <double,Object>
    //Assumes that probabilities[0] is the upper bound of the first range
    //If rnd is between the two bounds then return the upper one
    int i = 0;
    double lower, upper = 0;
    probabilities.Sort();
    foreach (double k in probabilities)
    {
        lower = i == 0 ? 0 : upper;
        upper = k;

        if (randy >= lower && randy <= upper)
```

[5] http://referencesource.microsoft.com/#mscorlib/system/random.cs

```
            {
                return k;
            }
            i++;

        }
        //For distributions with "other" then this will be the final key
        return 1.0;
    }
```

LISTING 6.2: Code for assigning distributions based on empirical data

There are many things for which there is no data, or no way to validate whether the data are an accurate approximation of the real world. That said, empirical data exists for distributions of browsers; CMSes, hosting providers, and countries. Regardless, it is hoped that the data are an accurate representation of the world, and that use could be drawn from it by tweaking the parameters.

6.3.2 Hosting Provider Population

Like with many phenomena associated with the Web, hosting providers follow a heavy tailed distribution both in terms of which country they reside in as well as their market shares (See Figures 6.5 to 6.7). As such, the distribution was selected according to both country and according to the market share within the country. The data for the hosting providers was obtained by running a Python script to scrape data from the `http://www.webhosting.info/` website, (see Appendix B.1). Aside from a small amount of exceptions, the website provided information for only the top 25 hosting providers. Given the distributions, it is unlikely that operators any smaller than the top 25 identified by `http://webhosting.info` would be involved in any legislative initiative.

6.3.3 Browser and CMS

Obtaining accurate data for browser market share is very difficult. An obvious solution has been to parse the `User-Agent` strings sent by the browsers as they visit certain Web pages, although this tends to be biased based on the demographic visiting the websites. Similarly, the `User-Agent` string has undergone many variations. The release of Internet Explorer 11, for example deliberately aimed to obscure the fact that it is Internet Explorer (Microsoft, 2013)[6]. There is nothing stopping impersonation of other `User-Agent` strings either. Analysis was conducted of an Apache access.log file of a major website, which revealed Internet Explorer 6 to be the browser with the highest market share! This seemed unlikely, so the data were discarded.

[6]The history of the User-Agent string in general also has a convoluted history as browsers attempt to impersonate each other, for a light-hearted take on this, see `http://webaim.org/blog/user-agent-string-history/`

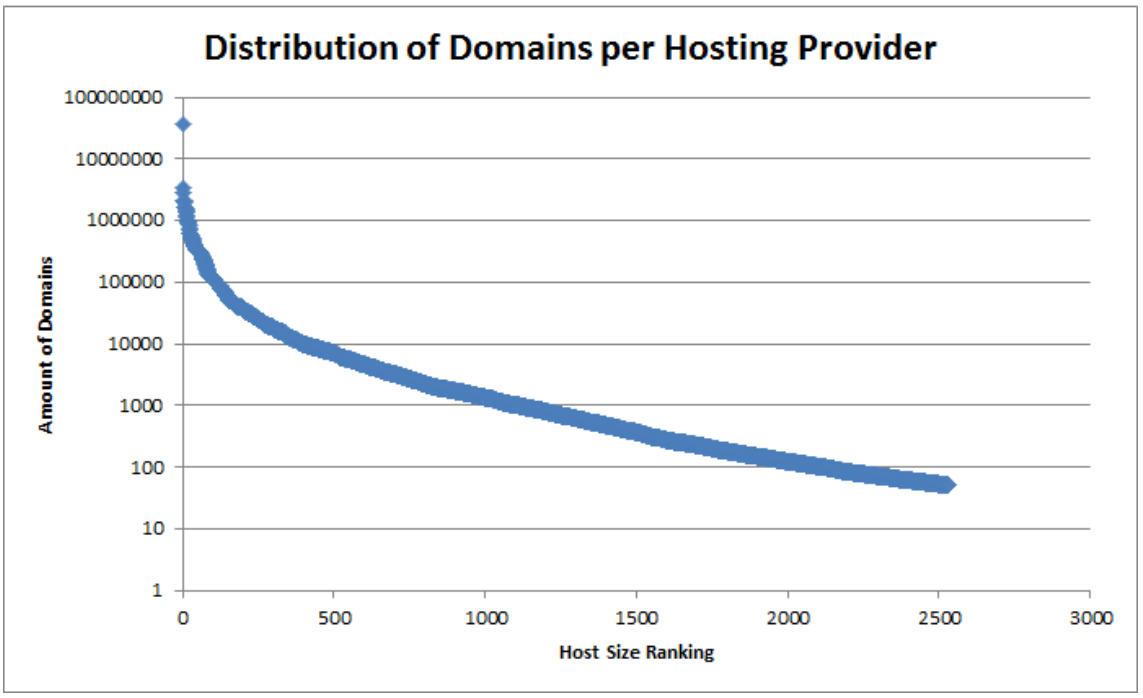

FIGURE 6.5: Graph demonstrating the worldwide distribution of domains per hosting provider (log scale)

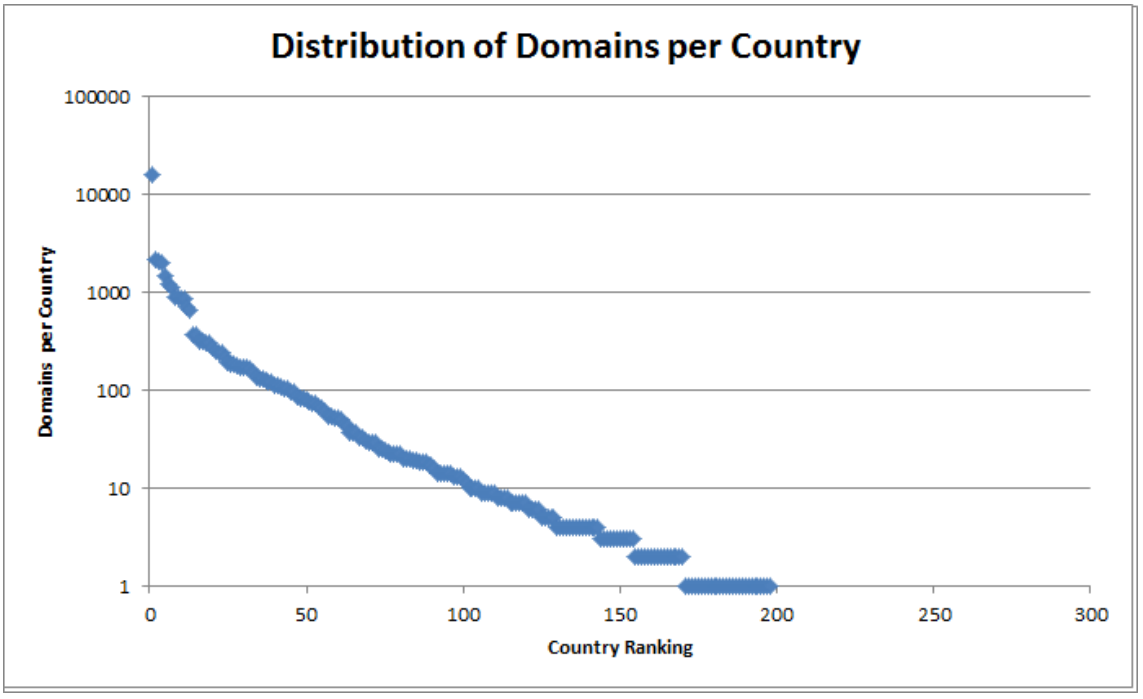

FIGURE 6.6: Graph demonstrating the distribution of domains per country (log scale)

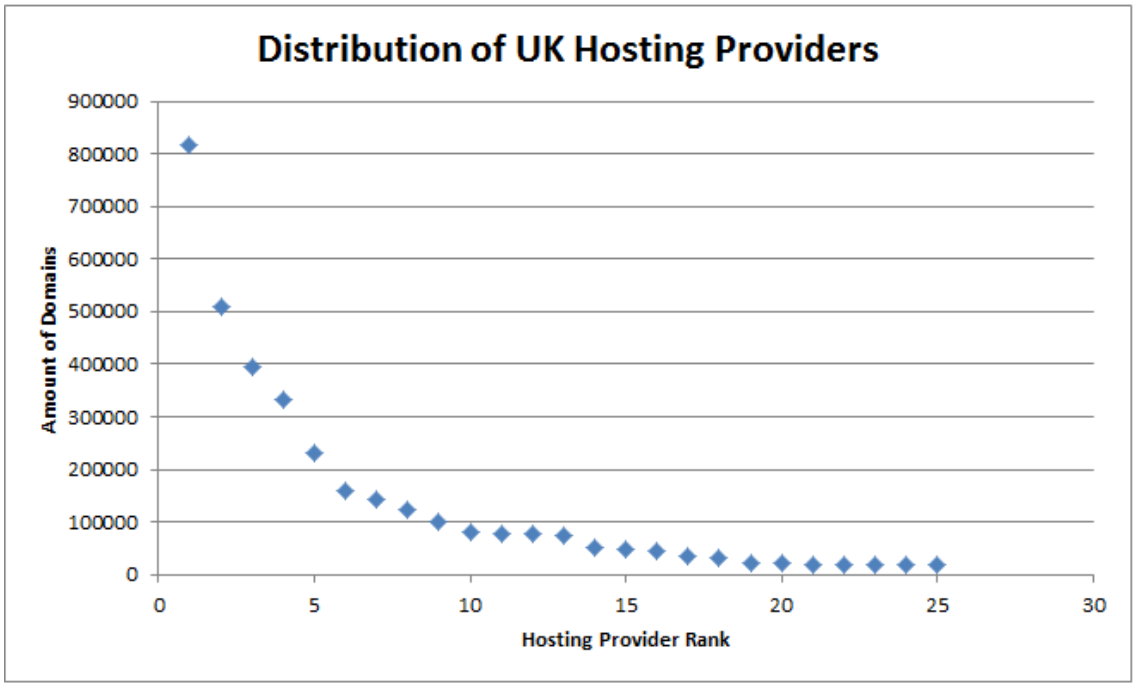

FIGURE 6.7: Graph demonstrating the distribution of domains for the UK

The "modern" approach is to user feature detection, so that customised content can be displayed to different browsers. This can be made more difficult by the fact that it is possible to "fingerprint" browsers so as to uniquely identify a particular user. This is a particular concern in regards to potential lack of privacy (Eckersley, 2010), and there have been some efforts by some browsers to hide as much version specific information as possible[7].

Differences in methodology or methods of counting mean that the distributions of browsers are likely to remain different. For example, it is possible that Chrome is over-represented due to its use of pre-rendering pages (Komoroske, 2011). Accepting these limitations, an effort was made to determine the distribution of different browsers from the Wikimedia Traffic Analysis Report (Wikimedia, 2014). From these data, the market share of the top five [desktop] browsers (those with a market share of over 1%) were extracted. To determine the initial probability of the user's browser being up to date, the latest two versions were counted as being up to date, due to new releases occurring in the middle of the month. Using Firefox and Chrome, the average appeared at approximately 85%. It is not possible to determine for Internet Explorer, since versions are supported on some platforms but not others, so an "old" version might in fact be fully supported. However, both Firefox and Internet Explorer have regular patch cycles, so the assumption will be that of the Internet Explorer browsers detected will be similar.

The CMS population is modelled over a distribution of the CMS population, taken from `w3techs.com` (W3techs.com, 2015). The market share of each CMS was determined

[7] See e.g. `https://wiki.mozilla.org/Fingerprinting`.

through scraping the website, and then each website in the model is assigned a CMS based on that. Within the distribution, it was considered that at the beginning of the model only 36.9% of the population used a CMS, so a CMS called "No CMS" was also added to account for 63.1% of the population. The amount of the websites which were regarded as having the latest version was determined by taking the mean of the values of the latest versions for WordPress and Drupal, which was 35%. The script to obtain this data can be seen in Appendix B.2.

The proportion of vigilant users was determined by the amount of browsers using the latest version (0.85), and the vigilant websites by the amount with the latest version of the CMS (0.35)

6.3.4 Vulnerabilities of Browsers and CMS

In the event that a piece of software becomes vulnerable, it is modelled by iterating through every agent using that software and updating their state to Vulnerable until such time as they choose to update. The frequency with which this happened was determined by analysing the frequency of CVEs (potentially) allowing RCE against popular browsers within a similar time frame. Data was extracted from the National Vulnerability Database from the three major browsers: Chrome, Internet Explorer, and Firefox for 2013. 2013 was chosen, because it was the last full year at the time the model was initially created, but also because the NVD changed the format of the XML towards the end of 2014 and it could not be so easily extracted. The mean of the three amounted to approximately 0.2. Internet Explorer had a slightly higher amount, probably due to the amount of different versions supported by Microsoft, and Chrome had very few. Together, that appeared roughly accurate.

The same methodology was attempted to determine the probability of a RCE vulnerability for server side software, although this turned out to be zero. By extending the period to search for, and lowering the requirements for the severity of the vulnerability, some results were returned, which was put at 0.003. There is less of a need to have RCE for server side attacks, since the goal is not necessarily to take over the entire server. A XSS attack can be used to embed malicious content for example, yet does not allow RCE on the server.

The script used to obtain the data can be seen in Appendix B.3.

6.3.5 Visiting a Website and Infection Probability

Every turn, a user will visit each website, with a probability from the ReachPerMillion value of the website, extracted according to the Alexa API[8]. From the population, the

[8] http://aws.amazon.com/awis/

Parameter	***Value***
Vigilant User Population	0.8
Vigilant Website Population	0.4
Vulnerable Website Infection	0.1
Vigilant User Recovery	0.9
Non Vigilant User Recovery	0.01
Vigilant Website Recovery	0.5
Blacklisted Website Recovery	0.8
Non Vigilant Website Recovery	0.1
Browser Vulnerability Discovered	0.1
Browser Attack Success	0.5
CMS Vulnerability Discovered	0.003
CMS Attack Success	0.4

TABLE 6.1: Default values for model parameters

top website is google.com, which would expect 435,000 people per million to visit it on any given day, so the parameter for visiting that website would be 0.435. Edwards et al. (2012) used the Alexa API as well, and showed it to have a power law distribution, and they also had a condition where the visits of websites was over a uniform distribution. Their aim was to look at the effects on variance in using the different distributions, but the aim for this simulation is to make use of the dynamics of the unequal distribution so as to require as few operators as possible to need to participate in an intervention.

6.3.6 Summary of Default Parameters

A summary of the values of the parameters which were used in the default condition of the simulation can be seen in Table **??**.

6.4 Results

The complexity of the simulation mean that there are many potential sets of results which could be extracted from the simulation. The majority of these have not been considered, because the aim of the simulation is to assess research question 4: "*Can actions by a single country, or group of countries, have a statistically significant effect on the worldwide prevalence of infections from drive-by downloads?*" It is only regarded as an exploratory, initial set of results to assess the efficacy of an intervention by a hosting provider. This section will be presented as follows.

Firstly, an examination of the default condition will be conducted, and this will be compared to the state of the real world. Following this, the relationship between the two populations will be discussed, in order to show that just because a large amount of websites are infected, it does not follow that a large amount of users will also be infected.

Vigilant Users	Non Vigilant Users	Vigilant Websites	Non Vigilant Websites
80,069	19,931	374	626

TABLE 6.2: The amount of each sub-population used in the simulations

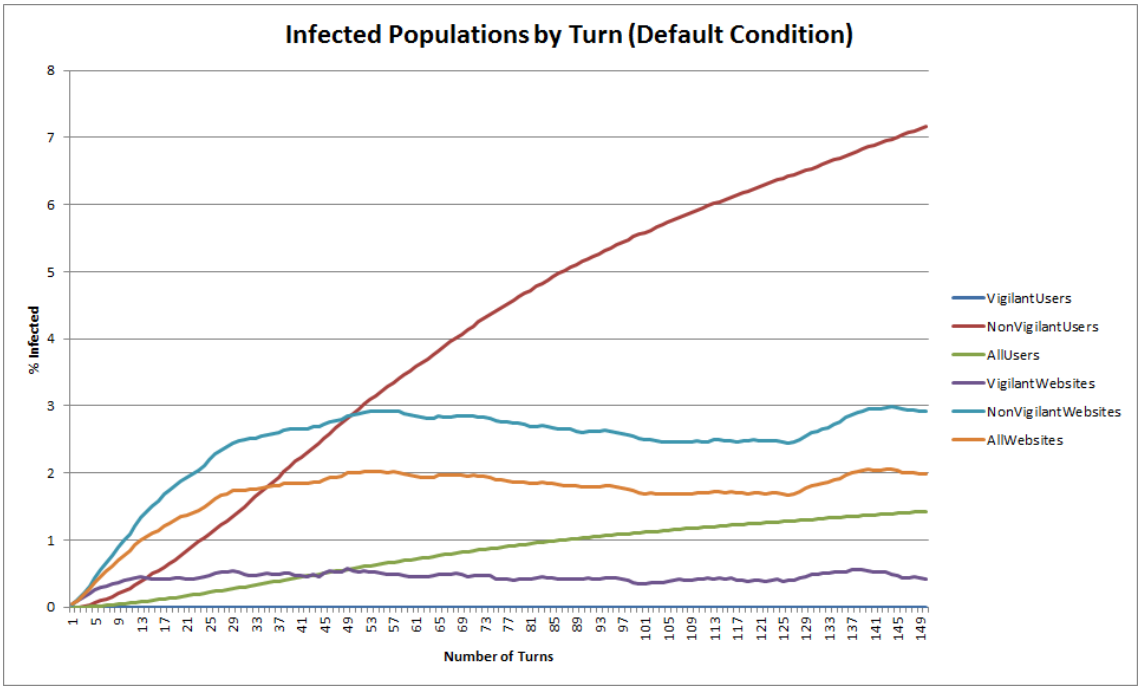

FIGURE 6.8: Turn infected websites v turn users infected

The efficacy of interventions by different states will be considered, at both ideal levels (100% compliance) and at other lower levels of compliance. The effects of modifying other parameters was also tested, in particular the effect of changing vigilance. At the end, some of the limitations of the simulation are considered, and the implications for policies discussed.

Each separate condition for the simulation was run over 150 turns, with 1,000 trials. The primary metric which was considered was the amount of the percentage of non-vigilant infected users at the end of the simulation. Although vigilant users can make a difference to the overall level given that they make up approximately 80% of the population by default (see Table 6.2).

Given the parameters described in Table 6.1, the probability of any one of this population becoming infected is incredibly unlikely. This does not mean, however, that there is no effect on the overall level of infected users. This can be seen in Figure 6.8, which tracks the percentage of the infected percentage over time (mean taken over 1,000 runs). The percentage of infected vigilant users is very close to zero (0.009), but the population is large enough, that it has an effect by the end, see table 6.4, where 22.8% of the total infected users are vigilant. An outline of the data for both absolute and percentage values can be seen in Tables 6.3 and 6.4.

	Mean	Median	Std. Deviation	Variance
Vigilant Users	0.01	0	0.0570	0.003
Non Vigilant Users	7.159	4.650	6.853	46.959
Vigilant Websites	0.428	0	1.13	1.278
Non Vigilant Websites	2.5	0.8	4.837	23.396

TABLE 6.3: Measures of central tendency of populations infected at 150 turns (%)

	Mean	Median	Std. Deviation	Variance
Vigilant Users	440.61	5	2734.059	7475077.809
Non Vigilant Users	1489.48	1029.5	1368.792	1873590.698
Vigilant Websites	125.81	89.5	101.061	10213.227
Non Vigilant Websites	246.75	167	205.572	42259.667

TABLE 6.4: Measures of Central Tendency of populations total amount infected

There have been few academic papers detailing the prevalence of drive-by downloads in the last couple of years, so the likes of Stone-Gross et al. (2011); Provos et al. (2007, 2008) are among the most recent. Given the speed at which security trends appear to change, these cannot be relied upon as an up to date source, so it was necessary to use reports by security companies which are more recent. Each of these are unreliable to a certain extent, because they have an inherent conflict of interest since they are selling products which claim to reduce the risks associated with insecurity. In addition, in many cases the methodology was not published either. A particular weakness of many of these, could also be the sample which is selected. For the client side, the results are limited to people who have actively chosen to install their security product; and agree to data sharing – likely characteristics of a vigilant population.

Some of the reports discussed in Section 2.2.1, WhiteHat Security report that some 86% of websites have a "serious" vulnerability (WhiteHat Security, 2013), and Checkmarx (2014) demonstrate a worrying level of insecurity in popular WordPress plugins. Although detailing the specific vulnerabilities which exist, and some overall figures about how much they protected, there are very few which provide information in a form which can be compared to. The Microsoft Security Intelligence Report Volume 17 offers some statistics for the amount of both client and server side (Batchelder et al., 2014). In the first half of 2014 (the latest statistics currently available) the report claims that there are 5.8 phishing websites per 1,000 hosts, and 12.1 hosting websites per 1,000 hosts.

This equates to 0.179% of websites which have some form of malicious content (whether malware or phishing). This does not conform strictly to drive-by downloads, but given that these are things which an attacker may choose to do after successfully compromising a website (recall the discussion in (Canali et al., 2013b)), this can be considered a useful guide for the simulation. The results in Table 6.3 show that when applied to the whole population of websites, then the level has an average of 1.66%. The difference

between the two is comfortably within a single standard deviation from the mean, and so represents a good approximation.

For client side statistics, the Microsoft report is used again. The relevant figure here is their Computers Cleaned per Mille (CCM), which is "an infection rate metric that is defined as the number of computers cleaned for every 1,000 unique computers that run the Malicious Software Removal Tool (MSRT)" (Batchelder et al., 2014). For the first half of 2014, the mean worldwide CCM was 9.0 (10.8 in Q1, 7.2 in Q2), which equates to 0.9%. Once again, taking the means of the two populations, the percentage of the population infected becomes 0.7156% which is, once again reasonably close. The figures which indicate incredibly high infection rates, e.g. Panda Security (2014), likely equates to the "encounter rate" in the Microsoft report, which is "the percentage of computers running Microsoft real-time security products that report a malware encounter" (Batchelder et al., 2014).

6.4.1 Relationships Between the Populations

The distribution of website reach appears to follow a heavy tailed distribution (See Figure 6.9). This is a documented phenomenon, Edwards et al. (2012) used a random sample of data from the Alexa rankings, and calculated a power law with exponent of $\gamma = -1.4$, which led to high degrees of variance in their experimental results – as they predicted.

Therefore, it was hypothesised that the mere fact that a lot of websites were infected did not follow that the next turn more users would get infected. The correlation for three different factors were assessed to test this effect:

- The websites infected on one turn v users infected on the same turn;
- Websites currently infected v users infected next turn;
- Total amount of websites infected v total amount of users infected.

The results from the three correlations can be summarised as follows, and the graphs can be seen in Figures 6.10, 6.11 and 6.12.

1. There was no evidence of a correlation between the amount of websites infected on a particular turn, and the amount of new users being infected next turn, $r = .057[.051, .064]$.
2. There was a moderate positive correlation between websites currently infected with the amount of users infected on the next turn, $r = .559[.555, .559]$.

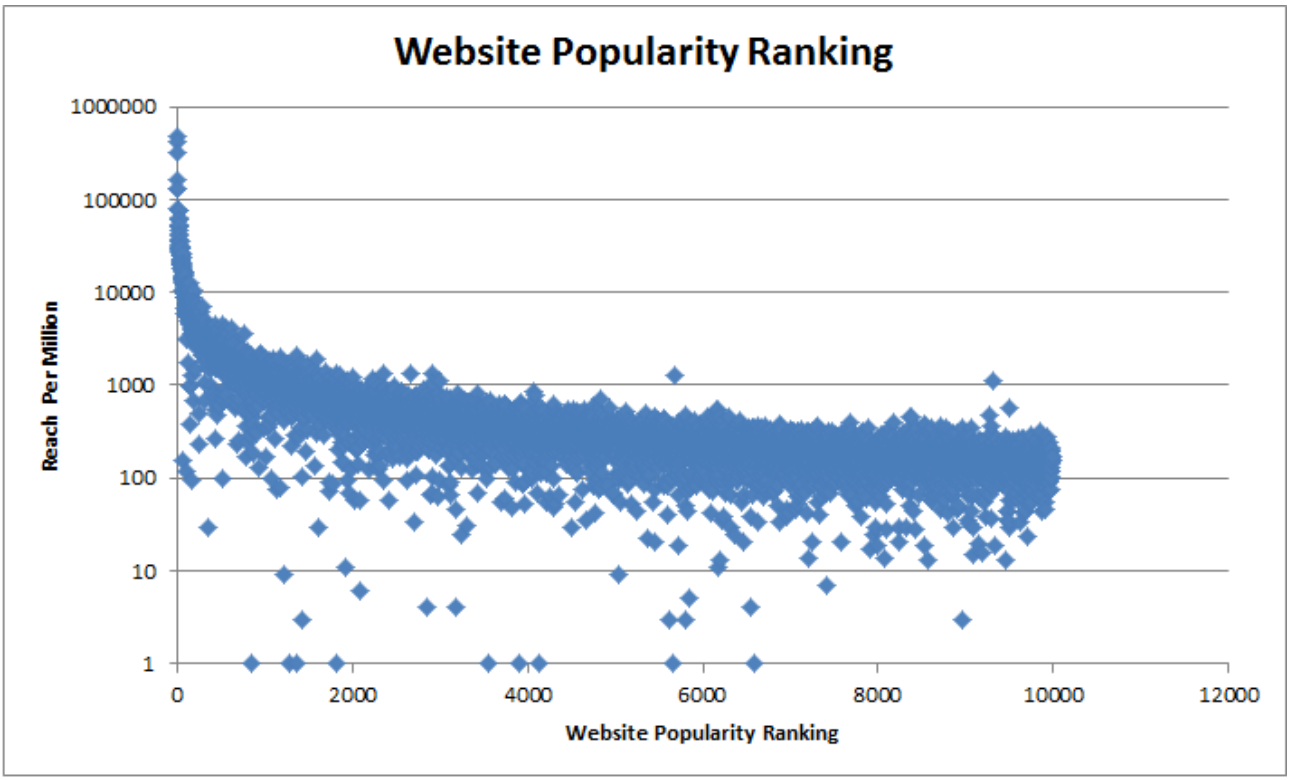

FIGURE 6.9: Ranking of Reach of the top 10,000 Websites

3. There was a strong positive correlation between the total websites infected and the total users infected $r = .862[.843, .879]$.

These results were mixed in terms of determining the degree of the relationship between website and user infections. The hypothesis could not be completely confirmed, since the results demonstrate that there is at least some relationship between the amount of websites infected and users infected. This can be explained by the fact that where large groups of websites are infected the probability is that at least one of them will be one of the popular ones. The lack of complete relationship in all cases can be seen from the amount of times zero users were infected at all, despite the fact that there were many websites infected. Similarly, even where the strong correlation exists, there are many outliers present.

6.4.2 Effects of intervention

Having established details relating to the "default" condition, the efficacy of various countries unilaterally adopting an intervention, which was enforced with a 100% compliance level was tested. Four different conditions were tested, interventions by the UK, EU, USA and World (i.e. all countries). This led to the following hypotheses being proposed:

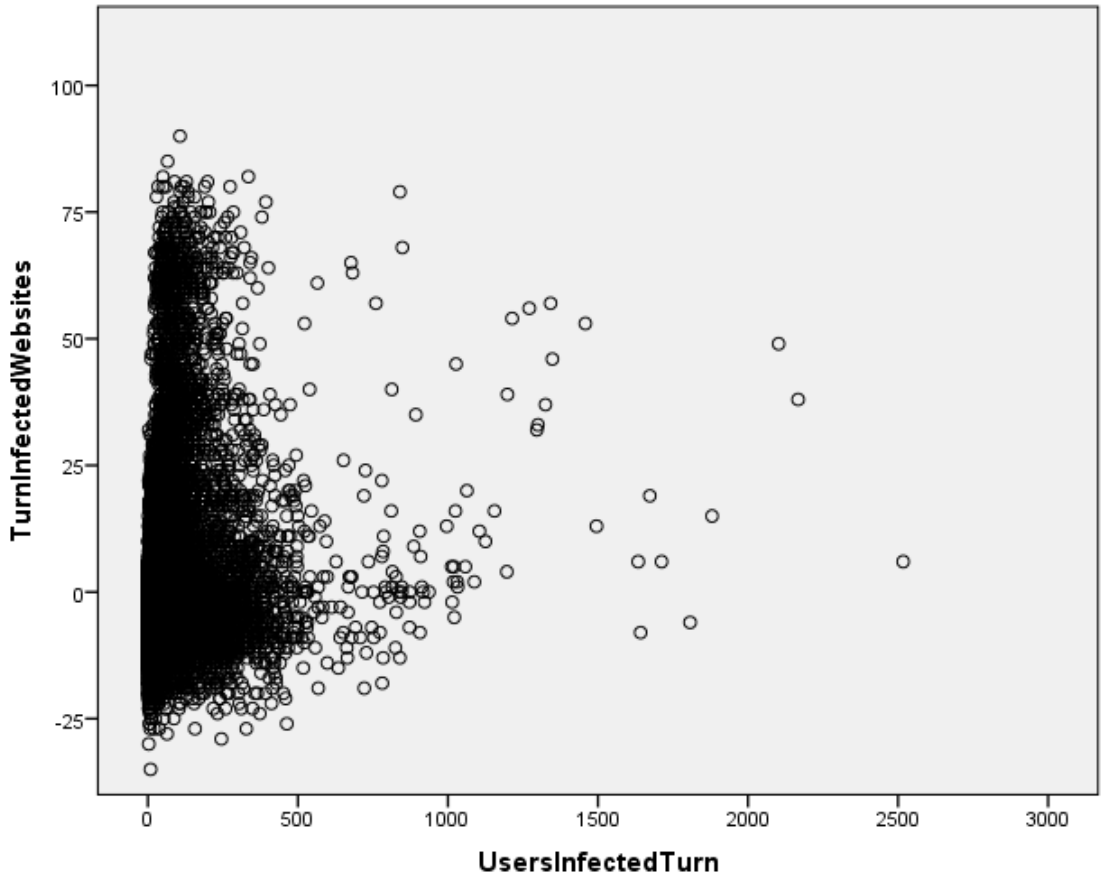

FIGURE 6.10: Turn infected websites v turn users infected

Condition	**Default**	**UK**	**EU**	**USA**	**World**
Non Vigilant Users Infected(%)	7.159	7.216	5.899	6.667	3.480

TABLE 6.5: % non-vigilant users infected at 150 turns, interventions with 100% compliance

1. Intervention by all countries in the world would have a significant effect compared to the default condition;
2. Intervention by the USA would have the same effect as the combined EU countries;
3. Intervention by either the EU or the USA would have a significant effect compared to no intervention;
4. Intervention by the UK would not be significant.

In order to assess these hypotheses, a one-way ANOVA test was performed between these five conditions, and there was a slight effect between intervention and the amount of non-vigilant infected users, $F(4, 4995) = 66.996, p < .001$,. The means can be seen in Table 6.5.

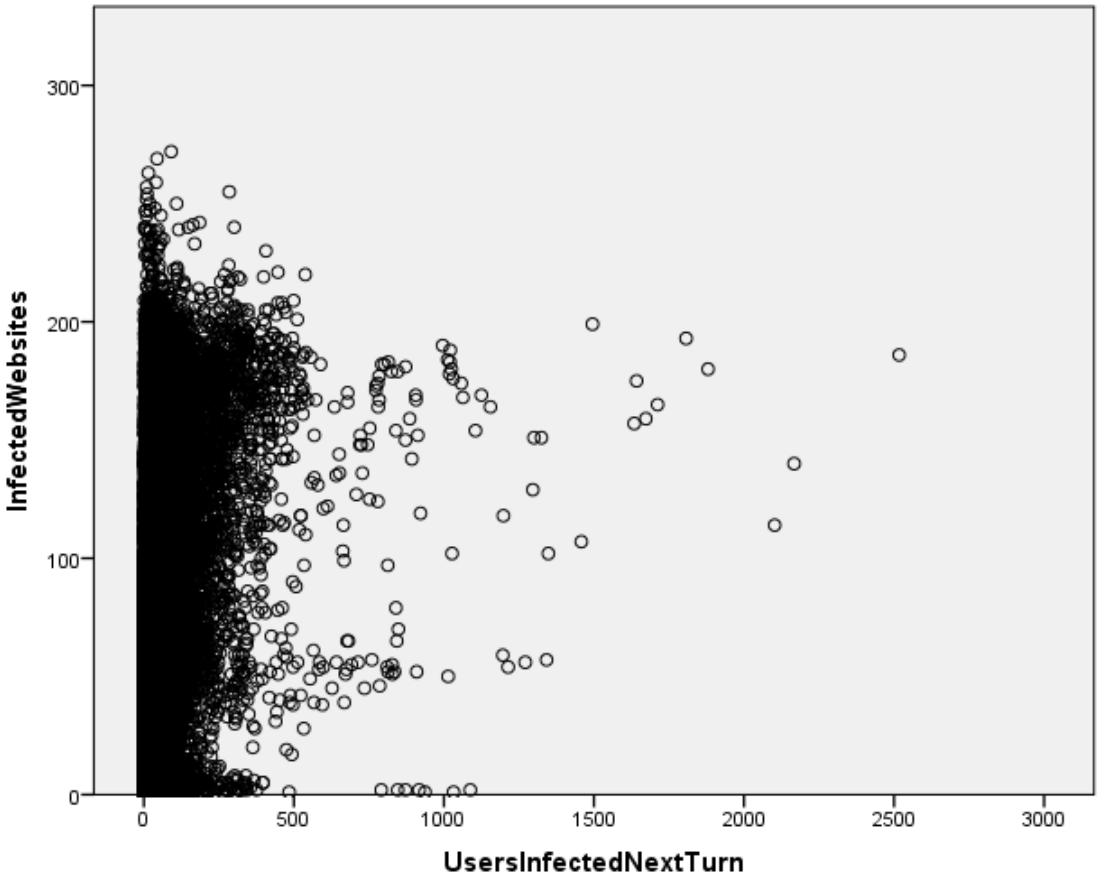

FIGURE 6.11: Amount of websites currently infected v amount of users infected next turn

Planned contrasts demonstrated that the World condition had an effect, with $t(3444.320) = 20.421, p < .001$, and that there was a difference between the EU and the USA $t(1976.326) = 2.8, p = .004$, so hypothesis 1 can be confirmed.

Games-Howell *post hoc* tests were conducted on the results, because homogeneity of variance could not be guaranteed ($p < .001$). Intervention by the UK was not significant ($p = 1.0$), or the USA ($p = .454$), but there was a significant difference in the EU condition ($p < .001$). This leads to Hypothesis 2 being rejected, and Hypothesis 3 being partly rejected. A further ANOVA test was run for the UK, where the notification period (i.e. the time between the host discovering the vulnerability, and then them blocking the website) which showed no significant difference with a shorter blocking period ($F(4, 4995), p = .261$), so Hypothesis 4 was confirmed.

6.4.2.1 Levels of Compliance

Whilst in an ideal scenario, every country which participated in the intervention would manage to get 100% compliance, unfortunately this is unlikely to occur in the real world. In order to assess the real world efficacy of the intervention, further tests were

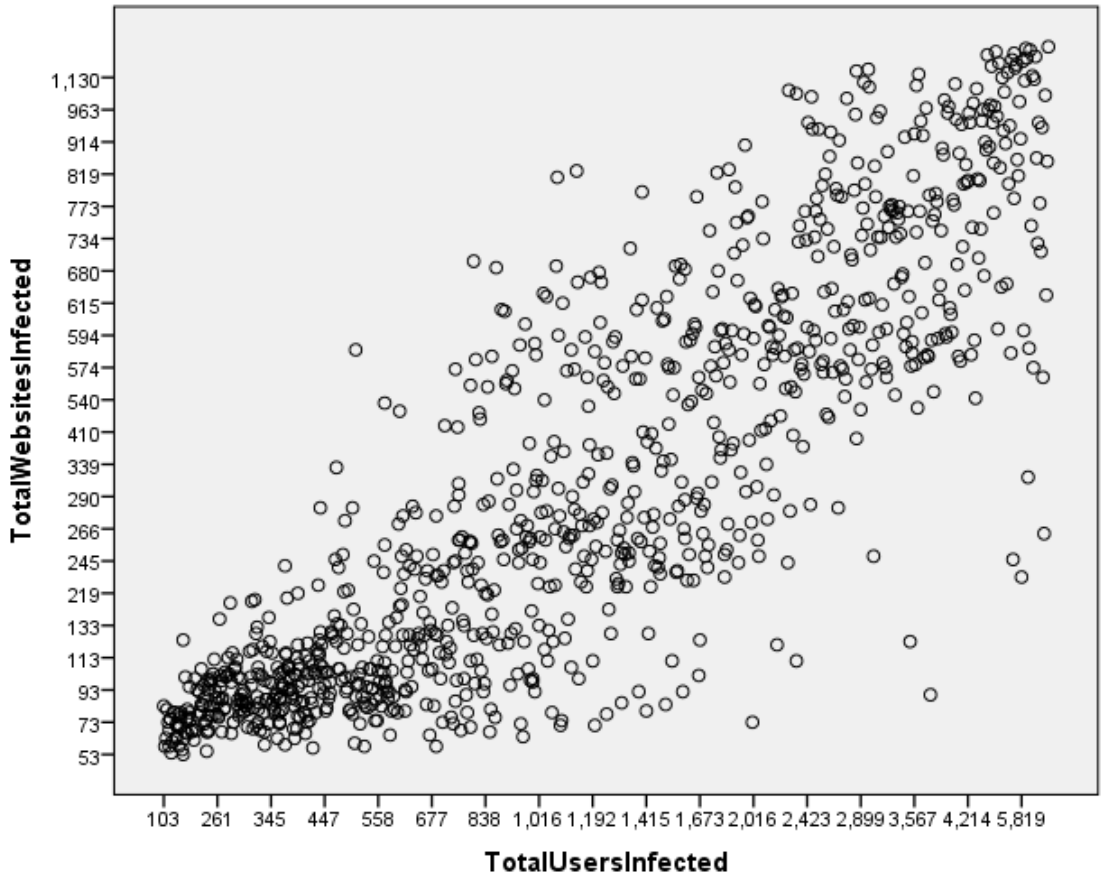

FIGURE 6.12: Total amount of websites infected v total amount of users infected

run in order to assess the performance under conditions of partial compliance. The two countries chosen for the intervention was the USA and the EU[9]. These were chosen since they both had a comparable share of the hosting providers, and both would require only one body to make the main legislation.

In addition, the analysis from the previous section had a surprising result that intervention by the USA would not be significant where the EU would. It was decided to investigate more closely the similarities between these two countries to get a better idea about whether the USA 100% condition was anomalous. Another potential difference which was considered worth exploring was the different distribution of website operators. Since the USA has the biggest operator by some margin, in a scenario where compliance is set at 80% then there will be occasions where this host is non-compliant. Given the worldwide market share (18.3%), it could conceivably have a greater effect. As such, the following hypotheses were proposed:

1. Any level of intervention will have an effect compared to no intervention;

[9]The EU is a supranational organisation rather than a country. However, given the legitimacy afforded to EU institutions to make law on behalf of all member states, the effect is arguably the same, and will be referred to as a country for the rest of this report. The effect of different countries within the EU having a different compliance level might be looked at in future work

2. Higher levels of compliance will have a greater effect than lower compliance;
3. As the compliance level gets higher, the difference will be less pronounced.
4. As the compliance levels get lower, intervention by the USA will have a smaller effect.

There was a significant effect of the level of compliance on the amount of infected non-vigilant machines at the end of the simulation, $F(10, 10989) = 5.819, p < .001$. Bonferroni *post hoc* tests were performed on compliance levels which indicated that compliance levels of 20%, 40% and 60% were not significant ($p = 1$ for 20% and 40%, and $p = .351$ for 60%) so Hypothesis 1 was rejected. Compliance levels of 80% and 100% were both significant (all ps $< .05$).

It can also be seen that higher levels have an effect compared to the lower levels of 20% and 40% (all ps $< .001$ for 80% and 100% compared to 0 – 40%). However, this does not continue, and no significant difference was found between 60% compliance and 100% compliance, (all ps $= 1.000$). This means that Hypothesis 2 can be confirmed, and Hypothesis 3 can be rejected.

In general, there was no significant difference between the countries either ($p = .944$), which suggests that the result for 100% compliance might have been an anomaly. It is more to be expected that the two countries would be more similar, and that the USA might have a better chance of success, given its higher market share.

There was no significant main effect between the countries ($p = .315$), which means that the rejection of Hypothesis 2 in the previous test should be treated with caution. Figure 6.13 also demonstrates the similarities between the two conditions, and shows some of the surprising results. The change between 80% and 100% compliance for the USA leads to a greater amount of infected hosts, which was unexpected, although the difference is not significant.

Despite expectations that the lower compliance rate would affect the USA more than the EU, the reverse seems to have happened. Although not significantly different for 20%, the mean drop in infected users by the EU from 40% to 60% is significant, and it is only at this point that the EU begins to perform better (although not significantly until the 100% mark.) Therefore Hypothesis 4 was rejected. This might be because of the fact that the distributions of the EU countries were spread out more, so that on each occasion there was more likely to be a chance of more than one moderately sized hosting provider not being compliant.

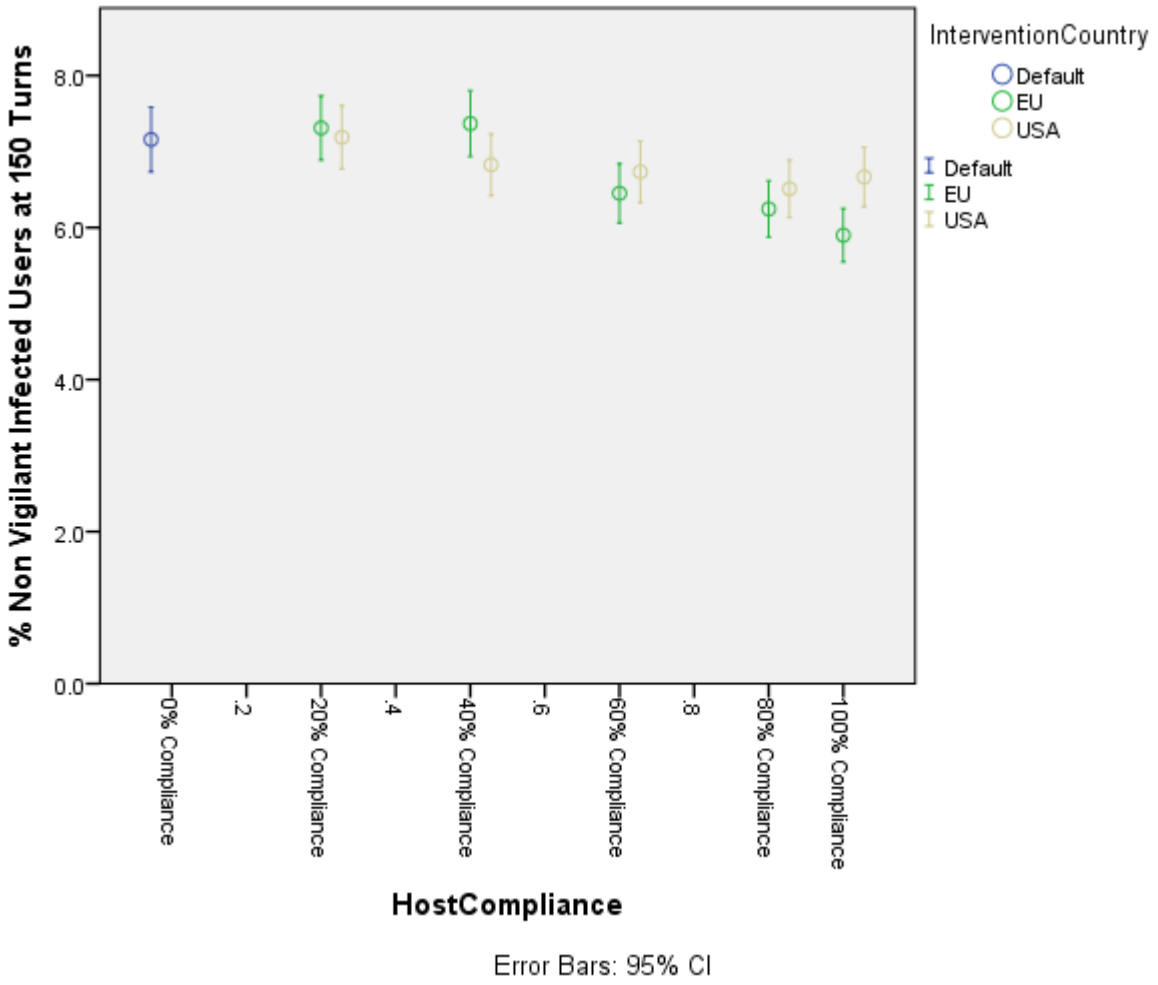

FIGURE 6.13: Percentage of infected vigilant users with different compliance levels

6.4.3 Effects of Modifying Parameters

Having found that there is a significant effect with intervention by different countries, and at different levels of compliance, the effect of altering other parameters were also tested in order to see the effect on the results in an otherwise default condition. For each condition, 1,000 simulations were run.

6.4.3.1 Website Population Count

It was expected that higher website populations would lead to a higher level of infected users, because even with small probabilities the cumulative amount makes it more likely that the conditions for a user infection will occur.

An ANOVA test was performed, where the population count was doubled each time (Figure 6.14. The test proved there was a slight effect between different website counts $F(5, 5994) = 35.195, p < .001$. Despite this, Games-Howell *post hoc* tests showed that

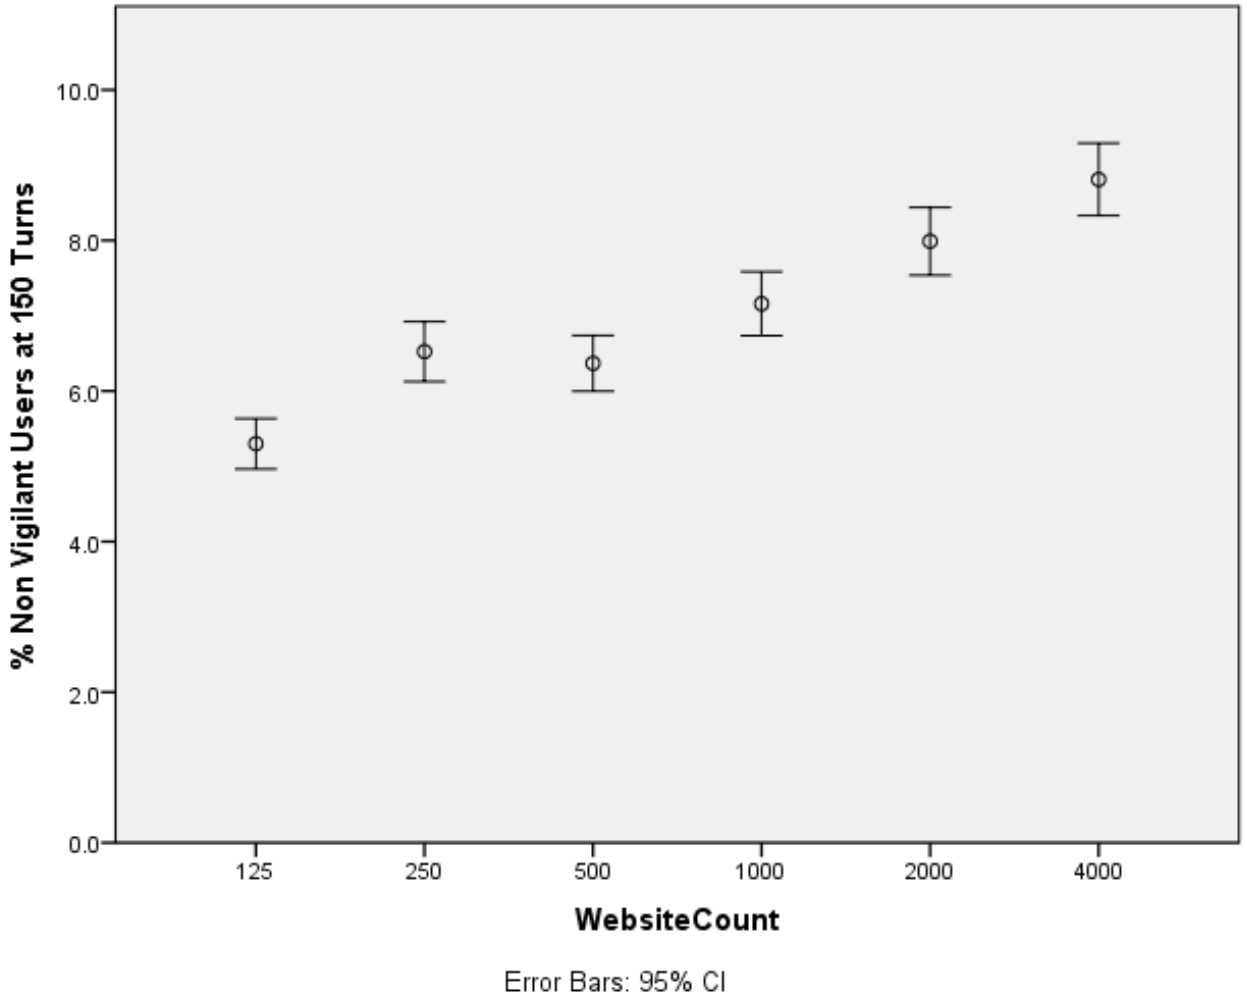

FIGURE 6.14: Difference in means between different sized website populations

the difference between conditions aside from 125 to 250 ($p < .001$, where 250 appears to be a slightly anomalous result, since it deviates from the otherwise consistent model. Doubling the website count twice does show significance in each case ($p < .001$).

6.4.3.2 Probability of Infection

Around the default level of 0.1 for the probability of a website to become infected once vulnerable. A one way ANOVA was performed, and it was revealed that there was a significant effect on the infected non-vigilant population of users $F(4, 4995) = 95.445$. Games-Howell *post hoc* tests revealed that a significant increase occurred with each raise up to 0.2, but that there was no significant difference between a probability of 0.2 and 0.4.

The results can be seen in Table 6.6.

Condition	0.025	0.05	0.1	0.2	0.4
Non Vigilant Users Infected(%)	3.824	5.352	7.159	8.419	8.29

TABLE 6.6: % non-vigilant users infected at 150 turns with different website infection probabilities

Website Vigilance (%)	0	20	40	60	80	100
% Non-Vigilant Users Infected	9.669	8.367	7.256	6.284	5.23	3.177

TABLE 6.7: % non-vigilant users infected at 150 turns with different website vigilance

User Vigilance (%)	0	20%	40%	60%	80%	100%
Total Users Infected	7819.94	6220.77	4672.84	3272.78	1881.64	446.82

TABLE 6.8: Total users infected at 150 turns with different user vigilance

6.4.3.3 Levels of Vigilance

As one of the major influences behind performing this simulation, it was decided to consider the effect of changing vigilance of first, the amount of vigilant websites; second, the amount of vigilant users; and finally the effect of both together.

An ANOVA test was carried out to assess the effect of website vigilance and user vigilance. For the user vigilance, a different metric had to be used for assessment rather than non-vigilant user levels, since the variable being changed is the vigilance level. Instead, this was assessed on the total amount of users infected.

There was a significant effect in the changing of website, $F(5, 5994) = 122.294$, and Games-Howell *post hoc* tests showed an increase between each increase (all ps $< .001$, except 40% to 60% where $p = .015$.

The effect of website vigilance can be seen in Table 6.7.

There was also a significant effect for the proportion of vigilant users in the population, $F(5, 5994) = 324.204$. Games-Howell *post hoc* tests revealed that each increase had a significant effect (all ps $< .001$). The effects can be seen in Table 6.8.

Both sets of results show similar patterns, and confirm that gains can continue to be made right up to the level of striving for 100% vigilance in one of the population. On the other hand, the sharp drop from 80% to 100% in the website vigilance levels could suggest that the parameters of the effectiveness of vigilant websites were too generous.

To assess the effect of increased vigilance levels in different degrees of vigilance in the other population, a two way ANOVA procedure was performed. Once again, this was done in terms of total users infected due to the inclusion of user vigilance as an independent variable in the test. Unsurprisingly, an effect was observed in both the effect of website vigilance $F(5, 35965) = 447.581$, and user vigilance $F(5, 35965) = 2224.695$,

and also for the interaction $F(24, 35965) = 45.133$. The effect can be seen in Figures 6.15 and 6.16.

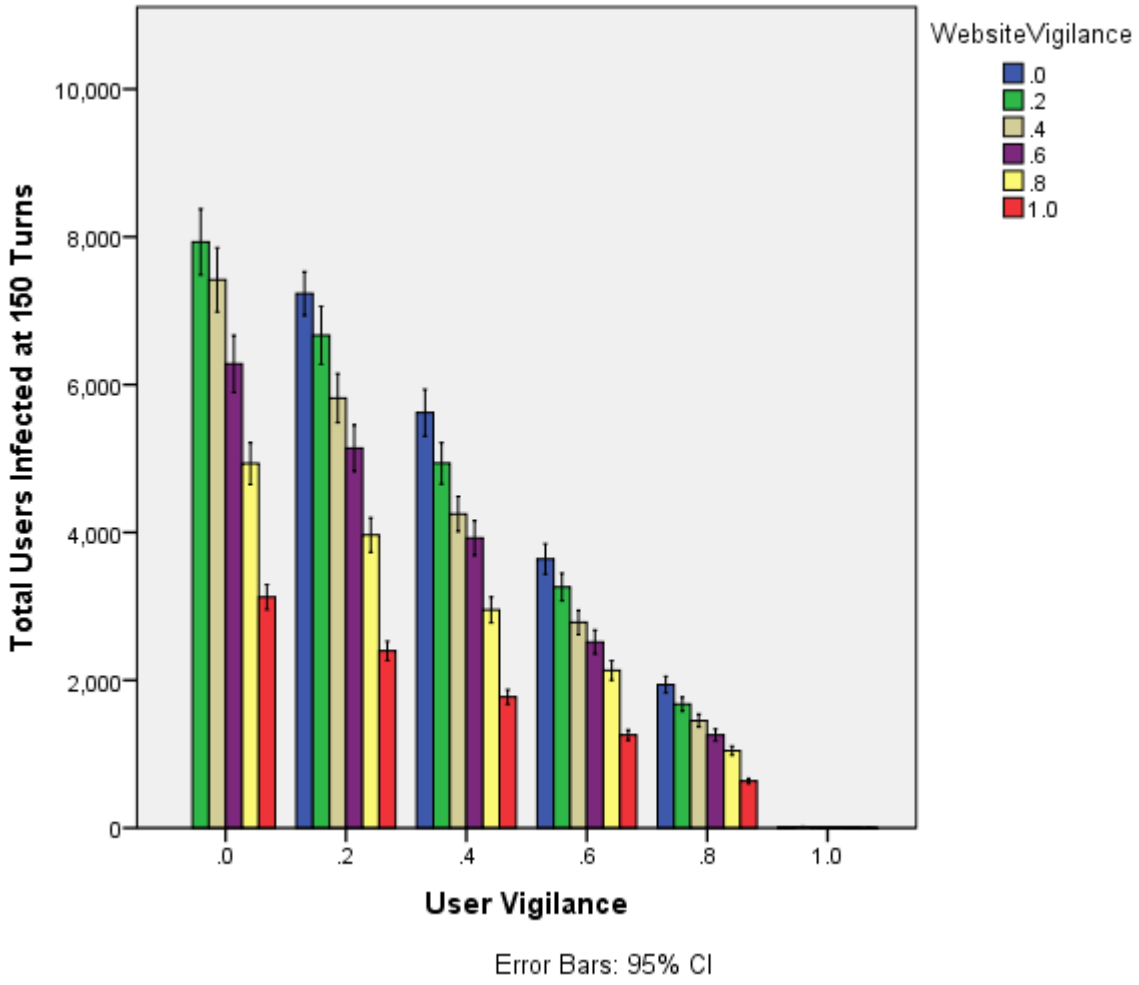

FIGURE 6.15: Effect of changing the proportion of vigilant users (absolute values)

This demonstrates that there is an alternative to blocking which can possibly even have more of an effect. This might be a policy which could be introduced, e.g. by providing information to ensure that websites do become more vigilant. However, those who are already causing the overall problem might be less inclined (or able) to increase their level of vigilance. The practicalities of such a system would have to be measured, although it could quite conceivably simply form a replacement to the hosts' obligations and merely require them to provide information or educational materials instead of implementing blocking procedures. It is beyond the scope of this research, however, and may be the subject of future work.

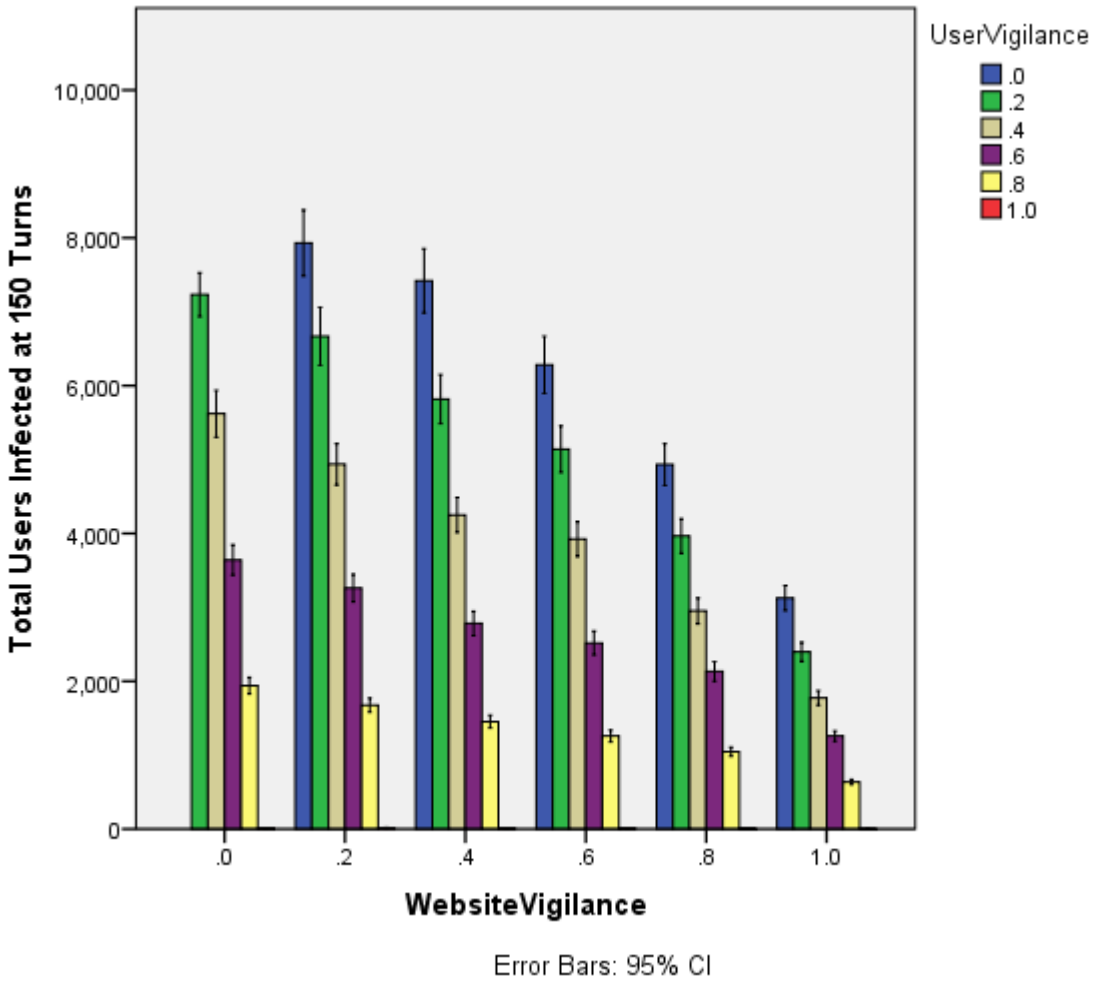

FIGURE 6.16: Effect of changing the proportion of vigilant websites

6.4.3.4 Conclusion

This section has analysed some of the parameters in the simulation which were not used when attempting to answer the main research question. In all cases, the output from the model has been largely as expected, supporting the idea that it works as designed and that it could be used to examine other scenarios. These could be done as future work, in particular work on the vigilance levels of the populations.

6.5 Discussion and Limitations

This chapter described an ABM which was used to analyse the efficacy of an intervention by hosting providers. It was assumed that upon noticing, the hosting providers would block in half the time it might otherwise take the website in question to recover from a vulnerability or infection. This was compared with a controlled default condition, and then with the effect of increasing vigilance of both users and websites.

The results showed that, in general, either approach would be successful in reducing the overall amount of infected machines on the Web in the event that this was enforced by the EU. This confirms the idea that it is not necessary to consider all the countries in the world when performing an intervention, provided that the countries who do participate have an appropriate worldwide market share. Somewhat surprising was the fact that the USA did not amount to anything significant, and this is something which requires further investigation. The USA had a higher market share, and as such, should have had more of an effect.

The surprise increase in infections between 80% and 100% with the USA intervention was not significantly different, but that it is increasing at all is definite reason to consider very carefully whether there might be a limitation to the accuracy of the simulation. The other anomalous result was the increase in user infections as the website population increased from 125 to 250, but then no significant change between 250 and 500 where the rest of that followed a consistent pattern. This means that possibly a larger N might be more appropriate for the amount of trials, which will be investigated more thoroughly in future work.

Whether this would be a worthwhile investment was not considered, and this would be far from a trivial task. The first thing to be established, is whether it would make an appreciable difference to the viability of cybercrime in the event of an intervention which made a statistically significant difference to overall infection rates. The research discussed in Section 2.3.3 provides some indications that attackers might fare badly in the event of an intervention like this. Both and in particular appeared to suggest that they would. Attackers who are discriminating in their targets, e.g. the Fake AV discussed by Stone-Gross et al. (2013) could well be affected more, whilst McCoy et al. (2012) believed the same would likely be true for spammers. On the other hand, the majority of the literature is by necessity speculative and those relying on figures from individual attackers inevitably do not have the complete picture.

As was established, the UK acting unilaterally failed to make a significant difference to global infection levels. That is not to say that this wouldn't have made a difference to machines in their own jurisdiction, but whether it is a worthwhile investment was left out of the scope. Obtaining figures to make an assessment such as this would be difficult. Firstly, it would be necessary to ascertain exact cost increases to hosting a website – and whether this might lead to either hosting providers or website operators moving to a different jurisdiction. In addition, establishing losses from cybercrime is generally regarded as a very difficult problem. The report produced by Detica (2011) came up with a figure an order of magnitude greater than that produced by Anderson et al. (2012).

Florêncio and Herley (2013) also made the point that many of the cybercrime surveys, or other methods, are very difficult to obtain results which would adequately extrapolate to

the population, even were they not methodologically poor. This is down to factors such as the distribution of losses, and the difficulties in obtaining sufficiently large samples. Some information is also lacking as well which could cause the case to be understated. Whilst the Cryptolocker ransomware reportedly asked for $300 in bitcoins to release the files (Kelion, 2013) the cost to the economy does not include the time spent working out how to get $300 into bitcoins, the costs of the data which was not recovered.

It might be that the government does not wish to consider in economic terms the costs of conducting an intervention, and operate on more ideological grounds. For example, the stated aim to make online trade safe in the UK, could be considered to be an aim based on principles since it . This is something which can still fall within the public health analogy, since it recognises that there remains issues outside of "rational" decisions.

Aside from a few results which might be anomalous, there are a few other minor limitations to the simulation. Firstly, as indicated above, the parameters selected for the vigilant users might have been overly generous of their ability to prevent compromise when compared to their non-vigilant counterparts. Educated guesses also had to be made for some parts as well, such as the likelihood of a vulnerable website getting compromised. It appears to be reasonably consistent with what might be expected, but there is no easy way to gain verifiable numbers for it.

The methods used to decide the websites visited by the end users could also be improved in the future. The Alexa rankings are made up of a series of fields including `ReachPerMillion`, `PageViewsPerMillion`, and `PageViewsPerUser`. In this instance, the only metric used was that of `ReachPerMillion`. Whilst this was probably the most appropriate, in some cases it may not accurately reflect the amount of views the websites could expect to get as the population was too small. Websites with a reach of 1 per million would be unlikely to have any traffic in this simulation, for example. In some cases, e.g. with social media websites, it might have been indicative to see the amount of page views as well, since participating in a social network can lead to additional Web based threats (Faghani et al., 2012). It is possible that the users who visited these websites might have been more susceptible if they spent more time on them. This is a subject for future work.

Similarly, the selection of websites was also done completely at random, without any consideration for local preferences. For example, a user visiting `www.google.co.uk` might also be assigned `www.google.es`, and where visiting a website, the local language was not considered. These appeared to be unnecessary complications to the model, and were excluded for this reason, but if a vigilance based solution were adopted in different countries, then this is something which might have to be considered.

Chapter 7

Conclusions and Future Work

This report set out to investigate the supplemental approaches to the law enforcement approach to combatting cybercrime, arguing that on their own the current approaches being used were inadequate. This led to legal and economic techniques being used alongside a computational simulation to prove the following hypothesis and research questions:

> The public health analogy remains a viable framework for combatting drive-by downloads, even though malware no longer spreads in a manner similar to a virus

1. How can the public health analogy be re-imagined given the current threat landscape?
2. Which stakeholder is the best to target in order to minimise the damage caused by drive-by downloads?
3. Are intermediary obligations to combat drive-by downloads appropriate?
4. Can actions by a single country, or group of countries, have a statistically significant effect on the worldwide prevalence of infections from drive-by downloads?

The focus of the literature in relation to the public health analogy has been on the notion of viruses, and propagation of malware. This thesis sought to analyse some additional elements which could provide more of a context, and how they could be used. More recent literature has presented the public health analogy in contexts different to the traditional compartmental models which have been used since (**?**), such as Mulligan and Schneider (2011) who compared the analogy in terms of its qualities as public goods.

The way that the public health analogy was reimagined was to remove the focus from epidemiological techniques and strategies, and consider the impact of any potential intervention in the context of efficacy and rights instead. Through analysis of the stakeholders, and literature from the security economics discipline, it was determined that there was little incentive in most cases for these stakeholders to intervene. Given this, government regulation was considered the only possibility, and the most appropriate means of guidance was that of health policy. Using the analogy in this way made it far more important to consider the impact upon rights and efficacy, because of the similarlity of the issues which existed.

In order to apply this to drive-by downloads, it was far more appropriate that it be considered in terms of the criminal as the parasite. This matches far better the analogy proposed, and allows the focus to be taken away from the spread of malware. The amount of machines which are infected is important, but we are not expecting that there be an actual epidemic. Rather, it is expected that a certain amount of non-vigilant users will continue to get infected and that this will make the model of the criminal economy viable.

Therefore, it is concluded that the public health analogy can be re-imagined in terms of efficacy and rights of interventions.

Stakeholders were analysed in Chapter 3 in terms of what they could do, and their motivations for doing so. In Chapter 5, the key points of efficacy and rights were applied to these same stakeholders according to the analogy. The combination of these factors demonstrated that the hosting providers would be best placed to deal with the issue of drive-by downloads. The simulation performed in Chapter 6 also proved that the application of this could have a significant impact.

The proposal suggested was that hosting providers pre-emptively block websites where they become vulnerable to infection, whether through outdated software or poor coding practices. This was deemed to be the best trade-off between efficacy and rights. ISPs were regarded as being excessively invasive on privacy or freedom of expression, but arguably more effective, whilst search engines lack the enforcement abilities and additional information to enable a better defence.

This does not preclude interventions by the other stakeholders, for example it is necessary that major software vendors continue to take care to fix vulnerabilities, and that search engines continue to do their best to keep malicious websites out of their lists. Similarly, future new interventions could be considered by other intermediaries, merely that according to the principles described within the public health analogy it is more appropriate to begin with intervention by hosting providers.

Hosting providers are the best placed stakeholder to minimise the damage caused by drive-by downloads

Intermediary obligations were analysed according to the analogy and immunities law within the EU. Although concerns are justifiably raised about the prospect of some interventions, such as content filtering, the obligations themselves are not necessarily a bad thing. To go after individual users or websites would be inefficient, and require the co-operation of the intermediaries in any case. As such, to use intermediary obligations are necesary, and doing so is done best with the trade-offs described. This is already alluded to in some of the cases requiring injunctions against ISPs (see generally *Cartier*).

By choosing hosting providers, it is necessary to consider current judicial thinking in relation to Article 14 of the E-Commerce Directive. The lack of protection for affirmative action (voluntary or otherwise) places hosting providers at risk in the event that they make a mistake. This is not the outcome which is desired at all, so protection akin to §230 Communications Decency Act in the USA would have to be introduced alongside the obligation.

The obligation proposed, to inspect for vulnerabilities of hosted websites, separated completely the user content from the technical content, sidestepping any data protection and privacy issues, as well as removing the "general" monitoring and being acceptable under Article 15 E-Commerce Directive.

Therefore, it it is concluded that intermediary obligations are appropriate to assist in solving the drive-by downloads problem.

The statistical output from the simulation revealed that it is possible for individual countries, or groups of countries, to make a statistically significant difference to the overall level of drive-by download prevalence worldwide. However, where this was applied to the UK, it was found that there was no statistically significant difference to the overall level. In the event that the UK were to unilaterally adopt such a solution, they would have to do it for their own reasons. Amongst these might be the government's stated aim of making the UK "one of the most secure places in the world to do business in cyberspace" (Cabinet Office, 2011).

However, an intervention by all the countries in the EU would have an effect, even in a realistic scenario where compliance was not at 100%. This makes it a plausible policy to pursue, although whether the impact of such a policy would be worth it remains to be proven.

A single country or group of countries can make a statistically significant difference to the overall level of drive-by downloads.

Given the answers to the research questions, it has been shown that the use of a modified version of the public health analogy can be useful in combatting drive-by downloads. As such, the hypothesis for the thesis is proved.

7.1 Limitations and Future Work

This is not a traditional thesis, confined to one subject, but rather is interdisciplinary. In order to be able to appreciate the issues surrounding malicious software and botnets appropriately, several distinct branches of both law and computer science need to be discussed. Therefore, greater emphasis was placed on the breadth of subjects rather than the depth. This has included elements of intermediary liability, security economics, public health law, all underpinned by general technological knowledge in relation to the Web and the methodology required to run a simulation. This is necessary, in order to understand the practicalities of implementing potential changes; and whether such a change would be likely to have an effect.

It is also acknowledged that many of the references contained within this report include references to blogs. Generally speaking, the only use for such references has been to reference news events which are not reported in more general, respected news outlets. Where this is not the case, then it is because there is simply no better source available. Journals are frequently out of date, particularly in a subject area as fast moving as the Web in general – and security in particular. Discoveries of vulnerabilities are often made by companies rather than academic researchers, and these companies will often have better data available, so where appropriate these will be used. Where they are, every effort has been made to limit it to practitioners who are well regarded, such as Brian Krebs and Bruce Schneier.

Firstly, the limitation of consideration to criminals motivated by profit represents only one facet of the global cybersecurity landscape. High profile denial of service attacks have led to considerable losses, yet, where these are done to make a political statement then any ideas about improving difficulty slightly do not apply to this group since they do not care about making a profit. There are also indications that the business models of criminals are changing slightly, the increased use of ransomware goes against the quiet, infiltration and making use of botnets. Similarly, on some occasions, the denial of service attacks have been used as a means of advertising services offered by the same group to protect against denial of service attacks.

Another area which would benefit from more research, is that of targeted attacks. In these, the adversary will expend greater effort in targeting a specfic victim for a greater payoff. In cases such as Stuxnet, the pay off was not financial, but rather the target was to cause damage to the Iranian nuclear facility. The attack against Sony Pictures in December 2014 was unprecedented in the damage it caused, and was reportedly the work of North Korean agents (although there is scepticism about this). Given the fallible nature of software, and adversaries with unlimited resources (in a practical sense), no company in the world is safe. Furthermore, given that it is necessary for Joe Public to rely on these companies to store personal or financial data, why are people still happy to continue to use these services?

Finally, more work on the public health techniques could also be a benefit in the future in order to determine exactly what a hosting provider (or other intermediary) should do in the event that they locate a website which they believe will become compromised. Machine learning techniques are beginning to emerge to predict the likelihood of a website becoming compromised, e.g. Soska and Christin (2014), and more work in this area is worthwhile.

Appendix A

Scanning for Vulnerable Websites

This appendix details the methodology and results for the application I wrote to detect which versions of WordPress were used by the top 10,000 websites. It was decided not to include this in the final thesis, because other work had done similar things (e.g. Vasek and Moore (2013)) and that these could also be detected by using the http://w3techs.com website. The text is as originally appeared in the upgrade report, in December 2013.

A.1 Approach

The approach of identifying the version of the CMS in use was to pre-compute hashes of common, static files which changed commonly between different versions. The presence of a file with a certain hash on the server would be able to indicate the version of the CMS installed. A prototype application was written in PHP, to test the top 10,000 websites from the Alexa list for the presence of WordPress installations. WordPress was chosen because it is the most popular CMS in use, and is said to power 17-18% of all websites, and the fact that it is open source allows the files to be easily analysed for comparison later.

The tasks of the application were as follows:[1]

1. Download the source code for all back versions of WordPress and extract
2. Iterate through the folders and compare to the next version through hashing each file
3. Store the hashes in an associative array in the format file => hash, and hash => version.

[1] The idea for this method came from http://dcid.me/texts/fingerprinting-web-apps.html – last accessed September 16 2013

4. Check for the presence of a meta tag indicating the version of WordPress. In some versions it was added by default in the meta tag in the format: `<meta name="generator" content="WordPress X.X.X" />`

5. Search the page for indications of it being a WordPress installation. This was done by searching for links with wp-content or wp-includes in the links to pages or images.

6. Identify if the WordPress installation might have a different root to the website itself, for example under blogs.domain.tld or domain.tld/blog. All common files must then be searched for from this root

7. Download the files from the website, hash them, and compare against the locally stored hashed files.

8. Store in the database whether the website is a WordPress site or not; and if so what version it was.

Manual analysis of the differences revealed that /wp-includes/js/tinymce/plugins/wordpress/editor_plugin.js changed amongst the most frequently, and it was possible to identify the version "family" from this. For example, 3.5.1 and 3.5.2 would be in the version 3.5 family. Analytics company w3techs estimate that 59% of the websites were on the latest version (3.5.1 at the time). The top 10,000 sites in the Alexa list were chosen as an initial sample for the application to run over.

It was hypothesised that the top ranking websites would have a higher percentage of updated sites because of the resources available to them and the potential business impact of having their site hacked. It was further hypothesised that top ranking websites would be less likely to use WordPress due to its history as merely a blogging platform, or the popularity making it more likely to be targeted in hacking attacks.

A.2 Results

The results can be seen in table A.2. The first observation is that the 59% figure suggested by w3techs is considerably higher than the figures observed by this application. This could be explained by the fact that higher ranking sites would be more inclined to want to run their own sites rather than use wordpress.com, which is a managed environment for WordPress sites. Lower sites using wordpress.com hosting would be more likely to use the latest version, thereby increasing the percentage with the latest version.

However, subsequent analysis of the wordpress.com files revealed that they contained an unknown hash. This could have been due to the fact that they were using version

3.52 which was released shortly afterwards. Alternatively, it could have been modified slightly to prevent exactly this kind of analysis. The hypothesis that the higher ranking sites would have greater security appeared to have been incorrect, since there was no noticeable improvement on the higher ranked sites compared to the lower ranked sites. This could also possibly be explained by the wordpress.com file being unrecognised, since 15% of the sites for the top 500 sites were unknown. Lower sites which were unknown could in fact be versions lower than version 2.0, which is the lowest version to be included in this study. There does appear to be a slight increase in sites older than 2 and 3 years as the sites get less popular, though the differences between them are fairly small.

This shows that there is a not insignificant proportion of websites which are running with known vulnerabilities. These websites represent a threat to users of the Internet, because of the variety of uses they can be put to. In relation to the public health analogy, there are a few ways in which they could be considered. In a standard epidemiological model, they could represent the susceptibles in a population able to pass on infections to anyone who comes into contact with them. A possibly better description would be to consider them to be like dirty water pumps since the act of visiting them (by a susceptible) could cause them to become infected. Like the water pumps, there is no immediately obvious way of telling which is safe and which isn't. Either way, despite the limitations discussed below, the study demonstrates that the analogy with public health does appear to hold.

Total Sites	WP Sites	WP %	Unknown	3.5.1	3.5	3.4.2	3.4.1	3.4	3.3.2	3.3.1	3.3	3.2.1	3.2	> 2 Years	> 3 years
500	13	2.6	15.38	30.77	7.69	7.69	0	0	7.69	7.69	0	0	0	15.38	0
1000	26	2.6	11.54	19.23	3.85	11.54	3.85	0	7.69	3.85	0	0	0	34.62	3.85
1500	48	3.2	10.42	18.75	2.08	16.67	6.25	0	6.25	4.17	0	0	2.08	31.25	2.08
2000	69	3.45	8.7	20.29	1.45	11.59	10.14	0	4.35	4.35	1.45	0	4.35	31.88	1.45
2500	95	3.8	7.37	21.05	1.05	12.63	8.42	0	5.26	5.26	1.05	0	4.21	32.63	1.05
3000	115	3.83	6.96	20	1.74	13.91	6.96	0	4.35	5.22	0.87	0	3.48	35.65	0.87
3500	144	4.11	6.94	22.92	1.39	13.89	5.56	0	4.86	4.17	0.69	0	3.47	35.42	0.69
4000	178	4.45	7.3	25.84	1.69	12.36	7.3	0.56	5.06	3.93	0.56	0	2.81	32.02	0.56
4500	215	4.78	6.98	26.51	1.86	12.56	6.98	0.47	4.65	3.26	0.47	0	3.26	32.56	0.93
5000	240	4.8	6.67	27.08	2.08	11.67	7.5	0.42	5	2.92	0.42	0	2.92	32.92	0.83
5500	276	5.02	6.16	31.16	1.81	10.14	6.88	0.36	4.35	2.9	0.36	0	3.26	31.52	0.72
6000	308	5.13	6.49	32.47	1.95	10.06	6.49	0.32	3.9	2.6	0.32	0	3.25	31.17	0.97
6500	338	5.2	5.92	35.21	1.78	10.65	5.92	0.3	3.85	2.96	0.3	0	2.96	29.29	0.89
7000	379	5.41	6.07	36.15	1.85	10.29	7.12	0.53	3.69	2.64	0.26	0	2.64	27.97	1.06
7500	407	5.43	5.65	37.59	1.97	10.32	7.13	0.49	3.44	2.7	0.25	0	2.46	27.27	1.47
8000	444	5.55	5.63	37.39	2.03	10.14	6.76	0.45	4.28	3.15	0.23	0	2.25	26.8	1.35
8500	483	5.68	5.18	38.51	2.07	9.73	6.42	0.41	4.35	3.11	0.21	0	2.48	26.5	1.24
9000	521	5.79	4.8	38.58	2.3	10.17	5.95	0.38	4.61	2.88	0.38	0.19	2.69	25.91	1.34
9500	562	5.92	4.8	38.97	2.31	10.14	5.69	0.36	4.63	2.67	0.36	0.18	2.49	25.8	1.42
10000	602	6.02	4.98	39.37	2.16	9.8	6.15	0.33	4.65	2.66	0.33	0.17	2.33	25.58	1.33

TABLE A.1: Incidences of different versions of WordPress installations. All columns are percentage values, except for the first two.

A.3 Limitations of the Study

This study was only a prototype, and it contained certain limitations.

The study only considered WordPress installations, from version 2.0 onwards. This represents a considerable percentage of both the Web and WordPress installs, but it is far from a complete picture, so future work would have been to identify other popular CMS software and what versions of that were installed. It was also observed that on many occasions, whilst the websites are up to date plugins are frequently not updated, so this would be expanded to include analysis of plugins associated with the CMS software. A similar approach could be used to determine the versions of plugins installed. A similar approach could be used for identifying the versions of popular plugins, because the source code is available at `http://wordpress.org/plugins/<<plugin name>>` and previous version numbers can be discovered from the "Compatibility" section of the page.

The sample was also a limitation. Whilst the top 10,000 sites provided an interesting insight, the popularity of a website does not matter if its hacked value is for DDoS or malicious hosting. It has been shown in the past that black hat SEO has been used to link malicious web pages to trending topics (Moore et al., 2011b), so an interesting addition would be to see how the top ranking sites for trending topics compare with the top website. Obtaining the current trending topics from Google is difficult, since access to the page requires the use of JavaScript and is not available by an API. It is also fairly small when considering the whole of the Web – it is entirely plausible that the results are different for the sites right at the top.

The lack of support for threads in PHP was a limiting factor in the size of the sample, because it did not scale. Even analysing 10,000 websites took nearly a week. A single threaded program can only perform one task at a time, and when there is a lot of waiting – both for responses and deliberate sleep to preserve the site's resources slows down the application considerably. It was decided to rewrite the version in Python which had better multi-threading support, along with improvements to the issues identified above.

Much effort was expended in improving the generalisability of the algorithm to get plugins, different CMSs and to better tell which version preceded which[2]. It turned out however, that an application which did this already existed: Blind Elephant[3] using the same approach. This minimised the amount of work needed to complete the application since it is open source and in Python and can simply be called to any website specified. A few modifications were made with the Blind Elephant code incorporated and it is currently being tested for WordPress, Drupal and Joomla websites.

[2]For example, consider parallel development of two version families. If done based purely on the modified date of the files, as in the original application, then this would not on occasions like this.

[3]`http://blindelephant.sourceforge.net`

As this report was being written, Joomla released a critical security update for a vulnerability[4] on July 31st. This was discovered some weeks after the release of the patch, and therefore possibly too late to consider analysis of patching patterns of top companies it does present an opportunity for a follow up study. By continuing to record the versions of the CMS software on sites on a regular basis, it should be possible to then retrospectively view the patching record in response to general patches, or patches in response to a CVE.

A final limitation which could be considered, is that although the application gives a general survey of what CMSs various sites use, it does not say why they chose not to upgrade their site given the risks. Future work is to identify the hosting providers through doing a WHOIS request on the IP address associated with the domain names to determine their policy about upgrading applications they host.

[4] http://developer.joomla.org/security/news/563-20130801-core-unauthorised-uploads

Appendix B

Data Gathered for the Simulation

In this appendix, the code used to generate the data which was used in the simulation is presented.

B.1 Hosting Provider Market Share

To determine the distribution of the hosts the following script was run to scrape information from the http://www.webhosting.info website on the 28th of May 2014. A manual check for the `robots.txt` page returned a 404 error, so no checks were neccessary. The first listing obtains the amount of hosts in each country, whilst the second obtains details about the distributions within the individual country.

```
for i in range(2,25):
        url = 'http://www.webhosting.info/webhosts/globalstats/?ob=HC\&oo=DESC\&pi=%d' % i
        html = ug.request(url)['html']
        soup = BeautifulSoup(html)
        rows = soup.find(id="AutoNumber7").find('table').findAll('tr')

        for tr in rows:
                if(tr.find('td') != None):
                        cells = tr.findAll('td')
                        country_url = cells[0].find('a')
                        country = country_url.text
                        #Having a > means that it's linking to the next page
                        #Don't want to try and add that into the database!
                        if "&lt;" in country:
                                break
                        amount_of_companies = int(cells[2].text.replace(',',''))
                        print "INSERT INTO Countries(Country, Url) \
                                VALUES('%s', '%s');" % (country, country_url['href'])
                        print "INSERT INTO HostPerCountry(CountryId, Amount)\n\
                        SELECT CountryId, %d \
                        FROM Countries WHERE Country = '%s'" % (amount_of_companies, country)
        #Robots was a 404 but lets be nice...
        print "About to sleep for 10s..."
```

```
with open('sql.txt', 'w') as f:
   for u in urls:
      i = 1
      while True: #We don't know how many pages there are going to be
         url = "%s?pi=%d" % (u[0], i)
         print url
         html = ug.request(url)['html']
         #print html
         soup = BeautifulSoup(html)
         rows = soup.find(id="AutoNumber7").find('table').findAll('tr')
         carry_on = False
         for tr in rows:
            if(tr.find('td') != None):
               cells = tr.findAll('td')#.find('a').text

               if len(cells[0].attrs) > 1:
                  if 'Next' in cells[0].text:
                     carry_on = True
                  break
               market_share = float(cells[2]   .text.replace('%',''))/100
               domains = cells[3].text
               query = "INSERT INTO HostingProviders\
                     (HostName,Country,MarketShare,Domains) \
                     VALUES('%s',%d,%f,'%s');\n" % (cells[1].text, u[1], market_share,domains)
               f.write(query)
         if carry_on:
            print "About to sleep for 5 seconds"
            time.sleep(5)
            i+=1
         else:
            break
```

B.2 CMS Market Share

The distribution of CMSes was obtained through scraping the w3techs website. The `algorithm()` function includes provision for getting the distribution of different versions for each different CMS. This was not used in the model.

```
from BeautifulSoup import BeautifulSoup
import sys
sys.path.insert(0, '..')
from URLGetter import URLGetter
import time
import os

ug = URLGetter()
def get_table_information(table, cms_param, filename = None):
   print "get_table_information"
   urls = []
   for tr in table:
      cms = ''
      th = tr.find('th')
      if th != None:
         if th.find('a') != None:
```

```
            a = tr.find('th').find('a')
            cms = a.text
            urls.append({cms:a['href']})
            cms = a.text

         #Graph table embedded inside the other table
         tables = tr.find('td').findAll('table')
         if len(tables) > 0:
            table = tables[0]
            share_text = table.find('tr').findAll('td')[1].text.replace('%','')
            share = 0
            if not 'less than' in share_text:
               share = float(share_text)
            if filename != None:
               with open('CMS/%s/%s.\
                  txt' % (cms_param.replace(':',''), filename.replace(':','')), 'a') as f:
                  query = "INSERT INTO CMSVersion(CMSId,Version,Share)\n\
                  SELECT CMSId, '%s', %f FROM CMSSoftware WHERE Name = '%s';\
                  \n" % (cms,share,cms_param)
                  f.write(query)
   return urls

def algorithm(urls, counter, already_done, cms, only_one_table = False):
   '''
   Recursively finds the pages relating to the different versions of each CMS
   Calls get_table_information to process the URL
   '''
   already_done.append(urls)
   for url in urls:
      filename = url.keys()[0]
      html = ug.request(url[filename])['html']
      soup = BeautifulSoup(html)
      table = None
      if soup.find(attrs={'class' : 'bars'}) == None:
         return
      if(only_one_table):
         table = soup.find(attrs={'class' : 'bars'}).findAll('tr')
      else:
         tables = soup.findAll(attrs={'class' : 'bars'})
         if len(tables) > 1:
            table = tables[1].findAll('tr')
         else:
            return
      urls = get_table_information(table, cms, filename)
      if urls in already_done:
         return
      counter += 1
      print "About to sleep for 10s"
      time.sleep(10)
      algorithm(urls, counter, already_done, cms, False)

def get_original_list():
   '''
   Returns a list of the CMS URLs according to w3techs
   '''

   cms_list = []

   html = ug.request('http://w3techs.com/technologies/\
```

```
      overview/content_management/all')['html']
   soup = BeautifulSoup(html)
   table = soup.find(attrs={'class' : 'bars'}).findAll('tr')
   for tr in table:
      cell = tr.find('th')
      if cell != None and cell.find('a') != None:
         a = cell.find('a')
         print a['href'], a.text
         cms_list.append({a.text : a['href']})

   return cms_list


def main():
   urls = get_original_list()
   for url in urls:
      #HACK
      u = [url]
      print "In the main list again..."
      already_done = []
      cms = url.keys()[0]
      os.makedirs(os.path.join('CMS',cms.replace(':','')))
      algorithm(u, 0, already_done, cms, True)

main()
```

B.3 NVD Scraper

The determination of how often a vulnerability would occur was obtained through using the mean of the average period in days between each vulnerability disclosure with a CVE score of ¿= 9.0 according to the National Vulnerability Database.

```
from BeautifulSoup import BeautifulSoup
import os
from datetime import datetime
import math



def parse_data(folder):
   #Manually substitute the software to search for
   listy = ['firefox', 'internet_explorer', 'chrome']

   for fi in os.listdir('CVEs'):
      running_total = 0
      if ('2013' in fi or '2012' in fi or '2014' in fi) and 'xml' in fi:
         for li in listy:
            with open(os.path.join('CVEs', fi)) as f:
               cve_dict = {}
               print "Parsing %s" % fi
               soup = BeautifulSoup(f.read())
               cves = soup.findAll('entry')
               for entry in cves:
                  cve_date = entry.find('vuln:published-datetime').text.split('T')[0]
                  products = entry.find('vuln:vulnerable-software-list')
```

```
                if products == None:
                    continue
                products = products.findAll('vuln:product')
                relevant_product = False

                for prod in products:
                    product = prod.text
                    split = product.split(':')
                    name = split[3]
                    if name == li:#in listy:
                        relevant_product = True
                        break
                severity = entry.find('vuln:cvss')
                if severity != None:
                    severity = severity.find('cvss:base_metrics').find('cvss:score').text
                    if float(severity) >= 8 and relevant_product:
                        if cve_date not in cve_dict:
                            cve_dict[cve_date] = 1
                        else:
                            cve_dict[cve_date] += 1
                            running_total += 1
            soup.decompose()#Think this gets around the MemoryError
            with open(os.path.join('CVEs', '%sCVEs.csv' % li), 'a') as f:
                for c in cve_dict.keys():
                    f.write("%s,%s\n" % (c, cve_dict[c]))


parse_data('')
```

Bibliography

Greg Aaron and Rod Rasmussen. Global phishing survey: trends and domain name use in 2H2009. *Anti-Phishing Working Group (APWG), Lexington, MA*, 2010.

Greg Aaron and Rod Rasmussen. APWG Global Phishing Survey: Trends and Domain Name Use in 2H2012, 2013.

Greg Aaron, Rod Rasmussen, and Aaron Routt. Global phishing survey: Domain name use and trends in 1h2014. Technical report, Anti-Phishing Working Group, 2014.

Kenneth S Abraham. *The forms and functions of tort law: an analytical primer on cases and concepts*. foundation Press, 1997.

Anne Adams and Martina Angela Sasse. Users are not the enemy. *Communications of the ACM*, 42(12):40–46, 1999.

George A Akerlof. The market for" lemons": Quality uncertainty and the market mechanism. *The quarterly journal of economics*, pages 488–500, 1970.

Devdatta Akhawe and Adrienne Porter Felt. Alice in warningland: A large-scale field study of browser security warning effectiveness. In *Proceedings of the 22th {USENIX} Security Symposium*, 2013.

Andy Alaszewski and Patrick Brown. *Making health policy: a critical introduction*. Polity, 2011.

Joaquin (European Commission) Almunia. Statement on the Google investigation, 2014.

R. Anderson. Open and closed systems are equivalent (that is, in an ideal world). *Perspectives on free and open source software*, pages 127–142, 2005.

Ross Anderson. 'Trusted Computing' Frequently Asked Questions.

Ross Anderson. Why information security is hard-an economic perspective. In *Computer Security Applications Conference, 2001. ACSAC 2001. Proceedings 17th Annual*, pages 358–365. IEEE, 2001.

Ross Anderson, Chris Barton, Rainer Böhme, Richard Clayton, Michel J G van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. Measuring the cost of cybercrime. In

Proceedings (online) of the 11th Workshop on the Economics of Information Security (WEIS), Berlin, Germany, 2012.

Ross Anderson, Rainer Böhme, Richard Clayton, and Tyler Moore. Security economics and european policy. *Managing Information Risk and the Economics of Security*, pages 55–80, 2009.

Efi Arazi. Pseudo random number generators in programming languages. Master's thesis, Efi Arazi School of Computer Science, 2011.

Andrew Ashworth and Jeremy Horder. *Principles of criminal law*. Oxford University Press, 2013.

Anushka Asthana and Tracy McVeigh. Government services to be online-only, 2010. Last accessed October 2015.

Sheena Asthana, Alex Gibson, Graham Moon, John Dicker, and Philip Brigham. The pursuit of equity in nhs resource allocation: should morbidity replace utilisation as the basis for setting health care capitations? *Social science & medicine*, 58(3):539–551, 2004.

Patrick Atiyah. *Vicarious Liability*. Butterworths London, 1967.

Albert-László Barabási, Réka Albert, and Hawoong Jeong. Scale-free characteristics of random networks: the topology of the world-wide web. *Physica A: Statistical Mechanics and its Applications*, 281(1):69–77, 2000.

Chris Baraniuk. Ashley madison: 'suicides' over website hack, 2015. Last accessed October 2015.

Dennis Batchelder, Nam Ng, Tim Rains, Joe Blackbird, Nial O'Sullivan, Jerome Stewart, Paul Henry, and Others. Microsoft security intelligence report volume 17 january through june 2014. Technical report, Microsoft, 2014.

BBC. More than 100 arrests, as fbi uncovers cyber crime ring. `http://www.bbc.co.uk/news/world-us-canada-11457611`, October 2010. Last Accessed March 2015.

BBC. Anonymous Wikileaks supporters explain web attacks. `http://www.bbc.co.uk/news/technology-11971259`, 2011.

BBC. China arrests thousands in latest internet crime crackdown. `http://www.bbc.co.uk/news/technology-18996811`, July 2012a.

BBC. Google's market share 'dips below 90%' in UK, 2012b.

BBC. Blackhole malware exploit kit suspect arrested. `http://www.bbc.co.uk/news/technology-24456988`, October 2013. Last accessed March 2015.

B.B.C. David cameron welcomes family-friendly internet filters. http://www.bbc.co.uk/news/uk-politics-25067051, November 2013. Last accessed March 2015.

BBC. Police hold 11 over ransomware scam 'affecting thousands'. http://www.bbc.co.uk/news/technology-21457743, February 2013.

B.B.C. Downing street presses isps over 'jihad reporting' button. http://www.bbc.co.uk/news/technology-30052211, 2014. Last accessed March 2015.

B.B.C. Lulzsec hacker helps fbi stop over 300 cyber attacks. http://www.bbc.co.uk/news/technology-27579765, 2014.

B.B.C. Sony cyber-attack: North korea faces new us sanctions. http://www.bbc.co.uk/news/world-us-canada-30661973, 2015.

Norman Begg, Mary Ramsay, Joanne White, and Zoltan Bozoky. Media dents confidence in mmr vaccine. *BMJ*, 316(7130):561, 1998.

David J Betz and Tim Stevens. Analogical reasoning and cyber security. *Security Dialogue*, 44(2):147–164, 2013.

Robert Beverly, Arthur Berger, Young Hyun, and Others. Understanding the efficacy of deployed internet source address validation filtering. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, pages 356–369. ACM, 2009.

D Bilar and E Filiol. On self-reproducing computer programs. *Journal in computer virology*, 5(1):9–87, 2009.

Leyla Bilge and Tudor Dumitras. Before we knew it: an empirical study of zero-day attacks in the real world. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pages 833–844. ACM, 2012.

Margaret Brazier and John Harris. Public health and private lives. *Med. L. Rev.*, 4:171, 1996.

J F Brenner. Nuisance Law and the Industrial Revolutions. *J. Legal Stud.*, 3:403, 1974.

British Telecom. Implementation of the online infringement of copyright (initial obligations)(sharing of ccost) order 2012, ofcom consultation published on 26 june 2012, bt response. http://stakeholders.ofcom.org.uk/binaries/consultations/onlinecopyright/responses/BT.pdf, September 2012.

Donna Buenaventura. Dutch ISPs Sign Anti-Botnet Treaty, 2009.

Robert K Burrow. Increased bank liability for online fraud: The effect of patco construction co. v. people's united bank. *NC Banking Inst.*, 17:381, 2013.

Juan Caballero, Chris Grier, Christian Kreibich, and Vern Paxson. Measuring pay-per-install: The commoditization of malware distribution. In *USENIX Security Symposium*, 2011.

Cabinet Office. The uk cyber security strategy protecting and promoting the uk in a digital world. 2011.

Davide Canali, Davide Balzarotti, and Aurélien Francillon. The role of web hosting providers in detecting compromised websites. In *Proceedings of the 22nd international conference on World Wide Web*, pages 177–188, 2013a.

Davide Canali, Davide Balzarotti, and Others. Behind the scenes of online attacks: an analysis of exploitation behaviors on the web. In *Proceedings of the 20th Annual Network & Distributed System Security Symposium*, 2013b.

Davide Canali, Leyla Bilge, and Davide Balzarotti. On the effectiveness of risk prediction based on users browsing behavior. In *Proceedings of the 9th ACM symposium on Information, computer and communications security*, pages 171–182. ACM, 2014.

Yannick Carlinet, Ludovic Mé, Hervé Debar, and Yvon Gourhant. Analysis of computer infection risk factors based on customer network usage. In *Emerging Security Information, Systems and Technologies, 2008. SECURWARE'08. Second International Conference on*, pages 317–325. IEEE, 2008.

Scott Charney. Collective Defense: Applying the Public-Health Model to the Internet. *Security & Privacy, IEEE*, 10(2):54–59, 2012.

Checkmarx. The Security State of WordPress' Top 50 Plugins. Technical report, Technical Report, 2014.

Shin-Ming Cheng, Weng Chon Ao, Pin-Yu Chen, and Kwang-Cheng Chen. On modeling malware propagation in generalized social networks. *Communications Letters, IEEE*, 15(1):25–27, 2011.

Eric Chien. Vbs.loveletter.var. http://www.symantec.com/security_response/writeup.jsp?docid=2000-121815-2258-99, 2001.

Nicolas Christin, Serge Egelman, Timothy Vidas, and Jens Grossklags. It's all about the Benjamins: An empirical study on incentivizing users to ignore security advice. *Financial Cryptography and Data Security*, pages 16–30, 2012.

D K Citron. Reservoirs of danger: The evolution of public and private law at the dawn of the information age. *S. Cal. L. Rev.*, 80:241, 2006.

Richard Clayton. Might Governments Clean-up Malware? In *Workshop on the Economics of Information Security (WEIS)*, 2011.

Richard Clayton, Steven J Murdoch, and Robert NM Watson. Ignoring the great firewall of china. In *Privacy Enhancing Technologies*, pages 20–35. Springer, 2006.

R H Coase. Problem of Social Cost, The. *Jl & econ.*, 3:1, 1960.

John Coggon. *What makes health public?: a critical evaluation of moral, legal, and political claims in public health*, volume 15. Cambridge University Press, 2012.

Alma Cohen and Liran Einav. The effects of mandatory seat belt laws on driving behavior and traffic fatalities. *Review of Economics and Statistics*, 85(4):828–843, 2003.

Fred Cohen. Computer viruses: theory and experiments. *Computers & security*, 6(1): 22–35, 1987.

Conficker Working Group. Conficker Working Group: Lessons Learned. Technical report, Tech. rep., Conficker Working Group–confickerworkinggroup. org, 2011.

Charlie Cooper. Cancer drugs fund: Life-extending drugs to be denied to nhs patients in england as fund overspends, 2015.

Council of Europe Treaty Office. Convention on cybercrime cets no.: 185. http://conventions.coe.int/Treaty/Commun/ChercheSig.asp?NT=185&CL=ENG, March 2015. Last Accessed March 2015.

Charlie Curtsinger, Benjamin Livshits, Benjamin G Zorn, and Christian Seifert. ZOZZLE: Fast and Precise In-Browser JavaScript Malware Detection. In *USENIX Security Symposium*, pages 33–48, 2011.

CYREN Security Blog. Update: Has the reported disruption of rustock affected spam levels. https://blog.cyren.com/articles/updated-has-the-reported-disruption-of-rustock-affected-spam-levels-1235.html, March 2011. Last Accessed March 2015.

Joe Davies(Microsoft). New Networking Features in Microsoft Windows XP Service Pack 2, 2004.

T Luis De Guzman. Unleashing a Cure for the Botnet Zombie Plague: Cybertorts, Counterstrikes, and Privileges. *Cath. UL Rev.*, 59:527, 2009.

Wouter de Vries. Hosting provider Antagonist automatically fixes vulnerabilities in customers websites. https://www.antagonist.nl/blog/2012/11/hosting-provider-antagonist-automatically-fixes-vulnerabilities-in-customers-websites/, 2012.

Department for Culture, Media & Sport and Ed Vaizey. Next steps to tackle internet piracy. https://www.gov.uk/government/news/next-steps-to-tackle-internet-piracy, 2012.

Department of Health. *Immunisation against infectious disease*. The Stationary Office under licence from the Department of Health, 2006.

Detica. The cost of cyber crime. Technical report, Cabinet Office, 2011.

Larry Dignan. Adobe, microsoft sync up patch schedule in overdue move. http://www.zdnet.com/article/adobe-microsoft-sync-up-patch-schedule-in-overdue-move/, November 2012. Last accessed March 2015.

Laura Donnely and Gregory Walton. 25 cancer drugs to be denied on nhs, 2015.

Peter Eckersley. How unique is your web browser? In *Privacy Enhancing Technologies*, pages 1–18. Springer, 2010.

Benjamin Edelman et al. Adverse selection in online'trust'certifications. In *WEIS*. Citeseer, 2006.

Benjamin Edwards, Tyler Moore, George Stelle, Steven Hofmeyr, and Stephanie Forrest. Beyond the blacklist: modeling malware spread and the effect of interventions. In *Proceedings of the 2012 workshop on New security paradigms*, pages 53–66. ACM, 2012.

Patrick Van Eecke. Online service providers and liability: A plea for a balanced approach. *Common Market Law Review*, 48(5):1455–1502, 2011.

Manuel Egele, Engin Kirda, and Christopher Kruegel. Mitigating drive-by download attacks: Challenges and open problems. In *iNetSec 2009–Open Research Problems in Network Security*, pages 52–62. Springer, 2009.

R A Epstein. Cybertrespass. *The University of Chicago Law Review*, 70(1):73–88, 2003a.

Richard A Epstein. In defense of the old public health-the legal framework for the regulation of public health. *Brook. L. Rev.*, 69:1421, 2003b.

European Commission. Green paper - damages actions for breach of the ec antitrust rules SEC(2005) 1732 /* com/2005/0672 final */, 2005.

European Commission. Green paper on consumer collective redress/* com/2008/0794 final */, 2008.

everything everywhere. Everything everywhere's response to ofcom's online infringement of copyright: Implementation of the online infringement of copyright (initial obligations) (sharing of costs) order 2012. http://stakeholders.ofcom.org.uk/binaries/consultations/onlinecopyright/responses/Everything_Everywhere.pdf, September 2012.

Jonathan I Ezor. Busting blocks: Revisiting 47 usc sec. 230 to address the lack of effective legal recourse for wrongful inclusion in spam filters. *Rich. JL & Tech.*, 17:1, 2010.

Mohammad Reza Faghani, Ashraf Matrawy, and Chung-Horng Lung. A study of trojan propagation in online social networks. In *New Technologies, Mobility and Security (NTMS), 2012 5th International Conference on*, pages 1–5. IEEE, 2012.

Nicolas Falliere, Liam O Murchu, and Eric Chien. W32. stuxnet dossier. *White paper, Symantec Corp., Security Response*, 2011.

FBI. Dnschanger malware. http://www.fbi.gov/news/stories/2011/november/malware_110911/DNS-changer-malware.pdf, November 2011. Last accessed March 2015.

Paul Ferguson. Network ingress filtering: Defeating denial of service attacks which employ IP source address spoofing. http://tools.ietf.org/html/rfc2827.html, 2000.

Dinei Florêncio and Cormac Herley. Sex, lies and cyber-crime surveys. In *Economics of Information Security and Privacy III*, pages 35–53. Springer, 2013.

Marc Fossi, Eric Johnson, Dean Turner, Trevor Mack, and Joseph Blackbird. Symantec report on the underground economy xii, 2008.

J Kevin French. Unauthorized and erroneous payment orders. *The Business Lawyer*, pages 1425–1445, 1990.

Huw Fryer, Roksana Moore, and Tim Chown. On the Viability of Using Liability to Incentivise Internet Security. In *Workshop on the Economics of Internet Security (WEIS)*, 2013.

Huw Fryer, Sophie Stalla-Bourdillon, and Tim Chown. Malicious web pages: what if hosting providers could actually do something... *Computer Law and Security Review*, 2015.

FTC. Statement of the Federal Trade Commission Regarding Googl's Search Practices In the Matter of Google Inc. FTC File Number 111-0163, 2013.

Joshua Gans, Stephen King, Robin Stonecash, and N Gregory Mankiw. *Principles of economics*. Cengage Learning, 2011.

Gian Franco Gensini, Magdi H Yacoub, and Andrea A Conti. The concept of quarantine in history: from plague to sars. *Journal of Infection*, 49(4):257–261, 2004.

Fiona Godlee, Jane Smith, Harvey Marcovitch, et al. Wakefields article linking mmr vaccine and autism was fraudulent. *BMJ*, 342, 2011.

J Goldberg. Twentieth Century Tort Theory. *Georgetown Law Journal*, 90, 2002.

Max Goncharov. Russian underground 101. 2012.

Dan Goodin. Fbi: Silk road mastermind couldnt even keep himself anonymous online. http://arstechnica.com/security/2013/10/silk-road-mastermind-unmasked-by-rookie-goofs-complaint-alleges/, October 2013.

Dan Goodin. Worm exploits nasty shellshock bug to commandeer network storage systems. http://arstechnica.com/security/2014/12/worm-exploits-nasty-shellshock-bug-to-commandeer-network-storage-systems/, December 2015. Last accessed March 2015.

Google. Using site speed in web search ranking. http://googlewebmastercentral.blogspot.co.uk/2010/04/using-site-speed-in-web-search-ranking.html, 2010.

Google. Finding more high-quality sites in search. http://googleblog.blogspot.co.uk/2011/02/finding-more-high-quality-sites-in.html, 2011.

Google. Google transparency report - making the web safer. http://www.google.com/transparencyreport/safebrowsing/, March 2015. Last accessed March 2015.

Lawrence Ogalthorpe Gostin. *Public health law and ethics: a reader*, volume 4. University of California Pr, 2010.

Scott Granneman. Infected in 20 minutes, 2004.

Frank L Greitzer, Andrew P Moore, Dawn M Cappelli, Dee H Andrews, Lynn Carroll, Thomas D Hull, et al. Combating the insider cyber threat. *Security & Privacy, IEEE*, 6(1):61–64, 2008.

Chris Grier, Lucas Ballard, Juan Caballero, Neha Chachra, Christian J Dietrich, Kirill Levchenko, Panayiotis Mavrommatis, Damon McCoy, Antonio Nappa, Andreas Pitsillidis, and Others. Manufacturing compromise: the emergence of exploit-as-a-service. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pages 821–832. ACM, 2012.

Chris Grier, Kurt Thomas, Vern Paxson, and Michael Zhang. @ spam: the underground on 140 characters or less. In *Proceedings of the 17th ACM conference on Computer and communications security*, pages 27–37. ACM, 2010.

Michael Joseph Gross. A declaration of cyber-war. http://www.vanityfair.com/culture/features/2011/04/stuxnet-201104, 2011.

Jeremiah Grossman. The web won't be safe or secure until we break it. *Communications of the ACM*, 56(1):68–72, 2013.

Garrett Hardin. The Tragedy of the Commons. *Science*, 162(3859):1243–1248, 1968.

Brian Haynes. Can it work? Does it work? Is it worth it?: The testing of healthcare interventions is evolving. *BMJ: British Medical Journal*, 319(7211):652, 1999.

Stephen E Henderson and Matthew E Yarbrough. Suing the Insecure?: A Duty of Care in Cyberspace. *New Mexico Law Review*, 32:11, 2002.

Stuart R Hene. Funds transfers under ucc article 4a: What is a commercially reasonable security system? *Consumer Fin. LQ Rep.*, 64:331–331, 2010.

Cormac Herley. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Proceedings of the 2009 workshop on New security paradigms workshop*, pages 133–144. ACM, 2009.

Cormac Herley. The Plight of the Targeted Attacker in a World of Scale. In *WEIS*, 2010.

Cormac Herley and Dinei Florêncio. A profitless endeavor: phishing as tragedy of the commons. In *Proceedings of the 2008 workshop on New security paradigms*, pages 59–70. ACM, 2009.

Cormac Herley and Dinei Florêncio. Nobody sells gold for the price of silver: Dishonesty, uncertainty and the underground economy. In *Economics of Information Security and Privacy*, pages 33–53. Springer, 2010.

Michael Hess and Bevan Rudge. Drupal core - highly critical - public service announcement - psa-2014-003. `https://www.drupal.org/PSA-2014-003`, October 2014. Last accessed March 2015.

Jack Hirshleifer. From weakest-link to best-shot: The voluntary provision of public goods. *Public Choice*, 41(3):371–386, 1983.

HM Government. *A Strong Britain in an Age of Uncertainty: The National Security Strategy*. The Stationary Office, London, 2010. ISBN 9780101795326.

HM Revenue & Customs. *2006–07 Accounts*. The Stationery Office, 2007.

Christopher Hodges. Collective redress in europe: The new model. *Civil Justice Quarterly, Forthcoming*, 2010.

Steven Hofmeyr, Tyler Moore, Stephanie Forrest, Benjamin Edwards, and George Stelle. Modeling internet-scale policies for cleaning up malware. *Economics of Information Security and Privacy III*, pages 149–170, 2011.

Petter Holme, Josh Karlin, and Stephanie Forrest. An integrated model of traffic, geography and economy in the internet. *ACM SIGCOMM Computer Communication Review*, 38(3):5–16, 2008.

Jenny Hope. Nhs to axe 21 cancer treatments as part of efforts to cut cancer drugs fund by £80 million, 2015.

Fu-Hau Hsu, Chang-Kuo Tso, Yi-Chun Yeh, Wei-Jen Wang, and Li-Han Chen. Browser-Guard: A Behavior-Based Solution to Drive-by-Download Attacks. *Selected Areas in Communications, IEEE Journal on*, 29(7):1461–1468, 2011.

Jeffrey Hunker. Us international policy for cybersecurity: five issues that won't go away. *J. Nat'l Sec. L. & Pol'y*, 4:197, 2010.

IIA. Internet Service Providers Voluntary Code of Practice for Industry Self-regulation in the area of Cyber Security, 2010.

Lukas Jeter and Shivakant Mishra. Identifying and quantifying the android device users' security risk exposure. In *Computing, Networking and Communications (ICNC), 2013 International Conference on*, pages 11–17. IEEE, 2013.

John P John, Fang Yu, Yinglian Xie, Arvind Krishnamurthy, and Martín Abadi. deSEO: Combating Search-Result Poisoning. In *USENIX Security Symposium*, 2011.

Vincent Johnson. Cybersecurity, Identity Theft, and the Limits of Tort Liability. *SCL Review*, 57:255, 2005.

Chris Kanich, Stephen Checkoway, and Keaton Mowery. Putting out a hit: crowdsourcing malware installs. In *Proceedings of the 5th USENIX Workshop on Offensive Technologies*, 2011.

Chris Kanich, Christian Kreibich, Kirill Levchenko, Brandon Enright, Geoffrey M Voelker, Vern Paxson, and Stefan Savage. Spamalytics: An empirical analysis of spam marketing conversion. In *Proceedings of the 15th ACM conference on Computer and communications security*, pages 3–14. ACM, 2008.

Louis Kaplow. Private versus social costs in bringing suit. *The Journal of Legal Studies*, pages 371–385, 1986.

Anestis Karasaridis, Brian Rexroad, and David Hoeflin. Wide-scale botnet detection and characterization. In *Proceedings of the first conference on First Workshop on Hot Topics in Understanding Botnets*, volume 7. Cambridge, MA, 2007.

Leo Kelion. Cryptolocker ransomware has 'infected about 250,000 PCs'. `http://www.bbc.co.uk/news/technology-25506020`, 2013. Last accessed March 2015.

Timothy Kelley and L Jean Camp. Online Promiscuity: Prophylactic Patching and the Spread of Computer Transmitted Infections. In *WEIS*, 2012.

Jeffrey O Kephart and Steve R White. Directed-graph epidemiological models of computer viruses. In *Research in Security and Privacy, 1991. Proceedings., 1991 IEEE Computer Society Symposium on*, pages 343–359. IEEE, 1991.

Jeffrey O Kephart and Steve R White. Measuring and modeling computer virus prevalence. In *Research in Security and Privacy, 1993. Proceedings., 1993 IEEE Computer Society Symposium on*, pages 2–15. IEEE, 1993.

William O Kermack and Anderson G McKendrick. A contribution to the mathematical theory of epidemics. In *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, volume 115, pages 700–721. The Royal Society, 1927.

Alex Komoroske. Prerendering in chrome. http://blog.chromium.org/2011/06/prerendering-in-chrome.html, June 2011. Last Accessed March 2015.

Brian Krebs. Silk road lawyers poke holes in fbis story. http://krebsonsecurity.com/2014/10/silk-road-lawyers-poke-holes-in-fbis-story/, October 2014. Last Accessed March 2015.

Brian Krebs. The internet of dangerous things. http://krebsonsecurity.com/2015/01/the-internet-of-dangerous-things/, January 2015a. Last accessed March 2015.

Brian Krebs. Ms update 3033929 causing reboot loop. http://krebsonsecurity.com/2015/03/ms-update-3033929-causing-reboot-loop/, March 2015b.

Stefan Kulk and Frederik J Zuiderveen Borgesius. Filtering for copyright enforcement in europe after the sabam cases. *European Intellectual Property Review, Forthcoming*, 2013.

Emily Kuwahara. Torts v. Contracts: Can Microsoft Be Held Liable to Home Consumers For Its Security Flaws? *S. Cal. L. Rev.*, 80:997–1433, 2007.

William M Landes and Richard A Posner. Positive Economic Theory of Tort Law, The. *Ga. L. Rev.*, 15:851, 1980.

Lawrence Lessig. *Code*. Lawrence Lessig, 2006.

John Leyden. BIGGEST DDoS ATTACK IN HISTORY hammers Spamhaus, 2013.

Zhen Li, Qi Liao, and Aaron Striegel. Botnet economics: uncertainty matters. In *Managing Information Risk and the Economics of Security*, pages 245–267. Springer, 2009.

Zhou Li, Kehuan Zhang, Yinglian Xie, Fang Yu, and XiaoFeng Wang. Knowing your enemy: understanding and detecting malicious web advertising. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pages 674–686. ACM, 2012.

Martin C Libicki. *Conquest in cyberspace: national security and information warfare*. Cambridge University Press, 2007.

D Lichtman and E Posner. Holding Internet Service Providers Accountable. *Sup. Ct. Econ. Rev.*, 14:221, 2006.

Nirmal Livingood, Jason and Mody. Recommendations for the Remediation of Bots in ISP Networks, 2012.

Nilly Madar, Tomer Kalisky, Reuven Cohen, Daniel ben Avraham, and Shlomo Havlin. Immunization and epidemic dynamics in complex networks. *The European Physical Journal B-Condensed Matter and Complex Systems*, 38(2):269–276, 2004.

Gregor Maier, Anja Feldmann, Vern Paxson, Robin Sommer, and Matthias Vallentin. An assessment of overt malicious activity manifest in residential networks. In *Detection of Intrusions and Malware, and Vulnerability Assessment*, pages 144–163. Springer, 2011.

Steve Marquess. Openssl needs corporate funding to avoid heartbleed repeat. http://www.zdnet.com/article/openssl-needs-corporate-funding-to-avoid-heartbleed-repeat/.

Alana Maurushat. Australia's accession to the cybercrime convention: Is the convention still relevant in combating cybercrime in the era of botnets and obfuscation crime tools. *UNSWLJ*, 33:431, 2010.

McAfee. Dissecting Operation High Roller. Technical report, Technical Report, 2012.

McAfee. Mcafee labs threats report june 2014. Technical report, Technical Report, 2014.

Stephen McCombie, Josef Pieprzyk, and Paul Watters. Cybercrime attribution: An eastern european case study. 2009.

Damon McCoy, Andreas Pitsillidis, Grant Jordan, Nicholas Weaver, Christian Kreibich, Brian Krebs, Geoffrey M Voelker, Stefan Savage, and Kirill Levchenko. Pharmaleaks: Understanding the business of online pharmaceutical affiliate programs. In *Proceedings of the 21st USENIX conference on Security symposium*, pages 1–1. USENIX Association, 2012.

Gary McGraw, Edward Felten, and Ryan MacMichael. *Securing Java: getting down to business with mobile code*. Wiley Computer Pub., 1999. ISBN 047131952X.

John PS McLaren. Nuisance law and the industrial revolution–some lessons from social history. *Oxford Journal of Legal Studies*, pages 155–221, 1983.

Jeffrey Meisner. Microsoft offers reward for information on rustock. http://blogs.microsoft.com/blog/2011/07/18/microsoft-offers-reward-for-information-on-rustock/, July 2011. Last visited March 2015.

Peter S Menell. A note on private versus social incentives to sue in a costly legal system. *The Journal of Legal Studies*, pages 41–52, 1983.

Damian Menscher. Notifying users affected by the dnschanger malware. http://googleonlinesecurity.blogspot.co.uk/2012/05/notifying-users-affected-by-dnschanger.html, 2012.

Microsoft. Understanding windows automatic updating. http://windows.microsoft.com/en-us/windows/understanding-windows-automatic-updating. Last accessed March 2015.

Microsoft. Microsoft Security Bulletin MS08-067 - Critical, 2008.

Microsoft. User-agent string changes. https://msdn.microsoft.com/en-us/library/ie/hh869301%28v=vs.85%29.aspx, 2013. Last accessed March 2015.

John Stuart Mill. *On liberty*. Broadview Press, 1999.

David Moore, Colleen Shannon, Geoffrey M Voelker, and Stefan Savage. Internet quarantine: Requirements for containing self-propagating code. In *INFOCOM 2003. Twenty-Second Annual Joint Conference of the IEEE Computer and Communications. IEEE Societies*, volume 3, pages 1901–1910. IEEE.

Tyler Moore and Richard Clayton. Evil searching: Compromise and recompromise of internet hosts for phishing. In *Financial Cryptography and Data Security*, pages 256–272. Springer, 2009.

Tyler Moore, Richard Clayton, and Ross Anderson. The economics of online crime. *The Journal of Economic Perspectives*, 23(3):3–20, 2009.

Tyler Moore, Nektarios Leontiadis, and Nicolas Christin. Fashion crimes: trending-term exploitation on the web. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 455–466. ACM, 2011a.

Tyler Moore, Nektarios Leontiadis, and Nicolas Christin. Fashion crimes: trending-term exploitation on the web. In *Proceedings of the 18th ACM conference on Computer and communications security*, pages 455–466. ACM, 2011b.

Nicola E Moran, S Blancafort, H Cowley, K Czabanowska, K Dierickx, Christian Munthe, Carlo Petrini, Elisabeth Petsetakis, Franz Piribauer, and Darren Shickle. Are compulsory immunisation and incentives to immunise effective ways to achieve herd immunity in europe? *Selgelid M, Battin M, Smith C (eds.), Ethics and Infectious Disease*, 2006.

Giovane Cesar Moreira Moura, Ramin Sadre, and Aiko Pras. Internet neighborhoods: the spam case. In *Network and Service Management (CNSM), 2011 7th International Conference on*, pages 1–8. IEEE, 2011.

Anuj Mubayi, Christopher Kribs Zaleta, Maia Martcheva, and Carlos Castillo-Chavez. A cost-based comparison of quarantine strategies for new emerging diseases. *Math Biosci Eng*, 7(3):687–717, 2010.

Deirdre K Mulligan and Fred B Schneider. Doctrine for Cybersecurity. *Daedalus*, 140 (4):70–92, 2011.

William H Murray. The application of epidemiology to computer viruses. *Computers & Security*, 7(2):139–145, 1988.

Satoshi Nakamoto. Bitcoin: A peer-to-peer electronic cash system. *Consulted*, 1(2012): 28, 2008.

Jose Nazario and T Holz. As the net churns: Fast-flux botnet observations. In *Malicious and Unwanted Software, 2008. MALWARE 2008. 3rd International Conference on*, pages 24–31, 2008.

Netcraft. March 2015 Web Server Survey, 2015.

netmarketshare.com. Desktop operating system market share. http://www.netmarketshare.com/operating-system-market-share.aspx?qprid=10&qpcustomd=0&qpsp=183&qpnp=1&qptimeframe=M, April 2014. Last accessed March 2015.

Mark EJ Newman, Stephanie Forrest, and Justin Balthrop. Email networks and the spread of computer viruses. *Physical Review E*, 66(3):035101, 2002.

Nuffield Council on Bioethics. *Healthy lives, healthy people: our strategy for public health in England*. Stationery Office, 2010.

Tobacco Advisory Group of the Royal College of Physicians et al. Going smoke-free: the medical case for clean air in the home, at work and in public places. *London: Royal College of Physicians of London*, 2005.

Anthony I Ogus. *Regulation: Legal form and economic theory*, volume 152. Clarendon Press Oxford, 1994.

Aleph One. Smashing the stack for fun and profit. *Phrack magazine*, 7(49):14–16, 1996.

Open Resolver Project. Open resolver project. http://openresolverproject.org/, October 2013.

Oracle. Critical patch updates, security alerts and third party bulletinsecurity alerts chicklet. http://www.oracle.com/technetwork/topics/security/alerts-086861.html, 2015. Last accessed March 2015.

Hilarie Orman. The morris worm: a fifteen-year perspective. *IEEE Security & Privacy*, 1(5):35–43, 2003.

Osterman ReGoogle. Best practices in email, web and social media security. http://www2.trustwave.com/rs/trustwave/images/Best_Practices_in_Email_Web_and_Social_Media_Security_Trustwave.pdf, January 2014. Last accessed March 2015.

OWASP. Top 10 2013–Top 10. https://www.owasp.org/index.php/Top_10_2013-Top_10, 2013.

Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The PageRank citation ranking: Bringing order to the web. 1999.

Panda Security. Quarterly report q3 2014. Technical report, Panda Labs, 2014.

H Van Dyke Parunak, Robert Savit, and Rick L Riolo. Agent-based modeling vs. equation-based modeling: A case study and users guide. In *Multi-agent systems and agent-based simulation*, pages 10–25. Springer, 1998.

Romualdo Pastor-Satorras and Alessandro Vespignani. Epidemic dynamics and endemic states in complex networks. *Physical Review E*, 63(6):066117, 2001.

PCI Security Standards Council. Payment card industry (pci) data security standard requirements and security assessment procedures version 3.0. Technical report, November 2013.

Sam Peltzman. The effects of automobile safety regulation. *The Journal of Political Economy*, pages 677–725, 1975.

Phillip Porras, Hassen Saidi, and Vinod Yegneswaran. Conficker C analysis. *SRI International*, 2009a.

Phillip Porras, Hassen Saïdi, and Vinod Yegneswaran. A foray into confickers logic and rendezvous points. In *USENIX Workshop on Large-Scale Exploits and Emergent Threats*, 2009b.

Jonathan B Postel. Rfc 821: Simple mail transfer protocol, 1982. *URL http://www. ietf. org/rfc/rfc0821. txt*, 15, 1982.

Niels Provos, Panayiotis Mavrommatis, Moheeb Abu Rajab, and Fabian Monrose. All Your iFRAMEs Point to Us. In *Proceedings of the 17th Conference on Security Symposium*, SS'08, pages 1–15, Berkeley, CA, USA, 2008. USENIX Association.

Niels Provos, Dean McNamee, Panayiotis Mavrommatis, Ke Wang, Nagendra Modadugu, and Others. The ghost in the browser analysis of web-based malware. In *Proceedings of the first conference on First Workshop on Hot Topics in Understanding Botnets*, page 4, 2007.

Hazhir Rahmandad and John Sterman. Heterogeneity and network structure in the dynamics of diffusion: Comparing agent-based and differential equation models. *Management Science*, 54(5):998–1014, 2008.

Greg Rattray, Chris Evans, and Jason Healey. American security in the cyber commons. *Contested Commons: The Future of American Power in a Multipolar World*, pages 137–176, 2010.

Eric Raymond. The cathedral and the bazaar. *Knowledge, Technology & Policy*, 12(3): 23–49, 1999.

Thomas Rid. Cyber war will not take place. *Journal of Strategic Studies*, 35(1):5–32, 2012.

Konrad Rieck, Tammo Krueger, and Andreas Dewald. Cujo: efficient detection and prevention of drive-by-download attacks. In *Proceedings of the 26th Annual Computer Security Applications Conference*, pages 31–39. ACM, 2010.

Uri Rivner. Anatomy of an Attack, 2011.

Susan Rose-Ackerman and Mark Geistfeld. The divergence between social and private incentives to sue: a comment on shavell, menell, and kaplow. *The Journal of Legal Studies*, pages 483–491, 1987.

David Rosenberg and Kathryn E Spier. Incentives to invest in litigation and the superiority of the class action. *Journal of Legal Analysis*, page lau006, 2014.

M.L. Rustad and T.H. Koenig. Tort of negligent enablement of cybercrime, the. *Berkeley Tech. LJ*, 20:1553, 2005.

V Sachin and N N Chiplunkar. SurfGuard JavaScript instrumentation-based defense against Drive-by downloads. In *Recent Advances in Computing and Software Systems (RACSS), 2012 International Conference on*, pages 267–272. IEEE, 2012.

Nathan A Sales. Regulating cyber-security? *Northwestern University Law Review*, 81: 1503, 2013.

Paul Sangster. Network endpoint assessment (nea): overview and requirements. *Network*, 2008.

Lisa Sattenspiel and D Ann Herring. Simulating the effect of quarantine on the spread of the 1918–19 flu in central canada. *Bulletin of mathematical biology*, 65(1):1–26, 2003.

Michael S. Schmidt and David E. Sanger. 5 in china army face u.s. charges of cyberattacks. `http://www.nytimes.com/2014/05/20/us/us-to-charge-chinese-workers-with-cyberspying.html?_r=0`, May 2014. Last visited March 2015.

Bruce Schneier. Information Security and Externalities. *ENISA Quarterly*, 2(4):3–4, 2007.

Bruce Schneier. The Plan to Quarantine Infected Computers, 2010.

Bruce Schneier. Cyberconflicts and national security'. https://www.schneier.com/essays/archives/2013/07/cyberconflicts_and_n.html, July 2013. Last accessed March 2015.

Bruce Schneier. We still don't know who hacked sony. http://www.theatlantic.com/international/archive/2015/01/we-still-dont-know-who-hacked-sony-north-korea/384198/, January 2015. Last accessed March 2015.

Guido Schryen. Is open source security a myth? *Commun. ACM*, 54(5):130–140, may 2011. ISSN 0001-0782.

Giuseppe Serazzi and Stefano Zanero. Computer virus propagation models. In *Performance Tools and Applications to Networked Systems*, pages 26–50. Springer, 2004.

S Shavell. Strict liability versus negligence. *NBER Working Paper*, (R0084), 1980.

Steven Shavell. Social versus the private incentive to bring suit in a costly legal system, the. *J. Legal Stud.*, 11:333, 1982.

Steven Shavell. Liability for harm versus regulation of safety, 1983.

Steven Shavell. The optimal structure of law enforcement. *Journal of Law and Economics*, pages 255–287, 1993.

Steven Shavell. The fundamental divergence between the private and the social motive to use the legal system. *The Journal of Legal Studies*, 26(S2):575–612, 1997.

Steven Shavell. The level of litigation: Private versus social optimality of suit and of settlement. *International Review of Law and Economics*, 19(1):99–115, 1999.

Steven Shavell. *Foundations of economic analysis of law*. Harvard University Press, 2004.

Seungwon Shin, Raymond Lin, and Guofei Gu. Cross-analysis of botnet victims: New insights and implications. In *Recent Advances in Intrusion Detection*, pages 242–261. Springer, 2011.

Kyle Soska and Nicolas Christin. Automatically detecting vulnerable websites before they turn malicious. In *Proc. USENIX Security*, 2014.

Pyda Srisuresh and Matt Holdrege. Ip network address translator (nat) terminology and considerations. 1999.

Sophie Stalla-Bourdillon. Sometimes one is not enough! securing freedom of expression, encouraging private regulation, or subsidizing internet intermediaries or all three at the same time: the dilemma of internet intermediaries liability. *Journal of International Commercial Law and Technology*, 7(2), 2012.

Sophie Stalla-Bourdillon. Online monitoring, filtering, blocking. what is the difference? where to draw the line? *Computer Law & Security Review*, 29(6):702–712, 2013.

Stuart Staniford, Vern Paxson, Nicholas Weaver, and Others. How to Own the Internet in Your Spare Time. In *USENIX Security Symposium*, pages 149–167, 2002.

Brett Stone-Gross, Ryan Abman, Richard A Kemmerer, Christopher Kruegel, Douglas G Steigerwald, and Giovanni Vigna. The underground economy of fake antivirus software. In *Economics of Information Security and Privacy III*, pages 55–78. Springer, 2013.

Brett Stone-Gross, Marco Cova, Lorenzo Cavallaro, Bob Gilbert, Martin Szydlowski, Richard Kemmerer, Christopher Kruegel, and Giovanni Vigna. Your botnet is my botnet: analysis of a botnet takeover. In *Proceedings of the 16th ACM conference on Computer and communications security*, pages 635–647. ACM, 2009.

Brett Stone-Gross, Marco Cova, Christopher Kruegel, and Giovanni Vigna. Peering through the iframe. In *INFOCOM, 2011 Proceedings IEEE*, pages 411–415. IEEE, 2011.

StopBadware and CommTouch. Compromised Websites: An Owner's Perspective. Technical report, 2012.

Symantec. Anonymous supporters tricked into installing zeus trojan. http://www.symantec.com/connect/blogs/anonymous-supporters-tricked-installing-zeus-trojan, 2012.

TalkTalk Group. Implementation of the online infringement of copyright (initial obligations)(sharing of ccost) order 2012, talktalk group subversion, non-confidential version. http://stakeholders.ofcom.org.uk/binaries/consultations/onlinecopyright/responses/TalkTalk_Group.pdf, September 2012.

Colin Tankard. Advanced persistent threats and how to monitor and deter them. *Network security*, 2011(8):16–19, 2011.

Ian Traynor. Russia accused of unleashing cyberwar to disable estonia. http://www.theguardian.com/world/2007/may/17/topstories3.russia, May 2007. Last accessed March 2015.

Trustwave. Trustwave 2013 Global Security Report. Technical report, Technical Report, 2013.

Trustwave. 2014 trustwave global security report. Technical report, Technical Report, 2014.

Lori Valigra. Online liability. 2013.

Michael Van Eeten, Johannes Bauer, Hadi Asgharia, Shirin Tabatabaie, and David Rand. The role of internet service providers in botnet mitigation an empirical analysis based on spam data. In *Proceedings (online) of the 9th Workshop on Economics of Information Security, Cambridge, MA*. TPRC, 2010.

Michel JG Van Eeten and Johannes M Bauer. Economics of malware: Security decisions, incentives and externalities. Technical report, OECD Publishing, 2008.

MJG van Eeten, H Asghari, JM Bauer, and S Tabatabaie. Isps and botnet mitigations: a fact-finding study on the dutch market. *TU Delft*, 2011.

Kenneth J Vandevelde. *Thinking like a lawyer: An introduction to legal reasoning*. Westview Press Colorado, 1996a.

Kenneth J Vandevelde. *Thinking like a lawyer: An introduction to legal reasoning*. Westview Press Colorado, 1996b.

Hal Varian. Managing online security risks, 2000.

Marie Vasek and Tyler Moore. Identifying Risk Factors for Webserver Compromise. 2013.

W3techs.com. Usage of content management systems for websites. `http://w3techs.com/technologies/overview/content_management/all`, March 2015. Last Accessed March 2015.

Gerhard Wagner. Collective redress–categories of loss and legislative options. *Law Quarterly Review*, 127, 2011.

Chenxi Wang, John C Knight, and Matthew C Elder. On computer viral infection and the effect of immunization. In *Computer Security Applications, 2000. ACSAC'00. 16th Annual Conference*, pages 246–256. IEEE, 2000.

Yang Wang and Chenxi Wang. Modeling the effects of timing parameters on virus propagation. In *Proceedings of the 2003 ACM workshop on Rapid Malcode*, pages 61–66. ACM, 2003.

Amalie M Weber. Council of europe's convention on cybercrime, the. *Berkeley Tech. LJ*, 18:425, 2003.

Lloyd L Weinreb. *Legal reason: The use of analogy in legal argument*. Cambridge University Press, 2005.

WhiteHat Security. Whitehat Website Security Statistics Report 2013. Technical report, Technical Report, 2013.

Wikimedia. Wikimedia traffic analysis report - browsers e.a. monthly requests or daily averages, for period: 1 sep 2014 - 30 sep 2014. http://stats.wikimedia.org/archive/squid_reports/2014-09/SquidReportClients.htm, September 2014. Last accessed March 2015.

Christopher Williams. Conficker seizes city's hospital network. http://www.theregister.co.uk/2009/01/20/sheffield_conficker/, January 2009. Last Accessed October 2015.

Anthony S Wohl et al. *Endangered lives: public health in Victorian Britain.* JM Dent and Sons Ltd, 1983.

Gilbert Wondracek, Thorsten Holz, Christian Platzer, Engin Kirda, and Christopher Kruegel. Is the Internet for porn? An insight into the online adult industry. In *Proceedings (online) of the 9th Workshop on Economics of Information Security, Cambridge, MA*, 2010.

WordPress. Updating WordPress, 2013.

Xiaofan Yang and Lu-Xing Yang. Towards the epidemiological modeling of computer viruses. *Discrete Dynamics in Nature and Society*, 2012, 2012.

Ting-Fang Yen, Victor Heorhiadi, Alina Oprea, Michael K Reiter, and Ari Juels. An epidemiological study of malware encounters in a large enterprise. In *Proceedings of the 2014 ACM SIGSAC Conference on Computer and Communications Security*, pages 1117–1130. ACM, 2014.

Ytosa. System.random serious bug. https://connect.microsoft.com/VisualStudio/feedback/details/634761/system-random-serious-bug, May 2011. Last Accessed March 2015.

Junjie Zhang, Christian Seifert, Jack W Stokes, and Wenke Lee. Arrow: Generating signatures to detect drive-by downloads. In *Proceedings of the 20th international conference on World wide web*, pages 187–196. ACM, 2011.

Table of Cases

Holomaxx Technologies Corporation v Yahoo! Inc 10-cv-04926 JF (PSG) (ND California United States District Court 2011)

Johnson v Arden 614 F 3d 785 (Court of Appeals, 8th Circuit 2010)

Jones v Dirty World Entertainment Recordings, LLC et al 840 F Supp 2d 1008 (Kentucky Dist Court, ED 2012), (*Jones v Dirty World Entertainment*)

ICI v Shatwell [1965] AC 656

Lanes Giftss & Collectibles and Others v Yahoo! and Others CV-2005-52-1 (Miller County, Arkansas Circuit Court 2005)

Leigh v Gladstone [1909] TLR 26

L'Oréal SA v eBay International AG [2009] EWHC (Ch) 1094

Case C-324/09 *L'Oréal SA v eBay International AG* [2011] ECR I–6011, (*L'Oréal v eBay*)

Patchett v Swimming Pool & Allied Trades Association Ltd [2009] EWCA Civ 717

Patco Construction Co, Inc v People's United Bank, d/b/a Ocean Bank No. 2: 09-cv-503-DBH

Patco Construction Co, Inc v People's United Bank, d/b/a Ocean Bank No. 11-2031, (2012) 684 F 3d 197

Perl (Exporters) Ltd v Camden London Borough Council (1983) 1 QB 342

R v Cambridge Health Authority, ex parte B (A Minor) [1995] EWCA Civ 49

Re C (Adult: Refusal of Medical Treatment) (1994) 1 All ER 819, (*Re C*)

Re R (A Minor: Wardship Consent to Treatment) (1991) 3 WLR 592, (*Re R*)

Re W (A Minor) (Medical Treatment: Court's Jurisdiction) (1992) 3 WLR 758, (*Re W*)

R(on the Application of British Telecommunications and TalkTalk) v Secretary of State for Business, Innovation and Skills and others [2011] EWHC (Admin) 1021

British Telecommunication and TalkTalk Plc v Secretary of State for Culture, Olympics, Media and Sport and others [2012] EWCA Civ 232

Rylands v Fletcher (1868) 3 HL 330

Case C-360/10 *Belgische Vereniging van Auteurs, Componisten en Uitgevers CVBA (SABAM) v Netlog NV* (2012) 2 CMLR, (*SABAM v Netlog*)

Case C-70/10 *Scarlet Extended SA v Société belge des auteurs, compositeurs et éditeurs SCRL (SABAM)* (2011) 2011 ECR I, (*Scarlet Extended*)

Smith v Trusted Universal Standards in Electronic Transactions (and others) 09-4567 (RBK/KMW), (*Smith v Trusted Universal Standards*)

St Helen's Smelting Co v Tipping (1865) 11 HL Cas 642

Stratton Oakmont, Inc v Prodigy Servs Co 94-031063

Twentieth Century Fox Film Corp and others v British Telecommunications Plc and others (2012) 1 All ER 806, (*Twentieth Century Fox v BT*)

Zango, Inc v Kaspersky Lab, Inc 568 F 3d 1169 (2009)

Zeran v America Online, Inc No. 97-1523, (1997) 129 F 3d 327

Table of Statutes

Domestic Law

Computer Misuse Act 1990

Copyright, Designs and Patents Act 1988

Digital Economy Act 2010

Fraud Act 2006

Police and Justice Act 2006

Public Health (Control of Diseases) Act 1984

Public Health (Infectious Diseases) Regulations 1988

European Law

Directive 98/34/EC of the European Parliament and of the Council of 22 June 1998 OJ L204/337

Directive 2000/31/EC of the European Parliament and of the Council of 8 June 2000 on certain legal aspects of information society services, in particular electronic commerce, in the Internal Market OJ L178/1, (E-Commerce Directive)

Directive 2001/29/EC of the European Parliament and of the Council of 22 May 2001 on the harmonisation of certain aspects of copyright and related rights in the information society OJ L167/0010

Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications) OJ L201/0037