The Impact of Incentives on Notice and Take-down

Tyler Moore and Richard Clayton

Computer Laboratory, University of Cambridge,
15 JJ Thomson Avenue, Cambridge CB3 0FD, United Kingdom.
firstname.lastname@cl.cam.ac.uk

Abstract We consider a number of notice and take-down regimes for Internet content. These differ in the incentives for removal, the legal framework for compelling action, and the speed at which material is removed. By measuring how quickly various types of content are removed, we determine that the requester's incentives outweigh all other factors, from the penalties available, to the methods used to obstruct take-down.

1 Introduction

Almost all schemes for the removal of undesirable content from the Internet are described as being a 'notice and take-down' (NTD) regime, although their actual details vary considerably. In this paper we show that the effectiveness of removal depends rather more on the incentives for this to happen, than on narrow issues such as the legal basis or the type of material involved.

It is impractical for Internet Service Providers (ISPs) to police the entirety of the content that their users place upon the Internet, so it is generally seen as unjust for ISPs to bear strict liability, viz: that they become legally liable for the mere presence of unlawful content. However, the ISPs are in an unrivalled position to suppress content held on their systems by removing access to resources – web-space, connectivity, file access permissions, etc. – from their customers. Hence many content removal regimes make ISPs liable for content once they have been informed of its existence, viz: once they have been put on 'notice'. If they fail to 'take-down' the material then sanctions against them may then proceed.

The ISP is often the only entity that can identify customers in the real world, and so they must necessarily become involved before the true originator can be held accountable for the presence of unlawful content. This gives rise to various complexities because the ISP may be bound by data protection legislation, or by common law notions of confidentiality, from disclosing the information haphazardly. Equally, ISPs are reluctant to drawn into acting as the plaintiffs' agent against their own customers – and at the very least demand recompense for their efforts, along with immunities when errors are made. Nevertheless, some benefits

Seventh Workshop on the Economics of Information Security (WEIS 2008), June 25–28 2008.
Created: 22 June 2008

do accrue from including the ISP in the process. They may be more familiar with the process than their customers, allowing them to reject flawed requests and assist in dealing with vexatious claims. The ISP's experience, along with their assessment of the standing of their customer, will enable them to assess the merits of the case, and perhaps advise their customer that the claim should be ignored. An ISP does not have any incentive to annoy a major commercial customer by suspending their website merely because of a dubious claim of copyright in a photograph it displays.

In fact, when we examine NTD regimes, we find that incentives are at the heart of the effectiveness of every process, outweighing the nature of the material or the legal framework for removal. Where complainants are highly motivated, and hence persistent, content is promptly removed. Where the incentives are weak, or third parties become involved with far less of an incentive to act, then removal is slow or almost non-existent.

In this paper we examine a number of notice and take-down regimes, presenting data on the speed of removal. We start by considering defamation in Section 2, which has an implicit NTD regime. In Section 3 we look at copyright which has, particularly in the United States, a very formalised NTD mechanism. In Section 4 we consider the removal of child sexual abuse images and show how slow this removal can be in practice. In Section 5 we contrast this with the removal of many different categories of 'phishing' websites, where we are able to present extensive data that illustrates many practical difficulties that arise depending upon how the criminals have chosen to create their fake websites. In Section 6 we consider a range of other criminal websites and show that their removal is extremely slow in comparison with phishing websites, and we offer some insights into why this should be so. In Section 7 we consider the issues surrounding the removal of malware from websites and from end-user machines, and discuss the incentives, such as they are, for ISPs to act to force their users to clean up their machines – and in particular to stop them from inadvertently sending out email 'spam'. Finally, in Section 8 we draw the various threads together to compare and contrast the various NTD regimes.

2 Defamation

Until very recently, access to mass media was limited by professionals (printers, newspaper editors, etc.) acting as gatekeepers, so that people could rarely express defamatory opinions to large audiences. Consequently, actions for defamation (sometimes technically distinguished by terms such as slander or libel) were also rare, and would typically involve a hand-written note (*Wilde v. Marquis of Queensbury 1895*), a letter (*Huff v. Huff 1915*), or a message on a golf club notice-

board (*Byrne v. Dean 1937*).[1] The defamation itself was usually obvious, and these cases were decided upon the interpretation of the law, or the unmasking of perjury by one of the protagonists.

Conversely, where books or mainstream journalism are involved, the gatekeeper has already formed an opinion that the material should be published, and may have edited the content to reduce the risk of action. These court cases typically revolve around 'justification' (whether the claim was in fact true) such as in *Irving v. Lipstadt 2000*, whether the court has jurisdiction (*Kroch v. Rossell 1937*), or the extent of the actual defamation (*Whistler v. Ruskin 1878*).

The Internet, and the ease of posting articles to Usenet (*Godfrey v. Demon Internet 1998*) or to web sites (*Totalise v. Motley Fool 2001*), has fundamentally changed the landscape to one in which there are no gatekeepers, but individuals are capable of rapidly propagating their words to a wide audience. This makes defamation more common – although since it is a civil matter, state aid for plaintiffs is seldom available, meaning that only the well-heeled, or those with the most open-and-shut cases, ever take action.

In the United Kingdom the law was revised by the Defamation Act 1996, when the Internet was already being widely used, albeit with 'user-generated content' still in its infancy. The Act enshrined the existing common law principle of 'innocent dissemination', and made it clear that distributors of defamatory material had a statutory immunity until they became aware of the nature of the material. In the context of Internet content this means that ISPs are not liable for a defamatory statement until they are put on notice of the existence of the material, and court actions will not succeed if they promptly take-down the material.

In 1999, in a pre-trial hearing of *Godfrey v. Demon Internet*, Morland J. set out the extent of the statutory (and common law) immunity in a lengthy judgment that struck out part of Demon Internet's defence (Morland 1999). The case was then settled out of court. This immediately led to a step-change in the number of notices being served upon Demon Internet – as people became aware of serving notices as a simple and effective method of removing defamatory material (Clayton 2000).

In the USA, defamation has some subtle differences from UK law. In particular if the statement relates to a public figure then it is necessary to prove 'actual malice'. More relevant to the present discussion, American ISPs also have a qualitatively different defence in that s230 of the Communications Decency Act 1996 gives an immunity where information is merely being transmitted and was originated by a third party. The leading case is *Zeran v. AOL* from 1997, which made it crystal clear that, in the US, serving a notice does not put an ISP under any obligation to take-down defamatory content.

Very few people can afford to pursue defamation cases to court, and those who do so are often very well-known. This leads to 'forum shopping' where cases are pursued in the British or Australian courts rather than in the United States. Con-

[1] In this paper we cite UK cases; similar developments occurred in the US, Australia and other jurisdictions.

versely, where a court action is unlikely, the practical effect of the different US and UK regimes is that when defamatory content is published on a UK web site it will be fairly promptly removed by the ISP. When it is republished on a US site, it then remains available for a considerable time.[2]

3 Copyright Violations

Rights holders have long complained about copyright violation by Internet users. As individuals have been able to access more and more bandwidth the focus of attention has moved from photographs, to songs, to feature films. Although most 'sharing' now takes place on peer-to-peer networks, the original mechanisms were the use of websites or Usenet articles. To deal with the dominance of client/server architectures of the time, the US passed Title II of the Digital Millennium Copyright Act 1998 (DMCA).

The DMCA gives an immunity (a 'safe harbor') to ISPs operating web or Usenet servers if they follow certain rules. They must provide a contact address, and if served with a valid notice alleging copyright infringement, they acquire 'actual knowledge' and must 'expeditiously' remove the material. However, if they are served with a 'counter notification' (a 'put-back' request) by their customer then they must restore the material 10 to 14 days later unless the matter has gone to court. The put-back notice has to identify the customer, who must submit to the court's jurisdiction, viz: they must firmly identify themselves as standing behind their claim that the take-down notice was mistaken.

Similar take-down provisions exist in the 2000 European Union 'Directive on Electronic Commerce', which gives a similar immunity to ISPs using very similar language – 'actual knowledge', 'expeditiously' etc. However, the Directive does not set out a put-back provision.

There have been many claims that the EU regime creates incentives for ISPs to remove items first without even bothering to ask questions afterwards. Two experiments have been performed to demonstrate this. In 2003 an Oxford research group posted material onto UK and US websites (Ahlert 2004). The material was an extract of John Stuart Mill's 1869 'On Liberty', discussing freedom of speech. The experimenters then wrote anonymously to the two hosting ISPs, falsely claiming that the material was still in copyright. The UK ISP removed the material, whereas the US ISP insisted upon the provision of the legally necessary 'on pain of perjury' declaration on the DMCA notice – which the researchers were not pre-

[2] It would be possible to give numerous instances of sites that have migrated to the US, however we have not provided any examples of this alternative type of 'forum shopping' because the authors are currently resident in the UK. In this jurisdiction there is some case law about providing pointers to defamatory material: in *Hird v. Wood 1894*, the court held that the defendant had defamed the plaintiff by merely standing on a road and mutely pointing out a path which, if followed, allowed one to view a notice on which a defamatory statement had been written, even though the authorship of that statement was never proved.

pared to make, so the material remained available. The researchers concluded that there was a substantial difference between the US and UK in how easily websites are removed, although one suspects that if they had perjured themselves, the difference would have disappeared.

In 2004 a similar experiment was performed by the Netherlands-based 'Multatuli Project' (Nas 2004). They placed some out-of-copyright material from a famous 1871 tract onto webspace provided by ten different Dutch ISPs. Their results were mixed, with some ISPs losing their first complaint and only acting on a follow-up message. By the end of the experiment, seven of the ten ISPs had removed the material, taking between 3 hours and 3 days to do so. However, in neither investigation did the customer protest the removal decision and suggest that the ISP taking the complaint at face value was incorrect. Hence the experiments do not necessarily represent the true situation, but merely show that ISPs are generally keen to avoid liability, do not establish the accuracy of complaints, and may need to be asked more than once before they act.

4 Child Sexual Abuse Images

Child sexual abuse images are often perceived as the most widely condemned form of Internet content, but this universality is relatively recent and remains inconsistent. For example, Japan did not pass its 'Child Prostitution and Child Pornography Prevention Law' until 1999.

Harmonisation of the laws in this area was one of the aims of the 2001 Convention on Cybercrime, but this has several optional aspects: the age limit should be 18, but can be as low as 16; simple possession need not be made a crime; and computer-generated material, no matter how realistic, may be tolerated. The last of these issues is a point of departure between the UK and the US. In the UK, child sexual abuse images generated on a computer are illegal if they are realistic enough to appear to be a photograph (viz: they are a 'pseudo-photograph'), whereas the US Supreme Court held in *Ashcroft v. Free Speech Coalition 2002* that since no real children were involved in creating this type of material, it was unconstitutional to ban it.

Notwithstanding these minor variations, the bulk of child sexual abuse images are illegal to distribute in all relevant jurisdictions, and hence it should be expected that any such material is promptly removed.

The Internet Watch Foundation (IWF) was founded in the UK in 1996 to operate a 'hotline' for reports of child sexual abuse images from the public. It employs trained staff to check these reports and pass them on to the UK police if illegal material is found. If the sites are in the UK then the police will act upon them directly, whereas if they are hosted elsewhere in the world then a report will be passed to the authorities in that country. Within the UK, the IWF will also pass the report directly to the ISP and, in the case of illegal images circulating on Usenet, will pass a report to all UK ISPs so that they can remove the article from their

servers. The IWF is a member of INHOPE,[3] and it will send a report to another INHOPE member if the material appears to be hosted in their country.

The IWF regularly publish statistics on where illegal images are hosted, but until recently they have not measured how long it takes to get them removed, and their figures still remain patchy. Anecdotally their view is that sites are generally removed in weeks rather than months, although they have much higher expectations for sites in the UK – where removal is expected within hours, a couple of days at most. Removal became so rapid that at one stage the IWF were reporting sites to the police first, so that evidence of what was on the site could be collected, and only two days later would they report the site to the relevant UK ISP.

The IWF have published a smattering of data on site longevity, which – since only fractions of a percent of all sites are now hosted in the UK – will in practice measure the speed of take-down of internationally hosted content. In mid-2006 they checked which sites were available at the start and end of a six week period and found that 287 (circa 20%) of sites had survived, including one dating from 1999 which had been the subject of 20 separate reports to the authorities (IWF 2006a). In 2007 (IWF 2007) they reported that 94 sites had been active for a year or more, 33 for two years and 32 for longer, and in late 2006 (IWF 2006b) that of the commercial sites they tracked (about half of the total) 62% were removed in a month, and 2% lasted more than a year.

To supplement the published reports, the IWF kindly provided us with sanitised data on the websites they track. They use an automated system which performs daily checks on whether the offending content remains available. Whenever the system detects removal, operators manually inspect the page to ensure it has been removed. The logs given to us include a pseudonym for the suspected URL, the date reported, the date removed and the date of reappearance (if observed).

We computed the lifetimes for websites reported during the calendar year 2007, which in some cases were already known to the IWF, but were mainly new. The total number of domains was 2585, although of course the number of individual pages with child sexual abuse images was much higher. We excluded 8 domains which had more than 100 individual reports, which we believe to be well-known 'free' web-hosting sites.

The lifetime of each website was calculated by comparing the date of first report to the date of first removal. Consequently, we do not consider how long it takes to remove any subsequent reappearance of images on the same website.[4]

The results are given in Table 1. Of the 2585 website domains reported to be hosting child sexual abuse images in 2007, nearly all (2531) of the websites removed images at least once on or before April 3, 2008. It took an average of 562 hours – over three weeks – to take down images hosted on these websites. The median lifetime is 264 hours, or 11 days. 54 websites reported in 2007 (2.1% of

[3] The International Association of Internet Hotlines: http://www.inhope.org

[4] The data provided by the IWF presents numerous difficulties whenever images reappear on the same website. Using the sanitised data they made available, it is impossible to distinguish between similar and distinct removals on the same website.

the total) have never had images removed. The average lifetime for these sites is 338 days (and growing). Combined together, the mean lifetime of all websites found to be hosting child sexual abuse images in 2007 is 719 hours (30 days).

Table 1. Lifetimes for Websites Hosting Child Sexual Abuse Images

-	Sites	Lifetime (hours)	
		mean	median
Removed websites	2531	562	264
Unremoved websites	54	≥8027	≥9216
Total	2585	719	288

While we have not measured the time to remove images when they reappear on websites, we have determined that within 24 weeks, images reappeared on 1070 sites, 41% of the total. Sometimes offenders reload new images onto free webspace, while at other times insecure websites are simply recompromised (we describe techniques for publishing illicit content below).

5 Phishing

Phishing is the term used when criminals entice people into visiting websites that impersonate the real thing, duping them into revealing passwords and other credentials, which will later be used for fraud. Many types of company are attacked in this way, from domain registrars, through auction sites and multi-user games to online merchants, but the vast majority of attacks are against financial institutions: banks, credit unions, credit card companies, online share brokers and so on.

In previous work we have identified wide variations in take-down time for different financial institutions and different types of attacker (Moore and Clayton 2007). We have subsequently determined that some of the variation can be ascribed to the company charged with removing the sites being unaware of its existence, viz: that no notice was issued, so no take-down occurred (Moore 2008).

In this paper we examine phishing attacks against a particular e-commerce company that conducts business using two very well-known brands, both of which are in the top 600 most visited websites in the world. We consider the data for attacks that were first reported during the month of January 2008. The lifetime figures we give are from the earliest point at which we know the site existed, to the last time that our monitoring system[5] indicates that it was hosting a fake page.

[5] For a detailed account of our 'feeds' of URLs of phishing websites, and our monitoring system we refer the interested reader to (Moore 2008; Moore and Clayton 2007; Moore and Clayton 2008). In the current context, the key point is that because we receive data from a number of disparate sources, we believe that our database of URLs is one of the most comprehensive available, and the overwhelming majority of phishing websites will come to our attention.

Where we can do no better, we use the timestamp from when we receive the URL, but almost all of our feeds, including the one from the brand owner we are considering, provide a timestamp from when they entered the URL into their internal systems – which we assume to be within a few minutes of when they start to verify the nature of the site and set their take-down processes into motion.

To avoid being traced, phishing attackers will not host the fake websites on their own personal machines. Some attackers use free webspace, where anyone can register and upload pages, but it is more common to encounter sites that are hosted on compromised machines; perhaps a residential machine, but often a server in a data centre. The hijacked machine will have come under the attacker's control either through a security vulnerability (typically unpatched applications within a semi-abandoned 'blog' or message-board), or because the user is running some malware, delivered by email or downloaded during a visit to a malicious website.

It is possible to distinguish these cases by examining the URL and by checking the IP address of the website. We now consider the various different types of hosting for phishing websites, and how NTD works in each case. The website lifetimes for the different phishing attack methods discussed in this section are listed in Table 2.

Table 2. Phishing Website Lifetimes by Attack Type

-	Sites	Lifetime (hours)	
		Mean	median
Free web-hosting			
all	395	47.6	0
brand owner aware	240	4.3	0
brand owner missed	155	114.7	29
Compromised machines			
all	193	49.2	0
brand owner aware	105	3.5	0
brand owner missed	155	103.8	10
Rock-phish domains	821	70.3	33
Fast-flux domains	314	96.1	25.5

5.1 Free Web-hosting

A typical URL for a website that has been set up at a free web-hosting provider is http://www.bankname.freehostsite.com/login, where the bankname is chosen to match or closely resemble the domain name of the financial institution being attacked. We compiled a list of known free web-hosting domains, and used this to determine which websites we were monitoring were hosted on free space. We also checked

the IP addresses of the websites against the IP address ranges used by the free providers.

To get the phishing website removed all that is necessary is to contact the webspace provider and draw their attention to the fraudulent site. They will then remove it and cancel the hosting account. In earlier work we commented on the wide disparity in take-down times between different providers, and upon a 'clued-up' effect, whereby when webspace providers were first exploited they would not know how to deal with the situation, but after a while they would acquire 'clue' and settle down to a steady-state removal time (Moore and Clayton 2007).

Some attackers favour hosting attacks on free webspace. In January 2008, for the two brands we are considering, we learnt of 395 phishing websites that were hosted on free webspace. The majority of these websites were removed before we could visit them, giving a median lifetime of 0 hours. However, there are a number of very long-lived websites, which dragged the mean lifetime up to 47.6 hours.

To understand why the mean is so much larger than the median, it is necessary to examine which websites were known about by the brand owner. Only 240 of the 395 free-hosting phishing websites impersonating the brands turned up in the company's own feed of phishing website URLs. This subset of websites was removed very quickly – 4.3 hours on average, with a 0 hour median. By contrast, the 155 websites that we learnt about from other sources, but the company remained ignorant of, had an average lifetime of 114.7 hours with a median of 29 hours.

5.2 Compromised Machines

For compromised machines, attackers may have restricted permissions, and are limited on where files can be placed. They add their own web pages within an existing structure, leading to URLs for their websites that have the typical form **http://www.example.com/user/www.bankname.com/** where, once again, the **bankname** is present to lend specious legitimacy should the user check which site they are visiting, yet fail to appreciate the way in which URLs are really structured.

The attacker may occasionally find that the existing DNS configuration permits URLs of the form **www.bankname.com.example.com**, but if this is not possible, and if the **example** part of the hostname makes it unlikely that the URL will be convincing, then the URL may use just the IP address of the compromised machine, perhaps encoded into hexadecimal to obscure its nature. Alternatively, to further allay suspicion, the fraudsters will sometimes go to the effort of registering their own domain name, which they will then point at either free webspace (as just discussed), which can often be configured to allow this to work, or to a compromised machine where they have sufficient control of the web server configuration so that it will respond to page requests. The domain names are usually chosen to be a variation on **bankname.com** such as **bankname-usa.com**, or they will use the bank's name as a subdomain of some plausible, but superficially innocuous domain, such

as bankname.xtrasecuresite.com. A half-way house to an actual domain name is the use of systems that provide domain names for dynamic IP address users, which results in the usage of domains such as bankname.dyndns.org.

In order to get a website removed from a compromised machine it is generally necessary to get in touch with the sysadmin who looks after it. In some cases that information can be gleaned from public records or from the rest of the website. In other cases it is necessary to work through the ISP to get a message delivered. Less commonly, where a domain name has been registered especially for the phishing attack, it is necessary to approach the appropriate domain name registrar and ask them to suspend the name.

We examined the attacks on the two brands by phishing websites that were hosted on compromised machines in January 2008 and found 193 websites[6] with an average lifetime of 49.2 hours and a 0 hour median, which is very similar to the lifetimes we measured for free web-hosting sites.

The similarities between compromised machines and free web hosts continue once we break down the lifetimes according to whether the brand owner was aware of the website. The 105 phishing websites hosted on compromised machines known to the company are removed within 3.5 hours on average (0 hour median). The 88 websites missed by the company remain for 103 hours on average, with a median of 10 hours.

Thus, for ordinary phishing websites, the main differentiator appears to be whether the organisation responsible for the take-down is aware of the site's existence. Free web-hosting companies and the administrators of compromised machines both appear to comply promptly with the take-down requests they received. However, the website administrators do need to be notified of the problem – phishing websites that brand owner did not know about, and so did not issue any notices for, remain up for considerably longer.

5.3 Rock-phish and Fast-flux Attacks

The 'rock-phish' gang operate in a completely different manner from the ordinary phishing attacks just described. This group of criminals perpetrates phishing attacks on a massive scale (McMillan 2006). The gang purchases a number of domains with meaningless names such as lof80.info. Their spoof emails contain a long URL of the form http://www.bank.com.id123.lof80.info/vr. Although the URL contains a unique identifier (to evade spam filters), all variants are resolved to a single IP address using 'wildcard DNS'. The IP address is of a machine that acts as a proxy,

[6] While our method for identifying compromised websites from the structure of phishing URLs has confirmed 193 websites, there are additional websites that we have not yet verified. Hence, the 193 websites should be viewed as a sample of a significantly larger population of compromised websites.

relaying web traffic to and from a hidden 'mothership' machine. If the proxy is removed, the DNS is adjusted to use another proxy, and so the only practical way to remove the website is to get the appropriate registrar to remove the domain name from the DNS.

A related form of attack is dubbed 'fast-flux'. The mechanism is similar to the one employed by the rock-phish gang, except that the domain name is resolved to many IP addresses in parallel (typically 5 or 10) and the IP addresses used are rapidly changed (sometimes every 20 minutes). For these attacks the only practical approach is to have the domain name suspended. We have identified several disjoint fast-flux networks. Interested readers can find more details of fast-flux in (Honeynet Project 2007), which gives many details about one of the networks that we also encountered, and about its use in phishing attacks in (Moore and Clayton 2007). Unlike rock-phish attacks, fast-flux networks are made available for hire as a type of 'bullet-proof' hosting. Hence, they are used for other types of attack in addition to phishing. We discuss the use of fast-flux domains for hosting online pharmacies in Section 6.3.

Besides using an innovative architecture, the rock-phish gang also attack multiple banks in parallel, with the URL path distinguishing between them. Since these bank 'micro-sites' generally appear and disappear together, we monitor the rock-sites generically, tracking whether the domain name remains active. For convenience, we track fast-flux sites in a similar manner, although they may attack only a single bank.

The rock-phish and fast-flux attack methods are not universally understood by the registrars who are asked to suspend domains. Splitting up the components of the attack (domains, compromised machines and hosting servers) obfuscates the phishing behaviour. Hence, each individual decision maker cannot easily recognise the nature of the attack – the domain registrar does not see an obviously impersonated domain name (e.g., barclaysbankk.com) and the ISP sysadmin does not find HTML for a bank site in a hidden sub-directory on a hijacked machine. Recent activities have highlighted the confusion domain name registrars are experiencing in addressing the threat from rock-phish attacks. Email-blacklist operator Spamhaus engaged in a public row with the Austrian domain registrar nic.at over the registrar's initial refusal to remove rock-phish domains (Spamhaus 2007).

The two brands we have studied so far have only been very briefly targeted by rock-phish and fast-flux attacks, so we instead examine all attacks of this type, irrespective of brand. The lifetime of the 821 rock-phish domains we monitored in January 2008 reflects the added difficulty faced during take-down procedures. The domains lasted 70.3 hours on average (median 33 hours), despite the additional attention rock-phish domains attract by impersonating many banks simultaneously. The lifetimes for the 314 fast-flux domains were similar, lasting 96.1 hours on average with a 25.5 hour median.

5.4 Common Features of Phishing Website Removal

As has been seen, phishing websites are generally removed fairly promptly. This might be viewed as quite remarkable given the multiple jurisdictions involved. The site may be in a different country than the bank, and the take-down company making the request may be in a third location. Furthermore, it is most unusual for the police or the courts to be involved in the procedure. There is no legislation anywhere prescribing the elements that need to be present on the notice – or indeed specifying what the penalties might be for ignoring the notice. In practice the vast majority of sysadmins have an understanding of what phishing is, they recognise the site as being part of a criminal enterprise, and they remove it. This was not always so – we have been told that when phishing was first starting in 2003 it was often more effective to point out the intellectual property infringements apparent on the website: the unauthorised use of logos, the similarity of design and text, and even in some cases, the unauthorised use of a particular rights-encumbered font.

Although the phishing sites are usually taken down it is, unfortunately, quite common for similar sites to reappear quickly. This occurs because the free web-hosting site does not have mechanisms to check for identical content being uploaded by a 'different' person; because the sysadmin for a compromised machine does not patch the security hole that led to the compromise; or because the registrar does not tighten up their procedures to prevent the purchase of domain names using the same *modus operandi* as the instance just suspended. Looking at comprehensive phishing data from October 2007 to March 2008 we found that approximately 22% of all compromised machines were recompromised within 24 weeks. A more detailed analysis of phishing-website recompromise and its causes can be found in (Moore and Clayton 2008). The recompromise rate for phishing is noticeably smaller than the 41% rate found for websites hosting child sexual abuse images. However, the comparison is somewhat inexact, since we cannot exclude all of the instances of free web-hosting from the analysis of child sexual abuse images due to our sanitised data source.

6 Fraudulent Websites

As the preceding discussion of the take-down of phishing websites shows, institutions being impersonated often have a very strong incentive to remove offending content. Consequently, miscreants have designed a number of scams that escape such scrutiny by creating websites for entirely fake institutions. While many types of fraudulent websites exist, we discuss three classes in this section: fake escrow agents, mule-recruitment websites and online pharmacies.

6.1 Fake Escrow Agents

One lucrative type of fraud is to set up fake escrow agents. Escrow agents serve as trusted intermediaries to facilitate large financial transactions between untrusted parties. For instance, someone buying a car on eBay might not want to pay the seller until she has received the car; likewise the seller might not want to ship the car until she has been paid. An escrow agent takes the money and the goods and completes the transaction once both parties have acted. In an escrow scam, a rogue seller offers an expensive item at a reasonable price. Once a buyer has been found, the seller suggests that they use an escrow agent of her choosing and points the buyer to the web page of the fake agent. The buyer sends her money to the fake escrow agent, only to later realise a fraud has occurred.

Fake escrow websites have been used extensively for the past few years. Because no organisation is being impersonated, no company tries to remove the websites. Only motivated volunteers, primarily acting through Artists Against 419 (AA419),[7] attempt to take down the websites. Initially, removing fake escrow websites took a very long time. Eventually, escrow website lifetimes diminished as volunteers developed efficient take-down procedures and established trust with ISPs. Meanwhile, the creators of the fake escrow agents have continued to turn the handle, creating 'new' companies with websites that borrow generously from prior incarnations. Hence, the battle between the fake escrow agents and volunteers has reached somewhat of a steady state.

We examined 696 fake escrow websites appearing between October and December 2007. On average, these websites remained for 222 hours, or over 9 days. The median lifetime was 24.5 hours, approximately one day. The volunteers are definitely making an impact, but given their limited resources they are certainly not as successful as the banks removing phishing websites.

Note that our analysis of lifetimes only includes escrow websites known to AA419, which is the only group we are aware of that are actively removing the websites.[8] It is undoubtedly the case that additional fake escrow websites exist. Most fake-escrow pages include curious phrases such as:

> "Thanks to our innovative view of courier transport and to our commitment to provide a competitive service, we soon were ahead in the sector, leaving all the traditional Trans companies behind."

While the names of the companies are changed frequently along with the URLs, the website content often remains the same. Using targeted web searches for 81 peculiar phrases repeatedly used on the escrow pages, we identified many more websites than those listed by AA419. Each web search found 9.8 domains on average, while approximately 2.4 of these domains were picked up by AA419.

---

[7] http://www.aa419.org

[8] Occasionally the legitimate escrow service **escrow.com** goes after fake sites that infringe upon their brand. Of course, additional volunteer groups may be operating, but we are unaware of any.

In all likelihood, these additional websites remain for much longer than those identified by the volunteers. So the fairest comparison between escrow and phishing websites is between the lifetimes of sites known to both removing parties. The average removal time of 4 hours (0 hour median) for the websites the brand owner knows about compares very favourably to the lifetimes of 222 hours (24.5 hour median) of escrow sites that are known to the volunteers.

6.2 Mule-recruitment Websites

One of the biggest challenges for phishing attackers is to 'launder' the proceeds obtained from victims. One method is to recruit 'money mules', who receive transfers of money from compromised phishing victim accounts, take a cut, and then forward the rest to third parties using non-revocable transactions, typically Western Union transfers. When the fraudulent transfers are detected, they are often reversed, and so it is often the case that the mule ends up out of pocket, rather than the original phishing victim.

Prospective mules are mainly recruited by sending spam email. Often the spam includes only an email address for correspondence. Other times there are links to a website of the purported company which is 'hiring' for jobs such as 'transaction processors', or 'sales executives'. Sometimes the mule-recruitment websites impersonate a legitimate business, but the company may be entirely fictitious. The existence of the website must be assumed to be important in engendering trust by the mule – who may even receive signed 'contracts of employment'. The apparent legitimacy makes the mules far more likely to ignore warnings given by Western Union against sending money to 'strangers'.

Table 3 Lifetimes of Mule Recruitment Websites

Company Name	Real?	Period	Sites	Lifetime (hours)	
				mean	median
Lux Capital	✓	Mar–Apr 2007	11	721	1050
Aegis Capital	✓	Apr–May 2007	11	292	311
Sydney Car Centre	✗	Jun–Aug 2007	14	171	170
Harvey Investment	✓	Sep–Oct 2007	5	239	171
Cronos Investment	✗	Oct–Nov 2007	12	214	200
Waller Truck	✓	Nov–Feb 2008	14	237	3
Overall			67	308	188

We tracked several sets of mule-recruitment websites during the course of 2007/8. These sites were clearly linked by the style of spam email sent and the

way in which a new set of sites commenced when an old set tailed off. Table 3 summarises the mule-recruiting companies, the number of different domain names that were used and for how long the websites remained available.

Where the websites impersonated existing companies we found that the usual response has been that a warning notice is placed onto the company's legitimate website, as in Figure 1. This appears to be intended to discourage correspondence, rather than to actively combat the money laundering. However, in one case the impersonated company does seem to have been more proactive. On the 7th and 8th of October 2007 the Draper Investment Company was impersonated under 7 different domain names. They were simultaneously removed around noon on the 9th of October (maximum lifetime 62 hours, mean 40 hours, median 39 hours). On 17th October the fictitious Cronos Investment Company made an entrance – the website design was identical to that of Draper Investment, except for the name. These sites clearly didn't engender the same reaction because they lasted considerably longer.

Figure 1. Warning Message on the Real Harvey Investment Web Page

Even though the money-laundering advertised by these sites directly harms banks attacked by phishing, none of the banks or take-down companies actively pursue the mule-recruitment websites. Although individual companies occasionally take action, we believe that in general only volunteer groups such as AA419 attempt to remove these sites. Even the volunteers treat these sites as less of a priority because the mules are seen as being complicit in phishing crimes. The lifetimes certainly reflect the lack of priority. The mule-recruitment websites we tracked had a lifetime of 308 hours (188 hours median). This is noticeably longer than for phishing, where the banks are actively seeking removal. It is also consid-

erably longer than for escrow websites, which may again reflect the priorities of the volunteers.

For one set of phishing websites, impersonating Waller Truck, we also made a collection of 256 incoming spam emails received by one of the authors between 27 Nov 2007 and 20 Feb 2008, which promoted 44 different URLs. By assuming that the sites were live at the moment each email was sent, we were able to calculate an alternative view of the lifetimes. The 33 URLs, for which more than one email was received, had a mean lifetime of 265 hours (median 124 hours).

We also applied 'capture-recapture' analysis to this data: AA419 and PhishTank[9] between them knew of 18 sites, of which 14 were in the emails received. Hence we estimate the total population to be 57.[10] We were also able to do a similar analysis for the Cronos websites. In this case we did not have the timestamps for the URLs, but the other sources knew of 12 sites, 10 of which overlapped with the email collection of 31 sites. In this case the overall population is estimated to be 37. Venn diagrams indicating the overlap are given in Figure 2.

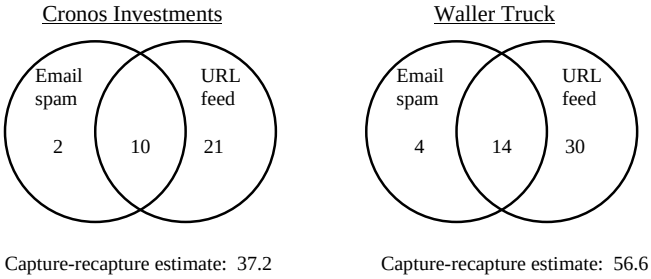

Figure 2. Venn Diagram Comparing Coverage of Mule Recruitment Websites from an Email Spam Source and the URL Feeds

The similarity of results from the email analysis and from the website lifetime measurements, and the relatively small number of sites that are likely to have been missed, means that even though relatively small numbers of sites have been tracked, the estimates of lifetimes seem reasonably robust.

[9] http://www.phishtank.com

[10] We use the standard formula for capture-recapture: $\frac{|sample_1| \times |sample_2|}{|overlap|}$ Our data does not satisfy all of assumptions necessary for this formula to hold – notably the population is dynamic, with sites appearing and disappearing. (Weaver and Collins 2007) computed the overlap between two phishing feeds and applied capture-recapture analysis to estimate the number of overall phishing attacks. They discuss how the capture-recapture assumptions can be accommodated for phishing. We leave deriving a more accurate estimate to future work.

6.3 Online Pharmacies Hosted on Fast-flux Networks

In Section 5.3 we described a particular category of phishing attacks called 'fast-flux'. Some online criminals have constructed a fast-flux network and made it available for hire. Clients include any group wishing to host material that is the target of NTD procedures. We have already presented lifetime figures for phishing websites hosted on fast-flux networks. We also examined 82 domains used by an online pharmacy from October to December 2007. The lifetimes of these domains is much longer than for the fast-flux domains used for phishing websites. The pharmacy domains remain for an average of 1370.7 hours, or over 8 weeks. The median lifetime is slightly longer at 1404.5 hours.

From these figures, it appears that almost no one is attempting to remove the online pharmacies, even though they are illegal and advertised through spam email. There is a large disparity between how long pharmacy websites remain available and the lifetimes of fast-flux phishing websites. This demonstrates that the longevity of the domains depends less on the hosting method used, and more on whether anyone is motivated to remove the offending content.

7 Spam, Malware and Viruses

The sending of unsolicited bulk email ('spam') has gone through considerable evolution over the past decade. The main senders are now compromised end-user machines on broadband links, which are operating as a part of a 'botnet'. The 'zombie' machines forming the botnet can be commanded to send spam, to scan other systems for security weaknesses, or to send large amounts of traffic as a part of a DDoS (distributed denial of service) attack. The operators of botnets hire out their services within a fairly sophisticated 'underground economy' (Franklin et al. 2007; Thomas and Martin 2006), sending spam for other criminals to promote their activities.

Until relatively recently, recruitment of machines into these botnets was done by the sending of email viruses – code that executed when the user was tricked into opening email attachments. The virus would then replicate itself by emailing everyone in the user's address book, while contacting a botnet controller for further instructions. A later refinement was to avoid the replication – which attracted attention – and only send out new copies of the code when the botnet was too small. The main infection vector at present appears to be placing malware onto websites in such a way as to infect visitors who have not applied security patches to their systems, or who can merely be inveigled into executing code supplied by the website.

Various volunteer organisations track these disparate activities. For example, Spamhaus[11] collates lists of IP addresses that are known to send spam email; Team Cymru[12] tracks various botnets by their characteristic scanning activities; and researchers such as (Enright 2007) and (Dagon et al. 2006) have developed ways of tracking particular botnets. In a slightly different realm, the StopBadware project[13] tracks websites that are infected with malware.

Some of these activities generate reports to the ISP whose customer's machine has been compromised. Additionally, many major ISPs (such as AOL[14] and MSN[15]) operate 'feedback loops', to let other ISPs know about incoming spam email to their servers. Finally of course, individual recipients of spam emails, or DDoS attacks, may generate reports of their own.

These reports are, in essence, notices of bad activity, and the expectation is that the ISP will pass the report on to their customer and the wickedness will be 'taken down'. The incentives for this take-down are weak, boiling down to implicit threats to block further traffic, or to name-and-shame ISPs that fail to act effectively. Some of the dynamics of blocking were discussed by (Serjantov and Clayton 2005), but there has been little other academic study.

The lack of incentives has been commented upon by legislators such as the UK House of Lords Science and Technology Committee, who recommended (House of Lords 2007) that after a short period the ISP should become legally liable for ongoing bad traffic (#3.69). However they found it hard to square the incentives with a wish to see ISPs being more proactive in monitoring (#3.68), and in any event, the UK Government totally rejected the proposal (United Kingdom Government 2007). Incentives to increase ISP participation in cleaning up compromised end-user machines also feature strongly in the recommendations made by (Anderson et al. 2008) in their report to the European Network and Information Security Agency (ENISA).

It would be instructive to measure the take-down times for this category of material because there are weak incentives on ISPs to act, but at the same time, there are strong incentives by complainants to see action taken. However, there are no published figures for the lifetimes of spam-sending (or DDoS-participating) machines. A key reason for this is that many of the end-user machines involved use dynamic IP addresses. Consequently, lifetimes may be artificially lowered (and the number of sources greatly exaggerated) by the problem machines regularly changing to new IP addresses – altering the only marker by which they can be distinguished.

[11] http://www.spamhaus.org

[12] http://www.team-cymru.org

[13] http://www.stopbadware.org

[14] AOL Feedback Loop Information: http://postmaster.aol.com/fbl/

[15] Microsoft Smart Network Data Services: https://postmaster.live.com/snds/

8 Comparing Take-down Effectiveness

We have just described many of the different categories of web content subject to NTD requests. For several categories we have obtained data on the associated websites' lifetimes. While the circumstances and assumptions for each category often vary, we can still draw useful comparisons. Table 4 summarises the lifetime data we have presented.

It is apparent that the presence of incentives to remove offending material has the greatest impact on website lifetimes. By far, phishing websites are removed fastest. Banks are highly motivated to remove any impersonating website because their continued appearance increases losses due to fraud and erodes customers' trust in online banking. Solid legal frameworks do not seem to matter as much. Courts almost never get involved in issuing orders to remove phishing websites. By contrast, other clearly illegal activities such as online pharmacies do not appear to be removed at all.

Table 4. Website Lifetimes by Type of Offending Content

-	Period	Sites	Lifetime (hours)	
			mean	median
Child sexual abuse images	Jan–Dec 2007	2585	719	288
Phishing				
Free web-hosting (two brands)	Jan 2008	240	4.3	0
Compromised machines (two brands)	Jan 2008	105	3.5	0
Rock-phish domains (all brands)	Jan 2008	821	70.3	33
Fast-flux domains (all brands)	Jan 2008	314	96.1	25.5
Fraudulent websites				
Escrow agents	Oct–Dec 2007	696	222.2	24.5
Mule-recruitment websites	Mar 07–Feb 08	67	308.2	188
Fast-flux pharmacies	Oct–Dec 2007	82	1370.7	1404.5

However, the banks' incentives are not perfectly aligned. Most banks remain narrowly focused on actively removing only those websites that directly attack their brand. Another key component of the phishing supply chain, mule-recruitment websites, is completely ignored and left to volunteers. Removing mule-recruitment websites is a collective-action problem: many banks are harmed by these websites, yet none takes action because they cannot be sure whether removing them will help themselves or their competitors. This lack of cooperation is somewhat surprising, given that there are numerous organisations within the financial sector whose remit includes tackling collective threats (e.g., the Financial

Services Technology Consortium (FSTC), the Financial Services ISAC in the US, APACS in the UK, and the Anti-Phishing Working Group (APWG)).[16]

Duped consumers are harmed most by fake escrow agents, yet they are in no position to remove the websites. Auction houses such as eBay do have a weak incentive to remove escrow websites, in that their continued existence undermines trust in online commerce. Volunteers are likely motivated by a sense of justice, but the figures demonstrate such an incentive is not entirely sufficient. Even the two experiments with spurious copyright infringement removal requests show what matters most is the perseverance of the requesters.

The technology chosen by the attacker does affect the speed of take-down, but the impact is much smaller than the incentive to remove. While the use of free web-hosting and compromised machines for hosting phishing websites exhibited similar lifetimes, the evasive techniques employed by the rock-phish gang and in fast-flux attacks leads to substantially longer lifetimes. However, online pharmacies using fast-flux techniques remain 14 times longer than phishing websites using the same approach. This provides further evidence that the defender's motivation for removal matters far more than the attacker's implementation strategy.

8.1 Lifetimes of Child Sexual Abuse Image Websites

The long lifetimes of websites hosting child sexual abuse images is particularly striking. In spite of a robust legal framework and a global consensus on the content's repulsion, these websites are removed much slower than any other type of content being actively taken down for which we have gathered data. An average lifetime of 719 hours is over 150 times slower than phishing websites hosted on free web-hosting and compromised machines. Since we are not privy to the hosting method used by child sexual abuse image websites, we do not know whether sophisticated techniques, such as those employed by the rock-phish gang, are used. Even here, the take-down time is around 10 times slower than for phishing. Take-down is more than twice as slow than for mule-recruitment websites, which are ignored by banks and only removed by volunteers. Only online pharmacies using fast-flux mechanisms are removed more slowly, and we have found no evidence that anyone is attempting to remove these websites at all!

The latest IWF Annual Report (IWF 2008) presents the website lifetimes using analysis that is similar to ours – they show 71% of websites removed within 50 days, and only 16 lasting all year. Data from earlier reports is much harder to directly compare. We therefore ran some additional tests to provide a more direct comparison to the previous IWF approach of checking which sites were still alive

[16] Financial Services Technology Consortium: http://www.fstc.org;
Financial Services Information Sharing and Analysis Center: http://www.fsisac.com;
Association for Payment Clearing Services: http://www.apacs.org.uk;
Anti-Phishing Working Group: http://www.antiphishing.org.

at the beginning and end of two particular four and six week periods. Note that this is not the same as counting the number of websites that are removed within six weeks. Using the IWF's approach, all websites alive on a starting date (whether they first appeared the day previously or three years before) are re-checked at the end of the period. This type of test tends to emphasise long-lived websites.

We conducted a similar test for phishing websites alive at midnight on October 1, 2007, counting the proportion of websites still alive four and six weeks later. The complete results are given in Table 5. Overall, a smaller proportion of phishing websites than child sexual abuse image websites are long-lived. 10.4% of all phishing websites were alive six weeks later, compared to 20% for websites hosting child sexual abuse images. 12.5% of all phishing websites were alive four weeks later, which is notably smaller than the 38% of commercial websites hosting child sexual abuse images.

Table 5. Proportion of Websites Still Alive After 6 and 4 Weeks Respectively

-	Sites > 6 weeks	Sites > 4 weeks	Sites
Child sexual abuse images	20.0%	38.0%	1400
Rock-phish domains	0.0%	0.0%	33
Fast-flux phishing	10.5%	15.7%	38
Ordinary phishing	24.0%	24.0%	25
All phishing combined	10.4%	12.5%	96

Examining the proportion of long-lived phishing sites broken down by type is instructive. Fast-flux and ordinary phishing sites suffered a few long-lived websites. By contrast, all rock-phish domains were removed within four weeks. This could be because rock-phish domains draw the attention of several banks, making it far less likely that they all let the domain slip through the cracks. Fast-flux and ordinary phishing attacks typically impersonate a single bank, making the occasional oversight possible.

However, it must be noted that exceptionally long-lived phishing websites are much less of a concern to the banks than long-lived child sexual abuse image websites are to groups such as the IWF. Phishing websites require spam advertisements to attract victims. If spam is no longer being sent on behalf of months-old websites, then little harm is being done. Long-lived websites hosting illicit pictures cause continued offence until their removal. Hence, the relatively poor performance in removing websites that host child sexual abuse images is especially troubling.

Applying the IWF's methods of analysis, child sexual abuse image websites fare worse than other types of offending content. But it remains difficult to construct a complete picture using these methods alone. Comparing the average website lifetime, as we have done in Table 4, unequivocally demonstrates that there is scope for hastening the removal of child sexual abuse images from the Internet.

An examination of the incentives can again shed light on why these websites are not removed more quickly. In the UK, the IWF works directly with the ISPs to remove offending websites. Websites hosted in the UK are claimed to be removed within 48 hours, which is believed to explain why only 0.2% of such websites are now hosted there (IWF 2006a). When the websites are hosted in other countries, the IWF notifies the appropriate law enforcement agency and perhaps a local INHOPE hotline, but then takes no further action.

Only 29 countries have hotlines that are members of INHOPE, and their policies vary on what to do with incoming reports. For example, the United States hotline 'CyberTipline' operated by the National Center for Missing and Exploited Children (NCMEC) states that they issue take-down notices to ISPs "when appropriate".[17] However, the IWF tells us that they only issue notices to members, which suggests that the incentive here is to use the notices as part of a 'carrot and stick' approach to growing their community.

Similarly, law enforcement responses also vary. Typically, reports are passed to a central agency operating at a national level. It is then up to this agency to pass the necessary information to the appropriate local jurisdiction, who then deal with passing information on to the responsible ISP. At any stage delays can be introduced, with further slowdowns triggered by evidence collection and assessment. The police are institutionally motivated to seek out the criminals, which is not always consistent with getting the material removed in the most timely manner.[18] Furthermore, law enforcement budgets are always very tight, and organisations may choose not to devote the necessary resources to process the reports quickly because they are not as highly motivated as INHOPE members.

Almost all the other types of material we have considered are dealt with on an international basis. While language can be a barrier to prompt action, borders are essentially immaterial to those seeking to have content taken down. However, because the police are made central to the process of dealing with child sexual abuse images, we can see a clear emphasis on jurisdiction since the police do not operate across national (or sometimes state or county) borders. The IWF told us that they would be "treading on other people's toes" if they contacted ISPs outside the UK, and that they "are not permitted or authorised to issue notices to takedown content to anyone outside the UK". The defamed, the rights holders, the banks, the takedown companies and the various groups of volunteers just do not think this way.

[17] http://www.ncmec.org/en_US/documents/CyberTiplineFactSheet.pdf

[18] In this paper we have not considered whether 'take-down' of child sexual abuse images is the optimal strategy. It could be argued that the correct approach is to locate the people behind the websites and that removing websites merely leads to a 'whack-a-mole' game that rapidly removes individual websites without decreasing the availability of the material. The attention that has recently been paid to site lifetimes in the IWF annual reports indicates that removal is now seen by them to be important. However (Callanan 2007) found that only 11% of all websites are reported to ISPs by member hotlines. They wish "not to interfere with any ongoing law enforcement investigation" and say that "depending on national legislation, the ISP sometimes prefers not to be informed about potentially illegal content." We do not understand this comment, unless it refers to the necessity, in some jurisdictions, for the ISP to make a report to the authorities.

9 Conclusion

In this paper we have examined a range of notice and take-down regimes. We have developed insights by comparing differing outcomes where underlying commonalities exist. The banks have adopted a narrow focus on phishing while overlooking mule recruitment. The evasive techniques of fast-flux networks appear unimportant, given that seemingly permanent online pharmacies and short-lived phishing websites use the same scheme.

The Internet is multi-national. Almost everyone who wants content removed issues requests to ISPs or website owners throughout the world, believing – not always correctly – that the material must be just as illegal 'there' as 'here'. Unexpectedly, in the one case where the material is undoubtedly illegal everywhere, the removal of child sexual abuse image websites is dealt with in a rather different manner. The responsibility for removing material has been divided up on a national basis, and this appears to lead directly to very long website lifetimes.

In sum, the evidence we have presented highlights the limited impact of legal frameworks, content types and attack methods on take-down speed. Instead, take-down effectiveness depends on how the responsibility for issuing requests is distributed, and the incentives on the organisations involved to devote appropriate resources to pursue the removal of unwanted content from the Internet.

Acknowledgments

We would like to thank the companies and voluntary groups who provide us with data about phishing website URLs. We are also extremely grateful to the IWF for their assistance in providing detailed statistics about website longevity. Tyler Moore is supported by the UK Marshall Aid Commemoration Commission and by US National Science Foundation grant DGE-0636782.

References

Ahlert, C., Marsden, C., and Yung, C. "How 'Liberty' Disappeared from Cyberspace: The Mystery Shopper Tests Internet Content Self-regulation", 2004. http://pcmlp.socleg.ox.ac.uk/text/liberty.pdf

Anderson, R., Böhme, R., Clayton, R., and Moore, T. "Security Economics and the Internal Market", ENISA, January 2008. http://www.enisa.europa.eu/doc/pdf/report_sec_econ_&_int_mark_20080131.pdf

Callanan, C., and Frydas, N.P. "2007 Global Internet Trend Report", INHOPE, September 2007. https://www.inhope.org/en/system/files/inhope_global_internet_trend_report_v1.0.pdf

Clayton, R. "Judge and jury? How "Notice & Take Down" gives ISPs an Unwanted Role in Applying the Law to the Internet", July 2000. http://www.cl.cam.ac.uk/~rnc1/Judge_and_Jury.pdf

Dagon, D., Zou, C.C., and Lee, W. "Modelling Botnet Propagation Using Time Zones", in *13th Annual Network and Distributed System Security Symposium (NDSS)*, San Diego, California, February 2006, pp. 235–249.

Enright, B. "Exposing Stormworm", October 2007. http://noh.ucsd.edu/~bmenright/exposing_storm.ppt

Franklin, J., Paxson, V., Perrig, A., and Savage, S. “An Inquiry into the Nature and Causes of the Wealth of Internet Miscreants”, in *14th ACM Conference on Computer and Communications Security (CCS’07)*, Alexandria, Virginia, October 2007, pp. 375–388.

Honeynet Project and Research Alliance. “Know Your Enemy: Fast-flux Service Networks, an Ever Changing Enemy”, July 2007. http://www.honeynet.org/papers/ff/fast-flux.pdf

House of Lords Science and Technology Committee. *Personal Internet Security, 5th Report of Session 2006–07*, The Stationery Office, London, August 2007.

Internet Watch Foundation. 2006 Half-yearly Report, IWF, July 2006. http://www.iwf.org.uk/documents/20060803_2006_bi-annual_report_v7_final4.pdf

Internet Watch Foundation. “IWF Reveals 10 Year Statistics on Child Abuse Images Online”, Press Release, IWF, October 2006. http://www.iwf.org.uk/media/news.archive-2006.179.htm

Internet Watch Foundation. “IWF Reports Increased Severity of Online Child Abuse Content”, Press Release, IWF, April 2007. http://www.iwf.org.uk/media/news.archive-2007.196.htm

Internet Watch Foundation. “2007 Annual and Charity Report”, IWF, April 2008. http://www.iwf.org.uk/documents/20080417_iwf_annual_report_2007_(web).pdf

McMillan, R. “‘Rock Phish’ Blamed for Surge in Phishing”, *InfoWorld* (12 December), 2006. http://www.infoworld.com/article/06/12/12/HNrockphish_1.html

Moore, T. “Cooperative Attack and Defense in Distributed Networks”, Tech Report UCAM-CL-TR-718, Computer Laboratory, University of Cambridge, June 2008.

Moore, T., and Clayton, R. “Examining the Impact of Website Take-down on Phishing”, in *Anti-Phishing Working Group eCrime Researcher’s Summit (APWG eCrime)*, Pittsburgh, Pennsylvania, October 2007, pp. 1–13.

Moore, T., and Clayton, R. “Evaluating the Wisdom of Crowds in Assessing Phishing Websites”, in *12th International Financial Cryptography and Data Security Conference (FC 2008)*, Tsudik, G. (Ed.), LNCS 5143, Springer-Verlag, Berlin, Germany, 2008, pp. 16–30.

Moore, T., and Clayton, R. “Evil Searching: Compromise and Recompromise of Internet Hosts for Phishing”, *in submission*, June 2008.

Morland J. “Laurence Godfrey v. Demon Internet Limited. Case No: 1998-G-No 30”, March 1999. http://www.hmcourts-service.gov.uk/judgmentsfiles/j932/godfrey2.htm

Nas, S. “The Multatuli Project: ISP Notice & Take Down”, in *SANE*, October 2004. http://www.bof.nl/docs/researchpaperSANE.pdf

Serjantov, A., and Clayton, R. “Modelling Incentives for Email Blocking Strategies” in *4th Workshop on the Economics of Information Security (WEIS)*, Cambridge, Massachusetts, June 2005.

Spamhaus. “Report on the Criminal ‘Rock Phish’ Domains Registered at nic.at”, Press Release, Spamhaus, June 2007. http://www.spamhaus.org/organization/statement.lasso?ref=7

Thomas, R., and Martin, J. “The Underground Economy: Priceless”, *USENIX ;login* (31:6), December 2006, pp. 7–16.

United Kingdom Government. “The Government Reply to the Fifth Report from the House of Lords Science and Technology Committee Session 2006–07 HL Paper 165 Personal Internet Security”, Cm7234, The Stationery Office, London, October 2007.

Weaver, R., and Collins, M. “Fishing for Phishes: Applying Capture-recapture to Phishing”, in *Anti-Phishing Working Group eCrime Researcher’s Summit (APWG eCrime)*, Pittsburgh, Pennsylvania, October 2007, pp. 14–25.