Journal of Economic Perspectives—Volume 23, Number 3—Summer 2009—Pages 3–20

The Economics of Online Crime

Tyler Moore, Richard Clayton, and Ross Anderson

The economics of information security has recently become a thriving and fast-moving discipline. This field was kick-started in 2001 by the observation that poorly aligned incentives explain the failure of security systems at least as often as technical factors (Anderson, 2001). As distributed computing systems are assembled from machines controlled by principals with divergent interests, microeconomic analysis and game-theoretic analysis become just as important for dependability as protocol analysis or cryptanalysis. The economic approach not only provides a more effective way of analyzing straightforward information-security problems such as privacy, spam, and phishing, but also gives insights to scholars of system dependability, conflict, and crime. An annual Workshop on the Economics of Information Security (WEIS) was established in 2002. The field now involves over 100 active researchers; the subject has drawn together security engineers, economists, and even lawyers and psychologists. For a survey of security economics in general, see Anderson and Moore (2006).

This paper will focus on the subject of online crime, which has taken off as a serious industry since about 2004. Until then, much of the online nuisance came from amateur hackers who defaced websites and wrote malicious software in pursuit of bragging rights. In the old days, electronic fraud was largely a cottage

■ *Tyler Moore is a Postdoctoral Fellow, Center for Research on Computation and Society (CRCS), Harvard University, Cambridge, Massachusetts. Richard Clayton is an Industrial Research Fellow and Ross Anderson is Professor of Security Engineering, both at the Computer Laboratory, University of Cambridge, Cambridge, England. Their e-mail addresses are ⟨tmoore@seas.harvard.edu⟩, ⟨Richard.Clayton@cl.cam.ac.uk⟩, and ⟨Ross.Anderson@cl.cam.ac.uk⟩, respectively.*

industry, local and inefficient: a typical card fraudster ran a vertically-integrated small business. For example, he might buy a card-encoding machine, get a job in a shop where he could copy customers' cards, and then go out at night to steal cash from automatic teller machines (ATMs). Similarly, electronic fraud might have involved a call-center employee collecting password data for use by an accomplice.

But now criminal networks have emerged—online black markets in which the bad guys trade with each other, with criminals taking on specialized roles (Thomas and Martin, 2006). Just as in Adam Smith's pin factory, specialization has led to impressive productivity gains, even though the subject is now bank card PINs rather than metal ones. As shown in Table 1, someone who can collect bank card and PIN data or electronic banking passwords can sell them online to anonymous brokers at advertised rates of $0.40–$20.00 per card and $10–$100 per bank account (Symantec, 2008). The information needed to apply for credit in someone else's name, such as name, social security number, and birthday, fetches $1 to $15 per set. The brokers in turn sell the credentials to specialist cashiers who steal and then launder the money.

A common modus operandi is for the cashier to transfer money from the victim's account to an account controlled by a "money mule." The mules are typically duped into accepting stolen money and then forwarding it. The cashiers recruit them via job ads sent in spam e-mails or hosted on websites such as Craigslist or Monster (Krebs, 2008a), which typically offer the opportunity to work from home as a "transaction processor" or "sales executive." Mules are told they will receive payments for goods sold or services rendered by their employer and that their job is to take a commission and forward the rest, using an irrevocable payment service such as Western Union. After the mule has sent the money, the fraud is discovered and the mule becomes personally liable for the funds already sent. Cashiers also use stolen money to pump up the price of a penny stock in which they have already invested; and they also launder money through online poker games and auctions.

The collection of bank passwords has also become specialized. "Phishermen" operate copies of genuine bank websites that encourage the unwary to log on so that their bank account numbers, passwords, and other credentials can be copied. These phishermen hire "spammers" to drive bank customers to their fake websites by sending e-mails that purport to come from their bank. Both the spammers and the phishermen use malicious software, or "malware," which is designed to infect the computers of people who run it; victims are duped into running it when they download a seemingly innocuous program or visit one of approximately three million infected websites (Provos, Mavrommatis, Rajab, and Monrose, 2008). The emergence of a profitable malware market means in turn that malware is no longer written by teenagers seeking to impress their peers but by specialist firms with budgets for R&D and testing. These firms in turn ensure that their products aren't detected by most antivirus software, and offer updates if they are (Schipka, 2007).

Table 1
Electronic Crime Figures

	Estimate	*Source*
Underground economy advertised unit prices		
Bank account credentials	$10–$100	Symantec (2008)
Credit cards	$0.40–$20	Symantec (2008)
Full identity (name, SSN, birthday, etc.)	$1–$15	Symantec (2008)
Online auction account credentials	$1–$8	Symantec (2008)
Number of compromised computers and websites		
Computers participating in botnets	5 million	Symantec (2008)
Computers infected with identity-theft malware	10 million	Panda Security (2009)
Websites hosting phishing (fake bank) pages	116,000	Moore and Clayton (2009)
Websites infecting visitors with malware	3 million	Provos et al. (2008)
Annual losses		
U.K. online banking fraud (6/2007–5/2008)	£36.5 million	APACS (2008)
U.S. direct identity theft losses (2006)	$2.8 billion	Gartner (2006)
European damages caused by malware (2006)	€9.3 billion	Computer Economics (2007)

One consequence is that while antivirus software previously detected most malware, it now detects only a minority of it. Infected computers may log their users' keystrokes to harvest credentials for electronic banking; they may also be recruited into botnets, which we discuss next. While estimates vary, the general consensus is that approximately 5 percent of computers worldwide are susceptible to malware infection at any given time (House of Lords Science and Technology Committee, 2007); one security provider has estimated that 10 million computers are infected with malware designed to steal online credentials (Panda Security, 2009).

With this new online crime ecosystem has come a new profession: the "botnet herder"—a person who manages a large collection of compromised personal computers (a "botnet") and rents them out to the spammers, phishermen, and other crooks. The computers in the botnet have been infected with malware that lets the botnet herder operate them by remote control, just as if they were robots. Nearly all e-mail spam is now sent by botnets. Many websites used by online criminals are hosted on botnets, ranging from online pharmacies through the fake banks used in phishing scams to the sham companies that "hire" money mules (Moore and Clayton, 2008a). Blackmailers also rent botnets and have threatened to overload bookmakers' websites with botnet traffic just before large sporting events; in 2004, three Russians were arrested after extorting several hundred thousand dollars in this way (Sullivan, 2004). A botnet was used to shut down parts of Estonia's infrastructure as a political protest (Lesk, 2007). Around five million computers participated in botnets unbeknownst to their owners in the second half of 2007 alone (Symantec, 2008).

Online criminals are also involved in many other types of fraud, from bogus lotteries through stock scams to advance-fee frauds, and in other crimes such as the

distribution of images of child sexual abuse. Some figures estimating costs of online crime are presented at the bottom of Table 1.

Thus far, the world has not coped well with the rapid growth and industrialization of online wickedness. Banks often hide fraud losses, or even blame their customers for fraud; and they hesitate to share information with other banks. Police agencies have also floundered.

Online crime has many similarities with the economics of conventional crime, the study of which has blossomed since Becker's (1968) seminal work. But there are some interesting differences, many of them driven by the global scale of online crime. We find a useful historical analogy two generations ago when criminals started using cars. Suddenly a burglar could break into several houses in a town where the police didn't know him and be home in time for breakfast. It took the police a generation to catch up, using national police networks and fingerprint databases. Online crime, like vehicle-assisted crime, will force big changes, both because it is transnational and also because it consists of a high volume of low-value offenses. Existing mechanisms for international police cooperation are designed for rare serious crimes, such as murder and terrorism, while online crime is petty crime committed on a global and industrial scale. Another difference is that conventional crime is generally committed by marginal members of society, especially by young men suffering multiple deprivation or abusing drugs or alcohol. In contrast, people with comparative advantage at online crime tend to be educated and capable, but they live in societies with poor job prospects and ineffective policing.

This paper begins by looking at the data on online crime. We then examine the collective-action aspects: people who connect infected personal computers to the Internet create negative externalities in that their machines may emit spam, host phishing sites, and distribute illegal content. The Internet's global, distributed architecture leads to security being dominated by the weakest link, which exposes the poor coordination among defenders both public and private. We present empirical evidence of how agile attackers shift across national borders as earlier targets wise up to their tactics, and discuss ways to improve law-enforcement coordination.

Finally, we will examine how defenders' incentives affect the outcomes. An interesting case study is to measure the average time required to remove different types of offending content from the Internet. Phishing sites that directly impersonate technology businesses, such as eBay, are typically knocked out within hours; sites that impersonate banks typically vanish within a few days; but money-laundering websites take far longer. They do not target a single bank, but spread their harm over the whole banking industry so that no one pursues them with much vigor. We will also show how the refusal of banks (and their contractors) to share information on phishing websites slows removal significantly and how a public-sector remedy is unlikely, at least in the short term.

Accurate Information

Hard statistics on losses from online crime are hard to come by in most countries. But without them, other figures—whether of vulnerabilities in computer systems, number of botnets, or size of spam flows—lack a grounded connection to the real economy.

One problem is that many of the statistics on security failure are collected by parties with an incentive to under- or overreport. For example, the anti-phishing group PhishTank has boasted about the large number of sites it identifies (OpenDNS, 2007), when accounting for duplicates would reduce its reported total several-fold (Moore and Clayton, 2007). The U.K. banking trade association APACS (Association for Payment Clearing Services) provides another example; it asserted a 726 percent increase in phishing attacks between 2005 and 2006, but only a 44 percent rise in losses (APACS, 2007). Governments, by contrast, often seek to minimize crime statistics. In a particularly egregious case, the UK government changed the rules so that fraud must be reported to the bank rather than to the police. This change caused the fraud figures to drop to near zero and was strongly criticized by a Parliamentary committee (House of Lords Science and Technology Committee, 2007). Internet service providers also have an incentive to undercount: they want to downplay the amount of bad traffic emanating from their customers, lest it affect their relationships with other Internet service providers.

But two approaches to measuring online crime do hold promise. Although individual banks are usually keen to keep data on fraud losses private, countrywide aggregation from banking data is possible. The Banque de France and APACS publish aggregated annual loss figures for the sums lost by banks to phishing attacks in France and the United Kingdom respectively, along with totals for theft through automatic teller machines and other financial fraud. As banks collect such statistics for operational, internal control, and audit purposes, aggregating them nationally can be straightforward. In a report we were commissioned to write on security economics and the internal market for the European Commission (Anderson, Böhme, Clayton, and Moore, 2008), we recommended that other countries follow the British and French lead in publishing nationally aggregated figures.

A different, and complementary, approach has emerged in many U.S. states and is now being considered by the European Parliament: security-breach reporting laws. Under a California law enacted in September 2002 (California State Senate, 2002), both public and private entities that conduct business in California must notify affected individuals when personal data under their control has been acquired by an unauthorized person. The law was intended to ensure that individuals are given the opportunity to protect their interests following data theft, such as when 45 million credit card numbers were stolen from T.J. Maxx's information

technology systems (Greenemeier, 2007). It was also intended to motivate companies to keep personal data secure; and Acquisti, Friedman, and Telang (2006) found a statistically significant negative impact on stock prices following a reported breach. Romanosky, Telang, Acquisti (2008) examined identity theft reports obtained from the Federal Communications Commission from 2002 to 2007. Using time differences in the adoption of state breach disclosure laws, they found a small but statistically significant reduction in fraud rates following statewide law adoption. Breach-disclosure laws also contribute data on security incidents to the public domain. The California law has inspired further laws in at least 34 other states, although their details vary.

Security breach notification could be improved with a central clearinghouse and some standardization of procedures. A clearinghouse would help to ensure that all reported breaches can be located by the press, investors, researchers, and sector regulators. Future U.S. or EU laws should also set minimum standards for notification; some U.S. companies have hidden notifications amongst lengthy marketing material. Finally, notification should include advice on what individuals should do; some notifications by U.S. firms have puzzled or terrified their recipients, rather than helped them with advice on risk reduction.

Some researchers have studied the new criminal markets directly. Researchers from the Internet security firm Team Cymru have long documented online crime (for example, Thomas and Martin, 2006). Franklin, Perrig, Paxon, and Savage (2007) monitored the public chat channels used by online criminals to contact each other, gathering extensive data on credit card fraud, spamming, phishing, and the sale of compromised hosts. Kanich et al. (2008) went a step further. They infiltrated a large botnet and altered the spam e-mails sent out so that they linked to a benign duplicate website under the researchers' control. They were able to provide the first independent answer to a long-standing question: how many people respond to spam? It turns out that 28 sales resulted from 350 million spam e-mails advertising pharmaceuticals—a conversion rate of 0.00001 percent. Such innovative surveillance will continue to be necessary. Unfortunately, in many countries (including the United Kingdom) such projects can only be carried out lawfully by law enforcement agencies—many of whom lack the technical expertise or simply don't care. Indeed, many countries have passed information security laws that probably do more to prevent security researchers from doing their jobs than to tackle actual crime.

Without accurate information on online crime, it is hard for private markets to provide incentives for more secure software. Akerlof's (1970) model of the "market for lemons" applies to many information security product and service markets. The security of software is hard enough for its developers to ascertain, let alone their customers; and consumers naturally refuse to pay a premium for quality they cannot assess.

Interdependent Security and the Difficulty of Coordination

In many contexts, security depends on the efforts of many interdependent principals. Hirshleifer (1983) told the story of Anarchia, an island whose flood defenses were built by individual landowning families and whose defense thus depended on the weakest link—that is, the laziest family. He compared this to a city whose protection against missile attack depended on the single best intercepting shot, and showed that best-shot provides better defense than weakest-link. Varian (2004) added a third case—where performance depends on the sum of the efforts of all the defenders, as in a democracy where we pay taxes and hire soldiers—and showed that this sum-of-efforts defense strategy is the best of all three.

Fixing online crime is hard because Internet security is often weakest-link. Millions of personal computers have been recruited to botnets, and attackers are spoilt for choice when selecting computers running out-of-date software to compromise. In addition, Internet insecurity is a bit like environmental pollution: someone who connects an insecure computer to the Net is creating a negative externality, as the computer can be used by others to attack third parties; there has even been a suggestion of a cap-and-trade system for software quality (Camp and Wolfram, 2004). The concentrated nature of the software industry, however, and its success to date in disclaiming liability, may mean that such solutions are some way off. Online malefactors also buy services from providers with the feeblest security policies, or which turn a blind eye to misbehavior.

The weakest-link nature of much Internet security brings us to the fundamental issue of where we can find, or create, effective control points. Responsibility for prevention could be placed on computer owners, Internet service providers, software suppliers, private security firms, law enforcement, or banks. Naturally, every one of these stakeholders wants someone else to fix the problem. Security is a classic collective-action problem. How can we get the actors in the system to respond to the social costs of online crime, not just the private costs?

A Central Role for Internet Service Providers

Internet service providers (ISPs) are unusually well placed to detect infection because evidence of a user's infection necessarily flows over an ISP's network. The approximately 4,000 ISPs in the United States range in size from mom-and-pop firms serving a few hundred customers in rural outposts to behemoths such as AT&T, Verizon, AOL, and Comcast, which each provide online connectivity to millions of households and businesses. Moreover, large ISPs have technical staff who can detect and clean up infected computers, while domestic users and small businesses are generally unable even to recognize when they are compromised. ISPs are also uniquely placed to limit the external impact of an infected computer: they control its Internet connection and can disconnect it if need be. Current best practice is less drastic: it is to quarantine infected computers into a "walled garden"

subnetwork from which they can access decontamination and software patches but not much else.

The market provides incentives for some Internet service providers to take action. An OECD study found that the strongest driver was the cost of customer support—one medium-sized ISP reported 1–2 percent of its total revenue was spent on handling security-related support calls (van Eeten and Bauer, 2008). Another incentive is that a provider whose customers emit too much bad traffic can find its peering arrangements with other companies damaged. However, very large ISPs are effectively exempt from peer pressure as others cannot afford to cut them off. Much of the world's bad traffic comes from the networks of these "too big to block" providers.

Some particularly negligent providers attract disproportionate amounts of bad business. For example, until 2007, a lot of the world's malware was hosted by the Russian Business Network, which ignored requests from international law enforcement (*The Economist,* 2007). After the Russian Business Network suddenly went offline in late 2007, malware distribution shifted elsewhere. In November 2008, a journalist from the *Washington Post* persuaded upstream bandwidth providers to shut off their connection to San Francisco-based McColo (Krebs, 2008b), which led to a temporary fall of almost 70 percent in the volume of spam worldwide. Apparently, many botnet herders had been using McColo to host their control computers. And recently EstDomains, which had served as the primary domain name registrar for malicious websites hosted by Russian Business Network (Krebs, 2008c), became the first domain registrar to have its accreditation terminated by the Internet Corporation for Assigned Names and Numbers (ICANN, 2008), the not-for-profit corporation that coordinates the Internet's naming system.

However, attackers can often exploit the next weakest link in the Internet far faster than defenders can knock them out. Moore and Clayton (2007) described how one leading group of online fraudsters, the rock-phish gang, operates. It registers many malicious web domain names to carry out attacks. Periodically, the gang picks a new domain name registrar and becomes an active customer, registering hundreds of web domains using false names and stolen credit cards. Because registrars will take down a domain name quickly if it looks abusive—such as a misspelling of a bank's name—the rock-phish gang chooses domains that do not appear to violate any trademark. It may take much longer to convince a naive registrar that an apparently innocuous domain is in fact being used for criminal purposes.

Figure 1 presents scatter plots of phishing website lifetimes based on data reported by Moore and Clayton (2007). The horizontal axis shows when the phishing site was reported; the vertical axis shows how long it lasted. The gang targeted Hong Kong domains first in March 2007, and then after the Hong Kong authorities wised up, they targeted Chinese domains in May 2007. Phishing sites on .hk (Hong Kong) domains and .cn (China) domains lasted much longer in the first months after the gang began targeting each domain type than in later months.

Figure 1
Scatter Plot of Phishing Site Lifetimes over Time Based on the Domain Targeted

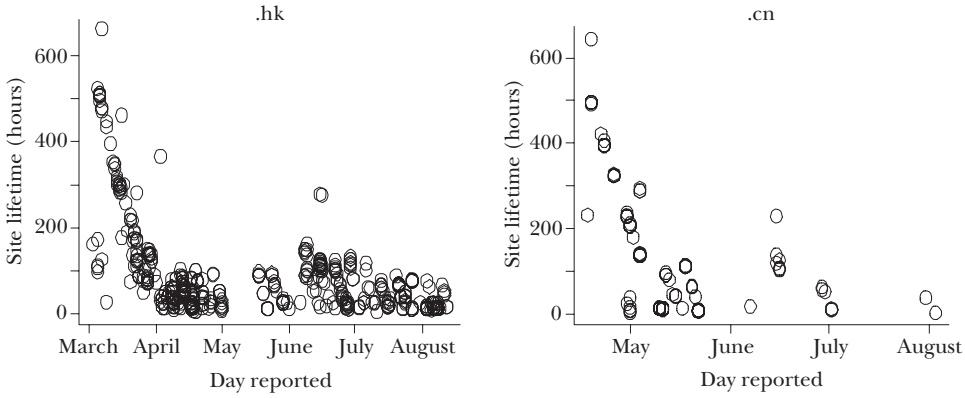

Source: Moore and Clayton (2007)

Other researchers also documented how websites hosting malware move from one registrar to the next (Day, Palmen, and Greenstadt, 2008).

Educating registrars is a work in progress. A few large firms perform most registrations (for example, GoDaddy, Network Solutions, register.com), but as with Internet service providers, there are thousands of smaller companies too. The Anti-Phishing Working Group's Internet Policy Committee (2008) has set itself the ambitious goal of educating all registrars about common threats, such as the rock-phish gang, before they are targeted.

Policymakers should also give Internet service providers, especially the big ones, a stronger incentive to stop infected computers attacking other users. For example, in our report for the European Commission, we proposed fixed statutory damages against an ISP that does not act within a fixed time period after being notified of an infected computer on its network (Anderson, Böhme, Clayton, and Moore, 2008) At present, takedown speed varies hugely: the best ISPs remove phishing sites in less than an hour, while others take weeks. Introducing a fixed minimum charge for not dealing with misbehaving websites after a reasonable notice period, say three hours, would provide a useful nudge.

Statutory damages of this general form have been used effectively in the airline industry, where the European Union has introduced them for airlines that deny passengers boarding due to overbooking, cancellations, or excessive delays. A passenger who can't fly can claim a fixed amount (typically 250 euros) from the airline, without having to produce hotel bills or other evidence of actual expenses. The airline may then sue other parties (such as airports or maintenance contractors) to recover these damages where appropriate. Similarly, we envision that Internet service providers would be able to recover damages from other negligent parties. Another of our recommendations is that vendors of network-attached

consumer equipment should have to certify that it is "secure by default," so if it turns out not to be, then ISPs could seek redress without being frustrated by the ubiquitous software liability disclaimers.

Movies and novels sometimes ascribe almost mythical abilities to computer hackers and especially to successful gangs like rock-phish. Yet our analysis and data suggest that the success of some of these gangs is as much strategic as technical: that is, successful attackers are not writing brilliant software, but instead are exploiting basic failings systematically. Ohm (2008) confirms this analysis in a discussion of the "myth of the Superuser"—people ascribe extraordinary capabilities to hackers when the reality revealed by empirical analysis is much more straightforward.

Sharing Security Data among Take-Down Firms

Collecting timely, comprehensive data on the latest online vulnerabilities and the currently compromised websites is essential for protecting consumers. However, many information security contractors that take down malicious websites on behalf of their clients keep such data to themselves, arguing that the information benefits their firm's competitive position. We have found that both the firms' customers and consumers would benefit if security contractors shared more data.

Back in the old days of the 1980s and 1990s, antivirus companies did not share virus samples; instead, they boasted of how comprehensive their lists were. Trade magazines published head-to-head comparisons of competing products, testing whether one antivirus company like Dr. Solomon caught more viruses than another one like Norton. In 1993, a series of press releases from the major companies claimed that some new virus was being overlooked by the competition, and it finally became clear that the overall effect of not sharing was damaging the industry. At that year's European Institute for Computer Antivirus Research (EICAR) conference, a meeting of the antivirus researchers led to an agreement that they would share samples of viruses with their competitors. This sharing continues to this day, improving the quality of protection available to consumers and businesses.

The anti-phishing industry has yet to learn this lesson. At its core lie specialist contractors, such as Cyveillance, RSA, and MarkMonitor, who are hired by banks to remove phishing websites and to suspend abusive domain names. These firms compile "feeds" of fishing sites—up-to-the minute lists of phishing page locations. Moore and Clayton (2008b) analyzed six months of feeds from multiple sources, including two anti-phishing contractors. In each case, many phishing websites were known somewhere in the industry, but not to the company with the take-down contract for the bank actually targeted by the site.

Figure 2 shows the findings for one large take-down company, called T_A, which removed phishing sites on behalf of 54 client banks during a six-month period. While T_A identified 9,008 of the sites impersonating its clients (represented by the left circle), a further 5,962 sites were known to its competitors (the right circle). As there could be other targets that neither T_A nor its competitors knew about, T_A missed at least 40 percent of its targets.

Figure 2

Venn Diagram Depicting Take-Down Company T_A's Awareness of and Response to Phishing Websites Impersonating Its 54 Client Banks during October 2007–March 2008

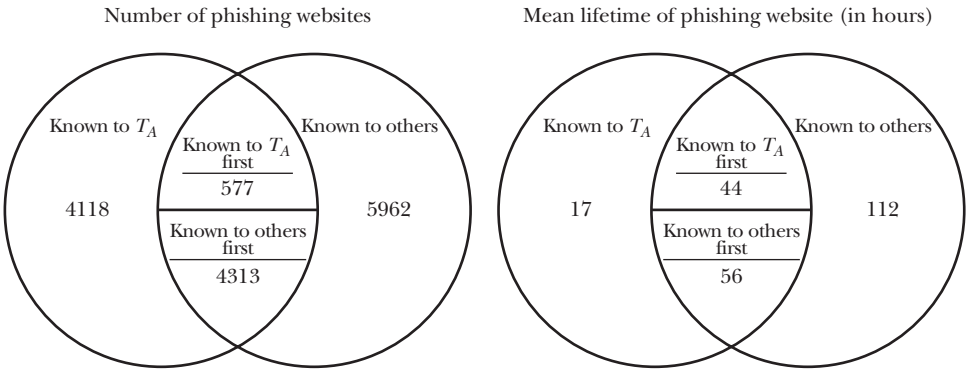

Source: Moore and Clayton (2008b).

Moore and Clayton approximate the costs of this situation relative to a more cooperative alternative by examining website lifetimes. The sites only T_A knew about are removed within 17 hours on average, while sites unknown to T_A last 112 hours on average, around four days longer. In addition, of the 9,008 websites T_A did know about, 4,313 were identified by the other sources first, at an average of 50 hours before T_A learned of them. The effect of such delays can be seen in the lifetime figures: these websites remain for 56 hours on average, 39 hours longer than sites known only to T_A. Phishing site lifetimes could be greatly shortened if take-down companies shared data, as the antivirus companies began to do 15 years ago. Their clients would gain directly, and the companies could also benefit—not only from the increased revenue opportunities of having more work to do, but also from being able to market a more valuable overall service. Moore and Clayton (2008b) calculated that website lifetimes could drop by half or more given information sharing and that for these two companies' clients alone—on some fairly rough estimates—around $330 million a year of fraud might be prevented by sharing.

As with the antivirus industry 15 years ago, sharing would have industry effects. Take-down companies compete on a number of factors, including price, customer service, speed of removal, and "feed comprehensiveness"—which is the industry term for how widely they search the web for offending sites. Sharing data would eliminate comprehensiveness as a competitive factor, which might have the negative effect of reducing incentives in this area (Olson, 2008). This objection could be dealt with by compensating those who provided more data to the common pool (Moore, 2008). Some well-established companies may view sharing as lowering the barriers to entry. However, competition would then move to speed of removal,

Table 2
Website Lifetimes by Type of Offending Content

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

Source: Moore and Clayton (2008a).

price, and service, at which incumbents would also have an advantage. And by increasing the number of sites that can be removed, the service itself would be worth more to the banks.

In our view, most existing take-down companies would benefit from sharing feeds—that is, the gains from the service being worth more would outweigh the loss of competition on comprehensiveness. The only likely losers would be the few companies that specialize primarily in producing feeds. For the banks that are customers of the take-down companies, greater feed sharing offers only benefits.

How the Incentives of Defenders Affect Take-Down Speed

Many different types of bad online content—from copyright violations to child sexual abuse images to phishing websites—are subject to take-down requests. Moore and Clayton (2008a) obtained data on the lifetimes of several types of websites, summarized in Table 2.

The lifetimes of questionable websites are heavily influenced by who has an incentive to find and complain about the offending material. Phishing websites are removed fastest: banks are highly motivated to remove any website that impersonates them. By contrast, other illegal activities such as online pharmacies do not appear to be removed at all.

However, most banks focus on removing only those websites that attack them directly. They ignore a key component of the phishing supply chain: mule recruitment. As described earlier, phishermen recruit "money mules," dupes who launder stolen money, typically using Western Union transfers. Because the fraudulent transfers are often reversed, the mule ends up out of pocket rather than the bank, and so banks lack an adequate incentive to crack down on mule recruitment. Their incentive is also dulled by a collective-action problem: it is hard to tell which bank will suffer from any given mule-recruitment campaign.

Thus, even though mule recruitment websites harm banks directly, not one of the banks or take-down companies actively pursues them. Typically, only vigilante groups such as "Artists Against 419" attempt any removal, and even they treat these websites as low priority because they see the mules as complicit in phishing. Finally, regulators may have trained banks to see money laundering as an issue of due diligence, rather than risk reduction; most money-laundering controls are aimed at crimes like drug trafficking in which the banks are not victims, and the incentives there steer them towards minimal and mechanical compliance. Moore and Clayton (2008a) found that mule-recruitment websites lasted 308 hours, far longer than phishing websites that directly impersonate banks (4 to 96 hours). This is an opportunity missed; the most rapid growth in spam late in 2008 has been for mule recruitment, which strongly suggests that mule shortage had become an important bottleneck in phishing operations.

Attack technology also affects take-down speed, but to a lesser extent. Naive crooks host their websites on free services or individual compromised web servers, which are easy for the contractors to take down; more sophisticated criminals such as the rock-phish gang mentioned earlier use evasive techniques such as fast-flux. Moore and Clayton (2007) describe this scheme: websites are hosted dynamically on a botnet, residing for just a few minutes on each computer and moving elsewhere before the removal service can do anything. But our figures show that the lifetime of an offending site is determined far more by the direct incentives that defenders have to take it down than by the attack technology. For example, Moore and Clayton (2008a) found that fast-flux phishing websites are removed in 96 hours, but fast-flux pharmacies are hardly removed at all (lasting nearly two months on average).

Coordination of Law Enforcement

There are tens of thousands of law enforcement agencies worldwide; many of them, even in developed countries, are fairly uninformed about computer crime. What is specifically illegal varies from one country to another: the leading global legal framework, the Council of Europe Convention on Cybercrime, has been ratified by the United States but has yet to be ratified by a majority of European Union member states.

Once nations have agreed on what is a crime, police forces will still have little incentive to work together on globalized volume crime. Suppose a phisherman in Russia sends out a million spams to random e-mail addresses. The largest police force in Britain, London's Metropolitan Police, might find ten thousand of these arriving in its area—London accounts for about 1 percent of global Internet use. The Met will be tempted to say "Oh bother, let the FBI deal with it," and focus on local street crime instead. Most local police forces prioritize crime-fighting by asking how many local citizens are victims, how many are perpetrators, and how serious is the damage locally. Using these criteria, it may be that few online attackers will seem worth pursuing, even if in aggregate they are having an enor-

mous effect. In the United Kingdom, for example, there are only two small police units specializing in online fraud (the Dedicated Cheque and Plastic Crime Unit and the Police Central e-crime Unit) and both rely on the banking industry for a lot of their funding.

The barriers to cooperation are further raised by the fact that online crime usually crosses national boundaries. Existing mechanisms for international police cooperation are expensive and slow—designed to catch the occasional fugitive murderer, but not for dealing with millions of frauds at a cost of a few hundred dollars each. The problem is compounded by sensitivities about national sovereignty: each individual case is reviewed by diplomats to ensure it isn't politically sensitive. Our suggestion here is that, following the precedent of SHAEF (the Supreme Headquarters Allied Expeditionary Force) in World War II and NATO today, countries should maintain liaison officers at a central command center that decides what tasks to undertake, whereupon the liaison officers relay requests to their own countries' forces. Such a permanent "joint operation" would deal with the glacial speed of current arrangements and the lack of international agreement on what to prioritize. The key is that countries must trust their liaison officers to assess which requests carry no political baggage and can be treated as straightforward police matters. We also need a mechanism to evolve a global strategy on cybercrime priorities. This will require both operational feedback and democratic accountability.

Public versus Private Action

The one cybercrime issue that really catches the attention of politicians and the popular media is websites that host images of child sexual abuse. Yet, curiously, we found that these websites are removed much more slowly than almost any other type of unlawful content. Their average lifetime of 719 hours is over 150 times that of normal phishing sites, and more than twice that of even the mule-recruitment websites. Why might this be?

During the 1990s, when the Internet came to public attention, policymakers established "child pornography" as the one Internet evil that all governments could agree to ban. In 29 countries, Internet service providers established hotlines to identify and take down offending material. In the United Kingdom, the Internet Watch Foundation (IWF) does this job and claims to remove child sexual abuse images hosted in Britain within 48 hours; only 0.2 percent of such sites are now hosted in the United Kingdom (Internet Watch Foundation, 2006). When the websites are hosted in other countries, the IWF notifies a local hotline or law enforcement agency, but then takes no further action.

Hotline policies and effectiveness vary and few, if any, are as effective as the Internet Watch Foundation. The U.S CyberTipline operated by the National Center for Missing and Exploited Children (NCMEC) states that they issue take-down notices to Internet service providers "when appropriate"; however, through October 2008, NCMEC apparently only issued notices to the subset of

ISPs that were actually members. A new U.S. law enacted in October 2008, the Protect Our Children Act, may fix this particular problem by making it compulsory for all ISPs to register with NCMEC. Law enforcement responses also vary. Typically, reports are passed to a national agency, which must then pass the information to a local jurisdiction, which then contacts the responsible ISP. Law enforcement budgets are always tight, and police forces will vary in their efforts in challenging such websites depending on how salient the issue is at that time in local politics.

Almost all other types of unlawful online material are dealt with on an international basis, and borders are essentially immaterial to a capable, motivated, private-sector firm seeking to have content taken down. However, the police have been given a monopoly on dealing with child sexual abuse images in many jurisdictions. In the United Kingdom, for example, simple possession of such material is a strict-liability criminal offense, effectively preventing the private sector from helping. (A company seeking to disable such material would likely possess it, in at least an on-screen image, at some stage of a takedown process.) Because the police have sole authority to pursue this material, jurisdiction becomes a significant stumbling block, for the police do not operate across national (or sometimes state or county) borders. The Internet Watch Foundation told us that they would be "treading on other people's toes" if they contacted Internet service providers outside the United Kingdom and that they "are not permitted or authorized to issue notices to take down content to anyone outside the UK." In contrast, with other kinds of online crime, banks, take-down companies, and even vigilantes show great flexibility in their willingness to pursue distant materials.

Police forces do however have a role in combating online crime. For example, we have noticed significant recent consolidation within the botnet and spam industries; as we noted, the takedown of McColo reduced spam worldwide by 70 percent when it broke the control that six large herders had over their botnets. In our view, police resources would best be concentrated on busting these large gangs, and the FBI with "Operation Bot Roast" is already moving in this direction. Ongoing tasks such as website take-down are better left to appropriately incentivized private contractors.

Concluding Remarks

Since about 2004, online crime has become organized and industrialized like no other crime, with the possible exception of the drugs trade. Many of the problems that banks and police forces face in controlling it already exist in traditional law enforcement but are made more acute online by factors ranging from network externalities to global scale. Unfortunately, crime looks set to be a permanent feature of the online ecosystem and to create significant costs for banks,

citizens, service providers, and others. We have presented the results of a number of recent research efforts that together explain how the online crime industry works, why current enforcement efforts are feeble, and how they could be improved.

With previous technology-driven crime innovations, from credit card fraud to the use of getaway cars in bank robbery, it took some time to work out the optimal combination of public and private security resources. Our analysis in this paper suggests that significant improvements are possible in the way we deal with online fraud. Criminal networks do have particular vulnerabilities—such as their money laundering operations. However, individual banks don't target money launderers because launderers attack the banking system as a whole, not any individual bank. Perhaps the banks' trade associations should target the laundrymen. Banks also fail to get their security contractors to share data on attacks where this could help them directly. This collective action problem is best dealt with by private-sector information sharing, as it was 15 years ago in the world of computer viruses. Finally, we suggest that the police should concentrate their efforts on the big phishing gangs.

To control online crime better, we need to understand it better. The key to this understanding is not so much technology, but gaining an economic perspective of the incentives faced by the different players.

■ *The authors are grateful to Allan Friedman, Steven Murdoch, and Andrew Odlyzko, who read this paper in draft and made helpful comments.*

References

Acquisti, Alessandro, Allan Friedman, and Rahul Telang. 2006. "Is There a Cost to Privacy Breaches? An Event Study." Paper presented at the International Conference on Information Systems (ICIS), Milwaukee, WI.

Akerlof, George A. 1970. "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics,* 84(3): 488–500.

Anderson, Ross. 2001. "Why Information Security is Hard—An Economic Perspective." *Proceedings of the 17th Annual Computer Security Applications Conference,* 358–65. IEEE Computer Society.

Anderson, Ross, Rainer Böhme, Richard Clayton, and Tyler Moore. 2008. "Security Economics and the Internal Market." European Network and Information Security Agency. http://www.enisa.europa.eu/doc/pdf/report_sec_econ_&_int_mark_20080131.pdf.

Anderson, Ross, and Tyler Moore. 2006. "The Economics of Information Security." *Science,* 314(5799): 610–13.

Anti-phishing Working Group Internet Policy Committee. 2008. "Anti-Phishing Best Practices Recommendations for Registrars." An APWG Industry Advisory. http://www.antiphishing.org/reports/APWG_RegistrarBestPractices.pdf.

APACS (Association for Payment Clearing Services). 2007. "Card Fraud Losses Continue to

Fall." Press release, March 14. http://www.apacs.org.uk/media_centre/press/07_14_03.html.

APACS (Association for Payment Clearing Services). 2008. "APACS Announces Latest Fraud Figures" Press release, September 25. http://www.apacs.org.uk/APACSannounceslatestfraudfigures.htm.

Becker, Gary. 1968. "Crime and Punishment: An Economic Approach." *The Journal of Political Economy,* 76(2): 169–217.

California State Senate. 2002. "Assembly Bill No. 700." http://info.sen.ca.gov/pub/01-02/bill/asm/ab_0651-0700/ab_700_bill_20020929_chaptered.pdf.

Camp, L. Jean, and Catherine D. Wolfram. 2004. "Pricing Security: A Market in Vulnerabilities." In *Economics of Information Security,* Vol. 12, *Advances in Information Security,* ed. L. Jean Camp and Stephen Lewis, 17–34. Boston: Kluwer Academic Publishers.

Computer Economics. 2007. "Malware Report: The Economic Impact of Viruses, Spyware, Adware, Botnets, and other Malicious Code." http://www.computereconomics.com/page.cfm?name=Malware%20Report.

Day, Oliver, Brandon Palmen, and Rachel Greenstadt. 2008. "Reinterpreting the Disclosure Debate for Web Infections." In *Managing Information Risk and the Economics of Security,* ed. M. Eric Johnson, 179–197. New York: Springer.

The Economist. 2007. "A Walk on the Dark Side." August 30.

Franklin, James, Adrian Perrig, Vern Paxon, and Stefan Savage. 2007. "An Inquiry into the Nature and Causes of the Wealth of Internet Miscreants." *Proceedings of ACM Conference on Computer and Communications Security (CCS),* 375–388. ACM Press.

Gartner. 2006. "Gartner Says Number of Phishing E-Mails Sent to U.S. Adults Nearly Doubles in Just Two Years." Press release, November 9. http://www.gartner.com/it/page.jsp?id=498245.

Greenemeier. 2007. "T.J. Maxx Parent Company Data Theft Is The Worst Ever." *Information Week,* March 29. http://www.informationweek.com/news/security/showArticle.jhtml?articleID=198701100.

Hirshleifer, Jack. 1983. "From Weakest-link to Best-shot: The Voluntary Provision of Public Goods." *Public Choice,* 41(3): 371–386.

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

Internet Corporation for Assigned Names and Numbers (ICANN). 2008. "Termination of Registrar EstDomains to Go Ahead." November 12. http://www.icann.org/en/announcements/announcement-12nov08-en.htm.

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

Kanich, Chris, Christian Kreibich, Kirill Levchenko, Brandon Enright, Geoffrey M. Voelker, Vern Paxson, Stefan Savage. 2008. "Spamalytics: An Empirical Analysis of Spam Marketing Conversion." *Proceedings of ACM Conference on Computer and Communications Security (CCS),* 3–14. ACM Press.

Krebs, Brian. 2008a. " 'Money Mules' Help Haul Cyber Criminals' Loot." *Washington Post.* January 25. http://www.washingtonpost.com/wp-dyn/content/story/2008/01/25/ST2008012501460.html.

Krebs, Brian. 2008b. "Major Source of Online Scams and Spams Knocked Offline." Blog titled "Security Fix." *Washington Post,* November 11. http://voices.washingtonpost.com/securityfix/2008/11/major_source_of_online_scams_a.html.

Krebs, Brian. 2008c. "EstDomains: A Sordid History and a Storied CEO." *Washington Post.* September 8. http://voices.washingtonpost.com/securityfix/2008/09/estdomains_a_sordid_history_an.html.

Lesk, Michael. 2007. "The New Front Line: Estonia under Cyberassault." *IEEE Security and Privacy,* 5(4): 76–79.

Moore, Tyler. 2008. "How Can We Co-operate to Tackle Phishing?" October 27. http://www.lightbluetouchpaper.org/2008/10/27/how-can-we-co-operate-to-tackle-phishing/.

Moore, Tyler, and Richard Clayton. 2007. "Examining the Impact of Website Take-down on Phishing." *Proceedings of the Anti-Phishing Working Group eCrime Researchers Summit,* 1–13.

Moore, Tyler, and Richard Clayton. 2008a. "The Impact of Incentives on Notice and Take-down." In *Managing Information Risk and the Economics of Security,* ed. M. Eric Johnson, 199–223. New York: Springer.

Moore, Tyler, and Richard Clayton. 2008b. "The Consequence of Non-cooperation in the Fight against Phishing." *Proceedings of the Anti-Phishing Working Group eCrime Researchers Summit,* 1–14. IEEE.

Moore, Tyler, and Richard Clayton. 2009. "Evil Searching: Compromise and Recompromise of Internet Hosts for Phishing." *Lecture Notes in Computer Science,* vol. 5628, pp. 256–72.

Ohm, Paul. 2008. "The Myth of the Superuser: Fear, Risk and Harm Online." *UC Davis Law Review,* 41(4): 1327–1402.

Olson, Eric. 2008. "A Contrary Perspective—Forced Data Sharing Will Decrease Performance and Reduce Protection." October 22. http://www.cyveillanceblog.com/phishing/a-contrary-perspective-%e2%80%93-forced-data-sharing-will-decrease-performance-and-reduce-protection.

OpenDNS. 2007. "OpenDNS Shares April 2007 PhishTank Statistics." Press release, May 1. http://www.opendns.com/about/announcements/14/.

Panda Security. 2009. "More than 10 Million Worldwide Were Actively Exposed to Identity Theft in 2008." March 10. http://www.pandasecurity.com/usa/homeusers/media/press-releases/viewnews?noticia=9602&sitepanda=empresas.

Provos, Niels, Panayiotis Mavrommatis, Moheeb Abu Rajab, and Fabian Monrose. 2008. "All Your iFRAMEs Point to Us." *Proceedings of the 17th USENIX Security Symposium,* 1–15. USENIX Association.

Romanosky, Sasha, Rahul Telang, and Alessandro Acquisti. 2008. "Do Data Breach Disclosure Laws Reduce Identity Theft?" Paper presented at the 7th Workshop on the Economics of Information Security, Hanover, NH. Available at SSRN: http://ssrn.com/paper=1268926.

Schipka, Maksym. 2007. "The Online Shadow Economy: A Billion Dollar Market for Malware Authors". MessageLabs White Paper. http://www.fstc.org/docs/articles/messaglabs_online_shadow_economy.pdf.

Sullivan, Bob. 2004. "Experts Fret over Online Extortion Attempts." *MSNBC.* November, 10. http://www.msnbc.msn.com/id/6436834/.

Symantec. 2008. *Symatec Global Internet Security Threat Report,* Vol. 13, *Trends for July–December 07.* http://eval.symantec.com/mktginfo/enterprise/white_papers/b-whitepaper_internet_security_threat_report_xiii_04-2008.en-us.pdf.

Thomas, Rob, and Jerry Martin. 2006. **"**The Underground Economy: Priceless." *USENIX ;login,* 31(6): 7–16.

van Eeten, Michael J. G., and Johannes M. Bauer. 2008. "The Economics of Malware: Security Decisions, Incentives and Externalities." OECD Science, Technology and Industry Working Paper No. 2008/1.

Varian, Hal. 2004. "System Reliability and Free Riding." In *Economics of Information Security,* Vol. 12, *Advances in Information Security,* ed. L. Jean Camp and Stephen Lewis, 1–15. Boston: Kluwer Academic Publishers.