The Economics of Spam*

Justin M. Rao
Microsoft Research

David H. Reiley
Google, Inc.

Keywords: *spam, externalities, email, arms race, screening*
JEL Codes: *D02, D23, D62*

The term "spam," as applied to unsolicited commercial email and related undesirable online communication, derives from a popular Monty Python sketch set in a cafe that includes the canned-meat product SPAM in almost every dish. As the waitress describes the menu with increasing usage of the word "spam," a group of Vikings in the cafe start singing, "Spam, spam, spam, spam, spam," drowning out all other communication with their irrelevant, repetitive song. The analogy to unsolicited commercial solicitations jamming one's inbox seems quite apt. Every day about 90 billion emails are sent to valid email addresses around the world; in 2010 an estimated 88 percent of this worldwide email traffic was spam (Symantec, 2010; MAAWG, 2011). Almost all of this spam is illegal under current laws.

How does spam differ from legitimate advertising? If I enjoy watching network television, using a social networking site or checking stock quotes online, I know I will be subjected to advertisements, many of which may be irrelevant or even annoying to me. Google, Yahoo!, Microsoft, Facebook, and others provide valuable consumer services, such as social networking, news and email, supported entirely by advertising revenue. While people may resent advertising, most consumers accept that advertising is a price they pay for access to valuable content and services. By contrast, unsolicited commercial email imposes a negative externality on consumers without any market-mediated benefit, and without the opportunity to opt out.

This negative externality makes spam particularly useful for teaching purposes. When asked for an example of an externality, most economists think of environmental pollution: groundwater toxins, acid rain, air pollution, global warming, and so on. Indeed, given the great linguistic generality of the term "pollution" (including noise pollution, light pollution, and others), it can be

*The authors completed this work while both were working at Yahoo! Research. We thank Yahoo! for giving us an unprecedented opportunity to pursue research with academic freedom and access to corporate data. We are grateful to our former Yahoo! colleagues Carlo Catajan, Raghav Jeyeraman, and Gareth Shue for taking time to help us understand the institutional details of this market. Anirban Dasgupta, Randall Lewis, Preston McAfee, Kunal Punera, Michael Schwarz, and Andy Skrzypacz provided helpful discussions. We thank JEP editors David Autor, John List, Chiang-Tai Hsieh, and especially managing editor Timothy Taylor, for advice and support in structuring and revising this paper.

difficult for economists to find an example of a negative externality that *cannot* be described as a form of pollution.[1] Spam offers an alternative example of an externality: distinct from pollution, familiar to students, and possessing a short, memorable name.

Of course, a similar externality has been present for decades in junk mail, telemarketing, and billboards. These intrusive activities also impose claims on consumer attention without offering compensation or choice. However, email spam is breathtakingly larger in magnitude, with quantities (if it went unfiltered) in the hundreds of emails per user per day—if our inboxes stood unguarded, they would quickly become totally useless. (In contrast, junk mail has not yet reduced our unguarded postal mailboxes to this fate.) One can purchase unsolicited email delivery on the black market for a price at least a thousand times less than that to send bulk postal mail. Spam has become such a widespread phenomenon that trademark holder Hormel finally stopped objecting to the use of the term to refer to unsolicited email (Templeton,).

Spam seems to be a truly extreme externality, in the sense that the ratio of external costs to private benefits is quite high. We estimate that American firms and consumers experience costs on the order of $20 billion annually due to spam. Our figure is more conservative than the $50 billion figure often cited by other authors, and we also note that the figure would be much higher if it were not for private investment in anti-spam technology by firms, which we detail further on. On the private-benefit side, based on the work of crafty computer scientists who have infiltrated and monitored spammers' activity (Stone-Gross et al., 2011; Kanich et al., 2008; Kanich et al., 2011; Caballero et al., 2011), we estimate that spammers and spam-advertised merchants collect gross worldwide revenues on the order of $200 million per year. Thus, the "externality ratio" of external costs to internal benefits for spam is greater than 100:1.

In this paper we start by describing the history of the market for spam, highlighting the strategic cat-and-mouse game between spammers and email providers. We discuss how the market structure for spamming has evolved from a diffuse network of independent spammers running their own online stores to a highly specialized industry featuring a well-organized network of merchants, spam distributors (botnets) and spammers (or "advertisers"). Indeed, email service provision has become more concentrated in part because the high fixed costs and economies of scale of filtering spam offer a significant advantage to large service providers. We then put the spam market's externality ratio of 100 into context by comparing it to other activities with negative externalities, such as pollution associated with driving an automobile, for which we estimate a ratio of 0.1–0.3, and for non-violent property crime such as automobile theft, for which we estimate a ratio of 7–30. Lastly, we evaluate various policy proposals designed to solve the spam problem, cautioning that these proposals often err in assuming away the spammers' ability to adapt.

[1] Editor David Autor points out traffic congestion as another useful teaching example that does not fall under the pollution blanket.

The History of Spam: Cat-and-Mouse Games

Email is sent via a “sender push” technology called Simple Mail Transfer Protocol (SMTP). Other examples of sender-push transfer include postal mail, text messaging and voice mail. (In contrast, the Hypertext Transfer Protocol (HTTP) used in web browsing is “receiver pull”—nothing shows up in your web browser until you specify a URL.) SMTP was designed in early 1980s when the trust level across what was then called the “Arpanet” was quite high. Accordingly, senders were not required to authenticate their emails. SMTP servers all over the world were programmed to cooperate in relaying messages. In many respects, SMTP replicates the “transfer protocol” of the U.S. Postal Service. Anyone in the United States can anonymously drop a letter in a mailbox and, provided it has proper postage, have it delivered, without any requirement for the sender to provide an authentic return address.

Spammers first developed technology to automate the sending of bulk email in the mid-1990s by opportunistically tapping into mail relay servers and anonymously floating a deluge of spam from phony domains (Goodman et al., 2007). In 1994, the attorneys Canter and Siegel hired a programmer to automate a posting to every USENET newsgroup in existence, so that thousands of discussion groups devoted to every topic from Star Trek to board games were inundated with advertisements for services to help immigrants apply for the Green Card lottery. This software soon evolved into the first automated bulk emailer (Zdziarski, pp. 10-13). In 1995, the first commercial “spamware,” aptly titled “Floodgate,” appeared for sale at a price of $100. Floodgate advertised its ability to harvest email addresses from a variety of sources including newsgroups, CompuServe classified ads, AOL Member Directory, and other sources. Then, via the included companion software Goldrush, it promised an ability to send out thousands of emails per hour (Zdziarski, p. 16, and Everett-Church, 1997). Such software, crude by today’s standards, enabled spammers to send email at a cost on the order of $0.0001 per message. Since then, the spam market has been shaped by the technological cat-and-mouse game between spammers and email service providers.

Anti-spam filtering techniques

As an early response to spam, Internet administrators developed authentication protocols: where previously one only had to type a password to collect one’s incoming mail, now most had to authenticate themselves by providing a password to send outgoing mail. To prevent domain spoofing—using the domain of a well-known company to make an email seem more legitimate—domain authentication routines check that the IP address listed in the Domain Name System matches the sending IP. However, many SMTP servers remained unauthenticated for a long time, and the default mail delivery protocol is still to deliver email no matter from where it is sent.

After authentication, the arsenal of filtering technologies consists of machine learning, crowdsourcing, and IP blacklisting. Such screening devices detect suspected spam messages and either reject them from being delivered, or send them to a junk-mail folder.

The machine-learning approach dates from the late 1990s (Sahami et al., 1998; Androutsopoulos et al., 2000). A typical machine-learning implementation uses “ground truth” data on a subset of observations to learn rules to classify the remaining data. With spam, the ground truth is given by human-labeled examples of spam and the algorithm is trained to recognize features of the email that predict whether it is spam. For example, one can have the classifier examine the predictive power of all words and word pairs in the subject lines of the emails, which might lead to dummy variables for the presence of “Viagra,” “Nigeria,” and “Free Money” being included as key predictors. Other examples of heavily weighted features include unusual punctuation common to spam, such as “!!!” and nouns associated with spam-advertised products (such as “Rolex”). Every URL contained in a message could also be treated as a possible predictor, because spam emails nearly always include the URL of the website to place orders for the advertised product. Any machine-learned filter can have false positives, that is legitimate mail that is filtered to the junk folder (for instance, it might be hard to converse legitimately about bank transfers in Africa).

Spammers responded with creative misspellings designed to avoid the filters, such as “V1agra,” created many unique URLs all mapping to the same order form, and included attachments of graphical images of text messages, which became popular with spammers when they realized that text-based classifiers could not find the text in the form of a GIF or JPEG image. Spammers also include irrelevant text passages, such as excerpts of news stories that are common in legitimate conversations, or create random permutations of words in each email to throw classifiers off track. Of course, over time the anti-spam classifiers continued to improve and adapt, too. Goodman (2007) presents a nice introduction to such anti-spam technology for non-experts.

Crowdsourcing collects additional data to improve the predictive power of machine-learning models. Large webmail providers such as Yahoo! Mail collect up-to-the-minute data when users press the “mark as spam” button to move email from the inbox to the junk-mail folder. Data from such marked spam can be used, as soon as that same day, to retrain the spam classifier. However, most users just delete irrelevant emails rather than marking them as spam. We took a random sample of six months of mail activity for 1.3 million active Yahoo! Mail users and found that only 6 percent of users ever marked any email as spam, but the vast majority deleted messages without reading them.

Spammers have developed a strategic response to the spam voting system. In addition to the “spam” button in the inbox, webmail services also provide a “not spam” button to mark false-positive messages in the junk-mail folder. In four months of 2009 Yahoo! Mail data, our former colleagues found that (suspiciously) 63 percent of all “not spam” votes were cast by users who never cast a single “spam” vote. After examining additional data on these accounts, such as IP address, position in the network of users, and repeatedly casting not-spam votes on a variety of emails that were receiving multiple “spam” votes from legitimate users, the authors concluded that the vast majority of these accounts were created by spammers to help their campaigns beat the spam filters (Cook et al., 2006; Ramachandran et al., 2011). The researchers discovered 1.1 million of these sleeper accounts and Yahoo! inserted a detection algorithm to mitigate this vote gaming.

The single most effective weapon in the spam blocking arsenal turns out to be "blacklisting" an email server (Cook et al., 2006; Ramachandran et al., 2007). In 2011, 80 percent of all emails received by Yahoo! Mail were rejected by their servers through IP blacklisting. Fortunately, just as the postmark from the sending post office limits the ability to spoof one's return address, Transmission Control Protocol (TCP) makes it impossible to spoof the IP address of the mail server from which the message was sent. Therefore, if email administrators noticed that their users were receiving tremendous amounts of mail from one server, they could "blacklist" such a server. Sharing blacklist information enables multiple organizations to shut down spam activity more quickly. For example, the Spamhaus Block List, founded in 1998 by Steve Linford, now protects nearly 1.8 billion email inboxes from spam.[2]

An unintended side effect of blacklisting occurs when a single user starts sending spam and causes their email server to be blacklisted. At that point, everyone else using the same email server will suddenly find their emails being blocked. This situation could arise within any large organization, like a college, or for a user of a small Internet service provider (ISP). Of course, information technology professionals can then sort out the problem, and organizations such Spamhaus strive to act quickly in unblocking any email server who was falsely accused or who corrects the problem with its users, but blacklists still routinely cause reliability problems for users trying to send email.

The larger email services such as Yahoo! Mail, Microsoft Hotmail, and Google Gmail have large, dedicated anti-spam and customer-support teams. The high fixed costs of anti-spam technologies and benefits of crowdsourced data have made it difficult for small email providers to compete, which has contributed to significant increases in concentration in email provision since the mid-1990s. We managed to obtain market-share data for the top 50 largest consumer web-based email services (including home ISPs) for the period 2006–2012. The data show that webmail provision has become increasingly concentrated in the "Big Three" of Hotmail, Yahoo! Mail and Gmail. The three-firm concentration ratio in this market has increased from 55% to nearly 85% over the last six years; we believe that spam is a significant contributor to this increase in concentration.

Botnets

Blacklists gradually made it impossible for spammers to use of their own servers (or others' open relay servers) with fixed IP addresses. Spammers responded with a "Whack-a-Mole" strategy, popping up with a new computer IP address every time the old one got shut down. This strategy was observed and named as early as 1996, and eventually became considerably cheaper with another major innovation in spam: the botnet.

[2] Source: http://www.spamhaus.org/organization/index.lasso, accessed 9 February 2012.

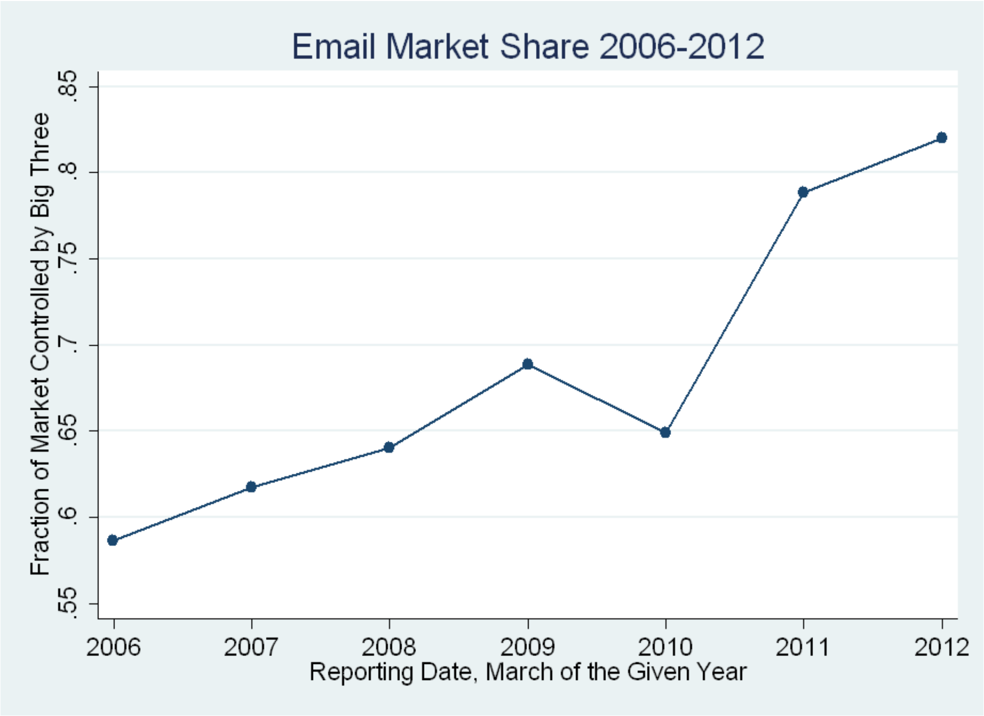

Figure 1: Webmail 3-firm concentration ratio, 2006–2012, *Source: comScore monthly reports from March of each year.*

A botnet is a network of "zombie" computers infected by a piece of malicious software (or "malware") designed to enslave them to a master computer. The malware gets installed in a variety of ways, such as when a user clicks on an ad promising "free ringtones." The infected computers are organized in a militaristic hierarchy, where early zombies try to infect additional downstream computers and become middle managers who transmit commands from the central "command and control" servers down to the front-line computers (John et al., 2009; Caballero et al., 2009; Cho et al., 2010).

The first spamming botnets appeared in 2003. Static blacklists are powerless against botnets. In a botnet, spam emails originate from tens of thousands of IP addresses that are constantly changing because most individual consumers have their IP addresses dynamically allocated by Dynamic Host Control Protocol (DHCP). Dynamic blacklisting approaches have since been developed; Stone-Gross et al. (2011) document that 90 percent of zombie computers are blacklisted before the end of each day. However, if the cable company assigns a zombie computer a new IP address each day, that computer gets a fresh start and can once again successfully send out spam.

In response to botnets, many ISPs, such as Comcast, began to prevent their users' computers from operating as send-mail servers. Unfortunately, this meant that individuals and small businesses could no longer run their own mail servers, as in the original, decentralized vision of the Internet, and now had to rely on larger commercial email vendors.

A second generation of botnets makes use of accounts at large commercial email providers. For example, a zombie could be programmed to sign up for hundreds of thousands of free email accounts at Gmail, and then send spam email through these accounts. Email providers have implemented sending thresholds designed to detect and prevent this sort of spamming. If a user exceeds these limits, the system may refuse to send out the email, or it may ask the user to solve a CAPTCHA (see next subsection). Such rules cut down on outbound spam, but also impose negative side effects on users who happen to be high-volume senders of legitimate email. In 2011, Yahoo! Mail experienced an average of 2.5 million sign-ups for new accounts each day. The anti-spam team deactivated 25 percent of these immediately, because of clearly suspicious patterns in account creation (such as sequentially signing up JohnExample1, JohnExample2,...) and deactivated another 25 percent of these accounts within a week of activation, due to suspicious outbound email activity.

In 2009, six botnets accounted for over 90 percent of botnet spam (Symantec, 2010; John et al., 2009). The largest botnet on record, known as "Rustock," infected over a million computers and had the capacity to send 30 billion spam emails per day, before it was taken down in March 2011. Microsoft, Pfizer, FireEye network security, and security experts at the University of Washington collaborated to reverse-engineer the Rustock software to determine the location of the command servers. They then obtained orders from federal courts in the United States and the Netherlands, allowing them to seize Rustock's command-and-control computers in a number of different geographic locations. (Microsoft financially supported the operation presumably because Rustock sent its emails through Windows Live Hotmail accounts, while Pfizer participated because a Rustock spam often advertised counterfeit versions of Pfizer's patent-protected Viagra.) If the servers had been located in less friendly countries, it is not clear whether the takedown could have been successful. The takedown of this single botnet coincided with a one-third reduction in email spam—and hence a one-quarter reduction in global email traffic (Thonnard and Dacier, 2011; Microsoft, 2011)—the efforts of these private firms produced a remarkably large *positive* externality.

CAPTCHA: Screening humans from bots

To avoid spammers setting up many commercial email accounts, services like Yahoo! Mail have implemented a screening device called a CAPTCHA, which is an acronym for "Completely Automated Public Turing test to tell Computers and Humans Apart." This test will be familiar to most readers as a set of twisty, distorted text characters. Spammers turned to visual-recognition software to break CAPTCHAs, and in response email providers have created progressively more difficult CAPTCHAS, to the point where many legitimate human users struggle to solve them.

However, the big breakthrough in CAPTCHA breaking arose when spammers figured out how to employ human labor to break CAPTCHAs for them. In this idea's first incarnation, a spammer would set up a pornography site, offering to display a free photo to any user who could successfully type in the text characters in a CAPTCHA image. In the background, their software had applied for a mail account at a site like Hotmail, received a CAPTCHA image, relayed it to the porn site, obtained text from by a user interested in free porn, and relayed this back to the Hotmail site (Kotadia, 2004).

More formal labor markets subsequently developed for CAPTCHA breaking (Motoyama et al. 2010). A market maker typically operates one website for interacting with buyers of CAPTCHA-breaking services, and another for interacting with workers who sell their labor. For example, one can purchase solutions from the "DeCaptcher" website, which transmits each CAPTCHA to a worker at the "PixProfit" website for breaking, then back to the customer at DeCaptcher. The customer may use a separate piece of software (such as GYCAutomator, which specializes in Gmail, Yahoo! Mail, and Craigslist CAPTCHAs) to transmit the CAPTCHA and its solution. The entire process takes less than 30 seconds. The market wage advertised for CAPTCHA-breaking laborers declined from nearly $10 per thousand CAPTCHAs in 2007 to $1 per thousand in 2009. These labor markets started with Eastern European labor and then moved to locations with lower wages: India, China, and Southeast Asia.

In February 2012, Kotalibablo.com was advertising to workers that they could earn wages starting at $0.35 per thousand. The same company operates the buyer-facing website Antigate.com, which at that time advertised a price of $0.70 per thousand to customers wanting to break CAPTCHAs. Motoyama et al. measured typical response times of around 10-15 seconds per CAPTCHA, with accuracy rates around 90 percent.(During one peak-load period, they experimentally measured a labor supply elasticity of approximately one: increasing their bid amount from $2 per thousand to $5 per thousand increased quantity solved from 8 to 18 per second.) Several websites can provide more than ten CAPTCHAs per second, putting total industry capacity (at a price of $1 per thousand) at over a million broken CAPTCHAs per day. These services market themselves as "Image to Text" providers and operate in the light of day—as of 2012 US law, there does not appear to be anything illegal about the services they offer.

CAPTCHAs are also used to authenticate senders in what are known as "challenge-response systems." Upon receiving an email from someone not in a preset contact list, an auto-reply is sent before the original message is delivered. The auto-reply requires that the sender solve a CAPTCHA, thus authenticating themselves as a human. Such systems have been available for at least seven years, but the market has for the most part rejected this technology, and with good reasons (Isacenkova and Balzarotti, 2011). First, the auto-reply "challenge" is often caught in a spam filter because it contains stock text, a link and is sent frequently from the same sender (all strong signals in machine-learned spam filters). Second, it requires that receivers maintain a continually updated contact list. Third, spammers can use the challenge-response system to bounce messages to unsuspecting "senders," which is known as "backscatter spam."

Hijacking accounts from legitimate users

Another recent strategy of botnets has been to hijack existing email accounts from legitimate users. (These same techniques can be used for even more nefarious purposes, such as hijacking a bank account; for more details, we refer readers to a recent JEP article about online crime (Moore et al., 2009).) Moore et al. 2009 in this journal that addresses this form of online crime in detail.) For example, "phishing" occurs when the culprit sends an email posing as a legitimate institution ("Hotmail user account services," often including the actual logo of the institution being spoofed) and asks the victim to visit a website to "verify your account password." In the practice of "keylogging," a type of malware records keystrokes and transmits information (especially suspected passwords) to the spammer. The practice of "packet sniffing" takes advantage of small companies and colleges who still transmit user passwords over the Internet in unencrypted text, and so a spammer "listening" at a login page can not only hijack that account, but also any other accounts (such as Yahoo! Mail) for which the user has conveniently chosen the exact same password. This technique was recently used to obtain access to 93,000 accounts on the Sony Playstation Network (Gross, 2011).

In 2005, an industry consortium established a technology standard called Domain Keys Identified Mail (DKIM) as a new weapon in the war against both spamming and phishing. Now adopted by a number of firms, including Yahoo! Mail, Gmail, PayPal, and eBay, this standard creates a digital signature that email senders can adopt. For example, if a phisher pretends to be PayPal asking a user to verify their account password, Yahoo! Mail will immediately notice that the message does not have the correct digital signature (based on public-key encryption) and will therefore reject the forged email without delivering it. Unfortunately, spammers have already responded to this strategy by trying to hijack the account of a corporate user that has been "whitelisted" via DKIM. In March 2011, a number of accounts became compromised at Epsilon, an email service provider who handles the sending of legitimate bulk email for a number of corporate clients, such as TiVo, Capital One, US Bank, and the Kroger grocery chain (Moyer, 2011).

On the whole, anti-spam efforts at large companies have mitigated the nuisance of spam to customers. However, the cat-and-mouse moves seem certain to continue.

Spammers and the *Field of Dreams*

From a spammer's perspective, any web platform delivering eyeballs is a natural target. In other words, as Kevin Costner's character in *Field of Dreams* famously heard, *If you build it, they will come.*

Spam is prevalent on social bookmarking sites (Krause et al., 2008) and online classifieds (Tran et al., 2011). On Twitter, spam takes the form of inserting a spammy link to an ongoing conversation between users (Yardi et al., 2009), using Twitter's "hashtag" feature. Twitter spam also occurs when an ostensible fan of a celebrity writes a message including the characters "@LadyGaga," in hopes of getting it exposed to her fans. Facebook suffers relatively less from spam because of the way it requires users to verify connections with each other, but spammers continue to invent new techniques, from malicious apps to friend requests from fictitious identities, that keep Facebook's anti-spam team quite busy (Warren, 2011; Cohen, 2012; Ghiossi, 2010). Text-messaging spam has become a serious problem in certain countries: one source estimated that 30 percent of text messages in China are now spam. However, in the United States the relatively high price of SMS messaging (often $0.10 per message, orders of magnitude higher than in China)) has kept text-message spam rates below 1 percent (Gómez Hidalgo et al., 2006). Text spam is aggressively filtered by cell phone providers and especially for text messages from a computer to a phone through a webmail client (Almeida et al., 2011). Providers of online instant-message software also struggle to block spam.

Next to email spam, the most prominent form of spam is known as "web spam" or "black-hat search-engine optimization." A typical web-spam implementation mines news feeds for headlines and automatically creates pages with snippets of popular stories. The article snippet is used under a Fair Use exception to copyright law and the remainder of the page is typically saturated with advertisements. Such web spam is deceptive, signal jams search engines and annoys users, but it is not illegal. It differs fundamentally from all the other forms of spam discussed in this paper in that it is not sender push (one only sees a web spam page if one voluntarily clicks). Web spam has been combated by machine learning about credibility of potential links, combined with downgrading low-credibility links in search results.((Caverlee and Liu, 2007; Zhou et al., 2007); see also (Ntoulas et al., 2006) and (Castillo et al., 2007).)

Market Structure

Most spam is illegal under the United States CAN-SPAM Act of 2003, which requires unsolicited emails to have valid return addresses and opt-out provisions. While many people use “spam” to refer to the (sometimes annoyingly frequent) messages they receive from businesses they have previously transacted with, for the purposes of this paper we define spam to be messages from economic agents who do not have a previous relationship with the customer and who do not offer opt-out provisions.

However, the spam market does have some similarities to the market for legitimate online advertising, whose institutions have been described in this journal by Evans (2009), in the sense that spam attempts to generate a sale. However, while in legitimate advertising the whole point is to promote awareness (of a firm or a product), spam typically uses obfuscation to get its message through. Spam-based advertising is dominated by “affiliate marketing,” in which a merchant recruits intermediaries known as affiliates (aka spammers) to advertise on its behalf, in return for a share of the final purchase amount (Levchenko et al., 2011; Samosseiko, 2009; Kanich et al., 2011; Kanich et al., 2008). Thus, a merchant advertising via spam generally shrouds its identity, hiding behind an array of cookie-cutter storefronts, in order to increase the chances of getting its offer through to users.

The supply (or “publishing”) side of the spam market has become dominated by botnets, as discussed earlier. Several teams of computer scientists have demonstrated that botnets are distinct economic entities from the merchants on the demand side of the spam market(John et al., 2009; Kanich et al., 2011; Stone-Gross et al., 2011). Major merchants are advertised by multiple botnets, and botnets compete with each other for clients (Thonnard and Dacier, 2011). A botnet may either rent out its services to independent spammers, or send its own spam while acting as an affiliate for a merchant. Both business models appear to be widely practiced (John et al., 2009; Kanich et al., 2011; Stone-Gross et al., 2011). The market structure appears to be an oligopoly (Zhao et al., 2009). The Stone-Gross team infiltrated the Cutwail botnet and documented its offerings, which range from a bare-bones rental of computation time on the compromised machines, to a user-friendly interface allowing a customer to create a mass mailing and test it against open-source spam filters before sending. Like publishers in the legitimate advertising market, botnets invest in significant fixed costs of ad serving, match advertisers with potential customers, and offer large reach.

To probe the demand side of the spam market, Levchenko et al. (2011), a team of 14 co-authors based at the University of California at San Diego and the University of California at Berkeley, used spam feeds to identify examples of spam, a web crawler to follow advertised URLs, and botnet infiltration and botnet detection algorithms (see also (John et al., 2009)) to monitor botnet activity. Table 1 presents statistics on the merchants tracked through this technique. The first row shows that shows that spam for only 45 merchants was sent from 350 million distinct URLs during the

Table 1: Breakdown of the spam supply chain

Stage	Pharmacy	Software	Replicas	Total
Unique URLs	346,993,046	3,071,828	15,330,404	365,395,278
Domains	54,220	7,252	7,530	69,002
Store-front styles	968	51	20	1,039
Merchants	30	5	10	45

Source: *Levchenko et al. (2011)*

data collection period. The second row of the table shows that there are more than 5000 times as many URLs as domain names used by spammers. For example, a spammer might register the domain pharma.com and then host thousands of identical pages with different URLs on the same domain: pharma.com/buy123.html, pharma.com/purchase01.html, and so on. There are also more than one thousand domain names per merchant. A merchant may be represented by several affiliate spammers, each of whom might register multiple domains. Large, reputable registrars generally reject applications for spammy-sounding domain names, such as those containing "med" or "pharm" (Kanich et al., 2008), but hundreds of registrars are willing to look the other way Levchenko et al. 2011). Row three gives the number of "store-front styles," which represent individual user interfaces, each with a distinct look and feel. When law enforcement tries to shut down illegal sales, they look for identical storefronts and try to take them down all at once, so store-front variation helps a merchant avoid complete shutdown. For each pharmaceutical merchant, there are approximately 30 distinct store fronts; this figure is much lower for software and replicas. The final row of the table shows the number of merchants anchoring the market. Despite the large numbers of domains, URLs, and store fronts, only 100 merchants had a measurable market share of spam activity, and fewer than ten merchants account for over 80 percent of the market (Levchenko et al., 2011; Kanich et al., 2011).

After tracking the merchants via the botnets, Levchenko et al. (2011) placed 120 orders for the advertised goods, spread across the 100 identified merchants. The affiliated spammer usually hosts the entire consumer storefront experience; that is, the spammer generally collects payment information and then hands the transaction to the merchant before credit-card authorization. Payment-processing services for these merchants are quite concentrated: a total of only seventeen banks serve the 100 merchants, with just three banks (from Latvia, Azerbaijan and St. Kitts & Nevis) processing the payments for more than 75 percent of the transactions. Postage stamps on the packages revealed the physical locations where the goods originated: nearly all the pharmaceuticals came from India, for example, while replica watches generally came from China.

Overall, while spammers have nearly free entry in registering domains and renting services from botnets, merchants appear to face more significant fixed costs, especially in obtaining payment-processing services. Only a small number of banks appear willing to take the risk of associating with gray-market merchants. This may explain why a relatively small number of merchants supply most of the market for these spam-advertised goods.

Assessing the externality

What are the costs of spam to users, and how does it compare with the return to spammers? A widely cited report from Ferris Research placed the worldwide cost of spam in 2005 at $50 billion; Ferris raised its estimate to $100 billion in 2007 and $130 billion in 2009.[3] However, the Ferris report did not describe how they estimated such key parameters as the amount of time per worker spent deleting spam; indeed, one of the authors of that report indicated to us that their work was "not a scientific survey," but that it attempted to be a lower-bound estimate. The most common estimate of profits of spammers seems to involve the phrase "millions of dollars a day," which in turn originated in a widely cited IBM press release.[4] In this section, we find these widely cited estimates of user costs and spammer profits are somewhat exaggerated, but of the right order of magnitude.

Measuring the diffuse costs of spam

The negative externalities imposed by spam include wasted time for consumers: both wading through irrelevant advertisements in one's inbox, and missing an important message that went to the junk-mail folder. They also include the costs of server hardware, which requires more than five times as much capacity as would be required in the absence of spam, as well as the costs of spam prevention services provided by firms to reduce the burden on end users.

The chief challenge in totaling up the social cost is credibly estimating the number of hours lost by people dealing with spam. Estimating the amount of spam that beats spam filters is difficult—after all, if we knew it was spam, we would have filtered it. But one approach is to look at the success of spam in achieving sales, and then scale up to infer how many spam messages must have gotten through. For example, Kanich (2008) et al. report that 382 million mailing attempts resulted in 28 sales. Internal Yahoo! data on similar "high ticket" items sold through display advertising (those with marginal profit greater than $50) indicate they typically have conversation rates of about 1 in 25,000. If the Kanich group's spam had conversion rates of this magnitude, then approximately 700,000 of the original 382 million emails—1.8 percent—evaded spam filers.

Our second approach considers several pieces of data: prices of spam services reported in Stone-Gross et al. (2011), spammer profit margins reported in Levchenko et al. (2011) and IP-blacklisting

[3] The reports are now available at http://www.email-museum.com.

[4] See http://gizmodo.com/354741/ibm-says-storm-worm-creators-making-millions-daily. Phishing for account information in order to steal money is a form of online crime representing less than 0.3 percent of all email traffic. Researchers at Microsoft found that conventional wisdom was an overestimate by 50 of the true profits to phishing (Herley and Florêncio, 2009).

rates from Kanich et al. (2008), and internal Yahoo! reporting. Take the cost estimate of $30 to deliver 5 million spam messages: four million of these will be bounced through blacklisting, but the remaining one million will be delivered (not necessarily to the inbox). Spammers net about $50 per actual sale (Levchenko et al. 2011). Thus, to break even a spammer needs one conversion per 8.3 million messages sent. The Kanich group found 1 conversion per 13 million emails sent, which is in the right ballpark. Dividing by our estimated conversion rate of 1/25,000 places the estimate of spam getting to the inbox at 3 percent. Internal auditing at Yahoo! finds a lower figure, but large webmail providers are presumably more effective at filtering spam than the average email service.

Thus, let's say that 1.8-3.0 percent of the 50 billion spam messages sent each day get through to the consumer. Suppose a consumer's average value for an hour of time is $25, and that on average, each piece of spam takes 5 seconds to deal with (including time spent dealing with false positives). Total end-user cost would then be approximately $12–20 billion per year. If a combination of better filtering and higher conversion meant only half as much spam had to reach the inbox to support the market, the our estimate of user costs falls to $6 billion, which we believe represents a conservative lower bound.

As to the costs of anti-spam technology, Ferris Research (2009) estimated the costs at approximately $6 billion in the United States, based on surveys of firms purchasing anti-spam solutions. This seems roughly correct, given that the largest anti-spam service provider, Symantec, had $6.2 billion in annual revenues in 2011. It is hard to know exactly how much of this revenue was due to spam as opposed to network security, as anti-spam services are often bundled with other services. Other firms providing anti-spam services to corporate clients include McAfee, Trend Micro, and Barracuda. It would also be necessary to add the labor costs of the IT staff who install and maintain the anti-spam solutions, and the costs of additional server capacity required by spam email. We believe $6 billion is a reasonable estimate for the total, which represents approximately $30 per user across 200 million users.[5]

If firms were not investing in anti-spam technology, end users would be receiving 300 times as much spam. That is, the anti-spam investment of $6 billion could be considered to save a total of $600 billion in user time. Realistically, a 300-fold increase in inbox spam would render email unusable for regular communication.

Taken together, the total costs of spam in the United States today appear to be between $18-26 billion. Our estimate is about half that of the widely cited Ferris Research number, because we use lower value of end-user time, and also because we use a lower estimate of the number of spams that reach user inboxes.

[5] For reference, Yahoo! Mail incurs anti-spam costs of approximately $55 million per year for 500 million active email accounts, a cost of $0.10 per user per year.

Table 2: Cost of spam advertising relative to other vectors

Advertising vector	CPM	Break-even conversion %	Break-even conversions
		w/ marginal profit=$50	per 100,000 deliveries
Postal direct mail	$250–1000	2-10%*	2000
Telemarketing	$50–250	0.1-0.5%	100
Legitimate online display ads	1–5$	0.002-0.006%	2
Retail spam†	$0.10-0.50	0.001–0.0002%	0.2
Botnet wholesale spam†	$0.03	0.00006%	0.06
Botnet via webmail ‡	$0.05**	0.0001%	0.1

Sources: †*Stone-Gross et al. (2011)*, ‡*Motoyama et al. (2010)*

* Direct Mailing Association reports 2.2%, **assuming botnet rental is delivery vector.

Measuring the private returns to spam

Researchers have used three tactics to estimate the revenues of botnets and merchants: 1) monitor botnet activity and infiltrate spot markets for spam services, 2) hijack a botnet to estimate the number of purchases generated by a merchant through a spam campaign, and 3) estimate order volume through periodically placing one's own orders and examining the gaps in the sequential order ID numbers.

As an example of the first approach, Stone-Gross et al. (2011) infiltrated the (then prolific) Cutwail botnet. They were able to monitor every advertising campaign run on the botnet, recording message volume, purpose, and associated merchants. Next, the team infiltrated a private web forum operated by the botnet masters as a market for spam services. The authors document two ways to publish spam through the Cutwail botnet. Retail spam services were offered at $100 to $500 per million emails in this market. "Wholesale" spam service involves separately acquiring email address lists and renting time on the botnet's spam-send infrastructure. A monthly rental, capable of pumping out 10 million messages per day, was priced at $10,000, or about $33 per million emails sent. A more premium wholesale spam product, sending all messages through webmail accounts (and therefore incurring the higher cost of having to break CAPTCHAs), cost about one-third more. The authors estimate that the Cutwail botnet earned $1.7–4.2 million in profit during the 14-month period of study.

In Table 2 we convert the cost estimates to the standard unit used in the advertising industry: cost per thousand impressions (CPM). To put the figures in perspective, we also include estimates for the cost of sending consumers messages via direct mail, telemarketing, or legitimate online advertising. We next suppose the average transaction, or "conversion," in online-advertising parlance, to produce profits of $50. Given this assumption, Column 3 gives the conversion rate necessary to break even on the each form of advertising. For legibility, Column 4 restates the break-even conversion rates in units of conversions per 100,000 ads.

Direct mail is the most expensive form of advertising, due to printing and postage costs; this medium thus requires high break-even conversation rates of at least 2 percent. For the case of $50 profit per sales, standard online display advertising can be profitable down to a conversion frequency of 2 per 100,000 ads, while "premium display" would require 10 per 100,000 ads. Retail spam is profitable down to 0.2 conversions per 100,000. Bulk spam through "wholesale" botnet rental is sustainable with a mere 0.06 conversions per 100,000 ads, or about 1 in 2,000,000. Clearly, spam can be orders of magnitude less effective than traditional forms of advertising and still remain profitable.

The second research technique, hijacking a botnet, appears in the influential 2008 "Spamalytics" paper (Kanich et al., 2008), in which the researchers co-opted control of a portion of the Storm botnet by carefully modifying the software instructions given to a set of downstream zombie computers. The modified instructions replaced the link to the spammer's storefront with a link to their own replica storefront. Users could place an order at the replica storefront, but would then receive an error message. The researchers could thus measure how many conversions would have been generated by the spam emails with their modified instructions.

In total, the group modified 345 million pharmaceutical emails sent from botnet zombies. Three quarters of these were blocked through blacklisting, and the remaining 82 million emails led to a scant 28 conversions, or about 1 in 3,000,000. This conversion rate is far lower than what could be profitable for a retail spam campaign. We suspect that the reason for this lack of success is that a large portion of this major spam campaign went to large e-mail providers like Yahoo! and Gmail and failed to evade their spam filters. We hypothesize that small-scale spammers can beat spam filters more easily, and can spend time crafting creatively targeted campaigns, while large-scale bulk campaigns spray email like a firehose, the vast majority of it blocked by filters.

The same research group also introduced the third estimation technique: placing sequential orders and drawing inferences from order ID numbers (Kanich et al., 2011). They began by making multiple purchases only a few seconds apart. Ten merchants were determined to use simple ascending rules for order IDs; for these merchants the researchers placed a series of orders spaced over a six weeks period. The order IDs fully revealed the quantity of other orders placed in the intervening time periods. The researchers also learned that one spammer hosted the images for his storefronts on a server belonging to someone else, which he hijacked through malware. The researchers notified the server's owner, who in turn gave them permission to monitor requests for the relevant image URLs, which provided reliable data on average order size and the basket of goods purchased. Each of these ten large spam-oriented merchants earned between $500,000 and $1.5 million per month in *revenue*—of course, profits would be lower. The researchers project that the entire spam-oriented market grosses revenues of about $180-360 million dollars annually. Over 80 percent of purchase dollars come from North American and Western European bank accounts.

We can check this revenue estimate using estimates of the prices and quantities of spam emails sent. Through mail-account auditing, Symantec and MAAWG both put the volume of spam directed at Western Europe and North America at about 60 billion attempted connections per day (Symantec, 2010; MAAWG, 2011). Of these, about 80 percent are blocked. Similarly, the Yahoo! Mail team told us that in October 2011, they received approximately 30 billion attempted connections per day, 80 percent of which were bounced, just under 10 percent of which went to the spam folder, and just over 10 percent went to a user's inbox. If the unblocked 20 percent of spam is priced at $50 per 1,000,000 ("premium bulk" rates), this would amount to $600,000 worth of spam being sent to Europe and North America each day—so perhaps $750,000 worldwide. This figure seems a bit high given our previous estimate of just under $1 million per day in revenues for the entire supply chain (which must also include the cost of goods sold), but it is of the right order of magnitude.

Overall, we feel comfortable with an estimate of total industry revenue for spam-advertised goods on the order of $300 million per year. One might, in principle, want to include consumer surplus in a calculation of the total benefits of spam. However, because consumers who wanted these goods would likely be able to find them via online searches in the absence of spam, we assume that the consumer benefits are less than the total revenues earned by the spam industry. Since we have estimated the revenues rather than the profits of the spam industry, and we know there are marginal costs to the goods sold, we will assume for convenience that the revenues represent approximately the total surplus generated by spam, including both producer and consumer surplus.

The "externality ratio" of spam put into context

Spam to end users located in the United States costs approximately $20 billion annually, compared with approximately $200 million in surplus generated by the spam to these same users. The ratio of the cost of this externality to society relative to the ratio of private benefits it generates is about 100:1.

To put this *extreme* externality into context, Table 3 provides estimates for the externalities associated with (1) the air pollution from driving a vehicle, and (2) the (non-violent) stealing of automobiles. For driving, we use a low value for the benefit accrued to a driver, a figure just above the operation cost per mile. In reality, people make many inframarginal trips, valued by the consumer well over the marginal cost. The cost estimate comes from Delucchi (1997), who does a nice job of accounting for the social cost of the various air pollutants emitted by an automobile (time congestion externalities are not measured so this estimate should be viewed as the cost of driving on an un-congested roadway). Delucchi's preferred estimate for the social cost per mile was $0.06; using this figure gives an externality ratio of only 0.1, three orders of magnitude less than the value we obtain for spam. By contrast, stealing automobiles has a much higher externality ratio, as demonstrated by Field (1993). The societal costs include uninsured losses to victims, insurance premiums, law enforcement patrol costs and the cost of prosecuting and incarcerating offenders

Table 3: Extracted revenue, imposed costs and externality ratios

Activity	Revenue/benefit	Cost	Externality ratio
Driving automobiles†	$0.60 per mile	$0.02–0.25 per mile	0.03-0.41
Stealing automobiles‡	$400-1200 million per year	$8–12 billion	6.7–30.3
Email spam	$160–360 million per year	$18-30 billion per year	56–187.5%

Source: †*Delucchi (1997),* ‡*Field (1993), FBI Uniform Crime Report (2010)*

who are caught. Adding it all up, the costs imposed on society by auto thieves are a whopping 7 to 30 times the revenue extracted from the vehicle theft.

In certain ways, non-violent auto theft turns out to be a fairly close analogue to spam. The costs of both auto theft and spam are high, and are distributed diffusely across the majority of the population (insurance rates and law enforcement costs account for the bulk of the costs of auto theft (Field, 1993)). Relative to other types of crime with poor insurance coverage, both have particularly diffuse costs. Unlike most crime, spam has no specifically identifiable victim, no especially wronged persons inspiring law enforcement to vigorously bring spammers to justice. In fact, the "victims" of spam, those who voluntarily make purchases from illegal advertising, arguably exert large negative externalities on the rest of society. Accounting for how much spam actually reaches the inbox, we estimate that only about 1 in 25,000 people needs to succumb to the temptation to make a grey-market purchase to make it profitable for spammers to inundate everyone with advertisements at current levels. From an economic perspective, one could imagine a law-enforcement system that provides disincentives to make such purchases.

We do note that while the externality ratio of spam is large, the cost comes in the form of attention and time, not lung disease and death, such as in the case of air pollutants. We are not aware of any estimates of the externality ratio of violent crime, but a quick back-of-the-envelope calculation puts the ratio of armed-robbery in the range of 10,000.[6] So there do exist examples of externality ratios even higher than that of spam, though these tend to have their harm concentrated in a small number of people. Various forms of air pollution are similarly diffuse to spam, and may have much larger social costs than spam, but their externality ratios are much smaller.

Policy Proposals

Considerable effort has gone into anti-spam measures. We already discussed many of the private (and cooperative) technological solutions that have been adopted by firms in an attempt to reduce the social cost of spam. Here we consider public-policy proposals from the legal and economic perspectives.

[6] If an armed robber stands to gain $100 (due to differences in the marginal utility of income) and the chance of a death in an armed robbery is 1/1000, then the ratio would be 10,000 if we use $10,000,000 as the value of a life.

Legal interventions

American spam legislation began in earnest with the Telephone Consumer Protection Act (TCPA) of 1991, which, as a response to rising fax-machine spam, required fax marketing to be opt-in.[7] The legislation also required phone telemarketers to offer an opt-out. In 2003, a consumer challenge to unsolicited email was unsuccessful; the Pennsylvania Superior Court ruled in *Aronson vs. Bright-Teeth Now* (2003 Pa. Super 187, 824 A.2d 320) that email transmission, without the tangible costs of paper and toner, was legally different from fax transmission (Court, 2003). The TCPA did little to stop telemarketing, especially with the *Aronson* decision, because opting out on a firm-by-firm basis was difficult and time consuming. However, the National Do-Not-Call Registry adopted in 2003 allowed consumers to opt out of all telemarketing (with some exemptions for nonprofits and politicians) by filling out a single form.

The first national legislation directed at email spam was the Controlling the Assault of Non-Solicited Pornography and Marketing Act of 2003. The cumbersome title created the catchy acronym "CAN-SPAM." The law requires unsolicited email to have a valid return address, to offer a simple opt-out option, and to identify itself as advertising in the subject line. The CAN-SPAM Act does not appear to have markedly impacted the illegal advertising market (Sipior et al., 2004). One reason is that much of spamming activity was already illegal, including the sale of counterfeit goods infringing on trademarks and intellectual property rights, or pharmaceuticals that are illegal to dispense without a prescription in many jurisdictions (or even to ship across state lines to a consumer with a valid prescription). In addition, jurisdictional boundaries hamper spam prosecutions. A spammer may be based in Latvia, work for a merchant in Moscow, send spam to the United States from a botnet with zombie computers all over the world, and have the final goods shipped from India. Governments around the world have not been willing to strain diplomatic relations with other countries over spammers.

A different legal tactic has been proposed by Levchenko et al. (2011). Recall that they found a potential choke point for spammers: the small number of banks willing to process payment for the merchants. American authorities might seek penalties for U.S. banks who transact with spammers in places like Azerbaijan, Latvia and St. Kitts & Nevis. Indeed, some of the basis for such legislation could come from the war on drugs, since a fair number of spam purchases are for controlled narcotic substances such as oxycodone.

Economic policy proposals

To correct problems created by a negative externality, the standard solution in the economist's toolkit is to levy a Pigouvian tax on the externality-causing activity. In the case of spam, the popular economic solution is to require a "postage stamp," costing perhaps a tenth of a cent, for

[7] Some illegal fax spam continued. Horror stories from recipients are documented at http://www.junkfax.org/fax/stories/Kirsch.html.

delivery of an unsolicited email advertisement, and transfer that postage amount to the receiver to compensate them for their attention (for example, (Kraut et al., 2002) and Bill Gates at the World Economic Forum in Davos, Switzerland, in January 2004.(Jesdanun, 2004)). However, pricing *all* email in order to disincentivize the irrelevant material is highly inefficient: many legitimate and useful emails, such as flight reminders and nonprofit newsletters, might well cease to exist.

A related option would be to levy penalties on consumers who purchase goods from spammers, on the grounds that every purchase goes a long way toward increasing the profitability of spam to U.S. consumers. However, enforcing such a law would be quite difficult without severe restrictions on privacy—like giving government the ability to monitor purchase receipts sent to webmail clients.

Instead, economic authors generally prefer the oft-discussed proposal of "attention bonds," a variation on the Pigouvian tax designed to put a nonzero price only on unsolicited emails (Loder et al., 2004). Given the approximately zero marginal cost of transmission, a zero price is efficient for those emails known to be desired by the recipient. Senders would post an attention bond (representing a price the sender is willing to pay for the recipient's attention) to be collected by the recipient upon receipt of the message. However, the recipient also has the option of "white-listing" a sender, which allows email to be sent without bonds. Under this system, any unbonded email not on a user's registered whitelist gets bounced. Ideally, consumers would have the ability to set individual reservation prices. A high-school student might be willing to look at any unsolicited email whose bond exceeded half a cent, while a busy lawyer might require $20.

While we admire the elegance of attention bonds, we wish to sound a note of caution. No ability currently exists to link email accounts with payment mechanisms. Should such adoption eventually become feasible, how might spammers respond? With attention bonds, a cyber-criminal could earn the size of the bond per email, say $0.05 (a figure often suggested, see for example (Van Alstyne, 2007)), by hijacking a legitimate account and sending mail to that account to collect the bond. Account hijacking is already a serious problem, and the incentives to hijack would increase by at least three orders of magnitude if one could steal $500 by sending 10,000 emails from a hijacked account. Of course, countermeasures could then be taken, but our point is that the attention-bond system will surely produce attempts to exploit the new system for profit. By the time one takes into account the transactions costs of setting up an attention-bond system, along with a much heightened incentive to hijack accounts, the overall welfare effects of such a change are unclear to us.

There are two key inefficiencies at work with the sender-push property right of SMTP. This paper has thus far focused on the first: unsolicited email imposes an externality on user attention. The second is that if spam has arguably created a stigma for legitimate email marketers, destroying potential surplus that could be created by legitimate players who would, in the absence of such stigma, offer some well-targeted emails to consumers who would mostly appreciate them. This inefficiency has presented an arbitrage opportunity for middlemen, such as "Daily Deal" sites like

Groupon and LivingSocial. A Daily Deal site collects email addresses via consumers opting in. If the deals turn out not to be of sufficiently high quality, consumers can easily opt out with a single action (much easier than opting out of unsolicited emails from hundreds of individual merchants). Merchants reach consumers through the transmission rights of the middleman, and pay a substantial fee to do so. At this writing in 2012, Groupon's market valuation exceeds $5 billion dollars, which gives an idea of the size of this second inefficiency.

In contrast to the high-level market-design interventions that have been proposed, we feel the most promising economic interventions are those that raise the cost of doing business for the spammers, which would cut into their margins and make many campaigns unprofitable. As mentioned, one fruitful avenue is to put legal pressure on domestic banks that process payment from foreign banks known to act on behalf of spam merchants. This could put downward pressure on conversion rates and with them, profits. Another proposal comes from our colleague Randall Lewis, who imagines "spamming the spammers" by identifying spam emails and placing fake orders on spam-advertised stores. This would increase the merchants' costs dramatically, as they would find it much more difficult to fill orders, and their banks may raise fees on them if they submit many invalid payment-authorization requests. Of course, an unintended consequence is that from time to time, a legitimate merchant will be inundated with bogus product orders.

Email-spam advertising has evolved over the past 15 years from a handful of independent "spam kings" to a well-organized, sophisticated market. The spam supply chain includes merchants at the top, affiliate spammers downstream, and a relatively concentrated market of botnets producing the majority of the spam emails. Nearly 50 trillion spam emails per year advertise a variety of products, including pharmaceuticals, gambling, counterfeit watches, gray-market job opportunities, pornography, software, and dating services. The costs of spam to consumers outweigh the social benefits by an enormous margin, on the order of 100:1. While we admire high-level economic proposals to introduce Pigouvian taxes on spam, our research on the cat-and-mouse games played by spammers lead us to be cautious about the possible unintended consequences of these proposals. Instead, we advocate supplementing current technological anti-spam efforts with lower-level economic interventions at key choke points in the spam supply chain, such as legal intervention in payment processing, or even "spam the spammers" tactics. By raising spam merchants' operating costs, such counter-measures could cause many campaigns no longer to be profitable at the current marginal price of $20-50 per million emails. These proposals are no panacea, but could bring about a significant reduction in spam.

References

Almeida, T., Gómez, J., and Yamakami, A. (2011). Contributions to the study of sms spam filtering: new collection and results. In *Proceedings of the 11th ACM Symposium on Document Engineering*, pages 259–262.

Androutsopoulos, I., Koutsias, J., Chandrinos, K., Paliouras, G., and Spyropoulos, C. (2000). An evaluation of naive bayesian anti-spam filtering. *Arxiv preprint cs/0006013*.

Caballero, J., Grier, C., Kreibich, C., and Paxson, V. (2011). Measuring pay-per-install: The commoditization of malware distribution. *Proceedings of the 20th USENIX Security Symposium*.

Caballero, J., Poosankam, P., Kreibich, C., and Song, D. (2009). Dispatcher: Enabling active botnet infiltration using automatic protocol reverse-engineering. In *Proceedings of the 16th ACM conference on Computer and Communications Security*, pages 621–634.

Castillo, C., Donato, D., Gionis, A., Murdock, V., and Silvestri, F. (2007). Know your neighbors: Web spam detection using the web topology. In *Proceedings of ACM SIGIR*, pages 423–430.

Caverlee, J. and Liu, L. (2007). Countering web spam with credibility-based link analysis. In *Proceedings of the Twenty-sixth Annual ACM Symposium on Principles of Distributed Computing*, pages 157–166.

Cho, C., Caballero, J., Grier, C., Paxson, V., and Song, D. (2010). Insights from the inside: A view of botnet management from infiltration. In *Proceedings of the 3rd USENIX conference on Large-scale Exploits and Emergent Threats: Botnets, Spyware, Worms, and More*, pages 2–2.

Cohen, D. (2012). Busted: Fake facebook friend requests.

Cook, D., Hartnett, J., Manderson, K., and Scanlan, J. (2006). Catching spam before it arrives: domain specific dynamic blacklists. In *Proceedings of the 2006 Australasian Workshops on Grid Computing and E-Research-Volume 54*, pages 193–202.

Court, P. S. (2003). Aronson vs. bright teeth now. pages 187–824.

Delucchi, M. and University of California, D. I. o. T. S. (1997). *The Annualized Social Cost of Motor-Vehicle Use in the US, 1990-1991: Summary of Theory, Data, Methods, and Results*. Citeseer.

Evans, D. (2009). The online advertising industry: Economics, evolution, and privacy. *Journal of Economic Perspectives*, 23(3):37–60.

FBI (2010). Fbi uniform crime report on motor vehicle theft.

Field, S. (1993). Crime prevention and the costs of auto theft: An economic analysis. *Crime Prevention studies*, 1:69–91.

Ghiossi, C. (2010). Explaining facebook's spam prevention systems.

Gómez Hidalgo, J., Bringas, G., Sánz, E., and García, F. (2006). Content based sms spam filtering. In *Proceedings of the 2006 ACM Symposium on Document Engineering*, pages 107–114.

Goodman, J., Cormack, G., and Heckerman, D. (2007). Spam and the ongoing battle for the inbox. *Communications of the ACM*, 50(2):24–33.

Gross, D. (2011). Again? sony's playstation network hit with another attack. In *http://articles.cnn.com/2011-10-12/tech/tec_gaming-gadgets_sony-playstation-network-attack_1_lulzsec-passwords-sony-pictures?_s=PM:TECH*.

Herley, C. and Florêncio, D. (2009). A profitless endeavor: phishing as tragedy of the commons. In *Proceedings of the 2008 Workshop on New Security Paradigms*, pages 59–70.

Isacenkova, J. and Balzarotti, D. (2011). Measurement and evaluation of a real world deployment of a challenge-response spam filter. In *Proceedings of ACM ICM 2011*.

Jesdanun, A. (2004). Is metered e-mail a viable anti-spam tactic?

John, J., Moshchuk, A., Gribble, S., and Krishnamurthy, A. (2009). Studying spamming botnets using botlab. In *Proceedings of the 6th USENIX Symposium on Networked Systems Design and Implementation*, pages 291–306.

Kanich, C., Kreibich, C., Levchenko, K., Enright, B., Voelker, G., Paxson, V., and Savage, S. (2008). Spamalytics: An empirical analysis of spam marketing conversion. In *Proceedings of the 15th ACM conference on Computer and Communications Security*, pages 3–14.

Kanich, C., Weaver, N., McCoy, D., Halvorson, T., Kreibich, C., Levchenko, K., Paxson, V., Voelker, G., and Savage, S. (2011). Show me the money: characterizing spam-advertised revenue. In *Proceedings of the 20th USENIX Security Symposium*, pages 8–12.

Krause, B., Schmitz, C., Hotho, A., and Stumme, G. (2008). The anti-social tagger: detecting spam in social bookmarking systems. In *Proceedings of the 4th International Workshop on Adversarial Information Retrieval on the Web*, pages 61–68.

Kraut, R., Morris, J., Telang, R., Filer, D., Cronin, M., and Sunder, S. (2002). Markets for attention: Will postage for email help? In *Proceedings of the 2002 ACM conference on Computer Supported Cooperative Work*, pages 206–215.

Levchenko, K., Pitsillidis, A., Chachra, N., Enright, B., Grier, C., Halvorson, T., Kanich, C., Kreibich, C., Liu, H., McCoy, D., et al. (2011). Click trajectories: End-to-end analysis of the spam value chain. In *Security and Privacy (SP), 2011 IEEE Symposium on*, pages 431–446.

Loder, T., Van Alstyne, M., and Wash, R. (2004). An economic answer to unsolicited communication. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 40–50.

MAAWG (2011). Email metric report. In *http://www.maawg.org/system/files/news/MAAWG_2010_Q3Q4_Metrics*

Microsoft (2011). Battling the rustock threat. In *Microsoft Security Intelligence Report, Special Edition.*

Moore, T., Clayton, R., and Anderson, R. (2009). The economics of online crime. *The Journal of Economic Perspectives*, 23(3):3–20.

Motoyama, M., Levchenko, K., Kanich, C., McCoy, D., Voelker, G., and Savage, S. (2010). Re: Captchas–understanding captcha-solving services in an economic context. In *Proceedings of the 19th USENIX Security Symposium*, volume 10.

Moyer, E. (2011). Breach exposes, chase, capital one, tivo customers. In *http://news.cnet.com/8301-1009_3-20050068-83/breach-exposes-chase-capital-one-tivo-customers.*

Ntoulas, A., Najork, M., Manasse, M., and Fetterly, D. (2006). Detecting spam web pages through content analysis. In *Proceedings of the 15th International Conference on World Wide Web*, pages 83–92.

Parry, I., Walls, M., and Harrington, W. (2007). Automobile externalities and policies. *Journal of Economic Literature*, 45(2):373–399.

Ramachandran, A., Dasgupta, A., Feamster, N., and Weinberger, K. (2011). Spam or ham? characterizing and detecting fraudulent "not spam" reports in web mail systems. In *Proceedings of ACM CEAS 2011.*

Ramachandran, A., Feamster, N., and Vempala, S. (2007). Filtering spam with behavioral blacklisting. In *Proceedings of the 14th ACM Conference on Computer and Communications Security*, pages 342–351.

Sahami, M., Dumais, S., Heckerman, D., and Horvitz, E. (1998). A bayesian approach to filtering junk e-mail. In *Learning for Text Categorization: Papers from the 1998 Workshop*, volume 62, pages 98–105. Madison, Wisconsin: AAAI Technical Report WS-98-05.

Samosseiko, D. (2009). The partnerka: What is it, and why should you care. In *Proceedings of Virus Bulletin Conference.*

Sipior, J., Ward, B., and Bonner, P. (2004). Should spam be on the menu? *Communications of the ACM*, 47(6):59–63.

Stone-Gross, B., Holz, T., Stringhini, G., and Vigna, G. (2011). The underground economy of spam: A botmaster's perspective of coordinating large-scale spam campaigns. In *USENIX Workshop on Large-Scale Exploits and Emergent Threats (LEET).*

Symantec (2010). Messagelabs intelligence: 2010 annual security report. In *http://www.messagelabs.com/mlireport/MessageLabsIntelligence_2010_Annual_report_FINAL.pdf*.

Templeton, B. The origin of the term "spam" to mean net abuse. *http://www.templetons.com/brad/spamterm.html*.

Thonnard, O. and Dacier, M. (2011). A strategic analysis of spam botnets operations. In *Proceedings of the 8th Annual Collaboration, Electronic messaging, Anti-Abuse and Spam Conference*, pages 162–171.

Tran, H., Hornbeck, T., Ha-Thuc, V., Cremer, J., and Srinivasan, P. (2011). Spam detection in online classified advertisements. In *Proceedings of the 2011 Joint WICOW/AIRWeb Workshop on Web Quality*, pages 35–41.

Van Alstyne, M. (2007). Curing spam: rights, signals & screens. *The Economists' Voice*, 4(2):4.

Warren, C. (2011). How to: Avoid and prevent facebook spam.

Yardi, S., Romero, D., Schoenebeck, G., et al. (2009). Detecting spam in a twitter network. *First Monday*, 15(1).

Zdziarski, J. (2005). *Ending spam: Bayesian content filtering and the art of statistical language classification*. No Starch Press.

Zhao, Y., Xie, Y., Yu, F., Ke, Q., Yu, Y., Chen, Y., and Gillum, E. (2009). Botgraph: Large scale spamming botnet detection. In *Proceedings of the 6th USENIX symposium on Networked Systems Design and Implementation*, pages 321–334.

Zhou, D., Burges, C., and Tao, T. (2007). Transductive link spam detection. In *Proceedings of the 3rd International Workshop on Adversarial Information Retrieval on the Web*, pages 21–28.