The Economics of Spam

by

Stephen Cobb, CISSP

Senior Vice President
Research & Education

ePrivacy Group

Date: February, 2003

The Economics of Spam

The rising tide of unwanted email threatens to swamp our inboxes and drown out the messages we do want. But why? The “junk mail” that the postal service delivers is not increasing at the same alarming rate. Some mail delivered by the postal service might be annoying, but the unwanted email we refer to as spam is way beyond annoying. All spam is irrelevant and intrusive but an alarming amount is offensive, fraudulent and illegal. And that's not all: spam is a burden to the person who gets it and the network of entities that deliver it.

The Parasitic Economics of Spam

So, if the people who receive spam don't want it, and the entities that deliver it would rather not, why does spam still exist? The answer lies in economics, more specifically, "the parasitic economics of spam." In a nutshell, the parasitic economics of spam means this: the act of sending a message costs the sender less than it costs all other parties impacted by the sending of the message.

A comparison with the traditional "snail mail" that the postal service delivers illustrates how spam-e-nomics works. The cost per message for snail mail declines as you increase volume, due to printing and handling efficiencies, bulk postage rates, and so on. But after a certain point it stops declining, because there are limits to the sliding scale for bulk mail postage. If you view this as a curve, as shown here, the curve flattens out.

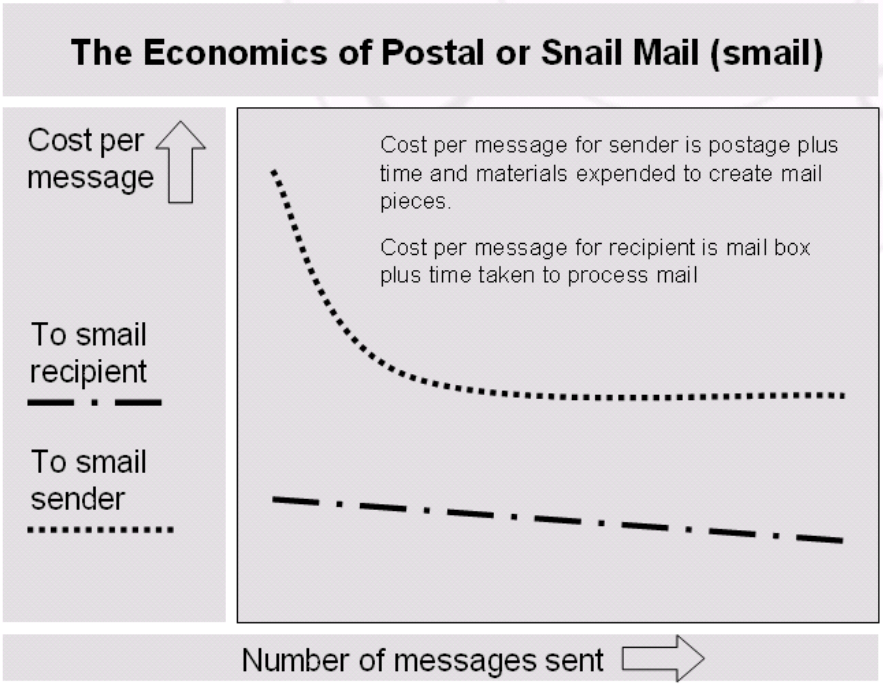

The per message cost of snail mail to the recipient is low to begin with, and relatively fixed. Most people check their snail mail once per day. Most junk mail is easy to identify and quickly sorted for disposal. By contrast, most people who use email check it multiple times per day, and spammers are constantly coming up with

new ways to make their messages look like legitimate email. This means that sorting is a nontrivial task performed multiple times per day.

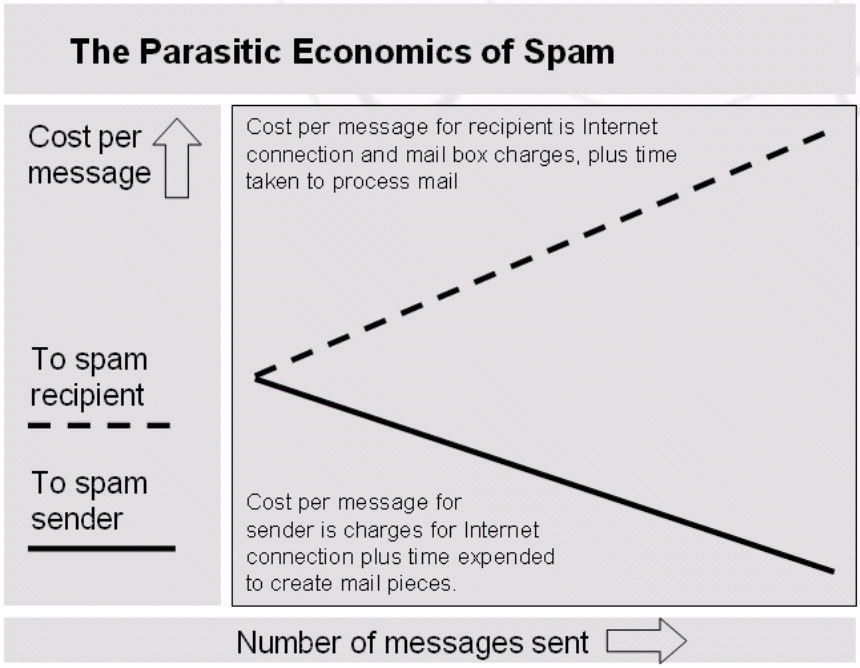

Conversely, per message costs for spammers start out low and decline rapidly with volume. Indeed, some spammers pay nothing for sending their messages, hijacking resources that belong to others. Software to automate spamming is cheap, as are mailing lists. Some spammers don't bother even bother with mailing lists and simply use a "brute force assault" on a target domain. The effect, as described by ISPs on the receiving end, is akin to that of a denial of service attack, wherein legitimate use of the ISPs services by those who pay for them is denied by a massive amount of illegitimate traffic (the term "brute force assault" is used because "brute force attack" has a meaning specific to computer intrusion and code-breaking-but the principle is the same in the same email is sent to all possible letter combinations at the target domain, as in aaaa@isp.net, aaab@isp.net, aaac@isp.net, and so on).

A Case Study

To see how the economics of spam work in real life, suppose you want to sell a new computer program that has a retail price is $49.95. You promise to pay a marketing company $19 for each sale they make. The traditional snail mail approach costs at least $1.00 for each brochure you mail to a prospective customer. Mail 5,000 brochures and you need a response rate better than 5% just to pay for the mailing (you need to sell 263 units at $19 commission each to pay the $5,000 mailing cost).

Take the bulk email route and the economics are quite different. A study by the Wall Street Journal (11/13/02) showed that a return rate as low as 0.001% can be profitable when using email. In one case cited, a mailing of 3.5 million messages resulted in 81 sales in the first week, a rate of 0.0023%. Each sale was worth $19 to the marketing company, so it took in $1,500. The cost to send the messages was minimal, probably

less than $100 per million messages (even a home computer connected to the Internet on a 56Kbps modem can send thousands of messages per hour). By the time the marketing company made the pitch to all of the 100 million addresses it had on file, it would probably have pocketed more than $25,000 on this project. Companies like this can make $50 for a mortgage lead, $85 for a cell phone sale.

Of course, some of the people who get these messages complain, which causes problems for the marketing company, and possibly the product or service vendor who has hired them. But not enough people complain to hurt the bottom line for either the sender of the marketing email or the entity on behalf of whom the marketing is being performed.

The parasitic economics of spam mean that the cost of handling and delivering the messages is born by someone other than the sender. Recent estimates put the level of spam at close to half of all email traffic, so theoretically, if there was no spam, ISPs could spend half as much as they currently do on facilities and bandwidth and still handle the same amount of legitimate traffic. Clearly, spam is a huge burden on ISPs and their costs are undoubtedly passed on to people who pay for Internet service (this is, consumers who pay for either dialup or broadband access at home, and enterprises that pay for Internet access for employees and business units).

Crossing the Line

The email marketing company profiled by the Wall Street Journal is relatively respectable. It pays for the bandwidth it uses and claims to honor opt-out requests from the people to whom it sends email. In short, the company does not believe that the messages it sends are spam. You could almost imagine this company agreeing to abide by a code of practice that reduces the extent to which it sends unwanted email. But to think that spam will go away if laws are passed to impose a code of practice would be a big mistake.

That's because the really big profits in spam come from breaking laws, such as those prohibiting deceptive business practices (in the United States, deceptive business practices were made illegal under the 1938 revision of the Federal Trade Commission Act, and they are also illegal under the laws of many states).

If you are prepared to break the law and lie about your product, you can make huge profits, like those raked in by the Arizona company that "herbal" pills at $2.50 per bottle and resold them via email for $59.95, labeled as penis enlargement pills. The millions of email messages that company sent might have been full of fabricated testimonials and bogus claims, but they produced results. When the company was busted by the State Attorney General, the seized profits included nearly $3 million in cash; a large amount of expensive jewelry; more than $20 million in bank accounts; 12 luxury imported automobiles (including 8 Mercedes plus assorted models from

Lamborghini, Rolls Royce, Ferrari and Bentley); an office building; and assorted luxury real estate in Paradise Valley and Scottsdale.

Deceptive claims are not the only way spammers break the law. Spamming is forbidden by most ISP service agreements and to enforce this, ISPs suspend or terminate accounts that are used to send spam. The result is that a lot of spammers send their messages by using someone else's account, without their permission. In other words, they steal Internet service.

Spam impacts the places to which the spam is sent, and the places via which it is sent, in other words, the inboxes of the recipients and the servers of Internet Service Providers. Internet connectivity costs money and part of the monthly fee that many consumers pay for access now funds the processing and disposition of spam.

Anti-Spam

Many ISPs now have elaborate blocking and filtering systems in place to prevent spam clogging their systems, but these take time and money to install and administer, and they are not without serious and potentially costly side-effects. If messages are blocked in error, not only the senders, but the intended recipients get upset. Many organizations regularly send messages-such as notices, statements, newsletters, and offers-at the request, or with the express permission, of the people to whom they are sending. These messages may be sent in large volumes, but they are not spam. If they are blocked, by an ISP filter that erroneously identifies them as spam, then both sender and recipient are likely to complain. Financial losses might be claimed. Lawsuits might arise.

A different type of blocking, which relies on lists of "bad" Internet addresses, can impact individuals as well as companies. If a lot of spam starts appearing from a specific address or range of addresses, those addresses might be placed on a "black list" (the term "block list" seems preferable, but is not as widely used). If those blocked addresses belong to your ISP you might not be able to send mail to people whose ISPs are blocking those addresses, even though you have done nothing wrong.

A lot of consumers get their email from a service provider such as Yahoo or Hotmail or AOL. As a service to users, these providers offer various forms of spam filtering at the individual level. Typically, these filters relegate spam to a separate folder, separating it from mail that appears to be legitimate. A similar service is available to any email user who wants to install a spam-filter program on their personal computer. Unfortunately, because spam filtering at the individual level is prone to false negatives and false positives, the user still has to weed some spam out of the inbox and some legitimate messages out of the spam bucket. Furthermore, the user's ISP still has to bear the cost of handling all of the mail, spam or not.

Another approach to spam filtering is the consensus model, whereby people who get messages that they consider to be spam report them as spam to some coordinating entity. A computer program is used to coordinate all of this input, but the input itself is human. If enough people say a particular message is spam, it is deemed to be spam and then blocked. Whether or not a particular message is spam is very hard to determine through purely automated processes. Computer algorithms to identify spam are far from perfect. So it might seem like a good idea to get human input in this way. Unfortunately, a consensus-based filtering system still has to learn how to evaluate input and could be skewed into false positives by accident or intent. There are numerous gray areas when defining spam, such as political messages, charities using email for fundraising, and so on. Does the fact that a person did not ask to receive a message automatically make it spam? If the message is relevant it might be welcome, and the recipient make not even think about whether or not they specifically requested it. Voting on spam could be hijacked by people who are upset with a particular sender.

Spam filtering can be refined by using white lists, which tell the filtering program to let certain email through, usually based on the source. Individual users can add their bank to their personal white list to prevent email from the bank being filtered out as spam. Some companies that offer filtering services also use white lists to allow all mail through from certain source deemed to be reputable. While white lists can help refine spam filtering, they are currently prone to spoofing, or falsification of email source data.

Every email contains information about where it comes from, but current email technology has no mechanism for making sure that the information in the header is correct. So if spammers discovers that all email from nicebank.com is being allowed through spam filters because it is on the white list, spammers can make their email look like it comes from that source, even though the content has nothing to do with banking. Some of the spoofing that occurs in spam is a really a form of corporate identity theft, which can eventually erode consumer trust in a brand or company.

Anti-spam laws have been widely discussed in recent years, but today their ability to reduce spam may be limited. For a start, many spammers have shown a willingness to breaking existing laws, wither they are anti-spam laws (such as those in Washington and California) or general business laws, like those which prohibit deceptive business practices (for example, any emailer that offers, but fails to honor, an option to opt-out of future mailings, is probably violating such laws).

While stronger laws with stiffer penalties may be worth passing, the fact remains that catching spammers is extremely hard, given the current state of email technology. As anyone who has tried to track a pieces of spam to their source will know, the trail often dead-ends at hijacked servers, of countries beyond the pale of feasible litigation.

Conclusion

Until something is done to alter the parasitic economics of spam and prevent spammers using and abusing resources paid for by others, the cost of spam will continue to be a burden on the Internet, retarding adoption and expansion, soaking up dollars that could be creating new levels of email service offering greater security, more intelligence, and even greater productivity improvements than we have seen so far.

ePrivacy Group has been studying the spam problem intently for a number of years and has been working on technology, called SpamSquelcher™, that alters the parasitic economics of spam. Such technology should become widely available in 2003, either under license or in product form, such as an appliance.

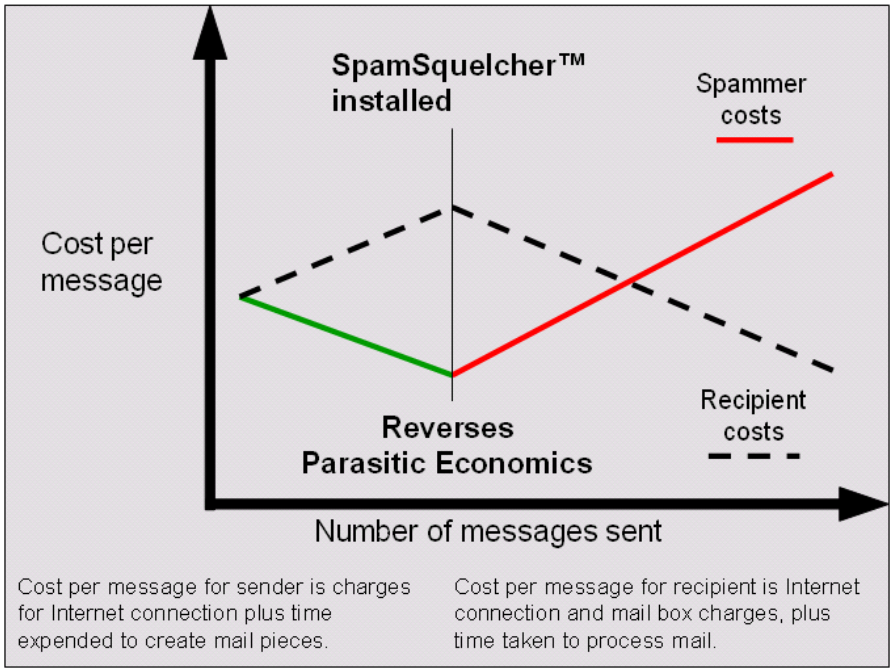

About ePrivacy Group:

ePrivacy Group is a trust technology company working to end spam by adding trust, privacy, and intelligence to email through initiatives like the Trusted Email Open Standard and patent-pending technology like Postiva™, which powers the Trusted Sender program overseen by TRUSTe and currently in use at MSN and other companies, and SpamSquelcher™ which dramatically reduces spam's impact on ISP and enterprise resources.

Founded by leading experts in privacy and security, ePrivacy Group is a privately held company based in Philadelphia, with offices in Los Angeles and Washington, D.C. For more information visit the website at http://www.ePrivacyGroup.com.

