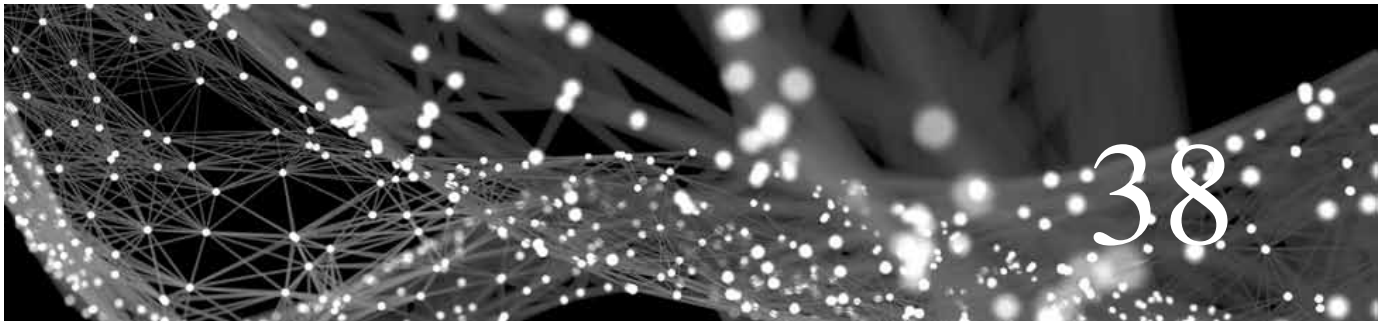

38

Spam

Finn Brunton

THE COUNTERHISTORY OF THE WEB

This chapter builds on a larger argument about the history of the Internet, and makes the case that this argument has something useful to say about the Web; and, likewise, that the Web has something useful to say about the argument, expressing an aspect of what is distinctive about the Web as a technology. The larger argument is this: spam provides another history of the Internet, a shadow history. In fact, following the history of 'spam', in all its different meanings and across different networks and platforms (ARPANET and Usenet, the Internet, email, the Web, user-generated content, comments, search engines, texting, and so on), lets us tell the history of the Internet itself entirely through what its architects and inhabitants sought to *exclude.* Identifying and describing spam, from the terminals of time-shared mainframes in the 1970s to the elaborate automated filtering systems of Gmail, meant having to talk about what the network is for, what the rules are, and who's in charge. And the first conversation, over and over again: what exactly is 'spam?'. Briefly looking at how this question got answered will bring us to the Web and what made it different.

Before the Web, before the formalization of the Internet, before Minitel and Prestel and America Online, there were graduate students in basements, typing on terminals that connected to remote machines somewhere out in the night (the night because computers, of course, were for big, expensive, labor-intensive projects during the day – if you, a student, could get an account for access at all it was probably for the 3 a.m. slot). Students wrote programs, created games, traded messages, and played pranks and tricks on each other. Being nerds of the sort that would stay up overnight to get a few hours of computer access, they shared a love of things like science fiction and the absurd comedy of *Monty Python's Flying Circus.* Alone at the terminals, together on the network, they would volley lines from *Python* sketches back and forth

– the dead parrot sketch, the dirty fork sketch, the spam sketch. This last was particularly popular because most of the dialogue was just the repetition of 'spam', whether sung by Vikings or shouted by the waitress, and it was therefore trivial to generate. You could write a simple program that, at the right spot in the dialogue, would post 'SPAM! SPAM! SPAM! SPAM! SPAM! SPAM! SPAM! SPAM!' over and over, relentlessly and without pause, filling the screen, killing the discussion, and often overloading the chat platform completely, kicking people offline. Jussi Parikka and Tony Sampson (2009), in the context of their larger analysis of spam, have shown that the sketch itself is built around a communications breakdown: the point where noise on the line overwhelms any particular signal. It was annoying, but playful and mischievous rather than malign, like unexpectedly blowing a vuvuzela in the middle of a conversation. This kind of noisy, frustrating behavior was dubbed 'spamming'.

The term came in useful in the ensuing decade-plus – though not with reference to advertising or commercial messages, which were their own category of etiquette violation. 'Spamming' remained the domain of noise and the indiscriminate, wasting time, attention, and bandwidth on redundant copies of messages, on overly verbose and off-topic postings, on tedious rants and cut-and-pasted slabs of text. Dave Hayes, prominent on the discussion system Usenet, wrote a manifesto in 1997 itemizing the forms of social misbehavior online – with 'commercial self-promotion' as a separate entry from 'SPAM', which meant precisely being a high-noise low-signal attention hog over a precious and expensive medium (Hayes, 1996). This definition shifted irrevocably in the spring of 1994, when two lawyers from Arizona posted a message across Usenet – that is, to computers around the world, to many thousands of users, indiscriminately – offering their services with the process of entering the US Green Card lottery. (We know this event best through the responses that quote it at the time – for instance (Larson, 1994).) The lottery was an initiative to simplify and speed up the process of getting papers to live and work in the United States as a foreign national: a subject of obviously narrow interest, which they had broadcast to computers from Singapore to Australia to the Netherlands – and, of course, throughout the United States, where the vast majority of recipients were citizens already. To enter the lottery, furthermore, only required sending in a postcard, but the message suggested that paid legal help would be needed – a sleazy commercial misrepresentation. Abuse of the network's many-to-many tools and global reach had been combined with a moneymaking scheme.

In the process of trying to describe this event, the global community on Usenet settled on *spam* as the term of art for their message and their action: the American lawyers had spammed the network. The word had jumped into the domain in which we identify today … or rather, it had come closer to our current understanding, in a way that is intimately intertwined with the development of the Web.

After assembling the whole history of spam, from an anti-Vietnam War message distributed on MIT's early time-sharing system in 1971, to a Digital Equipment Corporation ad on ARPANET in 1977, to present-day phishing, comment spam, 'spammy' posts and social network activity, clickbait and 419 ('Nigerian price') emails, I arrived at a working definition. What is 'spam'? *Spam is the manipulation of information technology infrastructure to exploit existing aggregations of human attention.* That is the meaning of 'spam' once all the technological particulars of search engine spamming or phishing campaigns have been worn away: following the term's broad application across the decades – both in English and as a loanword in other languages on the Internet – includes commercial and noncommercial activities, criminal and legitimate, with many different technologies and platforms, from email to

Twitter to content production. What remains consistent, I argue, is the model: spammers identify already existing collections of human attention, and imitate and manipulate their particular properties to extract value. I want to say a few more words explaining this before focusing on the Web in particular.

Spam is an information technology phenomenon. Across their many modes and domains, spammers push the properties of information technology to their extremes: the capacity for automation, algorithmic manipulation, and scripting; the leveraging of network effects and vast economies of scale; distributed connectivity and free or very low-cost participation. So many neglected blogs and wikis and other social spaces are out there on the Web: automatic bot-posted spam comments, one after another, will fill the limits of their server space. What this means – beyond characterizing spam as an activity – is that spammers take advantage of existing infrastructure in ways that make it difficult to extirpate them without making changes for which we would pay a high price. Indeed, in Geert Lovink's argument (2005), spam is akin to other network failures like identity theft in being inherent in the design – constitutional elements of yesterday's network architecture. Spammers partially survive by finding places where the potential value lost and effort expended in locking-down could exceed the harm they do – which reveals those places, and their value, to us.

More exactly, spammers find places where the open and exploratory infrastructure of the network hosts gatherings of humans, however indirectly, and where their attention is pooled. The use they make of this attention is exploitative not because they extract some value from it but because in doing so they devalue it for everyone else – that is, in plain language, they waste our time for their benefit. Recall the objections against the Green Card spam campaign on Usenet: it wasn't simply that the lawyers were acting commercially but that they didn't respect salience, barraging everyone indiscriminately with their lame message, treating the whole network as a passive audience whose time was theirs to spend. (Some of the complex distinctions inherent in spam, 'trash', 'junk', and 'waste' are considered in the analysis collected in Parikka and Sampson (2009), especially Galloway and Thacker's 'On Narcolepsy' (2008), and Gansing (2011).) Two consequences follow from studying spam in this light. First, we can see the history of networked computing as a thread in the history of the management and distribution of attention – Alessandro Ludovico (2005) has argued that spam is best seen as one instance in a long history, from traveling salesmen to personalized bulk postal mail to eye-catching billboards, of trying to interfere with our thoughts and provoke us into some form of consumer desire. Second, we can see in concrete terms how the nebulous shape of community online used spam to define itself. The intersection of these two topics brings us back to the distinctive history of the Web, the ways it gathered attention and formed communities, shown to us anew through spam's crass, hustling, inventive counterhistory. I will break this counterhistory up into four sections, which describe from spam's side how the Web became searchable, central, and social.

JUNK RESULTS

How, though, could spam have come to play a role in the Web? My brief sketch of spam's origins, above, is all social spaces: conversation threads on Usenet, chat on time-sharing computer networks, and of course email, where 'spam' as a concept and a business scaled up. The Web, though, was a kind of document navigation system at first – a markup language and set of protocols for authoring and exploring knowledge through hypertext. It was a project suited to a polyglot scientific community: developed by an English computer scientist, revised by a

Belgian, with the first site built in French, hosted on a Californian machine in the walls of a Swiss research institute, a knowledge presentation and navigation tool for one of the twentieth century's biggest capital-S Scientific communities. It's a context, and a technology, in which you can no more imagine spam thriving than you can imagine mold growing on a space station.

But mold does in fact grow in outer space, behind panels, on gaskets and insulation, on walls, under clothes. As human attention condensed, collected, and pooled on the Web, from Erwise to ViolaWWW to NCSA Mosaic, techniques began appearing to absorb and exploit it. Take a year, from the middle of 1994 to 1995, when we can see many different factors picking up speed: the global growth of users away from computer science professionals to the general population; the end of the noncommercial restrictions on the network in the United States; the shift in power from sysadmins to lawyers and entrepreneurs as social arbiters; and events like the Green Card lottery spam on Usenet (with subsequent publicity in newspapers – the first appearance of 'spam' in print – and major advertisers scenting blood in the water); Mosaic's booming download numbers; and the publication of *How to Make a Fortune on the Information Superhighway* – a cash-in book by those same Arizona lawyers, promising to teach readers how to market across the global network and get rich quick by exploiting the technology.

There were 20-odd websites in the fall of 1992, 10,000 by the end of summer in 1995, and millions by mid 1998. There were so many sites by then that finding what you were looking for – even knowing what was available to be found – was an enormous challenge, eloquently described in a paper published on April Fool's Day of 1998: 'The Anatomy of a Large-Scale Hypertextual Web Search Engine', by Sergey Brin and Larry Page. Others had already tried to solve the problem of abundance on the Web by developing search engines. The issue was, as Brin and Page put it, that 'some advertisers attempt to gain people's attention by taking measures meant to mislead automated search engines. … "Junk results" often wash out any results that a user is interested in' (2). Or, as a paper published on the very same day more bluntly put it: 'Some authors have an interest in their page rating well for a great many types of query indeed – spamming has come to the web' (Pringle et al., 1998: 1).

The form that spamming took reflects the unique particulars of the Web and search technology: it was designed not simply to dominate a conversation, as in chat, or flood a channel with messages, as with email and Usenet, but to make assertions of *relevance* and *salience.* We can see, through spam's development, how generations of search engines tried to model information on the Web in terms of what it meant for a user's query. The spiders that the first search engines sent out would go through the HTML source of a page, using the structure of the markup to assess the significance of words with greater or lesser degrees of importance and relevance to a search. A word in a URL (for uniform resource locator, the 'address' of the page) or in the first header tag – which is the markup for what the human reader would see as the 'title' of the page, as in <h1>My Homepage</h1> – was probably more important than one in the body text of a page and would be rated accordingly in the index. A set of elements called 'meta tags' were used in HTML specifically for the benefit of search engine spiders, with keywords listed for the page such that they would be invisible to the human reader but helpful to search indexing. Helpful in theory, anyway: though meta tag elements were popularized by early search engines such as AltaVista and Infoseek, they were so aggressively adopted by spammers that metadata was largely ignored by the turn of the century, with AltaVista abandoning the influence of meta tags on search results in 2002: 'the high incidence of keyword repetition and spam made it an unreliable indication of site content and quality' (Sullivan, 2002).

What precisely was the business plan of these early search spammers, and what were they putting in their web pages? Keywords were repeated in the meta tags and gathered in the page itself, hidden from the casual human reader's eye. One of the details that HTML can specify is the color of text, so the page's author could set the page's background to gray and make text the same shade of gray, invisible on the human reader's display while appearing to be normal text on the page as far as the spider was concerned. Innocuous pages with some form of spammy intent would have a mysterious gap at the bottom of the page. (Such techniques could also be used playfully or prankishly, part of the toolkit of the 'vernacular web' of bespoke HTML (Lialina, 2009).) The text on the page ended and there were no images, just a few inches of the gray background before the bottom. In that gap, in background-matching color and often minuscule font size, lay a magma flow of obscenity and pornography, product names, pop stars, distinctive phrases, cities and countries, odd terms seemingly plucked from Tristan Tzara's hat, selected because they happened to get good returns at that time. The text reads as though a Céline character worked for *Entertainment Tonight:* toyota ireland ladyboy microsoft windows hentai pulp fiction slut nirvana.

Such blocks of text illustrate a recurring theme in the development of spam on the Web and elsewhere: a matter-of-fact distinction between humans and machines, with different strategies for dealing with each. Almost every piece of spam, whether over email or in the context of spam blogs or comment spam, became *biface*, capable of being read in two ways with very different messages for the algorithm and for the human. (We will return to this distinction and its consequences for the Web at the end.) Spamming the early Web exacerbated this process, with techniques like 'cloaking'. Search engine spiders identify themselves by the way in which they request a page. This identification is part of the set of protocols that help to distinguish a normal Web browser from other platforms, like a phone, or a Braille display, making it possible to serve a compact page to the phone and text instead of images to the Braille device. This means you can serve one page to a spider, to be indexed and delivered as a search result, and an entirely different page to the user who clicks on the link. The signatures of spider requests, which trigger the cloak page, proved very difficult to disguise from spammers – which brings us back to Google's embryonic form in 1998: 'a prototype of a large-scale search engine which makes heavy use of the structure present in hypertext', created precisely to solve the problem posed by 'junk results', spammy web pages – and creating in turn a new way for spam to reshape the Web (Brin and Page, 1998: 1).

MUTUAL ADMIRATION SOCIETY

'The citation (link) graph of the web is an important resource that has largely gone unused in existing web search engines', wrote Brin and Page. 'These maps allow rapid calculation of a web page's "PageRank," an objective measure of its citation importance that corresponds well with people's subjective idea of importance' (3). Inspired by academic citation structure, they argued for reputation, essentially treating links as a measurable expression of social value. They were not the first to do this – earlier search engine projects, trying to beat the keyword-stuffing of the Web's first spammers, had tried to use numbers of links to roughly evaluate the meaningfulness of results. In response, spammers had started *link farms,* pages of nothing but links between spam sites, providing a cheap-and-easy boost to that metric. Part of Google's brilliance lay in the flaw in this strategy: spam pages are lonely. They may link to thousands of other sites, but the only *inbound* links, as a rule, come from other spam pages.

Links, in theory, carry an implicit endorsement, a vote of relevance made by a person. The spam-fighting question is: who is the person, and how much does their endorsement count for? Google's PageRank equation answered those questions with the behavior of the 'random surfer', an abstracted user of the late-1990s Web. This rather depressing model of a person starts on 'a web page at random and keeps clicking on links, never hitting "back," but eventually gets bored and starts on another random page' (4). The likelihood that this idle character, clicking ever forward along the link graph, will land on a given page defines PageRank. This means that other sites linking to your site matters – as does which sites link to those that link to you. It's a reputational model that works transitively, with links weighted differently by their significance: 'pages that have perhaps only one citation from something like the Yahoo! homepage are also generally worth looking at' (Yahoo! being, at the time, a directory of human-curated significant links.) What were spammers to do? Building on the algorithmic inference of social data, Google could make it 'nearly impossible to deliberately mislead the system'. The only workaround for spammers would be to build their own artificial societies.

A variety of strategies developed as Google's market share grew and other search engines around the world developed similar models. Websites with a high PageRank were transformed into kingmakers. A link from them could move a site onto the first page or top three returns of the different search sites, boosting attention and revenue. Sites took advantage of preexisting ideas for the human-curated Web, like 'Best of the Web' awards, 'Top 100 Sites' awards, and so forth; these awards included a badge, a little image, and a snippet of code to be copied into the winning site – a snippet that included a link to the award-giving site. The human user saw a little badge image, but the search engine spider saw an outgoing link: a digital endorsement. New habits of use and etiquette appeared among ordinary users of the Web: a comment in a blog post included the commenters' websites along with their names, to rack up another link. Posting something without including a 'via' link to the person you got it from – the 'via' being an additional outbound link as a kind of thanks for using their discovery – became increasingly rude, the sign of an uncouth person. 'Mutual admiration societies' arose, huge link-heavy sets of sites, each page linking to many of the others – all sites kept carefully unspammy, maintaining the pretense of legitimate use of the Web. Their business was not to produce spam sites themselves, but to charge for outbound links from the society. They were renting out their accumulated 'votes'. But even those had a characteristic shape: heavy cross-linking within a group of sites, all with only a few inbound links (because spam pages are lonesome), creating little islands of intense self-endorsement with no outside involvement. To analytic tools, it's a pattern as obvious as the newspaper ads taken out by vanity publishing houses for their new releases with the blurbs from friends and family – and easy for Google to discount accordingly.

In 1999, a company called Pyra Labs launched a service called Blogger. (Google would buy it in 2003.) Blogger's goal, as of so many related systems, from Flickr to Wikipedia, was to provide people with an intuitive means for publishing their content on the Web. It was remotely hosted, so you did not have to own a website domain name or pay for hosting; many of its processes were automated, so you did not have to design it or do any coding behind the scenes; and it had a useful and increasingly sophisticated Application Programming Interface (API) for connecting with other Web applications and automating processes. With the boom in weblog popularity and the peculiar chronological publishing model of blogs, came another three-letter acronym, RSS, 'Really Simple Syndication', which makes new posts or other changes on a site available in forms that are easy to use. (For the sake

of Web history completeness: RSS originally stood for 'RDF Site Summary', which highlights its relationship with the history of Web document formatting – but was retroactively changed to the more straightforward meaning.) Feed readers can gather the latest entries from RSS-enabled sites (which blogs soon were by default), material can be forwarded to mobile devices, and a page can feature the headlines or recent posts from other sites.

What does this have to do with Web spam? Consider the toolkit laid out by these developments: a content publishing system that can be easily automated (new accounts, posts on a prearranged schedule, modified settings), without the detailed work and paper trail of registering domain names and paying for Web hosting, and – with RSS – a faucet of other people's words, content that looked real and human because it actually was, unlike a lot of spam production. Hooked together with the right software tools, you can generate a new kind of mutual admiration and endorsement society with a network of spam blogs – or 'splogs'.

A splog production system will pull in RSS feeds from other blogs and news sources, chop them up and remix them, insert relevant links, and post the resulting material, hour after hour and day after day, with minimal human supervision. You can turn the machine on and leave the room while it makes money for you. With contextual advertising (including ads as a launching point for browser malware) you can make money through pageviews and the occasional click by running 'excerpt model' splogs, with fragments taken from other people's posts that are polling particularly well in Google's keyword metrics. A more ambitious system is 'full content' splogs, cross-linking in their hundreds and thousands to distort the shape of the Web. Each splog is assigned a set of keywords and feeds from which to pull related text, and in turn links to other splogs, which link to still more, forming an insular community on a huge range of sites – a kind of PageRank greenhouse that is not in itself meant to be read by people, but solely by search engine spiders. The splogs only work from a distance, appearing to be groups of people, the language and links functioning in aggregate. Taken in statistical total and algorithmic analysis, splogs resemble the patterns of a thriving community. Their posts are pitched at precisely the level of complexity the spider requires to accept their input as human, and they adapt human text for other machines to read and act on; affecting humans happens only indirectly – boosting the search ranking of a spammy appliance review site, for instance, that makes money through ads and affiliate links, or leading a human searcher into a fraudulent destination, whether a simple rip-off with ads and no meaningful content, or a site that might middleman a transaction, tacking on an additional fee or trying to force-download some adware.

This section opened with 'Google', a new-minted concept for a ranking algorithm, and ends with Google, a massively successful advertising company that runs a search engine. One question raised by this chronicle of spam's relationship to the Web and Web search is how complicit, or symbiotic, Google is with its own antagonist. Consider splogs: built and hosted on a platform Google owns, using text for content drawn from other sites hosted by Google, optimized to best fit Google's search engine algorithms, to boost the results for Google searches for sites that make money by hosting ads served through Google's affiliate advertising program (which, of course, also makes money for Google). Search engine spammers running their vast stables of spam blogs and sites are not anomalous, quantitatively or qualitatively; splogs now account for more than half of the total number of all blogs (Fetterly et al., 2004). They are the optimal users, from Google's perspective, constructing a system in which all the extraneous matter of people and conversation has been pruned away in favor of the automation of content production, search results, clicks, and ads served. This system in turn puts Google in the

contradictory position of having to analyze and expel many of their most dedicated customers: those who overexploit, and accidentally overexpose, the financial and attention economies and technologies that underlie the contemporary Web. Google is hardly alone in this problem, as we will see.

THE LANDING PAGE

Another shift in the dynamics of spam was developing, meanwhile, which reflects the Web's role as what Christian Sandvig (2013) calls an 'emerging essential'. It had been something that ran *on* infrastructure – running on top of the Internet, which ran in turn on top of the infrastructure of telephony – which became infrastructure itself, in the negative sense defined by Paul Edwards (2003: 187), 'those systems without which contemporary societies cannot function'. The Web became a key component for banking, health care, work and administration, and content creation and consumption, with browser standards and shared protocols becoming matters of urgent negotiation, monopolistic strategy, and even public safety (as in the push to standards like HTTPS) – with international implications from the top-level domains issued within the United States to legal decisions on hate speech in the EU (Goldsmith and Wu, 2008). In other words, it was not simply an aggregation of human attention around documents and content, navigated through search and hyperlinks, but a portal into many vulnerable and intimate parts of our working and personal lives. To understand how spam shifted accordingly, it helps to briefly look back for the last time to that Green Card lottery message in 1994.

What the lawyer-spammers were offering was, technically, an actual service (a misleading, borderline fraudulent one, but let that pass), whereas much of the spam we receive now has a very different agenda. You could call the real, working telephone number and schedule an appointment with the lawyers, just as, with much of the spam in the years following 1994, you could actually purchase the quack weight-loss pills, the deadstock toys, the counterfeit watches. Spam – spam of this era and this meaning – was loathed and despised, but it was also still somewhat legitimate, if only by accident. It thrived in the regulatory shadow of direct mail marketing, a powerful and moneyed interest that didn't want a legal precedent set that could close off a future advertising venue, and in the novelty of the increasingly popular and commercial Web, growing faster than legislation and defensive software could keep up. International guides to legal redress for spamming sprang up, their constantly updated confusion of potential laws – from CAN-SPAM in the United States to economic crime units in Norway to Canada's Department of Justice task force on pyramid schemes – highlighting the problem of figuring out where criminal lines were crossed (for example, Hollis, 2005). Many spammers were able to present themselves as brashly inventive promoters, with postal addresses and registered trademarks, seeking recognition in the classic tradition of entrepreneurial hustlers. What they produced is still what many people think of when they think of spam: the enthusiastic pitches full of mangled grammar and implausible stock photography, in the service of a recognizable, even traditional, class of dubious pleasures from timeshares and self-help books to diets and pornography.

But as the Web became an infrastructural system – and spam, for reasons too complex to go into here, became a progressively more embattled industry with less easy money – a new predatory spam technique took shape, to trick humans rather than machines. As far back as the mid 1990s, hackers and spammers alike had been finding ways to fool people into giving up their login information. Initially, the goal was to send spam from trustworthy-looking accounts, or within closed networks (whose members were often more naïve and easier to exploit).

In early 1996, in the Usenet newsgroup for the hacker magazine *2600*, the term for this technique makes its first appearance: 'phishing' ('mk590', 1996).

A representative spam business in the mid 1990s – who spelled it 'fishing' – used a simple ASCII picture, '<><', to note AOL accounts they'd captured to deluge people within the AOL network with spam ads (Brunton, 2013: 76). That was small change, though, compared with the uses to which phishing would be put. Targeted phishing messages used the same biface aspect of HTML – one side of the markup visible to people, the other for machines – to send email purporting to be from a bank, a credit card company, an email provider, or an employer, with a link whose innocuous text ('To resolve this block on your credit card, click here') disguises a suspicious URL (mastercard.l337haxx0r.ru, or whatever). Following the link reveals a careful – or sometimes not-so-careful – counterfeit of the original 'landing page' for the legitimate site. Careful study of spammer landing pages reveals how much HTML and CSS – the markup and styling vocabulary of early Web design – could convey the 'realness' of a particular online destination. Some crudely copy-and-pasted the HTML available through 'view source' commands for the sites they were pirating; others, faced with better-protected sites, reverse-engineered the design of their counterfeit, replicating color schemes, trying to duplicate the placement of images, and lifting or in some cases amusingly improvising the text.

Phishing sites are now often hosted on the compromised computers or servers that make up 'botnets', networks of many thousands of machines under the remote control of the spammer. These botnets generate the great bulk of the spam that we encounter (and much that we don't – spam can be upwards of 85% of *all* email on the Internet at peak times, most of it stopped by filters well before a person receives it (Messaging Anti-Abuse Working Group, 2011)). That verb, 'generate', is carefully chosen here: the botnet machines can run semi-autonomously, receiving command-and-control instructions for new spam campaigns and then spewing out messages in the millions, adjusting each one individually ('per-message polymorphism') and shifting strategy if they receive an unusually high number of rejections (Kreibich et al., 2008). In fact, a unique, targeted 'spear-phishing' attack, like the one that got John Podesta's Gmail login and disrupted the 2016 American election, is a flashback to a more artisanal, personal time in the Web. It had a carefully designed trick URL ('myaccount.google.com-securitysettingpage.tk') and a beautiful HTML email and login landing page, both mimicking Google's style to the pixel. (Similar care was taken with attempts to get internal email from the campaign of Emmanuel Macron in France – a tactic now so common that his staff prepared for it in advance.) That kind of attack is now the exception rather than the rule, the human touch for high-value targets. When we encounter a spam comment, a spam blog, a spam email, a message on Twitter @'d to you by a bikini avatar with a high-entropy name, we are very likely the first people to have ever seen it. It is the product of layers of wholly computational work, for which the humans merely set the parameters, assembled and passed around the world on a chain of mechanical writers and readers. This brings us to the last chapter of the Web's counterhistory, the closing act of those linkfarms and mutual admiration societies: the rise of the post-human social Web.

WE STILL BELIEVE THERE IS HUMAN INVOLVEMENT

The Web had always been social, of course, alight with cultures of linking, authoring, sharing, and citing, with forums, boards, comments, and 'virtual communities'. But as a matter of terminology, the Web became 'social' after it became searchable and

increasingly central, when Friendster, MySpace, Facebook, Orkut, LinkedIn, Bebo, Sina Weibo, Twitter, the *Marie Celeste* ghost ship that is Google+, and a million other social networks came to dominate much of the experience and use of the Web. This was an aggregation, a pooling, of human attention on a scale beyond a spammer's wildest dreams.

Furthermore, it was a model for aggregating human attention to which spamming came quite naturally. It was – and is – a terrain dominated by clickbait and linkbait, by eyeball-grabbing fake news (a technique pioneered by spam emails with links that launched malware downloads) and you'll-never-guess headlines bannered over the thinnest content, by the endlessly refilled candy bowl of meme culture, and advertisements indistinguishable from old-school spam come-ons (weight loss, penile enlargement, predatory home-loan scams, and other dragnets for dim fish). Even the legitimate human users developed spam-like approaches to their activity. Merlin Mann, a bemused witness to the dot-com scene, dubbed this activity on Twitter *personality spamming*, the work of arrogating attention for oneself, using social media to build an audience – often a very carefully quantified audience of 'followers' and 'rebloggers' – rather than a network of friends. It is the socially acceptable but aggressively eyeball-hungry work of those who would be, or act like, celebrities, 'influencers', or 'thought leaders'. From the Web 2.0 status culture analyzed by Alice Marwick (2013) to Whitney Phillips's media-savvy trolls (2015), to Limor Shifman's circulating memes (2013) and Sarah Jeong's 'Internet of Garbage' (2015), studies of the social Web capture how difficult it can be to distinguish from what its own users call spam.

Spam was a natural fit for this set of platforms and practices – so much so that it produced a similar paradox to that faced by Google, for which its most optimal customers were spammers. Spammers jumped into creating fake Twitter accounts and Facebook pages and YouTube accounts, serving up porn links and browser exploits and selling armies of Twitter followers, blocks of tens of thousands of Facebook 'likes' and YouTube views and upvotes. A huge portion of human time on the Web became devoted to interacting with and producing content that sought the illusion of salience through popularity – and if there was one thing spammers were good at, it was producing exactly that illusion. In a sweatshop model that recalls the early days of email spamming, employees of 'likefarming' firms will 'like' a particular brand or product for a fee. The going rate is a few US dollars for 1,000 likes (Schneider, 2004). Performed in narrowly focused bursts of activity devoted to liking one thing or one family of things, from accounts that do little else, this tactic is easy to spot, so they have to generate the appearance of casual use. They do this by liking pages recently added to the feed of Page Suggestions, which Facebook promotes according to its model of the user's interests – they behave, in other words, as ideal Facebook citizens, heavy users constantly clicking the thumbs-up and endorsing whatever Facebook's recommendation algorithm thinks they will endorse.

Twitter, likewise, has an enormous bot problem. It must regularly conduct sweeps to purge the bot accounts from its ranks, but the bots follow paying users in packs of thousands to make them look important and popular, as well as random humans to create the illusion of normalcy for their other activities. A too-successful purge is rewarded with outrage as users see their follower numbers plummet (and Twitter as a company sees its value drop, likewise, as the pool of active users shrinks). The same is true of buying views for your YouTube video, listens for your song on Soundcloud, or clicks on your ads. In a Web 2.0 version of Goodhart's Law ('When a measure becomes a target, it ceases to be a good measure', or, 'What gets evaluated, gets gamed'), any metric meant to describe human interest, esteem, or attention more generally will spur the development

of customized code to take it over for pay (Goodhart, 1981: 116). All of these vast social Web platforms have created models in which the spammers boost the metrics in exactly the way they're supposed to be boosted – just not as legitimate human users. (What these models say about the expectations for the humans, those hoof-clicking attention cattle, is left to the reader.)

I would like to close with a final problem, a ubiquitous mark on the structure of the Web left by spam, one that points towards the future: the humble CAPTCHA. The CAPTCHA system – the deformed letters on weird backgrounds that only humans can read, in theory, to verify their non-bot status – is meant to block automated posting, commenting, and account-creation tools, key components in the contemporary spammer arsenal. CAPTCHAs make it harder to start new Blogger blogs or open more free email accounts, and spammers have been working assiduously on different fronts to overcome them. In May 2008, the security company Websense documented a series of attacks on the account-creation process of email services. Many requests for accounts kept hitting the CAPTCHA stage, and most, but not all, failed (Whoriskey, 2008). The pace (replies in six seconds) and the failure rate (nine to one) suggested that computers were doing the solving. 'We still believe there is human involvement', said the company's statement. Botnet attacks on text recognition have improved enough since then that new forms of CAPTCHAs rely on identifying somewhat ambiguous visual information, like picking storefronts out of a set of pictures of buildings. To solve this, spammers have turned to automating humans with services like Captcha King, which retrieves the CAPTCHA images from things like the Twitter account-creation process for manual entry. An outsourced staff sits there all day banging out CAPTCHAs, with a guaranteed 'success rate of 95% with a response time of less than 90 seconds' (Krebs, 2012; Motoyama et al., 2010). Those poor souls, whose work makes regular data entry look exceedingly pleasant by comparison, are essentially being paid to be human – to exhibit a theoretically solely human characteristic.

In that labor, and in the statement 'We still believe there is human involvement', we can see a Web increasingly and finally dominated by the activity and content not of humans but of software, and humans directly responding to and directed by software. As Ben Light (2016) points out, human agency was always the centerpiece of Web 2.0 – but, observed more closely, it's clear we miss the real story of the contemporary Web if we fail to account for all the *nonhuman* agency and activity on it. This is an appropriately grim and paradoxical note for the close of this counterhistory of the Web: with Twitter accounts made by humans solving problems for machines, to provide other humans with the illusion of social activity, of 'Web presence'.

We have followed the work of drawing the line defining spam and non-spam through the history of the Web, from links and sites to blogs and search to the exquisite fakes that mislead users of the Web's infrastructure. Throughout, I have argued that the act of calling something 'spam' tells as much about what is being excluded as what is being identified: junk results and salient searches, personality spamming and meaningful social content, fake-out landing pages and the real thing, abusive advertising and legitimate applications. I hope I have also conveyed how blurry those categories can be. As we follow the movement of the word and the systems to which it is applied, spam exposes how vague and tricky the distinctions can be, and will continue to be. Spam exposes the failures and holes in models: how relevance is calculated, what a valuable collection of interlinked Web pages looks like, how people understand the pages they see and how their machines construe them, what constitutes approval, interaction, relationships, society, even humanness. In some of the cases I've described here, spam indicts the very systems it exploits; it produces optimal users, pushing

business models to their logical extreme. Following all these developments provides a different history of the Web – searchable, central, and social – through the ways each development created communities, attention, and the possibility of their own failure.

Which brings us back to the present and future of the Web, seen from spam's point of view: in which the humans involved have never been less important, mere fodder for content production and analytic stats, under the watchful eye of the platforms. The end of the line.

REFERENCES

Brin, S., and Page, L. (1998) 'The Anatomy of a Large-Scale Hypertextual Web Search Engine', *Computer Networks & ISDN Systems* 30(1–7): 107–117.

Brunton, F. (2013) *Spam: A Shadow History of the Internet*. Cambridge, MA: MIT Press.

Edwards, P. (2003) 'Infrastructure and Modernity: Force, Time, and Social Organization in the History of Sociotechnical Systems', in Thomas Misa, Philip Bray, and Andrew Feenberg (Eds.), *Modernity and Technology* (pp. 185–225). Cambridge: MIT Press.

Fetterly, D., Manasse, M., and Najork, M. (2004) 'Spam, Damn Spam, and Statistics: Using Statistical Analysis to Locate Spam Web Pages', *Proceedings of the 7th International Workshop on the Web and Databases* 67 (2004): 1–6.

Galloway, A.R., and Thacker, E. (2008) 'On Narcolepsy', in Jussi Parikka and Tony D. Sampson (Eds.), *The Spam Book: On Viruses, Porn, and Other Anomalies from the Dark Side of Digital Culture* (pp. 251–263). Cresskill, NJ: Hampton Press.

Gansing, K. (2011) 'Spamculture: The Informational Politics of Functional Trash', in Miyase Christensen, André Jansson, and Christian Christensen (Eds.), *Online Territories: Globalization, Mediated Practice and Social Space* (pp. 89–109). New York: Peter Lang.

Goldsmith, J., and Wu, T. (2008) *Who Controls the Internet?: Illusions of a Borderless World*. Oxford: Oxford University Press.

Goodhart, C. (1981) 'Problems of Monetary Management: The U.K. Experience', in Anthony S. Courakis (Ed.), *Inflation, Depression, and Economic Policy in the West* (pp. 111–146). Lanham, MD: Rowman & Littlefield.

Hayes, D. (1996) 'An Alternative Primer on Net Abuse, Free Speech, and Usenet' (http://www.jetcafe.org/dave/usenet/freedom.html).

Hollis, K. (2005) 'alt.spam FAQ or "Figuring out Fake E-Mail & Posts"' (Rev. 20050130, 30 January 2005) (http://digital.net/~gandalf/spamfaq.html).

Jeong, S. (2015) *The Internet of Garbage*. New York: Forbes Media.

Krebs, B. (2012) 'Virtual Sweatshops Defeat Bot-or-Not Tests' (Krebs on Security, 9 January 2012) (http://krebsonsecurity.com/2012/01/virtual-sweatshops-defeat-bot-or-not-tests/).

Kreibich, C., Kanich, C., Levchenko, K., Enright, B., Voelker, G., Paxson, V., and Savage, S. (2008) 'On the Spam Campaign Trail', *Proceedings of the 1st Usenix Workshop on Large-Scale Exploits and Emergent Threats* (http://cseweb.ucsd.edu/~savage/papers/LEETStormspam08.pdf).

Larson, W.L. (1994) 'Re: Green Card Lottery – Final One?' in news.admin.policy, 12 April 1994.

Lialina, O. (2009) 'A Vernacular Web 2', in Olia Lialina and Dragan Espenschied (Eds.), *Digital Folklore Reader* (pp. 58–69). Stuttgart: Merz Akademie.

Light, B. (2016) 'The Rise of Speculative Devices: Hooking Up with the Bots of Ashley Madison', *First Monday*, 21(6) (http://firstmonday.org/ojs/index.php/fm/article/view/6426/5525).

Lovink, G. (2005) 'The Principle of Notworking (Concepts in Critical Internet Culture)' (lecture, Hogeschool van Amsterdam), 24 February 2005.

Ludovico, A. (2005) 'Spam, the Economy of Desire' (Neural.it, 1 December 2005) (http://www.neural.it/art/2005/12/spam_the_economy_of_desire.phtml).

Marwick, A. (2013) *Status Update: Celebrity, Publicity, and Branding in the Social Media Age*. New Haven: Yale University Press.

Messaging Anti-Abuse Working Group (2011) 'Email Metrics Program: The Network

Operators' Perspective. Report #15 – First, Second and Third Quarter 2011' (http://www.maawg.org/sites/maawg/files/news/MAAWG_2011_Q1Q2Q3_Metrics_Report_15.pdf).

'mk590' (1996) 'AOL for Free?' in alt.2600, 28 January 1996.

Motoyama, M., Levchenko, K., Kanich, C., McCoy, D., Voelker, G., and Savage, S. (2010) 'Re:CAPTCHAs – Understanding CAPTCHA-Solving Services in an Economic Context', *Proceedings of the USENIX Security Symposium* (August 2010): 435–452.

Parikka, J., and Sampson, T.D. (2009) 'On Anomalous Objects of Digital Culture: An Introduction', in Jussi Parikka and Tony D. Sampson (Eds.), *The Spam Book: On Viruses, Porn, and Other Anomalies from the Dark Side of Digital Culture* (pp. 1–18). Cresskill, NJ: Hampton Press.

Phillips, W. (2015) *This Is Why We Can't Have Nice Things: Mapping the Relationship between Online Trolling and Mainstream Culture*. Cambridge, MA: MIT Press.

Pringle, G., Allison, L., and Dowe, D. (1998) 'What Is a Tall Poppy among Web Pages?', *Computer Networks & ISDN Systems* 30(1–7): 369–377.

Sandvig, C. (2013) 'The Internet as Infrastructure', in William Dutton (Ed.), *The Oxford Handbook of Internet Studies*. Oxford: Oxford University Press.

Schneider, J. (2004) 'Likes or Lies? How Perfectly Honest Business can be Overrun by Facebook Spammers' (TheNextWeb, 23 January 2004) (http://thenextweb.com/facebook/2014/01/23/likes-lies-perfectly-honest-businesses-can-overrun-facebook-spammers/).

Shifman, L. (2013) *Memes in Digital Culture*. Cambridge, MA: MIT Press.

Sullivan, D. (2002) 'Death of a Meta Tag' (*Search Engine Watch,* 30 September 2002) (http://searchenginewatch.com/article/2066825/Death-Of-A-Meta-Tag).

Whoriskey, P. (2008) 'Digital Deception', *Washington Post*, 1 May 2008.