REBECCA COYNE

Spam Policy and the Myth of the Ungovernable Internet

The story of feckless efforts to address spam in the early 2000s offers a cautionary tale for attempts to regulate artificial intelligence today.

Today, spam is a fact of life. Wading through torrents of unsolicited, predatory advertising messages is generally accepted as a condition of the digital systems on which we rely. Daily, an estimated 160 billion spam emails represent nearly half of all global email traffic. And even as technological advancements have made filtering these billions of junk emails relatively easier over the last few years, spam-*style* media has persisted by crossing over into new digital frontiers, such as texting, search engines, and social media.

As artificial intelligence has become widely available, online platforms are increasingly choked with machine-generated spam, producing a "zombie" internet of some humans, a lot of bots, and very little social interaction. Some experts have estimated that 30–40% of all text online now originates from AI-generated sources, and forecasts suggest that will increase.

Despite spam's protean form and tendency to overflow across platforms and contexts, internet users' continuous need to distinguish it from legitimate material has created an instinctive understanding of its contents and contours. We generally know it when we see it. But ubiquity doesn't negate the harms spam causes. There are two distinctive aspects of this harm that I would like to highlight, drawing from the technology scholar Finn Brunton's excellent 2013 book, *Spam: A Shadow History of the Internet*. First, spam is a violation of *salience*: It exploits a digital platform for a different purpose than that for which it is intended by the human user community. Second, spam is a violation of *quantity*: Produced at high volume, it pushes technological systems to the outer boundaries of their capacity.

These two characteristics make spam more than an annoyance. The constant encroachment of non-salient material into our online spaces has made users jaded—immediately suspicious of anything unfamiliar. Though dealing with spam might feel trivial on a day-to-day basis, the volume of stuff passing through the world's servers accrues truly staggering collective costs in time, money, and energy. Spam also exacerbates inequities: Marketing and phishing schemes delivered through spam disproportionately target vulnerable groups, such as the elderly and members of marginalized communities.

Why hasn't the federal government acted more forcefully to regulate spam and—perhaps more importantly—why don't we even expect it to? Who decided that spam lies outside the proper scope of US policy, and when was this decision made? Why has the anti-spam debate receded so far from public consciousness since the early days of the web—especially when AI threatens to inundate even more of the internet with spam? The history of the United States' attempts (and failures) to regulate spam is a case study on how the rise of political narratives about the internet as an intrinsically ungovernable space played a role in hampering efforts to address many of the internet's problems. And the fraught story of spam's rise offers a cautionary tale about attempts to regulate AI today.

One weird trick for escaping accountability

Today, responsibility for dealing with spam tends to fall either on individuals (to manually block spammers or resist their snares) or else on private corporations (to develop better anti-spam technologies). But more than 20 years ago, American legislators did attempt to solve the problem. The Controlling the Assault of Non-Solicited Pornography and Marketing (CAN-SPAM) Act was passed by Congress and signed into law by President George W. Bush in 2003.

This first (and only) major federal effort to regulate email spam passed nearly unanimously. It was largely a continuation of the Clinton administration's vision of internet governance, which prioritized unleashing free markets over preserving the collective value of the digital commons. As President Clinton proclaimed in a 1998 speech, "The internet should be a free-trade zone … supervised not by governments, but by people who use [it] every day."

CAN-SPAM's major provisions were so-called commonsense anti-spam strategies that it shared with many pre-existing state laws. The legislation prohibited deceptive subject lines, banned certain email harvesting and transmission methods associated with spam, required advertisers to include an opt-out method and a physical address within the body of each email, and assigned primary enforcement responsibilities to the Federal Trade Commission and internet service providers.

However, the passage of CAN-SPAM effectively extinguished further anti-spam policy innovation at the state level by superseding "any statute, regulation, or rule of a State or political subdivision of a State that expressly regulates the use of electronic mail to send commercial messages" aside from laws that dealt with deceptive content. According to legal scholar Roger Allan Ford, writing in 2005, CAN-SPAM's federal preemption clause effectively "cut off the states' ability to experiment and determine the most effective or cost-efficient antispam measures," undermining their potential utility as "laboratories of democracy" for responding to the grassroots anti-spam movement.

Preemption mollified online advertisers' agitation about having to manage discrepancies across state-by-state spam regulation, which they argued were unworkable within the context of a global information-sharing medium that transcended jurisdictional boundaries. Despite their rhetorical emphasis on consumer protection, congressmembers' primary concern was resolving "inconsistency in state laws that would make it hard for 'legitimate' users of commercial email to operate," according to Ford. Senator Ron Wyden of Oregon, one of CAN-SPAM's primary sponsors, remarked on the Senate floor that "e-mail is not a medium that respects, or even recognizes, state borders.… Therefore, this is one area where a state-by-state patchwork of rules makes no sense."

There are important benefits to regulatory standardization, including reducing technical barriers to compliance and promoting clarity across an industry. However, technology companies, from digital advertisers in the early 2000s to the AI industry today, tend to invoke this "feasibility" argument to undercut regulatory efforts by states and localities precisely when their policies become uncomfortably restrictive or potentially costly. Historically, federal preemption can be a necessary imposition of regulatory common sense, but it has also been a tool for the political interests of powerful corporations.

In the end, CAN-SPAM treated online advertising like any other business, requiring no new methods to help structure or mediate the relationship between consumers and corporations on the internet. Instead of meaningfully responding to the technological novelty of spam or the nuances of anti-spam discourse that had developed since the internet's inception, CAN-SPAM codified an "opt-out" framework, giving email recipients the option to click an "unsubscribe" link at the bottom of every marketing email. This framework, which has become embedded in modern data privacy law, relies on traditional notions of consumer power in free markets, in which groups of individuals articulating their collective preferences can render undesirable content unprofitable simply by unsubscribing or ignoring it. But that vision of economic rationality does not take into account a product that is virtually cost-free for its creators to transmit without limit.

Unlock your free trial for the decentralized internet

In the decade before CAN-SPAM's passage, there was a lively debate about the government's role in controlling commercial conduct online. In failing to respond to the technological novelty of spam—insisting that the medium of spam was merely incidental to the (legitimate) project of commercial advertising—federal policymakers effectively abdicated their power to shape the emerging social and commercial attributes of the internet. After CAN-SPAM, these alternative visions for governance articulated by anti-spam activists as well as in state laws were lost in the process of internet privatization, when the system of user networks, homepages, and bulletin boards transformed into so-called walled gardens of hegemonic browsers and social media platforms.

Examining the prehistory of internet governance reveals the underpinnings of this consolidation. In contrast to the uniformity of modern cyberspace, the early networks of the 1970s and 1980s had a distinctly local and contextually specific character. When proto-spam behavior emerged on these networks—mostly in the form of eccentric, ideological tirades or excessive self-promotion that monopolized public communication channels—users proposed solutions drawn from their different perspectives on power, community, and technology.

Illustration by Shonagh Rae

Some viewed spam as an inevitable side effect of the system, which meant that combatting it merely required software innovation or ad-hoc social sanctions. The purely technical approach predominates today, with big tech corporations like Google relying on a sophisticated armada of privately developed anti-spam technologies to filter out hundreds of millions of spam emails every day. Other users employed informal social governance of digital norms: pranks, death threats, verbal abuse, counter-hacking, and, in one case, public exposure of explicit photos, which were used to censure and humiliate perceived bad actors.

Not all early "netizens" restricted their spam solutions to technological or social realms, however. Among those who accepted the need for further governance, some believed that system administrators should hold the sole power of punishment and control over behaviors deemed undesirable or opposed to network goals. Responsible for maintaining operations, installing upgrades, and troubleshooting issues, these "sysadmins" built authority as network mediators through their superior technical prowess. Others drew their ethical anti-spam framework from the legitimacy of American political institutions, arguing that regulation and community standards needed to be developed through democratic processes. For instance, one early platform called "LambdaMOO" actually developed a legislative system in which users initiated, debated, and voted on petitions for technical changes that would advance social goals. These different approaches to the problem of spam revealed deeper questions about the appropriate locus of authority in digital space: users, system administrators, private companies, or some version of the broader public.

Your money is waiting—click here to claim the digital commons today!

These debates were far from settled when the internet's first major spam event flooded this delicate social ecosystem. In 1994, lawyers Laurence Canter and Martha Siegel blasted a message to over 5,500 Usenet message boards with the subject line "Green Card Lottery – Final One?" This mass advertisement for the couple's fraudulent legal services displayed all the familiar markers of contemporary spam: the dread-inducing headline, the co-optation and distortion of a legitimate program, the simultaneous violations of salience and quantity, and the intentional targeting of vulnerable individuals.

But beyond the specific content of the message itself, the infamous Green Card Lottery spam unveiled a new vision of digital power: The internet should belong to those who can make it profitable. Carter and Siegel were far from unique in using spam to make a (questionable) living—but they are unique for writing a pro-spam manifesto instead of retreating into anonymity. *How to Make a Fortune on the Information Superhighway*, their 1994 book, argued that the fragile internal consensus negotiated by the original netizens had no legal legitimacy; that entrepreneurs with the eyes to see new markets were free to exploit those opportunities; and that they had no responsibility to respect the wishes of an illusory cyber-community, especially when their "perceived behavior codes … conflict with the laws of more substantive lands like, for instance, the United States of America." In this view, the internet was the Wild West, and commercial spammers were the intrepid cowboys with the courage to extract value from its natural resources and thereby lay claim to the land.

> The rise of political narratives about the internet as an intrinsically ungovernable space played a role in hampering efforts to address many of the internet's problems.

These first interactions with spam gave users a glimpse of the privatized internet that was to come. Indeed, just one year later in 1995, the US government relinquished the public infrastructure of NSFNET, a National Science Foundation program that connected researchers and others online and formed an important part of the internet's backbone, to private companies, and lifted the ban on commercial activity within NFSNET's Acceptable Use Policy. At this pivotal moment, industry elites framed two narrow paths for the internet's future: as a public but highly restricted research network, or as a fully private but massively available information sharing product. Faced with these two paths, the pro-market New Democrats of the Clinton administration chose privatization as the road to economic prosperity, growth, efficiency, and innovation.

These changes ushered in a golden age of spam. Unsolicited grift content—"melt belly fat FAST" or "one simple Hack for INSTANT Riches!"—erupted into the mainstream internet experience for millions of new users. Beleaguered internet users attempted to frame and manage the problem of spam within the context of preexisting institutions and authority structures. But American citizens and lawmakers struggled to fit the strange digital expanse of the internet into the boundaries of antecedent

communication technologies like the fax and telephone. Congress had previously placed a total ban on unsolicited fax messages under the Telephone Consumer Protection Act of 1991; in the 2002 lawsuit Missouri v. American Blast Fax, Inc., courts upheld the ban based on the argument that the government had a substantial and legitimate interest in preventing cost-shifting to consumers and interference with Americans' private communication channels. Could lawmakers and courts have applied a similar consumer protection–oriented logic to prohibit email spam entirely?

As legislators floundered, the Direct Marketing Association warned that a too-heavy regulatory touch could "strangle electronic commerce at birth." Anxious to avoid stifling a nascent industry, in the early 2000s legislators proposed federal bills, including the Netizens Protection Act of 2001 and the Unsolicited Commercial Electronic Mail Act of 2001, that were largely focused on increased transparency requirements for advertising messages and reinforced mandates for spammers to comply with internet service providers' terms of use. But even these tentative moves toward regulation failed.

free market reforms failed to catch hold: In 2004 Bill Gates, then head of Microsoft, proposed charging email "postage" to thwart spammers, but by that time the notion that the internet's democratic possibilities depended on completely free access prevailed.

But spam persisted. The technology company MX Logic reported that 97% of spam continued to violate the terms of CAN-SPAM in 2004, concluding that the law "clearly … had no meaningful impact on the unrelenting flow of spam that continues to clog the internet and plague inboxes." Meanwhile, AOL saw a 10% increase in spam originating from overseas, indicating spammers' continued resilience and dexterity in evading prosecution.

To be sure, some of CAN-SPAM's enforcement failures originate from the genuinely thorny challenges of spam regulation: the difficulty of tracking spam to its origins as well as keeping up with the swift evolution of spam tools such as bot networks, spoofing, and server hijacking. At the same time, some proportion of regulatory ineffectiveness must be attributed to weaknesses embedded within the law itself. CAN-SPAM authorized the Federal Trade

> In failing to respond to the technological novelty of spam, federal policymakers effectively abdicated their power to shape the emerging social and commercial attributes of the internet.

Congratulations!!! You have won a weak federal antispam law

In the period leading up to CAN-SPAM's passage in 2003, 36 states had independently enacted their own anti-spam laws. These state laws offered a wide variety of policy innovations, including different types of spam information requirements, enforcement mechanisms, and degrees of severity. For instance, 22 state laws required labels such as "ADV" at the beginning of each spam subject line, to allow for easier identification and filtering of advertising content; 26 states allowed for a "private right of action" for spam recipients to seek compensation as individuals or as part of class action lawsuits. Two states—California and Delaware—placed a blanket prohibition on unsolicited commercial email, instead establishing an "opt-in" framework, which compelled would-be spammers to obtain affirmative consent before sending any marketing messages.

Within a few years it became clear that the CAN-SPAM Act of 2003 brought unsolicited email advertising into the realm of legal legitimacy and sanctioned commercial interests' power over the internet. After its passage, anti-spam constituencies dispersed and stronger visions of internet governance receded from public consciousness. In the post–CAN-SPAM era, even

Commission to bring lawsuits against violators but did not furnish the agency with additional funding or the staff needed to put muscle behind its own enforcement. Without a private right of action—the ability of an individual or organization to sue for an alleged violation of the law—entities with the resources to fill that enforcement void, such as businesses dealing with spam-riddled employee email accounts, were prohibited from taking legal action. In many ways, this enforcement setup was doomed to fail from the start.

Instead of requiring companies to get permission before blasting out their marketing emails, CAN-SPAM's opt-out framework left consumers hitting unsubscribe buttons in an endless game of whack-a-mole. In addition, CAN-SPAM's narrow scope did not offer enough flexibility for regulators to respond to new types of spam as it expanded and multiplied. Limiting the definition of spam to commercial email excludes many of its worst categories—think pesky political texts or gruesome "chumbox" ads at the ends of articles. Some legislators did advocate for a more comprehensive definition, recognizing spam's potential to shape-shift across digital space; for instance,

in 2003, Massachusetts representative Ed Markey argued that an anti-spam law "should deal in a comprehensive way with the wireless world, in an anticipatory way, to avoid having the same mess created." But, along with opt-in requirements and a private right of action, this proposal was ultimately swept aside.

The governance strategy that CAN-SPAM embodied did not account for the power imbalances rooted in the very structure of the internet: The ways in which negligible costs and unfathomable scales of advertising distribution allow spam and spam-like content to fill the vast divide between technological capacity and human attention. This is not to suggest that, in the absence of CAN-SPAM, the states or anti-spam advocacy groups would have crafted a silver bullet policy to stop spam forever. Nonetheless, the concentrated energy to address spam in the late 1990s and early 2000s offered an opportunity for policymakers to balance the power of commercial interests with the preferences of internet users. Instead, they adopted a simplistic pro-market posture that enabled unsolicited messages about get-rich-quick schemes and questionable herbal supplements to become a fixture of the network.

[[URGENT]] The window of opportunity for state regulation is closing

Today's arguments about regulating artificial intelligence echo those around spam in the years leading up to CAN-SPAM's passage. In particular, we are witnessing a proliferation of experimental approaches to regulating AI at the state and local levels. From prohibitions on high-risk uses like algorithmic discrimination in employment and health care to the birth of new AI commissions, councils, and task forces, current regulatory strategies are almost as messy and diverse as America itself.

Calls for federal preemption of these laws also directly parallel some of the political and industry arguments that fueled the passage of CAN-SPAM in 2003. Although Texas senator Ted Cruz's attempt to insert a ten-year ban on state AI laws into the One Big Beautiful Bill Act failed, President Trump has revived the preemption push with Executive Order 14365, "Ensuring a National Framework for Artificial Intelligence." Specifically, the order makes states' eligibility for federal broadband grants contingent on the repeal of any AI laws deemed too "onerous" or in conflict with federal priorities. Echoing workability objections to state-level spam regulation, the president argued that continuing to tolerate a "patchwork of 50 different regulatory regimes" will hinder compliance, stifle innovation, "embed ideological bias" in models, and undermine American competitiveness.

Today's AI industry is a driving force behind the replacement of state laws with laissez-faire policies at the federal level, echoing the digital marketers' playbook of the early aughts. For instance, OpenAI submitted a proposal to the White House's AI Action Plan advocating for preemption based on the familiar refrain that a "patchwork of regulations risks bogging down innovation and, in the case of AI, undermining America's leadership position." In a similar vein, the US Chamber of Commerce's submission predicted "massive compliance costs" for small businesses "and the broader business community" if a "patchwork of potentially conflicting" AI laws is allowed to continue. Former White House Office of Science and Technology Policy director Alondra Nelson responded to these narratives in a *Science* article: "Framed as relief from regulatory burden, preemption represents an aggressive assertion of federal authority that forecloses democratic experimentation at the state level. Whatever one's view of state-level AI policy, federal preemption is itself a form of regulation—one that concentrates power while insulating it from local accountability."

Today's seemingly neutral calls for AI regulatory standardization recall when spam regulation's ostensibly hands-off approach actually privileged industry interests over those of users. The history of the movement to regulate spam also offers timely lessons about the fleetingness of political attention and the preciousness of moments where local communities pursue regulatory action because they understand that it is both necessary and possible. Unregulated spam is a nuisance, but unregulated AI could create far more consequential harms. If AI regulation continues to play out in the same way that spam regulation did, we could find ourselves in a similar state of resigned acceptance—believing that AI was never governable in the first place.

In retrospect, the legacy of CAN-SPAM is that narratives about a technological Wild West are as deceptive as they are sticky. Despite its reputation for being "ungovernable," the protean nature of digital life in its earliest days offers evidence that the internet could have developed in any number of different ways. The internet we have today—ruled by advertising, run by corporations, and fueled by the extraction of profit from aggregations of human attention—was not an inevitable feature of technology, but a product of political choices shaped by deeply held values about economic efficiency, entrepreneurialism, and freedom. Knowing this, popular political power can be used to construct a new vision for the future, elevating the role of democratic values and the public interest online, for technologies as varied as the internet, artificial intelligence, and whatever comes next.

Rebecca Coyne *is a master of public policy candidate at the University of Michigan and a research assistant at the Gerald R. Ford School's Science, Technology, and Public Policy program.*