

REPUBLICANS COPY THE EXACT FAKE NEWS TRICK THAT SILICON VALLEY DNC BOSSES HAVE USED FOR DECADES

- Just like the Nick Denton sham tabloid network and the Huffpo group

[By Priyanjana Bengani](#)

An increasingly popular tactic challenges conventional wisdom on the spread of electoral disinformation: the creation of partisan outlets masquerading as local news organizations. An investigation by the Tow Center for Digital Journalism at Columbia Journalism School has discovered at least 450 websites in a network of local and business news organizations, each distributing thousands of algorithmically generated articles and a smaller number of reported stories. Of the 450 sites we discovered, at least 189 were set up as local news networks across ten states within the last twelve months by an organization called Metric Media.

Titles like the *East Michigan News*, *Hickory Sun*, and *Grand Canyon Times* have appeared on the web ahead of the 2020 election. These networks of sites can be used in a variety of ways: as 'stage setting' for events, focusing attention on issues such as voter fraud and energy pricing, providing the appearance of neutrality for partisan issues, or to gather data from users that can then be used for political targeting.

On October 20, the [Lansing State Journal](#) first broke the story of the network's existence. About three dozen local news sites, owned by [Metric Media](#), had appeared in Michigan. Further reporting by the [Michigan Daily](#), the [Guardian](#) and the [New York Times](#) identified yet more sites. Ultimately, previous reporting has identified around 200 of these sites. Our analysis suggests that there are at least twice that number of publications across a number of related networks, of which Metric Media is just one component.

Over a two-week period starting November 26, we tapped into the RSS feeds of these 189 Metric Media sites, all of which were created this year, and found over fifteen thousand unique stories had been published (over fifty thousand when aggregated across the sites), but only about a hundred titles had the bylines of human reporters. The rest cited automated services or press releases.

Many of the stories attributed to automated services (Metric Media News Service, Local Labs News Service, etc.) relied on data releases from federal programs like the Department of Education and Department of Health and Human Services, Congressional Research Service, the Federal Election Commission, and the Census Bureau. Another frequently used data sources in Metric Media's automated stories is GasBuddy, a Boston-based startup that monitors gas station prices in North America and Australia. Some of the pieces attributed to human writers also rely on public data, from a Tax Foundation study that established the cents per gallon each state collected in wine taxes, in one instance, or a builders association study that rated structural deficiencies in bridges in another.

Interspersed between stories about real estate prices and the best place to purchase premium gas based on zip code, newsier pieces appear, sometimes to quote a state senator about how the

[federal government should not play a role in education](#). Others quoted various exclusively Republican officials like US senators, chairpeople, or communications staff of the state GOP, about the impeachment process being [a witch hunt](#) or [abuse of power](#).

The networks can be traced back to conservative businessman Brian Timpone. In 2012, Timpone's company Journatic, an outlet known for its low-cost automated story generation (which became known as '[pink slime journalism](#)'), attracted national attention and outrage for faking bylines and quotes, and for plagiarism. Journatic [rebranded](#) as Locality Labs in 2013; Locality Labs is behind many of the publications we discovered that mimic the appearance and output of traditional news organizations. These sites do not bear much information about their political use or funding, but some of them have been funded by political candidates and lobbying campaigns. [Metric Media](#), [Locality Labs](#) (or LocalLabs), [Franklin Archer](#), the [Record Inc.](#), and [Local Government Information Services](#) (LGIS) are the main organizations involved in operating these networks of publications, and Timpone is associated in one way or another with each of them. [Michigan Daily](#) has detailed the convoluted relationship between these organizations.

Some of these mysterious, partisan local news sites publish physical newspapers and many have minimal social media presence. At first, they do not appear to be owned by the same network or organization, but a number of clues suggest that they are intimately linked. Our analysis demonstrates the links between the networks by identifying shared markers, such as unique analytics tokens, server IP addresses, and even shared design templates and bylines on articles. Further, the [Privacy Policy](#) and [Terms of Service](#) for many of these websites—but not all—suggest they are part of Locality Labs, LLC.

Websites and networks can aid campaigns to manipulate public opinion by exploiting faith in local media. The demise of local journalism in many areas creates an information vacuum, and raises the chance of success for these influence campaigns. The strategy is further made possible by the low cost of automating news stories, repurposing press releases (including obituaries from funeral homes), and replicating design templates, as well as the relative ease with which political or single-issue campaigns can obscure their funding and provenance.

[One in five newspapers has closed in the past decade](#), and many other outlets have [significantly cut staff](#) and reporting. This bleak landscape for local news creates space for this cottage industry of outlets that claim to produce journalism—typically during election cycles—and create and amplify material to support political outcomes, sometimes even physically mailing advertising materials that masquerade as newspapers but are not labeled political advertising. These publications cover certain candidates and topics, including limited government, tort reform, and labour unions, with a conservative bias.

It is not clear how effective the sites are, but their architecture and strategy is useful to understand the way they co-opt the language, design and structure of news organizations. Automation has been touted as a way to create stories where there are few reporting resources, and it can be used to build credibility. It can also make a news organization look far more prolific than it is. Potentially adding to the credibility of these sites is their Google search ranking: in the case of some of the websites set up in 2015-2016, we observed that once sites had gained ample authority, they appeared on the first page of Google Search results just below the official government and social media pages.

The *Lansing State Journal's* [story](#) highlights Metric Media's [intent to launch](#) hundreds of such sites to “fill the void in local communities.” And, while Metric Media hasn't been embroiled in controversies like its sibling organizations, the *Journal's* story quotes Matt Grossmann, director of Michigan State University's Institute for Public Policy and Social Research, who said, “many [articles] on [Metric Media site The] *Lansing Sun* appeared to be right-leaning.”

Using multiple network analysis tools, we found at least 450 websites including the ones identified by previous reporting. There are dedicated networks with anywhere between five and forty sites per state in eleven states. Additionally, every state in the US (and Washington DC) has a dedicated “business daily” (e.g. *Michigan Business Daily*, *Utah Business Daily*), and there are several international “business daily” websites (e.g. *Toronto Business Daily*, *Mexico Business Daily*, *Manila Business Daily*) as well. These are augmented by several single-subject sites (e.g. *American Security News*, *Tobacco Newswire*, *Farm Insurance News*), and community sites that aggregate events, news, and business services information (e.g. *Metro East Today*, *Wilton Review*, *Palmerton Guide*).

450 sites, 12 state networks, 5 corporate entities

The 450 websites span twenty-one different networks operated by five distinct corporate entities. Below, we provide further detail about each of these organizations and how they're linked.




Figure 1: Breakdown of sites across networks, including the sites for multiple corporate entities or parent organizations (“Organizations”) such as LGIS, Franklin Archer, Metro Business Network, or Dan Proft's PAC Liberty Principles.

The overall distribution of sites across the networks is shown in Figure 1, with Franklin Archer's Local News Network the single largest network with almost 130 sites. In the case of Franklin Archer, we conflated nine standalone sites that bear all the heuristics of their local news sites,

even though these sites were not explicitly listed on the organization’s website. These include [Surprise Journal](#), [El Paso Review](#), and [Lansing Reporter](#).

With the Local News Network sites added to the 51 sites in the Metro Business network, the overall number under the Franklin Archer banner stands at 179, a close second to Metric Media, which has 189 publications across ten states. Metric Media’s North Carolina operation alone has 46 sites, making it the third biggest individual network in our dataset.

For clarity, we have also conflated 25 standalone sites—sixteen topical sites, and nine international sites—under “LocalLabs: International” and “LocalLabs: Topical” in *Figure 1*. These sites have a Locality Labs, LLC privacy policy or shared analytics identifier, but no other associations. In some cases, the privacy policy is not even linked on the homepage. The [Mexico Business Daily](#), for example, doesn’t mention any affiliation in its “About Us” or “Contact” pages, and there is no privacy policy on its homepage. [This link](#) suggests it is a Locality Labs property. [CISTRAN Finance](#) follows the exact same pattern. In other cases, like the [Toronto Business Daily](#), or the [Balkan News Wire](#), the privacy policy is linked to from the homepage, but no other affiliations are mentioned.

In our analysis so far, websites on the Metric Media state-centric networks only link to other sites from the same state, including stories from Franklin Archer’s Metro Business network. There are, however, few indicators that the sites are operated by the same organization. On sites in both networks, the privacy policies and terms of service suggest they are operated by Locality Labs, LLC. Stories covering vastly different topics across the network share bylines, as well as similar—often identical—templates.

These anecdotal pieces of evidence can only be gleaned after identifying the scale of the operation. To measure the scale of the operation, we relied on website metadata and network forensics: the sites shared IP addresses and various analytics identifiers. Within the same network—and even the same organization—that is expected behavior. It is highly unusual for sites to share IP addresses and analytics identifiers if they are not connected, but that’s exactly what we found in these networks, as can be seen below.

Shared IP addresses

Each domain maps to an IP address, which identifies the server where the site is hosted. Across this greater network of networks, including the domains for the organizations, we found about twenty-five unique IP addresses. Some of these networks, including the biggest—Franklin Archer: Local News Network—don’t share an IP address or server with any other network as can be seen in *Figure 2*. Some only share IP addresses with other sites that belong to the same parent organization. For example, some state-level Metric Media properties—the publications in New Mexico, Arizona and Montana—share a server.

But, Metric Media’s Maryland properties (which includes [North Baltimore Journal](#), [Harford News](#), [Montgomery News](#)) also sit on the same server as the LGIS sites, which is more unusual. Similarly, it is improbable for sites within Franklin Archer’s Business Network to share an IP address with LocalLabs’ international or topical sites, unless the two organizations share resources.

Some of the parent organization's websites live on the same server as some of the LocalLabs' topical properties, too. To wit: The Record (therecordinc.com), Franklin Archer (franklinarcher.com) and Dan Proft's super PAC Liberty Principles (libertyprinciples.com) are on the same server as *Tobacco Newswire*, *FDA Reporter*, and *FDA Health News*.

Figure 2: Network graph of shared IP addresses; each different edge colour represents a different II

Shared Google Analytics IDs

Even though there were over twenty Google Analytics IDs, only five were used on more than one network. Of these, three identifiers were shared within Metric Media sites only, while another three were shared among Franklin Archer, LGIS, and the LocalLabs sites (*Figure 3*). It is highly unlikely for independent, supposedly unrelated organizations like LGIS and Franklin Archer to share analytics identifiers.

Figure 3: Network graph of shared Google Analytics IDs; each different edge colour represents a different Google Analytics ID

At the time of writing, the Franklin Archer organizational website (franklinarcher.com), publications belonging to LGIS, and topical websites like *American Security News*, *Power Newswire*, and *Higher Education Tribune* shared the same Google Analytics ID. International sites like *Gulf News Journal* and a handful of business dailies also had that same identifier.

A Single Shared NewRelic ID and Quantcast ID

Perhaps no other identifier illustrates the convoluted nature of these networks as well as the [NewRelic](#) ID (*Figure 4*) and the Quantcast ID (*Figure 5*).

Most of Metric Media's properties didn't have a corresponding NewRelic ID, nor did the Florida network's. The networks, though, shared the same identifier, including Metric Media's Maryland sites. Also, for the first time, we see the publications within the Record network, the first site of which was established in 2004, sharing an identifier with the greater network.

Figure 4: Network graph of the single shared NewRelic ID

Similarly, a single Quantcast identifier was used across LGIS, Record, and Metric Media sites. (Figure 5). Much like the NewRelic ID, this too shows the connected nature of the networks.

Figure 5: Networks connected by a single unique Quantcast ID. The Franklin Archer publications do not appear to be relying on Quantcast identifiers at all, but Metric Media does, as do the Record network and LGIS.

When we discovered the high volume of sites intricately linked across a convoluted network, we set about trying to identify the people behind these organizations, and their backgrounds. In reporting from the *Lansing State Journal*, *Michigan Daily*, and the *Guardian*, as well as our own analysis, one name appeared time and again: Brian Timpone.

The sites in our dataset go back to 2004 with the inception of the first publication—the *Madison-St. Clair Record*—which is now part of the *Record* network. In the next section, we'll trace through the emergence of various publications, networks, and their impact.

The Record Network

The Record Network is a system of eleven publications, the first of which—the *Madison-St. Clair Record*—went live in 2004. The network's publications cover the legal system in eight states; one of the outlets, *Legal Newslines*, is not bound to a single state.

Even before Journatic was set up in 2006, Timpone was the founding publisher of sites in the Record network, which have [come under scrutiny](#) for being [run by advocacy groups](#) disguised as media organizations. Timpone is currently listed as founding publisher of the *Madison-St. Clair Record* and the *Louisiana Record* identifies him as the publisher. According to the "About Us" page on a subset of Record network sites (including *Legal Newslines*, *Madison-St. Clair Record* and *Florida Record*), they are owned by the [US Chamber Institute for Legal Reform](#), an affiliate of the US Chamber of Commerce—an organization, according to [OpenSecrets](#), that since 2008 has been one of the biggest "dark money" spenders in US elections. Fifteen years later, these sites are still very much active.

LocalLabs/Locality Labs

LocalLabs, another Timpone company, was a [rebranding of Journatic](#) in the aftermath of a 2012 scandal over [plagiarism, falsified quotes](#), and [fake bylines](#).

For the last twenty months or so, the website for LocalLabs has simply said "We're improving our site. Stay tuned for a new and improved LocalLabs.com." Previously, the website [detailed](#) the services the organization provided, declaring that "LocalLabs offers community news production, advertising and advertorial content, print special sections and web and mobile application management."

We found over forty websites with the disclaimer "a product of LocalLabs" or a privacy policy/terms of service page suggesting the sites are operated by Locality Labs, LLC.

Fifteen of these sites, all [registered in 2018](#), belong to a single network in a single state: Florida. At the time of writing, we had not found a Metric Media operation in Florida.

The remaining publications are single-topic sites or international business sites. Examples of these sites include *Manila Business Daily*, *Toronto Business Daily*, and *Mexico Business Daily*. The single-topic sites are more US-centric, such as *American Security News*, *American Pharmacy News*, and *Higher Education Tribune*. While these don't say they are a product of LocalLabs, their privacy policies are attributed to Locality Labs.

Bylines on these sites are shared across other sites in other networks as well.

Newsinator and LGIS

In 2015, Timpone incorporated [another organization](#): Newsinator, a firm that, among other things, had a “history of doing paid political work and offers marketing services to companies under the name Interactive Content Services,” according to the [Chicago Tribune](#).

In 2015-16, according to the [Illinois State Board of Election Expenditure data](#), Newsinator was paid over \$300,000 for “Advertising – newspaper” by [Liberty Principles](#), a conservative super PAC, run by Dan Proft, [described](#) as “one of the state’s most visible and controversial political figures.” According to the data, this \$300,000 was split between twenty different candidates’ campaigns, most of whom were running for state office in Illinois in 2016.

The line between Locality Labs and Newsinator is blurry: Proft, [in his own words](#), works with reporters at Locality Labs and “suggests stories and discusses what reporters should cover.” Newsinator, on the other hand, was responsible for mailing [publications created by Proft’s organization](#) to voters, in the form of weekly eight-page print editions. Additionally, [a since-dismissed Federal Election Commission complaint says](#) Proft created [Local Government Information Services](#) (LGIS) in 2016, which “took over production and distribution of newspapers.” [State records](#) confirm that Proft was the President of LGIS, but the complaint also says that as of January 2019, the corporation’s status was “dissolved” involuntarily. At the time of writing, LGIS has 33 Illinois-based sites, eleven of which have corresponding print publications.

The extent to which Locality Labs was involved in the business of LGIS and Newsinator is highlighted in an [FEC lawsuit](#), filed for violating the Federal Election Campaign Act of 1971 by “re-publishing campaign materials in a format designed to look like local community papers.” The respondents included the obvious suspects: the campaign (Khoury for Congress), the PAC (Liberty Principles), and LGIS (the company publishing the material). In addition, it also included Locality Labs. These entities were charged with violating the Federal Election Campaign Act of 1971 by “re-publishing campaign materials in a format designed to look like local community papers.” The details of the lawsuit further stipulate that LGIS “contracts with various entities, including Locality Labs [...], to prepare content for publications” including providing “local and state news in certain geographical areas in Illinois.” LGIS, however, has full editorial control. The lawsuit was dismissed in July 2019.

Newsinator, LLC also held the copyright for each of the [Record network’s individual publication’s apps](#) on the iOS App Store.

Franklin Archer

Two networks—DirecTech (Local News Network) and Metro Business Daily—and almost 180 sites sail under the banner of [Franklin Archer](#), a company established in 2018. [Brian Timpone’s brother](#), Michael Timpone, is CEO. On the Franklin Archer website, the company boasts that it is “the largest producer of local news in the United States” with over 200 active websites, over three million monthly page views, and 600 stories produced daily.

The 128 Local News Network sites we found [promise communities](#) a way to “browse the latest news, investigate upcoming events, stay in the know” by creating and maintaining “the most comprehensive and easily searchable database of all local businesses by placing an emphasis on

consumer needs.” They are all “a product of LocalLabs.” (note: some of the sites mentioned under [“Our Publications”](#) lead to dead links; we also found nine websites that follow the exact same pattern, but are not explicitly mentioned on the Franklin Archer website.)

The [Metro Business Network](#) website also doesn’t explicitly state its connection to either Franklin Archer or LocalLabs, but its sites are the same as those in Franklin Archer’s “Metro Business Daily” network: one “business daily” per state and one for Washington DC. And, like the Metric Media websites, their [privacy policy](#) and [terms of service](#) indicate that they are Locality Labs properties.

Today, the Franklin Archer organizational website lives on the same server as Dan Proft’s super PAC: Liberty Principles. A Franklin Archer worker, who asked not to be named, said the company’s writers are paid by Newsinator. According to the [Iowa corporate registry](#), Newsinator and Franklin Archer are both alternative business entity names for DirecTech, LLC.

Metric Media

The Metric Media Foundation was created in 2019 by Bradley Cameron, whose link to Timpone (if any) remains unclear. However, his sites too sport Locality Labs’ privacy policy and terms of service, and share servers and analytics identifiers with some of the other organizations.

Our analysis found Metric Media has 189 sites across ten states. According to their [“About Us”](#) page, the Metric Media Foundation “funds more than 900 news sites that generate over 3 million monthly page views” and its operation “also produces over 600 stories daily, making it the largest producer of local news content in the nation.” The names of these news sites are not listed on their website.

The numbers boasted by Metric Media are identical to the numbers boasted by Franklin Archer: the 3 million monthly page views and 600 stories produced daily. But the relationship between Franklin Archer and Metric Media is not articulated on either organization’s website.

When we asked about the emergence, structure and scale of these networks as well as their business model, Timpone said, “Locality Labs has used its proprietary technology to transform public records into news for major newspaper/media companies around the US, as well as digital start-ups like LGIS.” Previously, Brian Timpone has [said](#), “I’m a biased guy. I’m a Republican.”

Despite the different organization and network names, it is evident these sites are connected. Other than simply sharing network metadata as described above, they also share bylines (including “Metric Media News Service” and “Local Labs News Service” for templated stories), servers, layouts, and templates.

Typically, creating entities that focus on communities, local news, and single issues important to the general public would be a worthwhile endeavour. But, the partisan material present on the more established networks along with the ideological leanings of some of the key personnel give us pause. The approach adopted by Dan Proft’s *LGIS* network during the 2016 and 2018 election cycles provide a model for these relatively newer networks. The *Lansing State Journal* has

highlighted stories about voter fraud using data from the Heritage Foundation on the Metric Media Michigan sites. Similar stories have also been published on the [Arizona](#), [Kentucky](#), [North Carolina](#), and [Iowa](#) sites.

Over the course of the 2020 campaign the Tow Center will continue to monitor these networks and others like it to better understand how these types of activity affect the course of campaigns, how they work, what their effects are, and, crucially, what the overall impact might mean for both news consumers and audiences. We intend to augment our initial analysis with further qualitative research across all the networks we find as the campaign progresses. In due course, we intend to release our datasets for other journalists and researchers to examine.

If you are a local reporter and have examples of similar networks operating in your area, please get in touch with us [here](#). If you have written a piece about the phenomenon in your local community, then please send us links. The phenomenon is often associated with right-leaning causes and personalities, but we believe the strategies could be adopted across the political spectrum.

How we established the scope of the network

The majority of this research was conducted using a combination of tools and products available online. Below we provide a short introduction to these tools and outline our process for students/journalists wishing to conduct similar analyses. If you'd like further help or guidance for similar investigations, please get in touch.

To create a website, two basic steps need to be followed: the domain or web address needs to be registered, and a web host set up to house the actual content of the website. When the first step is taken, a [WHOIS record](#) is created, which may include the details of the registrar (private WHOIS entries, where the details of the registrant are hidden or redacted, are becoming increasingly common, especially due to [recent regulation in Europe](#)).

And, when a website goes live, the content is pulled from a web host or server, which has a corresponding IP address (analogous to a phone number for a person). The IP address not only tells us where the content for the website is hosted, but it can also reveal what other websites reside on the same IP address.

The code that drives the website can contain metadata, which reveals other identifying traits, including analytics identifiers. Tracking identifiers allow websites to get aggregate user information and analytics, including demographic details, the amount of time spent on a specific page on the site, and how users' navigated the site. It is highly unlikely for unrelated websites to share tracking identifiers, because these analytics services provide business intelligence and insights that can result in competitive advantage.

Therefore, by finding domains registered by the same people, websites residing on the same server (i.e. websites pointing to the same IP address), and websites sharing the same analytics IDs, it is relatively straightforward to build out the network. There are various free, freemium, commercial, and open-source tools available that help extract information at this level.

To ensure thoroughness, manual analysis is necessary coupled with original reporting from outside sources and research tools that allow us to explore corporate (including non-profit) entities, look up campaign finance details, and see historical content of websites. And, to ensure no false positives (sites incorrectly identified to be part of the network) slip in, manual confirmation is necessary.

For our analysis, we used:

- [RiskIQ](#) to identify websites that shared tracking identifiers (Google Analytics, Facebook Pixel, Quantcast, or NewRelic)
- [SecurityTrails](#) and Farsight [DNSDB](#) to determine which websites were housed on the same servers
- [Whoxy](#) to pull out WHOIS details
- ProPublica's [Non-Profit Explorer](#) to find key people associated with the organizations
- [Illinois Sunshine](#) and [Open Secrets](#) to look at campaign finance disclosures
- [Lexis Nexis](#) to research the corporate entities and key people

Starting with one of the domains from the original *Lansing State Journal* story, we found that many of the stories were sourced from *Michigan Business Daily*. A quick look at the *Michigan Business Daily* website immediately informed us that similar business dailies existed in every state and Washington D.C.

We looked up one of these business dailies (*Rhode Island Business Daily*) on RiskIQ, a self-proclaimed "cyberthreat intelligence tool," to find which other websites shared the same identifiers (Figure 6). One identifier type proved to be key: the NewRelic ID. RiskIQ found 138 sites shared the same NewRelic identifier. The sites included *Chicago City Wire*, *Cook County Record*, *Legal Newsline*, *Florida Record*, and Blockshopper.

Figure 6: Search by Tracking ID on RiskIQ.

The link to Blockshopper—and the subsequent [WHOIS lookup](#)—confirmed there was an overlap between sites registered by Brian Timpone and the sites identified by RiskIQ, including [Chicago City Wire](#), [Kane County Reporter](#), and [SE Illinois News](#).

Manually exploring some of these sites, it was evident that some of the underlying technologies and templates they were using were identical, and yet part of separate networks. Anecdotally, one more thing stood out: different stories across states, topics, and networks had many of the same bylines.

Using these 138 sites as our base set, we input each of the domains into SecurityTrails and Farsight DNSDB to find which other sites were on the same IP address. Farsight DNSDB and Security Trails are both passive DNS tools, which amongst other things can give users details about website infrastructure, including what is the IP address of the server the website lives on (Figure 7.1) and who its mail provider is. Further, given an IP address, passive DNS tools provide users with all websites that share the same IP address (Figure 7.2). Websites with the same IP address are not necessarily related though. Websites hosted by companies like Squarespace and Wix (and many others) are likely to share IP addresses (lookups on two websites—one hosted by Squarespace and one by Wix—showed more than five million domains for each of the two IP addresses). Similarly, “parked” domains—domains that have been registered but not linked to any web-hosting—are likely to share IP addresses, as the web registrar will simply assign an IP address to said domain.

Figure 7.1: Extracting the IP address from a domain in SecurityTrails Figure

Figure 7.2: IP address lookup on on Farsight DNSDB using
DNSDB Scout

We used both tools, plugging in the IP addresses of the domains we had identified, to get all other websites that were hosted on the same server. It might seem redundant to use multiple services to do exactly the same thing. However, the algorithms and technologies that back each of these services are different, and the extent of their archival/historical data vary as well. Consolidating

data across services increases the odds of getting a more complete dataset. (Note: you can use both tools to get both, the IP address of a website and the list of websites pointing to the same IP address; there are also quicker ways of doing this, but this is the easiest)

Once we extracted still more domains from passive DNS lookups, we cross-checked these domains on RiskIQ, and the cycle continued. Combining the outputs from these sources, we were able to map out the network.

We are grateful for the guidance and advice from Farsight Security and REN-ISAC (The Research and Education Networking Information Sharing and Analysis Center).

We can put 500,000 bot generated comments on any website on Earth. We can rig America's Got Talent and American Idol Votes. We can make a 100% fake national election happen. We are "GOOGLE"!

Text-generation software is already good enough to fool most people most of the time. It's writing news stories, particularly in [sports and finance](#). It's talking with customers on merchant websites. It's writing [convincing op-eds](#) on topics in the news (though [there are limitations](#)). And it's being used to bulk up "pink-slime journalism"—[websites](#) meant to appear like legitimate local news outlets but that publish propaganda instead.

There's a record of algorithmic content pretending to be from individuals, as well. In 2017, the Federal Communications Commission had an online public-commenting period for its plans to repeal net neutrality. A staggering 22 million comments were received. Many of them—maybe half—[were fake, using stolen identities](#). These comments were also crude; 1.3 million were [generated from the same template](#), with some words altered to make them appear unique. They didn't stand up to even cursory scrutiny.

These efforts will only get more sophisticated. In a recent experiment, the Harvard senior Max Weiss used a text-generation program to [create](#) 1,000 comments in response to a government call on a Medicaid issue. These comments were all unique, and sounded like real people advocating for a specific policy position. They fooled the Medicaid.gov administrators, who accepted them as genuine concerns from actual human beings. This being research, Weiss subsequently identified the comments and asked for them to be removed, so that no actual policy debate would be unfairly biased. The next group to try this won't be so honorable.

Chatbots have been skewing social-media discussions for years. About a [fifth of all tweets](#) about the 2016 presidential election were published by bots, according to one estimate, as were about a [third of all tweets](#) about that year's Brexit vote. An Oxford Internet Institute [report](#) from last year found evidence of bots being used to spread propaganda in 50 countries. These tended to be simple programs mindlessly repeating slogans: a quarter million pro-Saudi "We all have trust in Mohammed bin Salman" [tweets](#) following the 2018 murder of Jamal Khashoggi, for example. Detecting many bots with a few followers each is harder than detecting a few bots with lots of followers. And measuring the effectiveness of these bots is difficult. The [best analyses](#) indicate that they did not affect the 2016 U.S. presidential election. More likely, they distort people's sense of public sentiment and their faith in reasoned political debate. We are all in the middle of a novel social experiment.

Over the years, algorithmic bots have evolved to [have personas](#). They have fake names, fake bios, and fake photos—sometimes [generated by AI](#). Instead of endlessly spewing propaganda, they

post only occasionally. Researchers can detect that these are bots and not people based on their patterns of posting, but the bot technology is getting better all the time, outpacing tracking attempts. Future groups won't be so easily identified. They'll embed themselves in human social groups better. Their propaganda will be subtle, and interwoven in tweets about topics relevant to those social groups.