GLOBAL

Russia's Election Interference Is Digital Marketing 101

The new Mueller indictment doesn't get at the root of the problem: the unchecked market power of social-media companies.

DIPAYAN GHOSH AND BEN SCOTT FEBRUARY 19, 2018

A Facebook like button is pictured at the Facebook's France headquarters in Paris, France, November 27, 2017. (REUTERS/BENOIT TESSIER)

Last Friday, the Justice Department charged 13 Russians with attempting to subvert the 2016 U.S. presidential elections. The case presented by Special Counsel Robert Mueller laid out an elaborate scheme of information operations, carried out primarily via the social media websites Facebook, Instagram, and Twitter. Through the Internet Research Agency, a so-called "troll factory" in St. Petersburg, the Russians created hundreds of fake accounts on these services, which then disseminated fake news and other misleading content about Democratic candidate Hillary Clinton to hundreds of thousands of users. They focused their campaign on topics that divide America—race, immigration, and religion—and targeted battleground states. According to figures reported by Facebook and Twitter, the

Russian campaign reached more than 125 million Americans on Facebook; over 675,000 people engaged with Russian trolls on Twitter. The Russians' effort is, of course, ongoing.

Thus far, the media coverage of Mueller's indictment has fixated on how all this could have happened, and probed whether the Trump campaign was involved. The answers to these questions will all emerge in time. The more troubling question is why it was so easy to make fools out of so many Americans.

Consider two things. First: While the Russians created fake accounts to pose as Americans on social media and buy ads, the technologies they deployed are all commonplace in the digital-marketing industry—this was no 007-style spycraft. Second: These days, Americans live in divisive, partisan information environments, chock-full of incendiary rhetoric. They have very low standards about the sources they accept as accurate, and yet aren't great at parsing fact from fiction on the Internet. Even "digital natives"— young people most at home in an online information environment—have proven inept at judging credibility. In other words, when the Russians set out to poison American politics, they were pushing on an open door.

How does a ready-made toolbox for digital manipulation already exist? For that, we have the digital-advertising industry to thank.

In a recent study on the digital-advertising industry that we published with New America and Harvard's Shorenstein Center, we analyzed how the tools of digital marketing can be readily repurposed by agents of disinformation. The basic idea is for advertisers to micro-target digital advertising at very specific demographic slices of social-media users to see how they respond. A disinformation operator could test hundreds of different messages, often aimed at thousands of different permutations of demographic groups on the advertising platforms of the most widely used social-media companies.

For example: A political advertiser (or communicator) might test a message about immigration in different cities across the country, or it might compare responses to that message based on age, income, ethnicity, education-level, or political preference. Because digital-media companies like Facebook collect vast amounts of data on their users, advertisers can parse based on age, income, ethnicity, political affiliation, location, education level, and many other consumer preferences that

indicate political interests. Once the ad buys indicate what messages get the biggest response from particular groups, the operator can organize its entire social-media campaign to reach those people and build out bigger and bigger audiences.

This is digital marketing 101. Start with a product to sell and test a variety of messages until the best one rises to the surface.

In the election-interference case, the "products" for Russian trolls were divisive political messages about issues like, say, religion. But just as with any other product, the ads ginning up fear and outrage about Islam in America benefited from Google and Facebook's machine-learning algorithms, which scan vast amounts of data and conduct tests on multitudes of political messages to determine the best way to find and engage an audience. Everybody makes more money if the ads work well—that is to say, if people click on them. The economic interests of advertisers and social media companies are essentially aligned. And while Facebook, Google, and Twitter are now taking steps to identify and block ads purchased by foreign agents and shut down these attempts to push fabricated news, the underlying machine of the ad tech market will, theoretically, accelerate users' consumption of all but the most egregious content.

When political advertisers—including purveyors of disinformation—get into the mix, the economics of audience segmentation and micro-targeted advertising start to produce what is known as a "negative externality" in the market, or an unintended outcome that harms the public. The system naturally organizes people into homogenous groups and feeds them more of what they want—typically, information that reinforces their pre-existing beliefs—and then ups the sensation-factor in order to hold people's interest for longer stretches of time.

A recent analysis of YouTube, for instance, showed that the videos in the "next up" queue were fed by an algorithm that prioritized keeping eyeballs glued on videos. The results predictably fed users content that matched previous preferences, or, failing that, just increased the level of sensationalism. In the wake of the Las Vegas shooting, users who watched at least one YouTube video questioning whether the shooting actually happened were then recommended more videos of the same sort—a dangerous example of how social-media algorithms can perpetuate and promote propaganda.

Today, even though hundreds of millions of people get their news and information from Google, Facebook, and Twitter, they are fragmented and polarized into a variety of isolated communities, ranging from the staunchly conservative to the hard left. In such an information environment, it's common for everyday users of social media to circulate incendiary content from dubious sources. So when the Russians inject streams of content suggesting that NATO is showering chemicals across Poland or that a Ukrainian policeman proudly donned a Nazi uniform, it doesn't seem so extraordinary for most of the audience.

Indicting 13 Russians, none of whom are likely to ever face justice, is not going to solve this systemic problem. Of course, to be clear: Russian agents should certainly be cut out of the market of persuasive messaging. But this won't fix things for good. A real solution will require a hard look at the relationship between information markets and democracy, and a focus on the public interest over the profit motive. By its very nature, the digital-information market fragments the political culture and separates people from a common set of facts that allows for functional self-government. Part of this problem lies in the economics of the market; proper regulation and new corporate practices can mitigate it.

In recent months, Congress and the leading internet companies have proposed various efforts to inject transparency into political advertising. But that can only do so much to limit the effects of the well-organized, nefarious disinformation operations of the future. Moving forward, social-media companies need to think about how to segregate the goals they implicitly share with disinformation actors.

In the near term, this means that companies like Facebook and Twitter could develop technology to weed out attempts at political disinformation—tools powered, perhaps, by the very same advanced algorithmic tools that enabled disinformation to spread on their platforms in the first place. But in the longer term, this is a problem of market power. Experts would do well to call for limits on the vast amounts of data available to a digital-advertising industry dominated by social-media and internet-platform companies by enforcing comprehensive privacy reforms. All of this demands a more concerted effort to curb the immense concentration of power enjoyed by the largest internet-platform companies.

As new operators—both foreign and domestic—learn from the Russians, this problem is bound to get worse before it gets better. Moving swiftly to reverse these

trends is vital.

We want to hear what you think about this article. Submit a letter to the editor or write to letters@theatlantic.com.