THE NEW ATLANTIS
A Journal of Technology & Society

The Ruin of the Digital Town Square

How Not to Regulate Social Media

Shoshana Weissmann

Regulation of social media has become a hot topic since the Russian meddling with the 2016 election and last year's widely misunderstood Facebook–Cambridge Analytica scandal. Legislators have debated various problems with social media and bots, with the only agreed-upon themes being that social media firms are up to something nefarious, and that something must be done about them. Unfortunately, policymakers are too confident that simply throwing more government at the problems—whatever those may be—will fix them, and are proposing a number of solutions that are well-intended but largely ill-conceived.

It's worth reconsidering, then, what exactly these controversies brought to light. Doing so will make clear that the laws proposed don't address the real problems and instead would trigger adverse unintended consequences.

At its core, the problem underlying the Facebook scandal was a lack of data security. Facebook provided Cambridge Analytica with data intended only for academic use, but it did not go as far as it should have to ensure that the data would in fact be used solely for this purpose.

Legislators had an opportunity to discuss data security issues during a congressional hearing with Facebook CEO Mark Zuckerberg last year. Yet many of the questions asked of Zuckerberg dealt not with data security, but with the fact that data was being collected in the first place. Some members of Congress even seemed to take issue with Facebook's use of that data for its targeted advertising. This was a bizarre criticism, as politicians routinely use these very same targeted-advertising tools during their campaigns. And although the Cambridge Analytica problem may have raised tangential concerns in legislators' minds, many of the concerns they raised in the hearing—like the suggestion that Facebook is censoring conservative viewpoints—were unrelated to the underlying data security issue.

The same dynamic applies to the problem of Russian bots' interference with the 2016 election. The core problem there was that political actors

Shoshana Weissmann *is the Digital Media Manager and a fellow at the R Street Institute.*

in Russia created large numbers of phony social media accounts designed to look like accounts of American citizens, then posted incendiary content from these accounts in order to sow social discord in the United States and influence the outcome of the election. Bots—social media accounts that are sometimes, though not necessarily, controlled by automation—were key to getting this content posted en masse. Yet Twitter data showed that the accounts controlled by humans creating similar content generated higher engagement than the automated bot accounts. Many elected officials, however, have ignored this, attempting to crack down on bot activity as if they were the main instruments of havoc. What's more, most officials misunderstand what bots are and how the Russians used them.

While the legislation mentioned below attempts to address concerns over data security and bots, the legislators drafting these bills fail to grasp what the real problems are—as is evident from the misguided solutions they propose.

Policymakers have proposed various legislative solutions to problems with social media privacy and bots. Are any of these proposals up to the job? Regrettably, the answer is no.

On the news site *Axios*, David McCabe helpfully breaks down the Social Media Privacy Protection and Consumer Rights Act of 2018. The bill, introduced by Senators Amy Klobuchar (D-Minn.) and John Kennedy (R-La.), mandates what social media companies are largely doing voluntarily: publishing simplified terms of service that more clearly explain how the platform will collect and use personal data. While disclosure requirements of this kind are common in many sectors—for example in health care and real estate—they are typically ineffective in actually ensuring informed consent, since there is tension between the legal need to disclose large amounts of information and the advantages of brevity and clarity. Law professors Omri Ben-Shahar and Carl E. Schneider make this case persuasively in their 2014 book *More Than You Wanted to Know: The Failure of Mandated Disclosure*.

Other provisions of the legislation could be applied much too broadly, with significant and negative unintended consequences. To deal with data breaches, for instance, the bill requires that an operator of an online platform send a notification to a user within 72 hours of learning that the user's personal data "has been transmitted in violation of the privacy or security program of the online platform." The legislation applies not only to platforms like Facebook and Twitter, but to any online platform that "collects personal data during the online behavior of a user of the online platform."

But operators of innumerable websites collect personal data, and many of them probably wouldn't realize that the legislation might apply to them. In fact, it is not even clear which cases this legislation would apply to. Consider a small business that creates a website using a service like Squarespace and collects email signups through that website. If Squarespace's data is leaked, then the law would presumably apply to Squarespace. But what if someone discovers the personal password of the small business owner and steals the email addresses? Would this law apply to the small business? It appears that it would, which means that countless small businesses would have to take on this new responsibility or face lawsuits.

The California Consumer Privacy Act similarly zeroes in on privacy and user data. Signed into law last June and going into effect in 2020, it gives users the right to have their data deleted, to forbid a platform from selling their data, and to request a copy of their data "to the extent technically feasible, in a readily useable format that allows the consumer to transmit this information to another entity without hindrance." The law applies only to those companies with "gross revenues in excess of twenty-five million dollars," but its effects would be broader than some realize, since those large companies serve smaller ones. For example, smaller companies who rely on larger companies' address lists to target ads to a narrow audience would no longer have access to as many addresses.

This legislation has a number of problematic features that would likely trigger potentially unwanted consequences. For instance, one implication of the text is that users have the right to switch their data to another platform offering the same service. But for much of the data in question, seamless transmission is just not how things would work. As it happens, Facebook already offers users the ability to download an ostensibly comprehensive dump of all the data it has on them. But for the ordinary user, there is no comparable entity to transmit that data to. If you've "liked" a page on Facebook, it's not possible to transmit that data to Twitter or any other major platform. Similarly, one cannot download all of one's Facebook posts and upload them to Pinterest. It wouldn't make sense to do so anyway, because different platforms have different functions and serve different purposes.

The law also allows the data to be delivered to the user electronically or by mail. If data is delivered by mail, it surely cannot be transmitted to a different online platform in that physical format. And the kind of user

data a platform could readily provide—the user's age, gender, location, occupation, interests, political leanings, and so forth—is probably not the kind of data one would be transmitting to another platform anyway. The user already knows this information, meaning it would be much easier to submit it to a new platform manually. Further, the law arbitrarily provides that a consumer cannot require the company to send over user data more than twice in a twelve-month period. Why twice rather than once or three times? Striking a balance between consumer interests and not overburdening companies is sensible, but it is unclear why two times is the magic number.

The law also has the potential to topple lots of businesses, as Dipayan Ghosh outlines in a July 2018 *Harvard Business Review* article. The danger it poses to social media platforms and search engines is obvious. For these companies, targeted ads are their main revenue sources, and they would be unable to target ads to consumers who demand that their data be deleted. But the bill would also harm businesses outside of the social media industry. Many businesses rely on the targeted-advertising capabilities of social media platforms to market their goods and services. And if they are limited in their ability to target ads to relevant audiences using social media platforms, they will cease to do so, which could cause them to go out of business. A 2018 survey found that 63 percent of African-American–owned and 35 percent of women-owned businesses were built on Facebook. Never before have small and new businesses had such powerful, or such cost-effective, ad-targeting capabilities, but this bill would severely undermine the advertising mechanisms that these small businesses depend on to flourish.

The risk could also extend to other companies that rely on the collection of user data. For example, Internet providers like AT&T and Verizon collect web-browsing data to help target advertisements. Other companies, as well as political campaigns, rent out or sell email lists to other like-minded companies and campaigns. People receiving unwanted emails can already unsubscribe from email lists, but this law grants the right to demand that one's email address not be made available for such lists in the first place.

The advantage to people who don't want to be bothered with emails is obvious. But it's worth keeping in mind that many businesses rely on gaining new contacts through third parties, and that this law would severely cripple their revenue streams. Multimillion-dollar data firms affected by the measure may not be the kinds of victims that elicit our sympathies. But what about the smaller firms and individuals that rely on

their data? Small congressional campaigns, for instance, rely on data from larger firms to reach out to voters for the first time.

If you are troubled by the concept of "targeted advertising," you should consider that it has been around since long before the advent of the Internet. Small businesses used to spend significant amounts of money on print or television advertisements to reach a small portion of potential customers located within fifty miles of the business. Now, thanks to online advertising, a small boutique can spend less money to target, say, thousands of women aged 18 to 35 who "like" Claire's on Facebook and live in the same town. Online targeted advertising is not something new, per se; it's simply a more efficient version of a much older practice.

Another California legislation, Senate Bill No. 1001, was signed into law last September and will go into effect in July. It purports to address the role of bots in fake news and advertising. But the law's definition of "bot" is too broad, and the method it uses is a poor match for the goal it seeks to accomplish.

The law makes it illegal to use online bots for commercial or political advertising in California unless the bot discloses its bot-hood in whatever communications it engages in. However, the law defines "bot" as "an automated online account where all or substantially all of the actions or posts of that account are not the result of a person." This definition likely applies to much activity that legislators may not have intended to regulate. For example, social media managers often use apps like Hootsuite and Buffer to automate the posting of social media content. These programs allow the manager to write a batch of tweets at one time and schedule them to post automatically over several weeks. Similarly, political campaigns and other organizations have programs for sending automated emails and text messages to appear as if they are from prominent political figures.

Many of these activities would likely fall under the legislation's definition of "bot." However, they are not inherently malevolent; they are simply efficient uses of technology to save labor. Of course, these technologies can be used nefariously, but the legislation has no means to ensure that it targets only bad actors. What it prohibits is using a bot "with the intent to mislead the other person about its artificial identity…in order to incentivize a purchase or sale of goods or services…or to influence a vote in an election." But "artificial identity" is ambiguous enough to apply to a great deal of automated online activity, much of which is perhaps misleading in some way but not actually harmful.

Furthermore, the legislation won't address the very problems it tries to solve: fake news and fraudulent advertising. To achieve provocative yet believable headlines, fake news has to be written by humans. The same is true of fraudulent advertising: It is humans, not robots, who create the fake ad. The Russian propaganda campaign was run by real people who wrote the content for the bots to post. In fact, a significant share of the actual posting was done by hand. Targeting automation does nothing to address the root problem here.

Moreover, one of the primary reasons legislators were concerned about advertising was that Russians also used bots to generate high "view" counts on videos in order to fool marketers into purchasing millions of dollars of video ads. But the legislation appears to apply only to bots that communicate with people, not to bots that create many views on videos, meaning that the legislation would not resolve this problem.

On the federal level, the proposed Bot Disclosure and Accountability Act of 2018, introduced by Senator Dianne Feinstein (D-Cal.), suffers from similar issues. It would require that anyone using software or other processes to "impersonate or replicate human activity" on social media "provide clear and conspicuous notice of the automated program" to other users. But this, again, is problematic because of the blurred distinction between human and automated activity.

The legislation also requires that platforms establish a process by which social media users can provide notice that their accounts use automated tools. This requirement is largely unnecessary, as one could simply add that information to one's biography on any platform that allows scheduling.

Furthermore, the law bans the use of automated social media tools for political advertising, even if the bot accounts are disclosed as such. Candidates and parties may not use "automated software programs or processes intended to impersonate or replicate human activity online to make, amplify, share, or otherwise disseminate any public communication." As with Senate Bill No. 1001, this prohibition might well apply to something as ordinary as a staffer scheduling a candidate's tweet to post during the evening or early morning, outside of normal working hours. This would hamper political upstarts and other candidates with limited resources, creating a further advantage for incumbent and wealthy candidates who can afford an all-hours social media team. Auto-sharing doesn't always result in the highest engagement, but it is a big help to cash-strapped campaigns. Automated targeted ads, which are widely used by political campaigns, would likewise be banned by this provision.

In each of these cases, there are varying degrees of mismatches between the intended aims of the legislation and the means to achieve them. These issues are complicated, and solutions are not easy. But the legislators who sponsored the bills appear unaware of the potential harms they could cause if enacted and enforced.

None of this is to say that solutions do not exist. They may exist, and they may highlight a more constructive role for government in resolving social media's dysfunctions. Reviving the federal Office of Technology Assessment would help. So would applying Yale law professor Jack Balkin's "information fiduciary" model to social media companies. Its basic idea is that our relationship to these companies and their use of our data could be regulated by similar principles as our relationship with doctors and lawyers, who need personal information about us to serve us well but who also have duties of confidentiality and loyalty toward us.

There are thoughtful and well-researched ideas to consider. But in order to succeed, the process must begin with a full understanding of social media technology as well as the business landscape that lawmakers are seeking to regulate.