Paper 1 | December 2020

The Weaponized Web

Tech Policy Through the Lens of National Security

Levers in the Digital Advertising Ecosystem

By Dipayan Ghosh, Lindsay Gorman, Bret Schafer, and Clara Tsao

alliance for securing democracy
G | M | F

HARVARD Kennedy School
MOSSAVAR-RAHMANI CENTER
for Business and Government

Introduction

Open societies have cultivated rapid technological advancement and market innovation—which have vastly outpaced democratic governance. Authoritarian powers have seized on the underlying opportunity to exploit the open standards of the democratically regulated online information environment and undermine democratic values and institutions while shoring up their own regimes. This poses a novel challenge for democracies, which must adapt to compete in this conflict over the data, architecture, and governance framework of the information space without compromising their principles. Effectively competing in this environment—and ensuring a democratic future for the online information space—will require policymakers to analyze technology and internet policy through the lens of national security.

This paper will form one part of a series of reports discussing policy proposals meant to address national security vulnerabilities in the online information ecosystem. Each report will be narrowly scoped in order to provide sufficient space to debate the relative merits of regulatory proposals. Although the authors make specific recommendations, the purpose of these reports is to foster deeper debate about potential policy options; as such, we present arguments for and against each regulatory option, with an emphasis on highlighting potential negative externalities. While intended for a global audience, the recommendations are decidedly U.S.-centric, due simply to the authors' deeper understanding of the U.S. regulatory landscape.

It is worth noting that these papers do not address the many underlying political and social issues that contribute to online harms. We also recognize that government regulation is not the only—and often not the best—tool to solve many issues in the digital domain. Our goal, therefore, is not to solve the myriad issues with how we produce, distribute, and consume information—it is to assist regulators and concerned stakeholders in thinking through legislative options to mitigate the national security concerns associated with malign foreign activity, interference, and alternative modes of governance online. This focus means that our recommendations are more narrowly scoped to the national security challenge than those advanced by others—including the Digital Innovation and Democracy Initiative, ASD's sister program at the German Marshall Fund.

Finally, it is our hope that these reports will move the tech policy conversation beyond empty platitudes, generalities, and well-intentioned but ultimately impractical proposals. Years after these problems first surfaced, it is our shared belief that it is time to stop admiring the problem and instead focus on concrete solutions.

Authoritarians and the Online Ad Ecosystem

Western democracies were first awakened to the potential for authoritarian exploitation of online political advertising in 2017, when Facebook revealed that the Russian government-linked Internet Research Agency had spent $100,000 on advertising on the platform in an attempt to influence the 2016 U.S. election.[1] Similar revelations from Google and Twitter contributed to concern as policymakers clamored for a response.[2] In the subsequent months and years, online platforms have taken numerous steps to improve transparency around digital advertising, principally in response to growing public and political pressure.[3] For its part, Congress has introduced several bills to mitigate the malign manipulation of online advertisements, though the issue has been mostly hamstrung by the politicization of foreign interference.[4]

The online advertising ecosystem presents a powerful medium for authoritarian actors and violent extremists to target specific demographics, build robust audiences, and even fund or profit from information operations. The same aspects that make online marketing a valuable industry for companies make it a useful tool for a range of bad actors. The ability to microtarget specific demographic and social groups allows marketers to segment and splice key audiences, but it also enables malign actors to target suppressive or divisive ads at those they deem susceptible—as the Internet Research Agency did in 2016.[5] Similarly, even as organizations—including journalistic outlets and public interest advocacy groups—rely on online ads to drive traffic to their site, authoritarians have adopted similar tactics. In 2019, Facebook revealed that employees of Russian state-controlled news agency and broadcaster Sputnik spent around $135,000 on ads as part of a coordinated operation to drive traffic to Sputnik content.[6]

Yet the focus on foreign-purchased political ads—while important—overlooks the more significant role that online advertising plays in funding the malign activities of both state and non-state actors online. According to a 2020 report from the Global Disinformation Index, nearly a quarter billion dollars' worth ($235 million) of advertising has been funneled to sites peddling disinformation.[7] Numerous other reports have outlined how ads from reputable brands have run on domains or social media platforms alongside content promoting violent extremism, hate speech, conspiracy theories, and state-backed propaganda.[8] American and European companies—and, at times, political parties[9]—have thus become the unwitting sponsors of content aimed at dividing democratic societies and undermining democratic institutions.

To address these concerns, online platforms have taken a number of steps to attempt to ensure the integrity of the online ad ecosystem and to stave off criticism. In large part, these efforts have focused on transparency through both the labeling of ads with purchaser information and the creation of online ad libraries.[10] Other policies have aimed to increase scrutiny of who is purchasing ads to prevent inauthentic manipulation.[11] Some platforms have restricted advertising altogether, most notably when Twitter banned all political ads in late 2019.[12]

While these efforts are an improvement on the pre-2016 online ad ecosystem, they are insufficient. Self-regulation has allowed companies to implement policies that are self-serving, contradictory, or unenforceable; and even well-intentioned measures have proved insufficient without robust deterrents. Moreover, authoritarian and extremist actors have rapidly adapted to changes on online platforms to circumvent these new policies, many of which have proved easy to manipulate.[13]

Policy Proposals

In this section, we outline a series of six legislative policy proposals for the online advertising ecosystem; provide background on the policy discussion; review arguments against and in favor of the proposal; and make a final recommendation for Congress.

1. Require greater transparency from AdTech companies by:

a. **Mandating that companies involved in the placement of programmatic ads (ads that are bought and sold through automated processes) provide advertisers with detailed disclosure reports identifying publishers (Site IDs) served by their ad buys; and**

b. **Legislating an aggregated and anonymized public disclosure requirement of ad spends by AdTech firms, similar to the SEC's quarterly reporting framework.[14]**

Background

The opacity of the online AdTech industry means that reputable brands are often the unwitting financial supporters of state-directed propaganda, extremist sites, and other outlets that can degrade democracy and potentially threaten national security. This is a byproduct of an enormously complicated online advertising ecosystem that obfuscates the money trail from advertisers to publishers, meaning advertisers and ad buyers do not necessarily know where their programmatic ads are placed.[15] Although this information is available to advertisers upon request, there are multiple layers that can obscure the true domain from advertisers. Resolving this opacity presents an opportunity for interest alignment between private sector advertisers who would like to see ads placed on sites to optimize sale conversion and public interest researchers and advocacy groups seeking to stop the funding of extremist content.

Arguments Against

Most ad networks—for example, Google AdSense—already offer advertisers the ability to see where their ads are placed, and true disclosure is often impossible due to the complexity of the advertising ecosystem. Without external pressure or oversight, transparency will not lead to more responsible advertising choices. Instead, requiring ad networks to publicly disclose ad spends would reveal proprietary marketing strategies that could be used by competitors. Finally, this is a problem best tackled by self-regulation and the market rather than government intervention, which is too slow to adapt to the rapidly changing online ecosystem.

Arguments in Favor

Without greater transparency into the publishers served by AdTech companies, ad dollars will continue to be funneled to sites that pose a national security risk. On walled garden platforms like Facebook and YouTube, advertisers of course know the domain their ads appear on, but not necessarily the content their brand appears alongside (for example, an ad running before a foreign propaganda video on YouTube). Advertisers have a right to know if they are unknowingly supporting content that threatens their brand reputation, particularly if they have used controls to avoid problematic content. Currently, however, responsible advertisers do not always have the requisite data to audit AdTech companies. Arming them with this knowledge would allow them to make more informed choices about their ad partners.

Civil society and advocacy groups have also traditionally played an important role in pressuring advertisers to curtail or boycott programming that promotes divisive, extremist, or propagandistic viewpoints. Watchdog groups, however, have far less insight into companies' ad spends than the companies themselves, and they are currently forced to scrape data from websites or screenshot individual instances of problematic placements. Public disclosures of the publishers served by programmatic ads would allow advocacy groups, like the Global Disinformation Index, to more effectively monitor and pressure advertisers who, in turn, can pressure ad networks that support extremist, junk, or propaganda websites.

Final Recommendation

Governments should require AdTech companies to disclose to advertisers the true domains and URLs served by their ads. Although most major AdTech companies provide this information upon request, the disclosures are often muddied by a lack of specificity or outright fraud (more on that below). AdTech companies should therefore be required to disclose to advertisers the publishers served by programmatic ads, and user-generated platforms should be required to disclose the specific URLs where ads appeared.

Legislation should also require ad networks to produce quarterly public reports to include the amount of money spent with individual publishers. These reports would allow both the advertising industry and watchdog groups to hold networks to account. However, this is an instance where there needs to be a balance struck between the benefits and harms of public disclosure. We therefore recommend that public reporting be anonymized and aggregated in a way that protects the marketing strategies of advertisers, while revealing which ad networks are supporting problematic content.

2. Appropriate funds for more rigorous investigations into and prosecution of digital ad fraud

Background

Central to the effort to improve transparency in the online space is the need to minimize various forms of ad fraud that not only harm legitimate business interests, but also siphon money to disinformation and extremist websites. Although ad fraud most commonly takes the form of fake clicks and manipulated traffic, it also includes the process by which websites with high traffic but low-quality content (namely pornography or state-sponsored propaganda) can be paired with one or more sites with low traffic but safe content. In 2019, for example, Uber alleged that its ad network partner, Phunware, committed wire fraud, racketeering, and common law fraud by defrauding Uber of $17 million by purchasing ad placements on behalf of Uber that were, according to the lawsuit, "not real ads, were illegitimate ads, and/or were prohibited ads, such as 'auto-redirects' or ads placed on prohibited sites such as pornographic websites."[16]

Although there are several ways malicious actors can commit this type of fraud, one of the easiest to prosecute is the intentional mislabeling of ad inventory, a process Check My Ads[17] has labeled "dark pooling."[18] To understand dark pooling, it is helpful to briefly (and simplistically) define the two different types of advertising account IDs. Every website has a number of unique ad IDs that allow them to be recognized on ad exchanges. These IDs are broken down into two types: direct and reseller accounts. A direct ID, as the name suggests, signifies to ad buyers that they are placing an ad *directly* on a specific website. A reseller ID indicates to advertisers that they are bidding on ad space across multiple websites, for example, all properties owned by a media conglomerate like the Gannett Company. This process is known as "pooling."

If the ad ID on an exchange is properly labeled as a reseller account, pooling is a perfectly legitimate practice, as advertisers are aware that they are bidding on ad space across multiple publishers. If, however, a direct ID is being shared across multiple sites, this is potentially a form of digital ad fraud because advertisers who believe they are bidding on one site, may, in fact, be bidding on a "dark pool" of sites, some of which may be less relevant to the advertiser or traffic in conspiracy theories or disinformation.

A report published in July 2020 details how several sites, including RT.com, share direct ID inventory with numerous other sites—essentially mislabeling reseller IDs as direct IDs.[19] This means that advertisers that have purchased a direct ad with a site that presumably carries less risk to their brand reputation, may, in fact, be supporting RT and other problematic sites that share that same direct ID. This form of profit sharing also makes it exceedingly difficult for advertisers to audit their ad spends, as a list of placements would not reveal direct IDs that have been fraudulently pooled.

Dark pooling and other ad fraud practices are likely already illegal, yet enforcement has been nearly nonexistent. This is due in part to the fact that some ad fraud schemes, like cross-domain spoofing,[20] are difficult to detect.

However, the Department of Justice has laid out a roadmap for the strategic prosecution of large-scale digital ad fraud rings, including a 2018 indictment of eight Russian and Kazakh cybercriminals accused of defrauding clients out of tens of millions of dollars.[21]

Arguments Against

Ad fraud is a problem best dealt with at the industry-level, and, indeed, the Association of National Advertisers and others have taken steps to combat widespread fraud. Plus, the specific tactics used to fraudulently funnel money into sites that pose a national security concern are already illegal, yet prosecutors have proven unwilling to bring cases. Given the scale, scope, and complexity of the problem, it would require a significant investment of prosecutorial and investigatory time, money, and personnel to tackle the problem, diverting resources from other law enforcement priorities.

Arguments in Favor

Ad fraud costs legitimate advertisers billions of dollars per year.[22] While most of the money earned in ad fraud schemes is not directed towards sites and actors that pose a national security risk, there is evidence that a range of malicious activity, from violent extremism to foreign disinformation operations and cybercrime, is financed, at least in part, through ad fraud. Plus, dark pooling and other deceptive measures that obfuscate the true recipient of ad revenue render transparency measures, like those outlined in the first recommendation, ineffective if not useless. Prioritizing the enforcement of existing fraud statutes would allow law enforcement to investigate and strategically prosecute cases that pose specific threats to national security.

Final Recommendation

Governments should appropriate funds to assist in the investigation and prosecution of entities engaged in fraud in the digital advertising space. While prosecuting each individual case of ad fraud is impractical, prioritizing funding for strategic enforcement against particularly egregious cases and actors would cut off a significant funding source for propaganda outlets and for-profit disinformation operations.

3. Create legal recourse for advertisers whose programmatic ads appear on sites that they have previously blacklisted

Background

Most AdTech companies offer advertisers the ability to blacklist (note that the use of this name, due to its racial connotations, is currently being revisited by the ad industry) sites or types of sites that they determine pose a risk to their brand reputation. Currently, however, if an ad network places an ad on a blacklisted site, advertisers have no real options for recourse, other than pulling future ads from the offending platform. Given the market dominance of Facebook, Google, Apple, and Amazon, this decision would potentially come with a huge cost to advertisers.

Arguments Against

Social media companies are already fined in countries like Germany (NetzDG) up to 5 million euros if terrorist content or hate speech are found on platforms and not deleted within 24 hours. It is more effective to push companies to focus on removing problematic content rather than being bombarded with lawsuits from aggrieved advertisers. Plus, given the enormously complex web of companies involved in the buying and selling of online ads, it would be challenging to determine liability. Government intervention in this case also feels like a form of industrial policy—taking a side in a dispute between two industries. Advertiser boycotts and other market-based measures have proven effective at motivating AdTech companies to invest more in platform security. For example, in July 2020, 1,000 companies participated in the Stop Hate for Profit campaign to boycott Facebook, dropping Facebook's share price by 8.3% and leading to promises of more robust enforcement of problematic content.[23] This suggests that additional legislation is unnecessary.

Arguments in Favor

While ad pull-outs by major brands have been somewhat effective at pushing large social media companies to enact changes, the market dominance by the major online advertising platforms limits the ability of advertisers to "break-up" with AdTech companies. Providing advertisers with a clear-cut legal mechanism—outside of tort law—to hold AdTech companies to account for violating their own advertising control policies would force greater due diligence from AdTech companies.

Final Recommendation

Congress should consider new regulatory policy to advance a system of recourse for digital advertisers. This legislation should be limited to instances where AdTech companies directly violate their own terms of service or the individual ad controls put in place by advertisers. This is particularly challenging, however, on advertising platforms dominated by user-generated content (for example, YouTube), where the screening of unsafe content is heavily reliant on automated processes that are prone to error. It is therefore our recommendation that this proposal needs further study to understand both the need for and the impact of legislation.

4. Require transparency from digital platforms: to publicly disclose the entities or individuals that pay for all digital targeted advertisements, to embed disclosures into digital ads themselves, and to perform due diligence to assess that the information provided by advertisers is accurate. This measure would expand upon existing proposals focused on political ads—most notably the *Honest Ads Act* introduced in the U.S. Senate—and apply to all digital advertisements placed on large platforms

Background

Targeted advertising in jurisdictions like the United States is gradually becoming the dominant form of corporate outreach to consumers, with this year—one defined by the coronavirus pandemic—marking the first time that marketing spent on digital advertising exceeded that spent on advertising in all traditional media formats.

A matter of public policy, however, is the growing discrepancy in transparency between the digital and traditional media environments when it comes to paid advertisements. Advertising in traditional media formats builds in certain forms of transparency, whether organically or through regulation. Such transparency measures are clearest in the case of political advertising. On broadcast television, for instance, consumers and the broader public can determine the provenance of the ads they see—that is to say, the entities that paid for and developed a broadcast political advertisement. Additionally, consumers have access to the aggregate spending and reach of a campaign's political advertisements over time because there is a record of where and when political advertisements have aired. Such circumstances and features of the traditional ad ecosystem contribute to a perception that political advertising on broadcast television is above board: There is accountability for an ad, consumers have some sense of where the ad has originated from, and there is a general sense that, should one wish to know how a campaign has spent its money in aggregate over broadcast networks during the course of an election cycle, such determinations are possible.

These protections do not exist in the digital landscape. Concerns over this discrepancy came into high resolution in the aftermath of the 2016 presidential election: A slow trickle of facts pried from Facebook, Google, and Twitter revealed that each platform had been infiltrated by agents of the Kremlin who covertly purchased and disseminated divisive ads targeting Americans in the lead-up to election day. The harms baked into the non-transparent digital space might extend to seemingly legitimate uses of digital advertising, as well: The Trump campaign's digital director for the 2016 election has revealed how he organized a contingency-based digital outreach program whereby seemingly endless configurations of ad content were automatically tested on various audience segments.[24] This meant the campaign could determine what kinds of ads worked with which communities of voters. But no law was in place to enable policymakers or the public to discern from the platforms which

political actors disseminated which ads to which audience segments in the lead-up to the election.

Such non-transparent political communication is harmful to the democratic process. As such, the *Honest Ads Act* was designed and introduced by Senators Mark Warner, Amy Klobuchar, and John McCain to force transparency on the platforms in the digital advertising realm, so that the online landscape could meet the standards in place for traditional media.[25] Most notably, the bill would establish that digital platforms hosting political ads would have to disclose the entity that publishes and finances the ad (that is, its provenance) and maintain a searchable online database of political ads so that the public and researchers could examine political ad campaigns. The bill has been described as a common-sense measure, and though it has not been pushed forward in a largely gridlocked and polarized Congress, under different political circumstances it could enjoy support.

Another important consideration is *how* ad disclosures are technically deployed. Efforts by platforms to provide contextual information on certain profiles and channels—namely, those operated by government officials and state media—have proven easy to abuse. For example, YouTube's state media labels do not appear when those videos are shared on Twitter. Requiring disclosures to be embedded in the ads themselves would help ensure consistency across platforms.

Arguments Against

Requiring platforms to collect and house data on all online ad spends would be expensive and onerous. It could also create a problem of big data overload—a process by which the public and watchdog organizations suffer from having too much data—making it difficult to effectively monitor and understand ad campaigns most relevant to democratic processes. Finally, paid-for ad disclosures in traditional media have done little to increase transparency or stem the flow of dark money into political or issue campaigns. Similar requirements online would be even more difficult to police, leading to rampant abuse by bad actors.

Arguments in Favor

A common criticism of the *Honest Ads Act* is that there is not a clear distinction between political and issue ads. Most of the ads purchased by Russia's Internet Research Agency, for example, were focused on issues rather than the candidates themselves, signaling that such a distinction would also hamper efforts to combat foreign interference. Requiring transparency on all ads, as recommended by Phil Howard of the Oxford Internet Institute,[26] as well as GMF's Digital Innovation and Democracy Initiative,[27] would remove this artificial distinction. Moreover, without these disclosures embedded into ads themselves, entities can find easy loopholes to skirt disclosure requirements by re-posting ads through third parties or sharing across different platforms.

Final Recommendation

We recommend that governments pass legislation requiring digital platforms whose monthly visitors exceed a certain threshold to disclose publicly the entities or individuals who pay for all targeted digital advertisements; to embed these disclosures into digital ads themselves; and to perform due diligence to assess that the information provided by advertisers is accurate. As Karen Kornbluh and Ellen Goodman have recommended, "Know your Customer" funding checks should be implemented to verify the completeness and accuracy of advertiser information, which could be modelled after the measure used by banks and financial institutions to combat money laundering.[28] Not only should platforms of a certain size be responsible for collecting data on advertisers (and ensuring that ads have appropriate disclaimers), but as the point of entry of this content into the digital information ecosystem, they should also be responsible for a level of due diligence to mitigate misrepresentation. Once verified, these disclosures should be embedded into ads themselves so that when a user retweets, shares, or forwards the ads, disclosures remain intact. This measure will guard against workarounds from malign actors that leverage third parties to spread undisclosed ads while raising the level of difficulty in sharing paid content of unclear provenance in the digital ecosystem.

5. Restrict advertisers' ability to target political ads beyond broad categories such as gender, age, and postal code

Background

As social media platforms amass ever-more precise data on user behavior to paint individualized pictures of consumers and fuel targeted advertisements, collateral impacts to the political process have accrued. In particular, microtargeted political advertisements can create opacity in civic discourse that undermines the information environment behind selecting political leaders. Armed with detailed information on what specific voters want to hear, campaigns and foreign influencers alike can target different and even conflicting messages to different slices of the electorate, absent accountability for those contradictions.

As the Senate Select Committee on Intelligence has assessed, "Advanced micro-targeting in the commercial sector is also rapidly becoming more effective. Propagandists will be able to continue to utilize increasingly advanced off-the-shelf capabilities to target specific individuals with highly targeted messaging campaigns."[29]

Pursuant to these concerns, several steps have been taken. FEC Commissioner Ellen Weintraub has called for the end of precise microtargeting practices by limiting targeting to one sub-level below the level of the election.[30] Twitter has banned political advertisements in their entirety. And in November 2019, Google announced that in cases of political ads, it will prohibit microtargeting beyond age, gender, and location.[31] Despite pressure, Facebook continues to enable microtargeting.[32]

Arguments Against

Microtargeting is not an inherently malign practice. Consumer and civil rights advocacy organizations that serve the interests of marginalized communities rely on precise microtargeting to reach audiences comprising their specific constituent voting demographics, such as people of color or the LGBTQ+ community, and to grow their platforms and voices. Restrictions on microtargeting thus may have collateral impacts that do not serve the public interest. A similar argument could apply to up-start politicians without large donor backing or name-recognition who rely on precise targeting for fundraising.

As a practical matter, such a proposal (at least in the United States) may face a steep uphill battle in Congress, given that politicians themselves benefit from and use microtargeting to reach constituents. In addition, a ban on microtargeting may face First Amendment challenges, meaning this action should be voluntary, as recommended by Federal Election Commissioner Ellen Weintraub.[33]

Finally, even if targeting were limited to location, gender, and age, this information could be enough, especially given advances in artificial intelligence and machine learning, to yield precision microtargeting even without further data collection.

Arguments in Favor

Restrictions on microtargeting would enhance the transparency and accountability of the political process and increase incentives for truthfulness in ads. Political advertisers would have to appeal to a wider audience, limiting their ability to harness and increase divisiveness.

For the big tech companies, limiting microtargeting of political ads would result in a practice similar to what Google is doing already, with a minimal effect on Twitter.

Finally, limitations on microtargeting are consistent with what most Americans desire. According to a Gallup poll funded by the Knight Foundation, 72 percent of Americans prefer not to be targeted at all, and another 6 percent are only comfortable with ad targeting that uses broader information such as gender, age, or ZIP code.

Final Recommendation

Congress should adopt regulatory policy that requires platforms to disclose information about the audience targeted by a given ad in the context of political advertising, as has been proposed in the *Honest Ads Act* introduced in the Senate,[34] as well as California legislation.[35]

Congress should consider regulatory policy that limits the amount and quantity of data that can be used in political ads; contemplates a minimum size for target audiences; and provides an opt-in choice for users to receive microtargeted ads.[36]

Such policy should include a carve-out for ad campaigns that facilitate the democratic process and simply promote non-partisan voter participation with accurate information about when and where to vote (for example, voter registration campaigns, and get-out-the-vote initiatives).

6. Pass legislation prohibiting foreign individuals and governments from purchasing any election-related ads, similar to the bipartisan *PAID AD Act*

Background

Although current law prohibits foreign governments and nationals from contributing directly to campaigns, they are still legally allowed to purchase ads supporting or opposing a candidate or political issue on social media at any time, and on TV up to 60 days before a general election or 30 days before a primary election.[37] Current restrictions on foreign political ads codified in the *Bipartisan Campaign Reform Act* of 2002 also do not restrain, according to the Supreme Court ruling in Bluman v. FEC, "foreign nationals from speaking out about issues or spending money to advocate their views about issues. It restrains them only from a certain form of expressive activity closely tied to the voting process—providing money for a candidate or political party or spending money in order to expressly advocate for or against the election of a candidate."[38] This would seemingly allow a foreign entity to purchase ads, transparently or otherwise, about key election issues, including ballot measures. In Maine, for example, Canadian government-owned electrical company Hydro-Quebec spent more than six million dollars during the 2020 election period in support of a ballot measure on a proposed transmission line through the state.[39] Currently, this ad spend is entirely legal, and likely represents a small fraction of foreign-purchased ads meant to influence American voters across the country.

Arguments Against

In the context of foreign disinformation campaigns, political ads purchased by foreign powers are a relatively minor concern compared to organic content. Moreover, a U.S. ban on foreign-purchased issue and political ads would likely inspire other governments to enact similar legislation (if such legislation does not exist); or, in a more likely scenario, would lead to tech platforms instituting a blanket ban on foreign-purchased ads during election periods. This could restrict the ability of U.S. or EU individuals, companies, or government actors to transparently support or oppose legislation in other countries. This could also have unintended consequences in the form of restricting genuine and civic-minded voter education efforts. A case study from the European Union is illustrative here. For a window during the 2019 EU Parliamentary elections, Facebook banned foreign-run political ads on its platform. The upshot was that organic get-out-the-vote efforts from Brussels were blocked in other member states due to their "foreign" origin. Such policies could therefore harm efforts to promote democracy, human rights, and civic participation in other countries.[40]

Arguments in Favor

Banning election-related ads would close a loophole exploited by the Internet Research Agency in 2016, when they purchased more than $100,000 in political and issue ads on Facebook. It would also prevent foreign governments and foreign entities from purchasing ads in support of or in opposition to ballot measures or other referendums. Currently, foreign governments, individuals, and corporate entities can legally purchase, with few restrictions, ads that directly or indirectly influence American voters. This is particularly true online, where rules

for digital ads continue to lag behind offline media.

Final Recommendation

Legislation that involves government restrictions on any ad with political implications is an overly blunt tool that could curtail the ability of multinational corporations or foreign individuals to transparently lobby for or against a range of issues. That said, there is a clear need for enhanced disclosures in political or issue ads purchased by foreign entities. We therefore recommend that such ads carry clear, prominently displayed disclosures specifically identifying whether a foreign entity sponsored the ad, in much the way that Twitter now clearly labels state media accounts. In addition, we recommend that the pre-election blackout period for foreign ads that exist offline be extended to the online space. Finally, this legislation should include a carve-out for foreign-purchased ads that promote democratic participation or truthful information about how, when, and where to vote.

About the Authors

Dipayan Ghosh, Ph.D. is the co-director of the Digital Platforms & Democracy Project at the Shorenstein Center on Media, Politics and Public Policy at the Harvard Kennedy School and faculty at Harvard Law School. A computer scientist by training, Ghosh previously worked at Facebook, where he led strategic efforts to address privacy and security issues. Prior, he was a technology and economic policy advisor at the White House during the Obama administration. Named to the Forbes 30 Under 30, he received a Ph.D. in electrical engineering & computer science from Cornell University, an MBA from the Massachusetts Institute of Technology, and completed post-doctoral work at the University of California, Berkeley.

Lindsay Gorman is the Emerging Technologies Fellow at the German Marshall Fund's Alliance for Securing Democracy and a consultant for Schmidt Futures. Lindsay has spent over a decade at the intersection of technology development and national security policy, including in the Office of U.S. Senator Mark Warner, the White House Office of Science and Technology Policy, and the National Academy of Sciences. In the latter post, she supported the Committee on International Security and Arms Control in track II nuclear and cyber security dialogues with Chinese and Russian experts. A physicist and computer scientist by training, she previously ran a technology consulting firm, Politech Advisory, advising start-ups and venture capital and has developed cybersecurity tools in Silicon Valley. Her research focuses on understanding and crafting a transatlantic response China's techno-authoritarian rise, from 5G and the future internet to information manipulation and censorship. Her technical expertise lies in artificial intelligence, statistical machine learning, and quantum materials. Lindsay holds an A.B. in physics from Princeton University, where she graduated magna cum laude, and a M.S. in applied physics from Stanford University.

Bret Schafer is the Alliance for Securing Democracy's Media and Digital Disinformation Fellow. As an expert in computational propaganda, he has appeared in the New York Times, USA Today, the Wall Street Journal, and the Washington Post, and he has been interviewed on NPR, MSNBC, CNN, Al Jazeera, and CBS and BBC radio. Prior to joining GMF, he spent more than ten years in the television and film industry, including stints at Cartoon Network and as a freelance writer for Warner Brothers. He also worked in Budapest as a radio host, in Berlin as a semi-professional baseball player in Germany's Bundesliga, and in Moscow as an intern in the Public Affairs Section at the U.S. Embassy in Russia. He has a BS in communications with a major in radio/television/film from Northwestern University, and a master's in public diplomacy from the University of Southern California, where he was the editor-in-chief of Public Diplomacy Magazine.

Clara Tsao is a non-resident fellow at the Alliance for Securing Democracy. She is an online disinformation expert and a civic tech entrepreneur, who recently co-founded the Trust & Safety Professional Association and the Trust & Safety Foundation to support the global community of professionals who develop and enforce principles and policies that define acceptable behavior and content online. Clara is also a non-resident senior fellow at the Atlantic Council's Digital Forensic Research Lab. Her previous roles include CTO at the U.S. Department of Homeland Security's Countering Foreign Influence Task Force and the interagency U.S. Countering Violent Extremism Task Force and Senior Advisor for Emerging Technology at the Cybersecurity Infrastructure Security Agency. She has spent a decade working in the technology industry across global teams at Microsoft, Apple, Sony PlayStation, AT&T, and also as a Google and Mozilla Technology Policy Fellow. Clara is also the Board Chair and President of the White House Presidential Innovation Fellows Foundation and a Senior Advisor at Tech Against Terrorism.

Special Thanks

The authors would like to thank the following individuals for their participation in an expert roundtable session in July 2020 and/or their feedback to earlier drafts of this paper:

David Agranovich

Claire Atkin

Dobromir Cias

Renee Diresta

Yael Eisenstat

Hala Furst

John Haigh

Nandimi Jammi

Julian Jarusch

Sam Jeffers

Ariel Fox Johnson

Karen Kornbluh

Sinead O'Sullivan

Robert Schaul

Danny Sepulveda

The authors would also like to thank Bradley Hanlon and Christina Revilla for their help organizing the roundtable.

Endnotes

1 Bradley Hanlon, "A Long Way To Go: Analyzing Facebook, Twitter, and Google's Efforts to Combat Foreign Interference," Alliance for Securing Democracy, December 19, 2018.
2 Daisuke Wakabayashi, "Google Finds Accounts Connected to Russia Bought Election Ads," New York Times, October 9, 2017.
3 "After Twitter, Google cracks down on political ads," CBS News, November 21, 2019.
4 Patrick B. Pexton, "Graham, Klobuchar introduce internet ads bill to boost transparency," Roll Call, May 8, 2019.
5 Scott Shane, "These Are the Ads Russia Bought on Facebook in 2016," New York Times, November 1, 2017.
6 Nathaniel Gleicher, "Removing Coordinated Inauthentic Behavior from Russia," Facebook, January 17, 2019.
7 "The Quarter Billion Dollar Question: How is Disinformation Gaming Ad Tech?" Global Disinformation Index, September 2019.
8 Paul P. Murphy, Kaya Yurieff and Gianluca Mezzofiore, "Exclusive: YouTube ran ads from hundreds of brands on extremist channels," CNN, April 20, 2018.
9 Mark Scott, "US presidential campaign ads run on YouTube content from Russia media, white supremacists," Politico, July 7, 2020.
10 Mark Sullivan, "Facebook expands ad transparency beyond politics: Here's what's new," Fast Company, March 28, 2019.
11 Bradley Hanlon.
12 Kate Conger, "Twitter Will Ban All Political Ads, C.E.O. Jack Dorsey Says," New York Times, October 30, 2019.
13 Jessica Brandt and Bradley Hanlon, "Online Information Operations Cross Platforms. Tech Companies' Responses Should Too," Lawfare, April 26, 2019.
14 "Exchange Act Reporting and Registration," U.S. Securities and Exchange Commission, accessed November 28, 2020.
15 Ian Leslie, "Advertisers trapped in an age of online obfuscation," Financial Times, February 28, 2017.
16 Allison Schiff, "Inside Uber's Fraud Suit Against Phunware," Ad Exchanger, August 26, 2019.
17 See checkmyads.org.
18 "So *that's* how Breitbart is still making money," Branded, Substack, July 22, 2020.
19 Zach Edwards, "Breitbart.com is Partnering with RT.com & Other Sites via Mislabeled Advertising Inventory," Medium, July 22, 2020.
20 See, for example, "Anti-Spoofing Protection in EOP," Microsoft 365, November, 17, 2020.
21 "Two International Cybercriminal Rings Dismantled and Eight Defendants Indicted for Causing Tens of Millions of Dollars in Losses in Digital Advertising Fraud," The United States Attorney's Office, Eastern District of New York, November 27, 2018.
22 Nicole Perrin, "Digital Ad Fraud 2019," Insider Intelligence, February 6, 2019.
23 Sheila Dang, "Exclusive: Facebook ad boycott campaign to go global, organizers say," Reuters, June 28, 2020.
24 Antonio Garcia Martinez, "How Trump Conquered Facebook—Without Russian Ads," Wired, February 23, 2018.
25 Tim Lau, "The Honest Ads Act Explained," The Brennan Center for Justice, January 17, 2020.
26 Phil Howard, "A Way to Detect the Next Russian Misinformation Campaign," New York Times, March 27, 2019.
27 Ellen Goodman, Karen Kornbluh, and Eli Weiner, "Safeguarding Democracy Against Disinformation," German Marshall Fund of the United States, March 24, 2020.
28 Goodman, Kornbluh, & Weiner.
29 "Report of the Select Committee on Intelligence, United States Senate, on Russian Active Measures Campaigns and Interference in the 2016 U.S. Election, Volume 2: Russia's Use of Social Media," United States Senate, October 2019.
30 Ellen L. Weintraub, "Don't abolish political ads on social media. Stop microtargeting." Washington Post, No-

vember 1, 2019.
31 Scott Spencer, "An update on our political ads policy," Google, November 20, 2019.
32 Kaili Lambe and Becca Ricks, "The basics on microtargeting and political ads on Facebook," Mozilla Foundation, January 14, 2020.
33 Bill Chappell, "FEC Commissioner Rips Facebook Over Political Ad Policy: 'This Will Not Do,'" NPR, January 9, 2020.
34 "The Honest Ads Act," Website of U.S. Senator Mark R. Warner, May 2019.
35 "Assembly Bill 2885 – False campaign speech and online platform disclosures," California State Assembly, July 30, 2020.
36 Julian Jaursch, "Why EU must limit political micro-targeting," EU Observer, July 22, 2020.
37 52 U.S. Code § 30121.Contributions and donations by foreign nationals
38 Benjamin Bluman, Et al v. Federal Election Commission, The Supreme Court, September 1, 2011.
39 Robbie Feinberg, "Maine Lawmakers Call On Hydro-Qubec[sic] To Stop Campaign On CMP Transmission Line Ballot Referendum," Maine Public, July 29, 2020.
40 Mark Scott, Laura Kayali, and Maia De La Baume, "Facebook to Cave to EU Pressure after Row Over Political Ad Rules," Politico EU, April 18, 2019.