



UvA-DARE (Digital Academic Repository)

Platform ad archives: promises and pitfalls

Leerssen, P.; Ausloos, J.; Zarouali, B.; Helberger, N.; de Vreese, C.H.

Published in:
Internet Policy Review

DOI:
[10.14763/2019.4.1421](https://doi.org/10.14763/2019.4.1421)

[Link to publication](#)

Creative Commons License (see <https://creativecommons.org/use-remix/cc-licenses/>):
CC BY

Citation for published version (APA):
Leerssen, P., Ausloos, J., Zarouali, B., Helberger, N., & de Vreese, C. H. (2019). Platform ad archives: promises and pitfalls. *Internet Policy Review*, 8(4). <https://doi.org/10.14763/2019.4.1421>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

UvA-DARE is a service provided by the library of the University of Amsterdam (<http://dare.uva.nl>)



Platform ad archives: promises and pitfalls

Paddy Leerssen

Institute for Information Law (IViR), University of Amsterdam, Netherlands, p.j.leerssen@uva.nl

Jef Ausloos

Institute for Information law (IViR), University of Amsterdam, Netherlands

Brahim Zarouali

Department of Communication Science, University of Amsterdam, Netherlands

Natali Helberger

Institute for Information Law (IViR), University of Amsterdam, Netherlands

Claes H. de Vreese

Department of Communication Science, University of Amsterdam, Netherlands

Published on 09 Oct 2019 | DOI: 10.14763/2019.4.1421

Abstract: This paper discusses the new phenomenon of platform ad archives. Over the past year, leading social media platforms have installed publicly accessible databases documenting their political advertisements, and several countries have moved to regulate them. If designed and implemented properly, ad archives can correct for structural informational asymmetries in the online advertising industry, and thereby improve accountability through litigation and through publicity. However, present implementations leave much to be desired. We discuss key criticisms, suggest several improvements and identify areas for future research and debate.

Keywords: Advertising, Micro-targeting, Platforms, Transparency, Politics

Article information

Received: 30 Apr 2019 **Reviewed:** 16 Sep 2019 **Published:** 09 Oct 2019

Licence: Creative Commons Attribution 3.0 Germany

Funding: This work was funded by the RPA Personalized Communications at the University of Amsterdam, which is a joint initiative between the Institute for Information Law (IViR) and the School of Communication Research (ASCoR).

Competing interests: The author has declared that no competing interests exist that have influenced the text.

URL: <http://policyreview.info/articles/analysis/platform-ad-archives-promises-and-pitfalls>

Citation: Leerssen, P. & Ausloos, J. & Zarouali, B. & Helberger, N. & de Vreese, C. H. (2019). Platform ad archives: promises and pitfalls. *Internet Policy Review*, 8(4). DOI: 10.14763/2019.4.1421

INTRODUCTION

In 2018, the online platforms Google, Facebook and Twitter all created political ad archives: publicly accessible databases with an overview of political advertisements featured on their services. These measures came in response to mounting concerns over a lack of transparency and accountability in online political advertising, related to illicit spending and voter manipulation. Ad archives have received widespread support in government and civil society. However, their present implementations have also been criticised extensively, by researchers who find their contents to be incomplete or unreliable. ¹ Increasingly, governments and civil society actors are therefore setting up their own guidelines for ad archive architecture – in some cases even binding legislation. Ad archive architecture has thus rapidly gained relevance for advertising law and policy scholars, both as a tool *for* regulation and as an object *of* regulation. ²

This article offers an overview of the ad archive governance debate, discussing the potential benefits of these tools as well as pitfalls in their present implementations. Section two starts with a basic conceptual and legal framework which describes the basic features of archives and applicable regulations, followed by a normative framework which discusses the potential benefits of ad archives in terms of transparency and accountability. Section three reviews the shortcomings of current ad archive initiatives, focusing on three core areas of ongoing debate and criticism. Firstly, we discuss scoping: ad archives have faced difficulty in defining and identifying, at scale, what constitutes a “political advertisement”. Secondly, verifying: ad archives have proven vulnerable to inauthentic behaviour, particularly from ad buyers seeking to hide their true identity or the origin of their funding. Thirdly, targeting data: ad archives do not document in meaningful detail how ads are targeted or distributed. We propose several improvements to address these shortcomings, where necessary through public regulation. Overall, we argue that both legal scholars and communications scientists should pay close attention to the regulation of, and through, this novel and potentially powerful tool.

PROMISES: THE CASE FOR AD ARCHIVES

CONCEPTUAL FRAMEWORK: WHAT ARE ‘AD ARCHIVES’?

This paper focuses on ad archives, which are systems for the automated public disclosure of advertisements via the internet. The key examples are Facebook’s Ad Library, Google’s Advertising Transparency Report and Twitter’s Ad Transparency Center. These systems document the advertisement messages sold on the platform, as well as associated metadata (e.g., the name of the buyer, the number of views, expenditure, and audience demographics). These archives are public, in the sense that they are available without restriction to anyone with a working internet connection.

In practice, the major ad archives have focused on documenting political advertisements, rather than commercial advertisements. Beyond this, they differ in important respects. Firstly, they differ significantly in how they define “political” advertising in order to determine what ads are included in the archive. The major archives also differ in how they verify their contents – particularly the identity of their ad buyers – and in terms of the metadata they publish related to ad targeting. Section three considers these questions of scoping, verifying and targeting in further detail.

The major ad archives went live in 2018. Facebook's archive was first announced in October 2017 and went live the next year in May 2018. Google and Twitter followed soon after. They initially focused exclusively on the United States, but they have since gradually expanded their efforts. Facebook and Twitter's archives now offer worldwide coverage, although certain functions are still regionally restricted. Google covers only the US, the European Union and India (Google, 2019a).

In theory, ad archives can be created not only by platform intermediaries but also by a range of other actors, including advertisers, academics or NGOs. For instance, political parties can maintain their own online database documenting their political advertisements, as has been proposed in the Netherlands (Netherlands Ministry of the Interior, 2019). As early as 2012, Solon Barocas argued for a centralised non-profit database, or 'clearing house', for political ads (Barocas, 2012). The London School of Economic's Truth and Trust Commission proposes that the government administer a central database, or "political advertising directory" (Livingstone, 2018). The investigative journalists of ProPublica have maintained a public database of Facebook ads which they crowd-sourced from a group of volunteers (Merrill & Tobin, 2019). While we do not discount these approaches, our discussion focuses on platform-operated archives, since these have recently attracted the most widespread traction in policy and practice.

LEGAL FRAMEWORK: WHY ARE PLATFORMS BUILDING ARCHIVES?

Formally speaking, the major platform ad archives are self-regulatory measures. But they emerged in response to significant public pressure from the ongoing "techlash" (Smith, 2018; Zimmer, 2019). These "voluntary" efforts are therefore best understood as an attempt to stave off binding regulation (Wagner, 2018). Indeed, platforms have no immediate commercial incentive to offer transparency in their advertising practices. The role of public regulation, or at least the threat thereof, is therefore essential in understanding the development of ad archives (see Vlassenroot et al., 2019). Below we offer an overview of key policy developments.

Both platforms and policymakers present ad archives as a means to improve accountability in online political advertising (e.g., Goldman, 2017; Warner, 2017). Political advertising in legacy media has historically been regulated in various ways, to prevent undue influence from concentrated wealth on public discourse. Online advertising is placing new pressure on these legacy regimes. In many cases, the language of existing law has simply not been updated to apply online. Furthermore, online political micro-targeting has unique affordances that can enable new types of harms demanding entirely new regulatory responses. For instance, platform advertising services lower the barrier to buying ads across borders, and to buy ads under false or misleading identities. Furthermore, micro-targeting technology, which enables advertisers to target highly specific user segments based on personal data analytics, can enable novel methods of voter deception, manipulation and discrimination (Borgesius et al., 2018; Chester & Montgomery, 2017). For instance, targeted advertising can enable politicians to announce different or even conflicting political programmes to different groups, thereby fragmenting public discourse and making it more difficult to hold politicians accountable to their electoral promises (Bodó, Helberger, & de Vreese, 2017; Borgesius et al., 2018). Targeted advertising can also enable discrimination between voter groups, both intentionally through advertisers' targeting decisions and unintentionally through undocumented algorithmic biases (Boerman, Kruikemeier, & Borgesius, 2017; Ali et al., 2019).

These concerns about online advertising are compounded by the fact that the online advertising ecosystem is difficult to monitor, which undermines efforts to identify, diagnose and remedy potential harms (Chester & Montgomery, 2017). This opacity is due to personalisation:

personalised advertisements are invisible to everyone except the specific users they target, hiding them from observation by outsiders (Guha, Cheng, & Francis, 2010). As Benkler, Faris and Roberts observe, this distinguishes online advertisers from mass media advertisers, who necessarily acted “in the public eye”, thus “suffering whatever consequences” a given message might yield outside of its target audience (Benkler, Faris, & Roberts, 2018, p. 372). As a result, the online advertising ecosystem exhibits structural information asymmetries between, on one side, online platforms and advertisers, and on the other, members of the public who might hold them accountable. Researchers can potentially resort to data scraping methods, but these suffer from severe limitations and are vulnerable to interference by the platforms they target (Bodó et al., 2018; Merrill & Tobin, 2019). Accordingly, targeted advertising creates structural information asymmetries between advertisers and their publics.

These concerns over online political advertising took centre stage in the “techlash”, which followed the unexpected outcomes of the 2016 Brexit referendum and US presidential elections. In the UK, the Vote Leave campaign was accused of deceptive messaging, and violations of data protection law and campaign spending law in their political micro-targeting activities (Merrick, 2018; Waterson, 2019a). In the US, ad spending from Russian entities such as the Internet Research Agency raised concerns about foreign election interference. In both countries, Facebook shared selected advertising data sets in response to parliamentary investigations (Lomas, 2019; Shane, 2017). But these came well over a year after the events actually took place – driving home the general lack of transparency and accountability in the advertising ecosystem. Similar controversies have also played out subsequent elections and referenda, such as the Irish abortion referendum of 2018 which drew an influx of foreign pro-life advertisements (Hern, 2018). The actual political and electoral impact of these ad buys remains debatable (e.g., MacLeod, 2019; Benkler, Faris, & Roberts, 2018). But in any case, these developments drew attention to the potential for abuse in targeted advertising, and fuelled the push for more regulation and oversight in this space.

Ad archives have formed a key part of the policy response to these developments. The most prominent effort in the US is the Honest Ads Act, proposed on 19 October 2017, which would require online platforms to “maintain, and make available for online public inspection in machine readable format, a complete record of any request to purchase on such online platform a qualified political advertisement” (Klobuchar, Warner, & McCain, 2017, Section 8(a)(j)(1)(a)). This bill has not yet passed (Montellaro, 2019). But only several days after its announcement, Facebook declared its plans to voluntarily build an ad archive, which would largely conform to the same requirements (Goldman, 2017). Google and Twitter followed suit the next year.

Since 2018, governments have started developing binding legislation on ad archives, often with resistance from platforms. Canada’s Elections Modernization Act of December 2018 compels platforms to maintain public registers of political advertising sold through their service. Facebook and Twitter have sought to comply with these measures, but Google instead responded by discontinuing the sale of political advertisements in this jurisdiction altogether (Cardoso, 2019). Similarly, the State of Washington’s Public Disclosure Commission attempted to regulate ad archives by requiring advertisers publicly disclose political ads sold in the state (Sanders, 2019). In this case, both Google and Facebook have refused to comply with the disclosure rules and instead banned political advertising in this region (Sanders, 2019). Citing federal intermediary liability law, the Communications Decency Act of 1996, Facebook contended it was immune to any liability for political advertising content (Sanders, 2019). Some reporters also claim that Facebook has lobbied to kill the Honest Ads Act, despite publicly claiming to support regulation and implement its requirements voluntarily (Timmons &

Kozłowska, 2018).

Europe is also poised to regulate ad archives. In the run-up to the EU elections of May 2019, the European Commission devised the Code of Practice on Disinformation, which is not a binding law but rather a co-regulatory instrument negotiated with major tech companies including Google, Facebook, Twitter, Microsoft and Mozilla. ³ By signing the Code, these companies have committed to a range of obligations from fact-checking and academic partnerships to the creation of ad archives (European Commission, 2018, Section II.B.). Furthermore, leaked documents from the European Commission show that political advertisements will receive particular attention in the upcoming reform of digital services rules (Fanta & Rudl, 2019). Member states are also exploring the regulation of ad archives. In the UK and the Netherlands, parliamentarians have expressed support for further regulation in, respectively, a [parliamentary resolution](#) and a [committee report](#) (Parliament of the Netherlands, 2019; House of Commons Select Committee on Digital, Culture, Media and Sport, 2019). France has passed a [binding law](#) requiring the public disclosure of payments received for political advertisements – if not a comprehensive regulation of ad archives *per se* (Republic of France, 2018).

Ad archives exist alongside a number of other proposals for regulating targeted advertising. One popular measure is installing user-facing disclaimers, intended to inform audiences about e.g., the identity of the advertisers, the source of their funding, and/or the reason why they are being targeted. Another approach is to regulate funding, e.g., through spending limits, registration requirements, or restrictions on foreign advertising. Finally, targeting technology and the use of personal data can also be regulated. Some combination of these measures is found in, *inter alia*, the US Honest Ads Act, the EU's Code of Practice, Canada's Elections Modernization Act, and France and Ireland's new election laws. The EU's General Data Protection Regulation (GDPR) is also a highly relevant instrument, since it grants users' information rights, and constrains the ability for advertisers to use personal data for ad targeting purposes (Bodó, Helberger, & De Vreese, 2017).

Of course, present ad archive initiatives are far from uniform. Definitions of e.g., the relevant platforms, disclosure obligations and enforcement mechanisms all differ. An exhaustive comparative analysis of these differences would exceed the scope of this paper. The second half of this paper discusses how these policy initiatives differ on some of the key design issues outlined above (scoping, verifying, and targeting data), and how the major platforms have responded to their demands. First, we discuss the policy principles driving this new wave of regulation.

NORMATIVE FRAMEWORK: WHAT ARE THE POLICY GROUNDS FOR AD ARCHIVES?

Ad archive initiatives have typically been presented in terms of 'transparency and accountability', but these are notoriously vague terms. The concrete benefits of ad archives have not been discussed in much depth. To whom do ad archives create accountability, and for what? The answer is necessarily somewhat abstract, since ad archives, being publicly accessible, can be used by a variety of actors in a variety of accountability processes. Indeed, this diversity is arguably their strength. Other advertising transparency measures have focused on particular groups of stakeholders, such as user-facing disclaimers, third party audits or academic partnerships. Ad archives, by contrast, can enable monitoring by an unrestricted range of actors, including not only academics but also journalists, activists, government authorities and even rival advertisers – each with their own diverse capacities and motivations to hold advertising accountable. In this sense, ad archives can be seen as recreating, to some extent, the public

visibility that was inherent in mass media advertising and is now obfuscated by personalisation (see above). Broadly speaking, this public visibility can be associated with two types of accountability: (a) accountability to the law, through litigation, and, (b) accountability to public norms and values, through publicity.

Ad archives can contribute to law enforcement by helping to uncover unlawful practices. Although online political advertising is not (yet) regulated as extensively as its mass media counterparts, it may still violate e.g., disclosure rules and campaign finance regulations. And, as discussed previously, new rules may soon be coming. Commercial advertising, for its part, may be subject to a range of consumer protection rules, particularly in Europe, and also to competition law, unfair commercial practice law and intellectual property law. Ad archives can allow users to proactively search for violations of these rules. Such monitoring could be done by regulators, but importantly also by third parties including commercial rivals, civil rights organisations, consumer protection organisations, and so forth. These third parties might choose to litigate independently, or simply refer the content to a competent regulator. Indeed, regulators often rely on such third party input to guide their enforcement efforts, e.g., in the form of hotlines, complaints procedures and public consultations. In most cases, litigation is likely to be straightforward and inexpensive, since most platforms operate notice and takedown procedures for the removal of unlawful advertising without the need for judicial intervention.⁴ Platforms can also remove advertising based on their own community standards, even if they do not violate any national laws. In this light, ad archives can contribute to enforcement in a broad sense, including not only public advertising laws but also platforms' private standards, and relying not only on public authorities but on any party with the time and interest to flag prohibited content.

In addition to litigation, ad archives also facilitate publicity about advertising practices, which can serve to hold platforms accountable to public norms and values. Journalists, researchers and other civil society actors can draw on archives to research and publicise potential wrongdoings that might previously have flown under the radar. For instance, the US media has a strong tradition of analysing and fact-checking television campaign ads; ad archives could help them do similar coverage of online political micro-targeting. Such publicity may encourage platforms and/or advertisers to self-correct and improve their advertising standards, by raising the threat of reputational harm. And failing such a private ordering response, publicity can also provide an impetus for new government interventions. In these ways, ad archives can contribute not only to the enforcement of existing laws, but also to informed public deliberation, and thus to the articulation and enforcement of public norms and values (see Van Dijck, Poell, & de Waal, 2018). Such publicity effects may be especially important in the field of online political advertising, since, as discussed, this space remains largely unregulated under existing laws, and raises many novel policy questions for public deliberation.

In each case, it is important to note the factor of deterrence: the mere *threat* of publicity or litigation may already serve to discipline unlawful or controversial practices. Even for actors who have not yet faced any concrete litigation or bad publicity, ad archives could theoretically have a disciplinary effect. In this sense, a parallel can be drawn with the concept of the Panopticon, as theorised in surveillance studies literature; subjects are disciplined not merely through the fact of observation, but more importantly through the pervasive possibility of observation (Foucault, 1977; Lyon, 2006). Put differently, Richard Mulgan describes this as the *potentiality* of accountability; the possibility that one "may be called to account for anything at any time" (Mulgan, 2000, p. 564). Or, as the saying goes: The value in the sword of Damocles is not that it drops, but that it hangs (e.g., *Arnett v. Kennedy*, 1974).

Of course, these accountability processes depend on many other factors besides transparency alone. Most importantly, ad archives depend on a capable and motivated user base of litigators (for law enforcement effects) and civil society watchdogs (for publicity effects). For publicity effects, these watchdogs must also be sufficiently influential to create meaningful reputational or political risks for platforms (see Parsons, 2019; Wright & Rwabizambuga, 2006). These conditions can certainly not be assumed; which researchers are up to the task of overseeing this complex field, and holding its powerful players to account? This may call for renewed investment in our public watchdogs, including authorised regulators as well as civil society. Ad archives might be a powerful tool, but they rely on competent users.

Finally, of course, the above analysis also assumes that ad archives are designed effectively, so as to offer meaningful forms of transparency. As we discuss in the following section, present implementations leave much to be desired.

PITFALLS: KEY CHALLENGES FOR AD ARCHIVE ARCHITECTURE

Having made the basic policy case for the creation of ad archives, we now discuss several criticisms of current ad archive practice. Firstly, we discuss the issue of scoping: which ads are included in the archive? Second, verifying: how do ad archives counteract inauthentic behaviour from advertisers and users? Third, targeting: how do ad archives document ad targeting practices? Each of these issues can create serious drawbacks to the research utility of ad archives, and deserve further scrutiny in future governance debates.

Ad archive architecture is very much a moving target, so we emphasise that our descriptions represent a mere snapshot. Circumstances may have changed significantly since our time of writing. Accordingly, the following is not intended as an exhaustive list of possible criticisms, but rather as a basic assessment framework for some of the most controversial issues. For instance, one important criticism of ad archives which we do *not* consider in detail is the need for automated access through application programming interfaces (APIs). When ad archive data is exclusively available through browser-based interfaces, this can make it relatively time-consuming to perform large-scale data collection. To enable in-depth research, it is clear that ad archives must enable such automated access. Until recently, Facebook did not offer public API access to their ad archive data (Shukla, 2019). And once the API was made publicly accessible, it quickly appeared to be so riddled with bugs as to be almost unusable (Rosenberg, 2019). As noted by Laura Edelson, these API design issues are not novel or intractable from a technical perspective, but eminently “fixable”, and thus reflect sub-standard implementation on the part of Facebook (Rosenberg, 2019). In response, Mozilla, together with a coalition of academics, has drafted a list of design principles for ad archive APIs (Mozilla, 2019). Such public, automated access can be seen as a baseline condition for effective ad archive policy. What then remains are questions about the *contents* of the archive, which include scoping, verifying and targeting.

SCOPING: WHAT ADS ARE INCLUDED IN THE ARCHIVE?

A key design question for ad archives is that of scope: what ads are included in the archive? First, we discuss the concept of “political” advertising, which is the central scoping device in most existing initiatives and has led to many implementation challenges. Second, we discuss the attempts to exempt news reporting from political ad archives.

“Political” ad archives: electoral ads v. issue ads v. all ads?

Ad archive initiatives, both self-regulatory and governmental, have emphasised “political” advertising rather than commercial advertising. However, their precise interpretations of this concept differ significantly. Below we discuss these differing approaches and relevant policy trade-offs.

The main dividing line in existing political ad archives is between issue ads and electoral ads (or “campaign ads”). “Election ads” explicitly reference an election or electoral candidate, whereas “issue ads” reference a topic of national importance. Google focuses exclusively on election ads, whereas Facebook and Twitter also include issue ads in certain jurisdictions, and even non-political ads. Most public policy instruments also focus on issue ads, including the US Honest Ads Act and the EU Code of Practice. There is good reason to include issue ads, since they have been central to recent controversies. During the 2016 US election, for instance, foreign actors such as the Russian-controlled Internet Research Agency advertised on divisive issues such as racial politics, sexual politics, terrorism, and immigration, in an apparent attempt to influence the election (Howard et al., 2018). An approach which focuses on election ads would fail to address such practices.

However, the drawback of “issue ads” as a scoping device, is that the concept of a political “issue” is broad and subjective, and makes it difficult for archive operators to develop actionable definitions and enforce these in practice. Google, in its implementation reports for the EU’s Code of Practice, reported difficulties in developing a workable definition of a “political issue” (Google, 2019). The European Commission later lamented that “Google and Twitter have not yet reported further progress on their policies towards issue-based advertising” (European Commission, 2019). In Canada, where the Election Act also requires the disclosure of issue-based ads, Google has claimed that they are simply unable to comply with disclosure requirements (Cardoso, 2019). These difficulties might explain why the company announced plans, as discussed previously, to ban political advertising entirely for Canadian audiences during election periods.

Yet these attempts to ban political advertising, as an alternative to disclosure, beg the question whether platforms can actually enforce such a ban. After all, the platforms themselves admit they struggle to identify political ads in the first place. Simply declaring that political ads are prohibited will not guarantee that advertisers observe the ban and refrain from submitting political content. Could platforms then still be liable for a failure to disclose? Here, a tension emerges between ad archive regulation and intermediary liability laws, which typically immunise platforms for (advertising) content supplied by their users. Canada, Europe and the US all have such laws, although their precise scope and wording differ. Indeed, Facebook has argued that it is immunised against Washington State’s disclosure rules based on US federal intermediary liability law – the Communications Decency Act of 1996 (Sanders, 2018a). Similarly, the EU’s intermediary safe harbours, which prohibit “proactive monitoring obligations” imposed on platforms (e-Commerce Directive 2000/31/EC, Article 15). Such complex interactions with intermediary liability law should be taken into account in ongoing reforms.

Compared to Google, Facebook is relatively advanced in its documentation of issue ads. But that company too has faced extensive criticism for its approach. The company employs approximately 3,000-4,000 people in reviewing ads related to politics or issues, using “a combination of artificial intelligence (AI) and human review”, and is estimated to process upwards of a million ad buyers per week in the US alone (Matias, Hounsel, & Hopkins, 2018).

Facebook’s website offers a list of concrete topics which they consider “political issues of national importance”, tailored to the relevant jurisdiction. The [US list of political issues](#) contains 20 entries, including relatively specific ones such as “abortion” and “immigration”, but also relatively broad and ambiguous ones such as “economy” and “values” (Facebook, 2019). The [EU list](#) contains only six entries so far, including “immigration”, “political values” and “economy” (Matias, Hounsel, & Hopkins, 2018).

Despite these efforts, research suggests that Facebook’s identification of political issue ads is error-prone. Research from Princeton and Bloomberg showed that a wide range of commercial ads are at risk of being mislabeled as political, including advertisements for e.g., national parks, veteran’s day celebrations, and commercial products that included the words “bush” or “clinton” (Frier, 2018; Hounsel et al., 2019). Conversely, data scraping research by ProPublica shows that Facebook failed to identify political issue ads on such topics as civil rights, gun rights, electoral reform, anti-corruption, and health care policy (Merrill & Tobin, 2019). These challenges are likely to exacerbate as platforms expand their efforts beyond the United States to regions such as Africa and Europe, which contain far greater political and linguistic diversity and fragmentation. Accordingly, further research is needed to determine whether the focus on issue ads in ad archives is appropriate. It may appear in future that platforms are able to refine their processes and identify issue ads with adequate accuracy and consistency. But given the major scaling challenges, the focus on issue ads may well turn out to be impracticable.

In light of the difficulties with identifying “issue ads”, one possible alternative would be to simply include all ads *without* an apparent commercial objective. In other words, a definition *a contrario*. This approach could capture the bulk of political advertising, and would avoid the difficulties of identifying and defining specific political “issues”. Such an approach would likely be more scalable and consistent than the current model, although this might come at the cost of increased false positives (i.e., a greater overinclusion of irrelevant, non-political ads in the archive).

Another improvement could be to publish *all* advertisements in a comprehensive archive, regardless of their political or commercial content (Howard, 2019). This would help third parties to independently evaluate platforms’ flagging processes for political ads, and furthermore to research political advertising according to their *own* preferred definitions of the “political”. This is what Twitter does in its Ad Transparency Center: the company still takes steps to identify and flag political advertisers (at least in the US), but users have access to all other ads as well (Twitter, 2019a). However, only political ads are accompanied by detailed metadata, such as ad spend, view count, targeting criteria, et cetera. Facebook, in an update from 29 March 2019, also started integrating commercial ads into its database (Shukla, 2019). Like Twitter, however, these ads are not given the same detailed treatment as political ads. In this light, Twitter and Facebook appear to be moving towards a *tiered* approach, with relatively more detail on a subset of political ads, and relatively less detail on all other ads.

Of course, a more fundamental advantage of comprehensive publication ads is that it extends the benefits of ad archives to *commercial* advertising. Commercial advertising has not been the primary focus of ad archive governance debates thus far, but here too ad archives could be highly beneficial. A growing body of evidence indicates that online commercial ad delivery raises a host of legal and ethical concerns, including discrimination and manipulation (Ali et al., 2019; Boerman, Kruikemeijer, & Borgesius, 2017). Furthermore, online advertising is also subject to a range of consumer protection laws, including child protection rules and prohibitions on unfair and deceptive practices. With comprehensive publication, ad archives could contribute to

research and reporting on such issues, especially if platforms abandon their tiered approach and start publishing more detailed metadata for these ads.

Platforms may not be inclined to implement comprehensive ad archives since, as discussed, their commercial incentives may run counter to greater transparency. But from a public policy perspective, there appear to be no obvious drawbacks to comprehensive publication, at least as a default rule. If there are indeed grounds to shield certain types of ads from public archives – though we see none as of yet – such cases could also be addressed through exemption procedures. The idea of comprehensive ad archives therefore warrants serious consideration and further research, since it promises to benefit the governance of both commercial *and* political advertising.

Exemptions for news reporting

Some ad archive regimes offer exemptions for news publishers and other media actors. News publishers commonly use platform advertising services to promote their content, and when this content touches on political issues it can therefore qualify as an issue ad. Facebook decided to exempt news publishers from their ad archive in 2018, following extensive criticism from US press industry trade associations, who penned several open letters criticising their inclusion in ad archives. They argued that “[t]reatment of quality news as political, even in the context of marketing, is deeply problematic” and that the ad archive “dangerously blurs the lines between real reporting and propaganda” (Carbajal et al., 2018; Chavern, 2018). Similar exemptions can now also be found in Canada’s Elections Modernization Act and in the EU Code of Practice (Leathern, 2019). However, the policy grounds for these exemptions are not particularly persuasive. There is little evidence to suggest, or reason to assume, that inclusion in ad archives would meaningfully constrain the press in its freedom of expression. Indeed, ad archive data about media organisations is highly significant, since the media are directly implicated in concerns about misinformation and electoral manipulation (Benkler, Faris, & Roberts, 2018). Excluding the media’s ad spending is therefore a missed opportunity without a clear justification.

VERIFYING: HOW DO ARCHIVES ACCOUNT FOR INAUTHENTIC BEHAVIOUR?

Another pitfall for ad archives is verifying their data in the face of fraud and other inauthentic behaviours. One key challenge is documenting ad buyers’ identities. Another is the circumvention of ad archive regimes by “astroturf”, sock puppets and other forms of native advertising. More generally, engagement and audience statistics may be inaccurate due to bots, click fraud and other sources of noise. As we discuss below, these pitfalls should serve as a caution to ad archive researchers, and as a point of attention for platforms and their regulators.

Facebook’s archive in particular has been criticised for failing to reliably identify ad buyers (e.g., Edelson et al., 2019). Until recently, Facebook did not verify the names that advertisers submitted for their “paid for by” disclaimer. This enabled obfuscation by advertisers seeking to hide their identity (Albright, 2018; Andringa, 2018; Lapowsky, 2018; O’Sullivan, 2018; Waterson, 2019). For instance, ProPublica uncovered 12 different political ad campaigns that had been bought in the name of non-existent non-profits, and in fact originated from industry trade organisations such as the American Fuel & Petrochemical Manufacturers (Merrill, 2018). Vice Magazine even received authorisation from Facebook to publish advertisements in the name of sitting US senators (Turton, 2018). More recently, Facebook has therefore started demanding proof of ad buyer identity in several jurisdictions, such as photo ID and notarised forms (Facebook, 2019b). Twitter and Google enforce similar rules (Google, 2019b; Twitter,

2019b). The Canadian Elections Modernization Act now codifies these safeguards by requiring platforms to verify and publish ad buyers' real names.

Such identity checks are only a first step in identifying ad buyers, however. Ad buyers wishing to hide their identity can still attempt to purchase ads through proxies or intermediaries. In theory, platforms could be required to perform even more rigorous background checks or audits so as to determine their ultimate revenue sources. But there may be limits to what can and should be expected of platforms in this regard. Here, ad archive governance intersects with broader questions of campaign finance regulation and the role of "dark money" in politics. These issues have historically been tackled through national regulation, including standardised registration mechanisms for political advertisers, but many of these regimes currently do not address online advertising. Platforms' self-regulatory measures, though useful as a first step, cannot make up for the lack of public regulation in this space (Lapowsky, 2018; Livingstone, 2018). Even Facebook CEO Mark Zuckerberg has called for regulation here, arguing in a recent op-ed that "[o]ur systems would be more effective if regulation created common standards for verifying political actors" (Zuckerberg, 2019).

Another weak spot for ad archives is that they fail to capture "native advertising" practices: advertising which is not conducted through social media platforms' designated advertising services, but rather through their organic content channels. Such "astroturfing" strategies have seen widespread deployment in both commercial and political contexts, from Wal-Mart and Monsanto and from Russian "troll farms" to presidential Super PACs (Collins, 2016; Howard et al., 2018; Leiser, 2016). Ad archives do not capture this behaviour, and indeed their very presence could further *encourage* astroturfing, as a form of regulatory arbitrage. Benkler, Faris, and Roberts suggest that ad archive regulation should address this issue by imposing an independent duty on advertisers to disclose any "paid coordinated campaigns" to the platform (Benkler, Faris, & Roberts, 2018). One example from practice is the Republic of Ireland's Online Advertising and Social Media Bill of 2017, which would hold ad buyers liable for providing inaccurate information to ad sellers, and also prohibit the use of bots which "cause multiple online presences directed towards a political end to present as an individual account or profile on an online platform" (Republic of Ireland, 2017). Enforcing such rules will remain challenging, however, since astroturfing is difficult to identify and often performed by bad actors with little or no interest in complying with the law (Leiser, 2016).

For ads that are actually included in the archive, inauthentic behaviour can also distort associated metadata such as traffic data. Engagement metrics, including audience demographic data, can be significantly disturbed by click fraud or bot traffic (Edelman, 2014; Fulgoni, 2016). Platforms typically spend extensive resources to combat inauthentic behaviour, and this appears to be a game of cat-and-mouse without definitive solutions. In light of these challenges, researchers should maintain a healthy scepticism when dealing with ad archive data and, where necessary, continue to corroborate ad archive findings with alternative sources and research methods (see, generally: Vlassenroot et al., 2019).

The above is not to say that *all* information supplied by ad buyers should be verified. There may still be an added value in enabling voluntary, unverified disclosures by ad buyers in archives. Facebook, for instance, gives advertisers the option to include "Information From the Advertiser" in the archive. Such features can enable good faith advertisers to further support accountability processes, e.g., by adding further context or supplying contact information. It is essential, however, that such unverified submissions are recognisably earmarked as such. Ad archive operators should clearly describe which data is verified, and how, so that users can treat

their data with the appropriate degree of scepticism.

TARGETING: HOW IS AD TARGETING DOCUMENTED?

Another key criticism of ad archives is that they are not detailed enough, particularly in their documentation of ad targeting practices. Micro-targeting technology, as discussed previously, is the source of many public policy concerns for both political and commercial advertising, including discrimination, deception, and privacy harms. These threats are relatively new, and are both undocumented and unregulated in many jurisdictions – particularly as regards political advertising (Bodó et al., 2017). Regrettably, ad archives currently fail to illuminate these practices in any meaningful depth.

At the time of writing, the major ad archives differ significantly in their approach to targeting data. Google’s archive indicates whether the following targeting criteria have been selected by the ad buyer: age, location, and gender. It also lists the top five Google keywords selected by the advertiser. Facebook’s Ad Library, by contrast, does not disclose what targeting criteria have been selected, but instead shows a demographic breakdown of the actual audience that saw the message - also in terms of age, location and gender. Twitter offers *both* audience statistics and targeting criteria, and covers not only the targeting criteria of age, location, and gender, but also their preferred language. These data vary in granularity. For instance, Google’s archive lists six different age brackets between the ages of 18 and 65+, whereas Twitter lists 34. For anyone familiar with the complexities of online behavioural targeting, it is apparent that these datasets leave many important questions unanswered. These platforms offer far more refined methods for ad targeting and performance tracking than the basic features described above.

For better insights into ad targeting, one helpful rule of thumb would be to insist that ad archives should include an equivalent level of information as is offered to the actual ad buyer – both in terms of targeting criteria and in terms of actual audience demographics (Mozilla, 2019). For some targeting technologies, full disclosure of targeting practices might raise user privacy concerns. For instance, Facebook’s Custom Audience feature enables advertisers to target users by supplying their own contact information, such as email addresses or telephone numbers. Insisting on full disclosure of targeting criteria for these custom audiences would lead to the public disclosure of sensitive personal data (Rieke & Bogen, 2018). Anonymisation of these data may not always be reliable (Ohm, 2010). In these cases, however, Facebook could at a minimum still disclose any additional targeting criteria selected by the ad buyer in order to refine this custom audience. Furthermore, ad performance data, rather than ad targeting data, can also provide some insight into targeting without jeopardising the custom audience’s privacy (Rieke & Bogen, 2018). Other platforms’ advertising technologies might raise comparable privacy concerns, demanding a case-by-case assessment of relevant tradeoffs. These exceptions and hard cases notwithstanding, however, there are no clear objections (either technical or political) that should prevent platforms from publicly disclosing the targeting methods selected by their advertisers.

In light of such complexities, designing appropriate disclosures will likely require ongoing dialogue between archive operators, archive users and policymakers. The first contours of such a debate can already be found in the work of Edelson et al., Rieke & Bogen, and Mozilla, who have done valuable work in researching and critiquing early versions of Google, Twitter and Facebook’s data sets (Edelson et al., 2019; Mozilla, 2019; Rieke & Bogen, 2018). For the time being, researchers may also choose to combine ad archive data with other sources, such as Facebook’s Social Science One initiative, or GDPR data access rights, in order to obtain a more detailed understanding of targeting practices (Ausloos & Dewitte, 2018; Venturini & Rogers,

2019). For instance, Ghosh et al. supplemented ad archive research with data scraped with ProPublica's research tool, which gave insights into ad targeting that were not offered through the ad archive (Ghosh et al., 2019). Along these lines, ad archives can help to realise Pasquale's model of "qualified transparency", which combines general public disclosures with more limited, specialist inquiries (Pasquale, 2015).

CONCLUSION

This paper has given an overview of a new and rapidly developing topic in online advertising governance: political ad archives. Here we summarise our key findings, and close with suggestions for future research in both law and communications science.

Ad archives can be a novel and potentially powerful governance tool for online political advertising. If designed properly, ad archives can enable monitoring by a wide range of stakeholders, each with diverse capacities and interests in holding advertisers accountable. In general, ad archives can not only improve accountability to applicable laws, but also to public opinion, by introducing publicity and thus commercial and political risk into previously invisible advertisements.

Public oversight will likely be necessary to realise these benefits, since platforms ostensibly lack the incentives to voluntarily optimise their ad archives for transparency and accountability. Indeed, our analysis here has already identified several major shortcomings in present ad archive policies: scoping, verifying, and targeting. To realise the full potential of ad archives, these issues will require further research, critique, and likely regulation. Our review suggests that major advances can already be made by comprehensively publishing *all* advertisements, regardless of whether they have been flagged as political; revoking any exemptions for media organisations; requiring basic verification of ad buyers' identities; documenting how ad archive data is verified; and disclosing all targeting methods selected by the ad buyer (insofar as possible without publishing personal data).

Looking forward, ad archives present a fruitful research area for both legal and communication sciences scholars. For legal scholars, the flurry of law making around political advertising in general, and transparency in particular, raises important questions about regulatory design (in terms of how relevant actors and duties are defined, oversight and enforcement mechanisms, etc.). In future, ad archives also deserve consideration in commercial advertising governance, in such areas as consumer protection, child protection, or anti-discrimination.

The emergence of ad archives also has important implications for communications science. Firstly, ad archives could become an important *resource* of data for communications research, offering a range of data that would previously have been difficult or impossible to obtain. Although our paper has identified several shortcomings in this data, they might nonetheless provide a meaningful starting point to observe platforms' political advertising. Secondly, ad archives are an interesting *object* of communications science research, in terms of how they are used by relevant stakeholders, and how this impacts advertising and communications practice. Further research along these lines will certainly be necessary to better understand ad archives, and to make them reach their full potential.

ACKNOWLEDGEMENTS

The authors wish to thank Frédéric Dubois, Chris Birchall, Joe Karaganis and Kristofer Erickson for their thoughtful reviewing and editing of this article. The authors also wish to thank Frederik Zuiderveen Borgesius and Sam Jeffers for their helpful insights during the writing process, as well as the participants in the ICA 2019 Post-Conference on the Rise of the Platforms and particularly the organisers: Erika Franklin Fowler, Sarah Anne Ganter, Dave Karpf, Rasmus Kleis Nielsen, Daniel Kreiss and Shannon McGregor.

REFERENCES

- Albright, J. (2018, November 4). Facebook and the 2018 Midterms: A Look at the Data – The Micro-Propaganda Machine. Retrieved from <https://medium.com/s/the-micro-propaganda-machine/the-2018-facebook-midterms-part-i-recursive-ad-ccountability-aco90d276097>
- Ali, M., Sapiezynski, P., Bogen, M., Korolova, A., Mislove, A., & Rieke, A. (2019). Discrimination through optimization: How Facebook’s ad delivery can lead to skewed outcomes. *Arxiv [Cs]*. Retrieved from <https://arxiv.org/pdf/1904.02095.pdf>
- Angelopoulos, C. J., Brody, A., Hins, A. W., Hugenholtz, P. B., Leerssen, P., Margoni, T., McGonagle, T. van Hoboken, J. V.J. (2015). Study of fundamental rights limitations for online enforcement through self-regulation. Institute for Information Law (IViR). Retrieved from https://pure.uva.nl/ws/files/8763808/IVIR_Study_Online_enforcement_through_self_regulation.pdf
- Andringa, P. (2018). Interactive: See Political Ads Targeted to You on Facebook. *NBC*. Retrieved from <http://www.nbcсандiego.com/news/tech/New-Data-Reveal-Wide-Range-Political-Actors-Facebook-469600273.html>
- Arnett v. Kennedy, 416 U.S. 134 (Supreme Court of the United States, 1974).
- Parliament of the Netherlands. (2019). *Motion for Complete Transparency in the Buyers of Political Advertisements on Facebook*. Retrieved from <https://www.parlementairemonitor.nl/9353000/1/j9vvij5epmj1eyo/vkvudd248rwa>
- Ausloos, J., & Dewitte, P. (2018). Shattering one-way mirrors – data subject access rights in practice. *International Data Privacy Law*, 8(1), 4–28. doi:10.1093/idpl/ipy001
- Barocas, S. (2012). The Price of Precision: Voter Microtargeting and Its Potential Harms to the Democratic Process. *Proceedings of the First Edition Workshop on Politics, Elections and Data*, 31–36. doi:10.1145/2389661.2389671
- Benkler, Y., Faris, R., & Roberts, H. (2018). *Network Propaganda: Manipulation, Disinformation, and Radicalization in American Politics*. Oxford: Oxford University Press.
- Bodó, B., Helberger, N., & Vreese, C. H. de. (2017). Political micro-targeting: a Manchurian candidate or just a dark horse? Towards the next generation of political micro-targeting research. *Internet Policy Review*, 6(4). doi: 10.14763/2017.4.776
- Boerman, S. C., Kruikemeier, S., & Zuiderveen Borgesius, F. J. (2017). Online Behavioral Advertising: A Literature Review and Research Agenda. *Journal of Advertising*, 46(3), 363–376. doi:10.1080/00913367.2017.1339368
- Borgesius, F.J., Moller, J., Kruikemeier, S., Ó Fathaigh, R., Irion, K., Dobber, T., Bodo, B., & de Vreese, C. (2018). Online Political Microtargeting: Promises and Threats for Democracy. *Utrecht Law Review*, 14(1). doi:10.18352/ulr.420
- Carbajal, A., Kint, J., Mills Wade, A., Brooks, L. T., Chavern, D., McKenzie, A. B., & Golden, M. (2018). Open Letter to Marck Zuckerberg on Alternative Solutions for Politics Tagging. Retrieved from https://www.newsmediaalliance.org/wp-content/uploads/2018/06/vR_Alternative-Facebook-Politics-Tagging-Solutions-FINAL.pdf

- Cardoso, T. (2019, March 4). Google to ban political ads ahead of federal election, citing new transparency rules. *The Globe and Mail*. Retrieved from <https://www.theglobeandmail.com/politics/article-google-to-ban-political-ads-ahead-of-federal-election-citing-new/>
- Chavern, D. (2018, May 18). Open Letter to Mr. Zuckerberg. *News Media Alliance*. Retrieved from <http://www.newsmediaalliance.org/wp-content/uploads/2018/05/FB-Political-Ads-Letter-FINAL.pdf>
- Chester, J., & Montgomery, K. C. (2017). The role of digital marketing in political campaigns. *Internet Policy Review*, 6(4). doi:10.14763/2017.4.773
- Collins, B. (2016, April 21). *Hillary PAC Spends \$1 Million to 'Correct' Commenters on Reddit and Facebook*. Retrieved from <https://www.thedailybeast.com/articles/2016/04/21/hillary-pac-spends-1-million-to-correct-commenters-on-reddit-and-facebook>
- House of Commons Select Committee on Digital, Culture, Media and Sport. (2019). *Disinformation and 'fake news': Final Report*. Retrieved from https://publications.parliament.uk/pa/cm201719/cmselect/cmcomeds/1791/179103.htm#_idTExtAnchor000
- Edelman, B. (2014). Pitfalls and Fraud In Online Advertising Metrics: What Makes Advertisers Vulnerable to Cheaters, And How They Can Protect Themselves. *Journal of Advertising Research*, 54(2), 127–132. doi:10.2501/JAR-54-2-127-132
- Edelson, L., Sakhuja, S., Dey, R., & McCoy, D. (2019). An Analysis of United States Online Political Advertising Transparency. *ArXiv [Cs]*. Retrieved from <http://arxiv.org/abs/1902.04385>
- European Commission. (2018, September 26). *Code of Practice on Disinformation*. Retrieved from <https://ec.europa.eu/digital-single-market/en/news/code-practice-disinformation>
- European Commission. (2019). *Third monthly intermediate results of the EU Code of Practice against disinformation*. Retrieved from <https://ec.europa.eu/digital-single-market/en/news/third-monthly-intermediate-results-eu-code-practice-against-disinformation>
- Fanta, A. & Rudl, T. (2019, July 17). Leaked document: EU Commission mulls new law to regulate online platforms. *Netzpolitik.org*. Retrieved from: <https://netzpolitik.org/2019/leaked-document-eu-commission-mulls-new-law-to-regulate-online-platforms/>
- Facebook. (2019a). Issues of national importance. Retrieved from <https://www.facebook.com/business/help/214754279118974>
- Facebook. (2019b). Ads about social issues, elections or politics. Retrieved from <https://www.facebook.com/business/help/208949576550051>
- Frier, S. (2018, July 2). Facebook's Political Rule Blocks Ads for Bush's Beans, Singers Named Clinton. *Bloomberg*. Retrieved from <https://www.bloomberg.com/news/articles/2018-07-02/facebook-s-algorithm-blocks-ads-for-bush-s-beans-singers-named-clinton>

- Fulgoni, G. M. (2016). Fraud in Digital Advertising: A Multibillion-Dollar Black Hole: How Marketers Can Minimize Losses Caused by Bogus Web Traffic. *Journal of Advertising Research*, 56(2), 122. doi:10.2501/JAR-2016-024
- Foucault, M. (1977). *Discipline and Punish The Birth of the Prison*. New York: Pantheon Books.
- Ghosh, A., Venkatadri, G., & Mislove, A. (2019). *Analyzing Political Advertisers' Use of Facebook's Targeting Features*. 7. Available at <https://www.ieee-security.org/TC/SPW2019/ConPro/papers/ghosh-conpro19.pdf>
- Goldman, R. (2017). Update on Our Advertising Transparency and Authenticity Efforts. *Facebook Newsroom*. Retrieved from <https://newsroom.fb.com/news/2017/10/update-on-our-advertising-transparency-and-authenticity-efforts/>
- Google (2019a). *Implementation Report for EU Code of Practice on Disinformation*. Retrieved from https://ec.europa.eu/information_society/newsroom/image/document/2019-5/google_-_ec_action_plan_reporting_CF162236-E8FB-725E-CoA3D2D6CCFE678A_56994.pdf
- Google (2019b). Verification for election advertising in the European Union. Retrieved from <https://support.google.com/adspolicy/answer/9211218>
- Guha, S., Cheng, B., & Francis, P. (2010). Challenges in measuring online advertising systems. *Proceedings of the 10th Annual Conference on Internet Measurement - IMC '10*, 81. doi:10.1145/1879141.1879152
- Hansen, H. K., Christensen, L. T., & Flyverbom, M. (2015). Logics of transparency in late modernity: Paradoxes, mediation and governance. *European Journal of Social Theory*, 18(2), 117–131. doi:10.1177/1368431014555254
- Hounsel, A., Mathias, J. N., Werdmuller, B., Griffey, J., Hopkins, M., Peterson, C., ... Feamster, N. (2019). Estimating Publication Rates of Non-Election Ads by Facebook and Google. Retrieved from <https://github.com/citp/mistaken-ad-enforcement/blob/master/estimating-publication-rates-of-non-election-ads.pdf>
- Howard, P. N., Ganesh, B., Liotsiou, D., Kelly, J., & François, C. (2018). *The IRA, Social Media and Political Polarization in the United States, 2012–2018* [Working Paper 2018.2]. Oxford: Project on Computational Propaganda. Retrieved from <https://comprop.oi.ox.ac.uk/research/ira-political-polarization/>
- Howard, P. (2019, March 27). A Way to Detect the Next Russian Misinformation Campaign. *The New York Times*. Retrieved from <https://www.nytimes.com/2019/03/27/opinion/russia-elections-facebook.html?module=inline>
- Keller, D. & Leerssen, P. (in press). Facts and where to find them: Empirical research on internet platforms and content moderation. In N. Persily & J. Tucker (eds), *Social Media and Democracy: The State of the Field*. Cambridge: Cambridge University Press.
- Klobuchar, A., Warner, R., & McCain, J. (2017, October 19). *The Honest Ads Act*. Retrieved from <https://www.congress.gov/bill/115th-congress/senate-bill/1989/text>
- Kuczerawy, A. (2019, in press). Fighting online disinformation: did the EU Code of Practice forget about freedom of expression? In E. Kuźelewska, G. Terzis, D. Trottier, & D. Kloza (Eds.),

Disinformation and Digital Media as a Challenge for Democracy. Cambridge: Intersentia.

Lomas, N. (2018, July 26). Facebook finally hands over leave campaign Brexit ads. *Techcrunch*. Retrieved from: <https://techcrunch.com/2018/07/26/facebook-finally-hands-over-leave-campaign-brexit-ads/>

Lapowsky, I. (2018). Obscure Concealed-Carry Group Spent Millions on Facebook Political Ads. *WIRED*. Retrieved from <https://www.wired.com/story/facebook-ads-political-concealed-online/>

Leathern, R. (2019). Updates to our ad transparency and authorisation efforts. Retrieved from: <https://www.facebook.com/facebookmedia/blog/updates-to-our-ads-transparency-and-authorisation-efforts>

Leiser, M. (2016). AstroTurfing, ‘CyberTurfing’ and other online persuasion campaigns. *European Journal of Law and Technology*, 7(1). Retrieved from <http://ejlt.org/article/view/501>

Livingstone, S. (2018). *Tackling the Information Crisis: A Policy Framework for Media System Resilience* [Report]. London: LSE Commission on Truth, Trust and Technology. Retrieved from <http://www.lse.ac.uk/media-and-communications/assets/documents/research/T3-Report-Tackling-the-Information-Crisis-v6.pdf>

Lyon, D. (2006). *Theorizing Surveillance: The Panopticon and Beyond*. Devon: Willan Publishing.

Macleod, A. (2019). Fake News, Russian Bots and Putin’s Puppets. In A. MacLeod (Ed.), *Propaganda in the Information Age: Still Manufacturing Consent*. London: Routledge.

Matias, J. N., Hounsel, A., & Hopkins, M. (2018, November 2). We Tested Facebook’s Ad Screeners and Some Were Too Strict. *The Atlantic*. Retrieved from: <https://www.theatlantic.com/technology/archive/2018/11/do-big-social-media-platforms-have-effective-ad-policies/574609/>

Merrick, R. (2019, December 25). Brexit: Leave ‘very likely’ won EU referendum due to illegal overspending, says Oxford professor’s evidence to High Court. *The Independent*. Retrieved from: <https://www.independent.co.uk/news/uk/politics/vote-leave-referendum-overspending-high-court-brexit-legal-challenge-void-oxford-professor-a8668771.html>

Merrill, J. B. (2018). How Big Oil Dodges Facebook’s New Ad Transparency Rules. Retrieved 22 April 2019, from ProPublica website: <https://www.propublica.org/article/how-big-oil-dodges-facebooks-new-ad-transparency-rules>

Merrill, J. B., & Tobin, A. (2019, January 28). Facebook Moves to Block Ad Transparency Tools — Including Ours. *ProPublica*. Retrieved from <https://www.propublica.org/article/facebook-blocks-ad-transparency-tools>

Montellaro, Z. (2019). House Democrats forge ahead on electoral reform bill. *POLITICO*. Retrieved from <https://politi.co/2GO4eJ8>

Mozilla (2019, March 27). Facebook and Google: This is What an Effective Ad Archive API Looks Like. *The Mozilla Blog*. Retrieved from:

<https://blog.mozilla.org/blog/2019/03/27/facebook-and-google-this-is-what-an-effective-ad-archive-api-looks-like>

Mulgan, R. (2000). Comparing Accountability in the Public and Private Sectors. *Australian Journal of Public Administration*, 59(1), 87–97. doi:10.1111/1467-8500.00142

Ohm, P. (2010). Broken Promises of Privacy: Responding To The Surprising Failure of Anonymization. *UCLA Law Review* 57, 1701–1777. Retrieved from <https://www.uclalawreview.org/pdf/57-6-3.pdf>

Netherlands Ministry of the Interior. (2019). *Response to the Motion for Complete Transparency in the Buyers of Political Advertisements on Facebook*. Retrieved from: <https://www.tweedekamer.nl/kamerstukken/detail?id=2019Z03283&did=2019D07045>

O’Sullivan, D. (2018). What an anti-Ted Cruz meme page says about Facebook’s political ad policy. *CNN*. Retrieved from: <https://www.cnn.com/2018/10/25/tech/facebook-ted-cruz-memes/index.html>

Parsons, C. (2019). The (In)effectiveness of Voluntarily Produced Transparency Reports. *Business & Society*, 58(1), 103–131. doi:10.1177/0007650317717957

Pasquale, F. (2015). *The Black Box Society: The Secret Algorithms That Control Money and Information*. Cambridge: Harvard University Press. Retrieved from <https://www.jstor.org/stable/j.ctt13x0hch>

Republic of France. (2018). *LoI n° 2018-1202 du 22 décembre 2018 relative à la lutte contre la manipulation de l’information*. Retrieved from <https://www.legifrance.gouv.fr/affichTexte.do?cidTexte=JORFTEXT000037847559&categorieLien=id>

Republic of Ireland (2017), *Online Advertising and Social Media (Transparency) Bill 2017*. Retrieved from: <https://data.oireachtas.ie/ie/oireachtas/bill/2017/150/eng/initiated/b15017d.pdf>

Rieke, A., & Bogen, M. (2018). *Leveling the Platform: Real Transparency for Paid Messages on Facebook*. UpTurn Report. Retrieved from <https://www.upturn.org/static/reports/2018/facebook-ads/files/Upturn-Facebook-Ads-2018-05-08.pdf>

Rosenberg, M. (2019, July 25). Ad Tool Facebook Built to Fight Disinformation Doesn’t Work as Advertised. *The New York Times*. Retrieved from: <https://www.nytimes.com/2019/07/25/technology/facebook-ad-library.html>

Sanders, E. (2019, May 9). Washington Public Disclosure Commission Passes Emergency Rule Clarifying That Facebook and Google Must Turn Over Political Ad Data. *The Stranger*. Retrieved from <https://www.thestranger.com/slog/2018/05/09/26158462/washington-public-disclosure-commission-passes-emergency-rule-clarifying-that-facebook-and-google-must-turn-over-political-ad-data>

Sanders, E. (2018, October 16). Facebook Says It's Immune from Washington State Law. *The Stranger*. Retrieved from: <https://www.thestranger.com/slog/2018/10/16/33926412/facebook-says-its-immune-from-wa>

shington-state-law

Shane, S. (2017, November 1). These are the Ads Russia Bought on Facebook in 2016. *The New York Times*. Retrieved from <https://www.nytimes.com/2017/11/01/us/politics/russia-2016-election-facebook.html>

Shukla, S. (2019, March 28). A Better Way to Learn About Ads on Facebook. *Facebook Newsroom*. Retrieved from <https://newsroom.fb.com/news/2019/03/a-better-way-to-learn-about-ads/>

Singer, N. (2018, August 16). ‘Weaponized Ad Technology’: Facebook’s Moneymaker Gets a Critical Eye. *The New York Times*.
<https://www.nytimes.com/2018/08/16/technology/facebook-microtargeting-advertising.html>

Timmons, H. & Kozlowska, H. (2018, March 22). Facebook’s quiet battle to kill the first transparency law for online political ads. *Quartz*. Retrieved from:
<https://qz.com/1235363/mark-zuckerberg-and-facebooks-battle-to-kill-the-honest-ads-act/>

Turton, W. (2018, October 30). We posed as 100 senators to run ads on Facebook. Facebook approved all of them. *VICE News*. Retrieved from:
https://news.vice.com/en_ca/article/xw9n3q/we-posed-as-100-senators-to-run-ads-on-facebook-facebook-approved-all-of-them

Twitter. (2019a). *Implementation Report for EU Code of Practice on Disinformation*. Retrieved from
http://ec.europa.eu/information_society/newsroom/image/document/2019-5/twitter_progress_report_on_code_of_practice_on_disinformation_CF162219-992A-B56C-06126A9E7612E13D_56993.pdf

Twitter. (2019b). How to get certified as a political advertisers. Retrieved from
<https://business.twitter.com/en/help/ads-policies/restricted-content-policies/political-content/how-to-get-certified.html>

van Dijck, J., Poell, T., & de Waal, M. (2018). *The Platform Society: Public Values in a Connective World*. Oxford: Oxford University Press.

Van Til, G. (2019). Zelfregulering door online platforms: een waar wondermiddel tegen online desinformatie? [Self-regulation by online platforms: a true panacea against online disinformation?]. *Mediaforum*, 1(13). Retrieved from
https://www.ivir.nl/publicaties/download/Mediaforum_2019_1_vanTil.pdf

Vandor, M. (2018, 18). Indexing news Pages on Facebook for the Ad Archive. *Facebook Media*. Retrieved from: <https://www.facebook.com/facebookmedia/blog/indexing-news-pages-on-facebook-for-the-ad-archive>

Venturini, T., & Rogers, R. (2019). “API-Based Research” or How can Digital Sociology and Journalism Studies Learn from the Facebook and Cambridge Analytica Data Breach. *Digital Journalism*, 7(4), 532–540. doi: 10.1080/21670811.2019.1591927

Vlassenroot, E., Chambers, S., Di Pretoro, E., Geeraert, F., Haesendonck, G., Michel, A., & Mechant, P. (2019). Web archives as a data resource for digital scholars. *International Journal of Digital Humanities*, 1(1), 85–111. doi:10.1007/s42803-019-00007-7

Wagner, B. (2018). Free Expression?: Dominant information intermediaries as arbiters of internet speech. In M. Moore & D. Tambini (Eds.), *Digital Dominance*. Oxford: Oxford University Press.

Warner, R. (2017). *The Honest Ads Act (primer)*. Retrieved from <https://www.warner.senate.gov/public/index.cfm/the-honest-ads-act>

Waterson, J. (2019, January 14). Obscure pro-Brexit group spends tens of thousands on Facebook ads. *The Guardian*. Retrieved from <https://www.theguardian.com/politics/2019/jan/14/obscure-pro-brexit-group-britains-future-spends-tens-of-thousands-on-facebook-ads>

Waterson, J. (2019, April 3). Facebook Brexit ads secretly run by staff of Lynton Crosby firm. *The Guardian*. Retrieved from: <https://www.theguardian.com/politics/2019/apr/03/grassroots-facebook-brexit-ads-secretly-run-by-staff-of-lynton-crosby-firm>

Zuckerberg, M. (2019, March 30). The Internet needs new rules. Let's start in these four areas. *The Washington Post*. Retrieved from https://www.washingtonpost.com/opinions/mark-zuckerberg-the-internet-needs-new-rules-lets-start-in-these-four-areas/2019/03/29/9e6f0504-521a-11e9-a3f7-78b7525a8d5f_story.html

FOOTNOTES

1. E.g. Mozilla, 2019; Mattias, Hounsel, & Hopkins, 2019; Merrill, 2018; Rieke & Bogen, 2018; Edelson et al., 2019; Andringa, 2018; Lapowsky, 2018; O'Sullivan, 2018; Waterson, 2019; Albright, 2018; Howard, 2019. See Section three for further discussion.

2. Parallel to the more general distinction between the governance *of* platforms and the governance *by* platforms (Gillespie, 2018).

3. The Commission describes the Code as a 'self-regulatory' instrument. However, given the Commission's involvement in its development and oversight, we consider 'co-regulatory' a more apt description (Kuczerawy, 2019; more generally see Angelopoulos et al., 2015).

4. Installing such notice and takedown for unlawful content is a requirement under EU law. In the US, notice and takedown procedures are only required for copyright and trademark claims, and the majority of takedown occurs on a strictly voluntary basis. In practice, much of the content removed under these regimes is assessed on the basis of platforms' voluntary standards (Keller & Leerssen, 2019).