Ad Delivery Algorithms: The Hidden Arbiters of Political Messaging

Muhammad Ali*
mali@ccs.neu.edu
Northeastern University
Boston, MA, USA

Piotr Sapiezynski*
p.sapiezynski@northeastern.edu
Northeastern University
Boston, MA, USA

Aleksandra Korolova
korolova@usc.edu
University of Southern California
Los Angeles, CA, USA

Alan Mislove
amislove@ccs.neu.edu
Northeastern University
Boston, MA, USA

Aaron Rieke
aaron@upturn.org
Upturn
Washington, DC, USA

ABSTRACT

Political campaigns are increasingly turning to targeted advertising platforms to inform and mobilize potential voters. The appeal of these platforms stems from their promise to empower advertisers to select (or “target") users who see their messages with great precision, including through inferences about those users’ interests and political affiliations. However, prior work has shown that the targeting may not work as intended, as platforms’ ad delivery algorithms play a crucial role in selecting which subgroups of the targeted users see the ads. In particular, the platforms can selectively deliver ads to subgroups within the target audiences selected by advertisers in ways that can lead to demographic skews along race and gender lines, and do so without the advertiser’s knowledge. In this work we demonstrate that ad delivery algorithms used by Facebook, the most advanced targeted advertising platform, shape the political ad delivery in ways that may not be beneficial to the political campaigns and to societal discourse. In particular, the ad delivery algorithms lead to political messages on Facebook being shown predominantly to people who Facebook thinks already agree with the ad campaign’s message even if the political advertiser targets an ideologically diverse audience. Furthermore, an advertiser determined to reach ideologically non-aligned users is non-transparently charged a high premium compared to their more aligned competitor, a difference from traditional broadcast media. Our results demonstrate that Facebook exercises control over who sees which political messages beyond the control of those who pay for them or those who are exposed to them. Taken together, our findings suggest that the political discourse’s increased reliance on profit-optimized, non-transparent algorithmic systems comes at a cost of diversity of political views that voters are exposed to. Thus, the work raises important questions of fairness and accountability desiderata for ad delivery algorithms applied to political ads.

*These authors contributed equally to this research.

WSDM ’21, March 8–12, 2021, Virtual Event, Israel

ACM ISBN 978-1-4503-8297-7/21/03.
https://doi.org/10.1145/3437963.3441801

ACM Reference Format:
Muhammad Ali, Piotr Sapiezynski, Aleksandra Korolova, Alan Mislove, and Aaron Rieke. 2021. Ad Delivery Algorithms: The Hidden Arbiters of Political Messaging. In *Proceedings of the Fourteenth ACM International Conference on Web Search and Data Mining (WSDM ’21), March 8–12, 2021, Virtual Event, Israel.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3437963.3441801

1 INTRODUCTION

Political campaigns spend millions of dollars on advertising to get their messages out to voters. This spending has been increasingly migrating from traditional broadcast media (e.g., television and newspapers) to the internet. For example, at the U.S. state level, ten times as many candidates advertise on Facebook than on TV [13]. The reason for the shift is that online advertising platforms promise to lower the cost of advertising and increase the efficacy of campaigns through detailed *targeting*, where advertisers can specify the users they would like to reach using *attributes*. Common targeting attributes include demographic characteristics, interests, visited locations, phone numbers, relationship status, and wealth. Platforms are known to determine attributes in a variety of ways, including by relying on those provided by the users explicitly, algorithmic inferences by the platforms based on users’ activity (both on the platforms and elsewhere), or augmentation with data supplied by third-party data brokers [8, 10, 12, 32].

The implications of the *ad creation and targeting* phase of the advertising process, where the advertiser uploads their ad creative and selects their desired target users, have been subject to recent scrutiny and policy debate. Targeting attributes offered by the ad platforms have been shown to enable advertisers to prevent certain ethnic groups from seeing ads [5, 29]. For example, in 2016, the Trump campaign used these techniques in an attempt to lower turnout among young women and Black voters [17], and there is evidence that Russian organizations used these tools with the goal of influencing the 2016 U.S. presidential elections [27, 30]. Recently, U.S. Federal Election Commission (FEC) Chair Ellen Weintraub argued that platforms should limit political advertisers’ ability to narrowly target ads to ensure that “a broad public can hear the speech and respond” [33]. Shortly after, Google announced that it would significantly limit election ad targeting in order to “promote increased visibility of election ads” [16].

Scientific attempts to rigorously measure the effects of online political targeted advertising outside of the controlled lab environments [22] have been limited by the challenge of controlling for the optimization decisions made by platforms in their *ad delivery* phase, or the process by which platforms select which ads get shown to which users [9].

Given the platforms' desire not to overwhelm the users with too many ads (especially those with potentially upsetting content), the finite budget of the advertiser, a large potential audience of ad recipients, and the competition from other advertisers for those recipients, platforms need to select a subset of the targeted users who will actually be shown the ad. This selection is commonly performed through auctions, where the outcome is determined not only based on the advertisers' willingness to pay, but also on the platform's long-term business and growth goals, such as the platform's desire for its users to see *relevant* ads (and, therefore, maintain its user base) and the platform's desire for its advertisers to achieve their desired outcomes (and, therefore, achieve the platform's revenue and advertiser growth goals). We call the algorithmic approach that the platform uses to balance these goals during the selection *ad delivery optimization*.

Prior work has showed that ad delivery optimization, and its reliance on algorithmically inferred "relevance" of an ad to a user, can lead to troubling results in the context of life-opportunity ads [4]. Specifically, it demonstrated that ads targeting the same gender- and race-balanced audiences for various jobs were delivered to vastly different groups of users: cashier job openings were shown predominantly to women, taxi driver job openings were shown predominantly to Black users, and artificial intelligence and lumber jobs openings were shown to majority white and male audiences [4]. These results were subsequently extended with controls for user job qualifications [18] and reproduced in several European countries [19]. Further prior work showed that the existing ad ecosystem provides little support for advertising to a demographically representative cohort [14].

In this work, we investigate the impact of ad delivery optimization in real-world advertising platforms on a different arena: political discourse. We focus on Facebook because of its critical importance to today's digital political advertising and its pioneering role in targeted advertising. Specifically, we seek to answer: *Is a political campaign advertising on Facebook able to reach all of the electorate?* Or, *is Facebook preferentially delivering ads to users who it believes are more likely to be aligned with the campaign's political views?* Additionally, *to what extent does Facebook vary ad pricing based on its hypothesized match between the target audience's and campaign's political views?*

The answers to these questions are particularly urgent and salient in light of the debate unfolding over the "microtargeting" of political ads for at least two reasons. *First*, skews resulting from ad delivery can raise similar concerns to those raised about narrow targeting: an electorate who cannot "hear and respond" to political speech. *Second*, ad delivery algorithms might counteract the goals of restricting microtargeting by redirecting ads according to the choices of the platforms (in spite of broader target audiences). In other words, limiting targeting options transfers more power with regards to political message delivery to the platforms; and therefore, makes the investigation of their delivery algorithms and their implications for political discourse even more important.

To rigorously answer the questions posed, we became a political advertiser and spent over $13,000[1] to run political ads under controlled conditions, and observed how Facebook's algorithms delivered them. We used Facebook's ad reporting features, combined with proxies, to understand who our ads were delivered to and how our budget was split across users with different political leanings.

1.1 Main Findings

The results of our analyses offer the following contributions:

First, we show that, despite identical targeting parameters, budgets, and competition from other advertisers, the content of a political ad alone can significantly affect which users Facebook will show the ad to. For example, when we run two campaigns, each targeting the same audience comprised of an equal number of registered Democratic voters and Republican voters, we find that our ad for a Democratic candidate delivers to an audience that is 70% Democratic, while the audience reached by the ad for a Republican candidate is only 40% Democratic.

Second, we find that it can be difficult and more expensive for political campaigns to have their content delivered to those who Facebook believes are not aligned with the campaign's views. For example, we find that when targeting an audience of conservative users, in the first day of the ad campaign, Facebook delivers our liberal-leaning ad to only 4,772 users, while our conservative-leaning ad to 7,588 users.[2] We find that the underlying reason for the differences in delivery is that our liberal-leaning ads targeting conservative users are charged significantly more by Facebook than our conservative-leaning ads ($15.39 versus $10.98 for 1,000 impressions), despite being run from the same ad account, at the same time, and targeting the same users.[3]

Third, we observe these effects persist for ads that do not prompt user engagement, which suggests that the ad delivery decisions made by Facebook are not driven exclusively by user reactions to the ad but instead are made at least partially by Facebook itself.

Taken together, our results indicate that Facebook preferentially shows users political ads whose contents Facebook predicts are aligned with their political views. This ad delivery choice has negative implications for both users and campaigns. For users, such delivery limits users' exposure to diverse viewpoints unbeknownst to the users, especially if the predictions about the alignment are based on an algorithmic analysis of users' activity and third-party data, rather than information explicitly provided by the users. For campaigns, such delivery may inhibit them from reaching beyond their existing "base" on Facebook, as getting ads delivered to users the platform believes are not aligned with their views may become prohibitively expensive. Furthermore, unlike in traditional media, this may imply that campaigns of equal financial means are not equal in their ability to reach a particular audience[4], with the price differential decided exclusively and non-transparently by Facebook.

[1] Throughout the paper we refer to prices in U.S. Dollars.
[2] We find a similar, but flipped, effect if we target an audience of liberal users.
[3] Again, we see a similar, flipped effect when targeting liberal users.
[4] The work of [23] hypothesizes differential pricing using a different methodology.

Importantly, these effects may be occurring without users' or campaigns' knowledge or control.

Stepping back, our findings raise serious concerns about whether Facebook and similar ad targeting platforms are, in fact, *amplifying* political *filter bubbles* by economically dis-incentivizing content they predict is not aligned with users' political views. Put simply, Facebook is making decisions about which political ads to show to which users based on its own priorities, such as user engagement or financial growth. Although Facebook's role was not entirely unpredictable given the previous work on delivery optimization in the context of job ads [4], we confirm it extends to political advertising, a context in which Facebook's choices may have significant negative externalities on political discourse in society at large.

Our investigation presents a new example of an empirical study of a black-box algorithmic system, and the challenges of pursuing such a study. It thus raises questions of the accountability desiderata for ad delivery optimization in the context of political advertising.

Ethics All of our experiments were conducted with careful consideration of ethics. *First*, we obtained Institutional Review Board review of our study, with our protocol being marked as "Exempt". We did not collect any users' personally identifying information from Facebook, and did not collect any information about users who visited our site after clicking on our ads. *Second*, we minimized harm to Facebook users when running our ads by only running "real" ads, i.e., if a user clicked on one of our ads, they were brought to a real-world page not under our control that was relevant to the topic of the ad. In the few cases where the ads pointed to a domain we controlled, the visiting users were automatically and immediately redirected to a real page that we did not control. *Third*, we minimized harm to Facebook itself by participating in their advertising system as any other advertiser would and paying for all of our ads. We registered as an advertiser in the area of "Social Issues, Elections or Politics" [2], meaning our ads were subject to the same review as the ads of other political campaigns. *Fourth,* we minimized the risk of altering the political discourse through careful choices of the ad content, and running approximately the same number of copies of ads for Republican and Democratic candidates, with the same budgets. The total amount we spent on political advertising while collecting data for this paper was minuscule compared to the ad budgets of real campaigns in the same period (likely in the millions of dollars [11]).

2 METHODOLOGY

Our experiments consist of four stages: audience creation, ad creation, collection of data on delivery, and statistical analysis. Here, we briefly describe the decisions made at each stage.

2.1 Audience creation

Facebook's advertising interface allows us to target users based on their inferred political interests. We created two audiences this way, selecting a geographic region centered around a town and Facebook's inferred characterization of interests such as "Likely engagement with US political content (Conservative)" and "Likely engagement with US political content (Liberal)". We further narrowed the targeting by specifying additional required characteristics such as those who are, according to Facebook's characterization, "interested in" topics such as "Donald Trump for President", "Make America Great Again", or "Bernie Sanders". We aimed to approximately match the sizes of liberal- and conservative-leaning audiences for the region by adjusting the targeting radius around a chosen location until the Estimated Daily Reach provided by Facebook was close to matching.

2.2 Ad copy creation

We registered as political advertisers on Facebook (which required confirming our identity and residence in the United States). In our experiments we ran two types of ads: **generic** and **real**.

The **generic** ads did not feature any candidate or a political stance. Instead, they showed an image of the American flag, and the ad copy encouraged the viewers to register to vote, see Figure 1e. All ads appeared to point to our domain `psdigital.info`, but if any user clicked the ad, they would be redirected to the official `fec.gov` webpage.

The majority of the **real** ads replicated the ads run by official political campaigns that we obtained from the Facebook Ad Library [11], see Figure 1a-d. Ads for Bernie Sanders' merchandise store were the only exception, as his campaign had not advertised merchandise on Facebook; we created the ad creative for this ad. Whenever the replicated ad was written in the first person, we changed it to be a third person reference to the name of the candidate. We chose Donald Trump and Bernie Sanders for the ads because at the time of experiment design (early July 2019), they had spent most on Facebook advertising among the major candidates of each party [11].

Each of our experiments consisted of two ad campaigns: one with a copy and/or linked content that is liberal-leaning and the other – that is conservative-leaning. Each of these ad campaigns, in turn, featured two ads that looked exactly the same to the users, but targeted different audiences: one targeting a liberal-leaning audience, and another – a conservative-leaning one. Therefore, each experiment consists of four ads in total.

It is worth noting that using two candidates/parties is sufficient for studying political ads in the U.S. because of the primarily two-party system (Republicans vs. Democrats). Our proposed methodology can be extended to multiple political parties to understand skews in other countries where several parties might be competing in an election.

2.3 Performance optimization and statistics

When creating ad campaigns, advertisers on Facebook are asked to specify their *objective*, or what they are trying to achieve, and the *optimization* that Facebook should use to achieve the objective. Unless stated otherwise, all of our campaigns ran with the "Reach" objective and "Reach" optimization, which according to Facebook's documentation [26], means Facebook would allocate the campaign's specified budget to maximize the number of unique users to whom the ad is shown, rather than to maximize, for example, engagement or "Traffic" (showing the ad to the users most

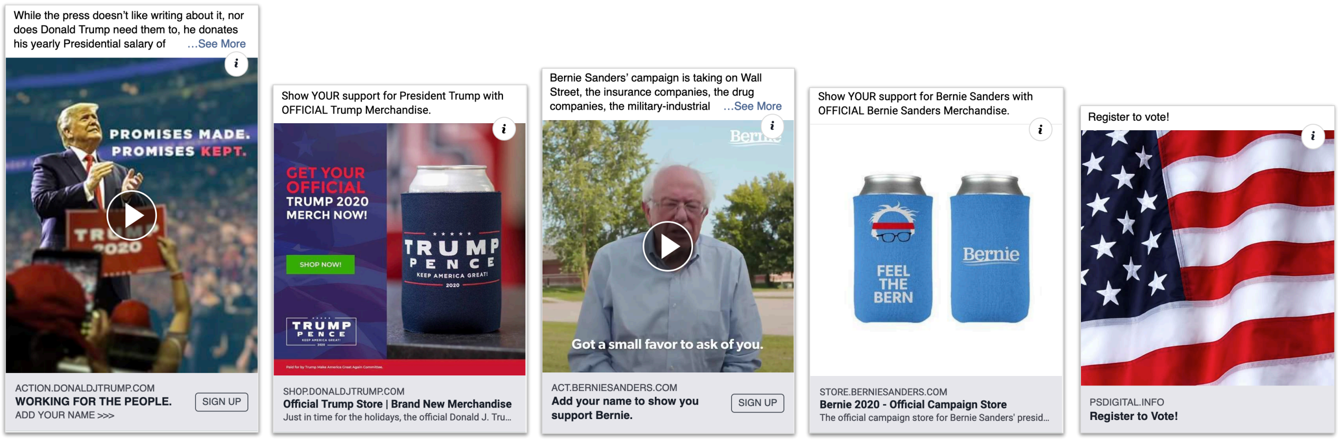

a) Trump - Issues b) Trump - Merch c) Sanders - Issues d) Sanders - Merch e) Generic

Figure 1: Ads used in our experiments concerning political issues and promoting candidates' merchandise.

likely to click). Consequently, our campaigns were charged for impressions rather than for clicks. Although we cannot verify this as Facebook's political ad archive does not reveal the optimization objectives of political ad campaigns, we hypothesize that political campaigns aiming to get their message out to as many people as possible and / or aiming to reach an ideologically diverse audience are likely to use this combination of objective and optimization.

After choosing the audiences and selecting the ad copies, we submitted our campaigns for review. Upon acceptance, the ad campaigns started presenting the ads to users. Using the advertising interface, we tracked the number of users reached by each ad, as well as the cost Facebook charged us for impressions every five minutes over the entire lifetime of the ads.

2.4 Statistical analysis

In the course of this work we compare the fractions of Democrats (or Republicans) among the users exposed to two ads that differ in their content. The comparison process consists of two steps and is based on previous work [4].

First, we estimate the fraction of Democrats in each ad, and the 99% confidence interval around that estimate using the method recommended by Agresti and Coull [3], shown in Equation (1):

$$
\begin{aligned}
L.L. &= \frac{\hat{p} + \frac{z_{\alpha/2}^2}{2n} - z_{\alpha/2}\sqrt{\frac{\hat{p}(1-\hat{p})}{n} + \frac{z_{\alpha/2}^2}{4n^2}}}{1 + z_{\alpha/2}^2/n}, \\
U.L. &= \frac{\hat{p} + \frac{z_{\alpha/2}^2}{2n} + z_{\alpha/2}\sqrt{\frac{\hat{p}(1-\hat{p})}{n} + \frac{z_{\alpha/2}^2}{4n^2}}}{1 + z_{\alpha/2}^2/n},
\end{aligned}
\tag{1}
$$

where $L.L.$ is the lower confidence limit, $U.L.$ is the upper confidence limit, $\hat{p}$ is the observed fraction of Democrats in the audience, n is the total size of the audience exposed to the ad. To obtain the 99% interval we set $z_{\alpha/2} = 2.576$.

Second, we compare whether the fractions in two scenarios are statistically significantly different. Since in the vast majority of our results the confidence intervals do not overlap (easily judged visually from the subsequent figures), the difference is statistically significant.

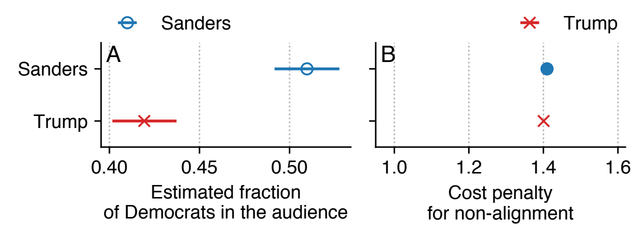

Figure 2: Delivery statistics for ads that look identical to users, but appear partisan to the Facebook classification mechanism. (A) Ads that appear to promote the Democratic candidate are shown more to liberal users and vice-versa. (B) The financial penalty for trying to show an ad that Facebook deems non-aligned; reaching the same number of people in the same audience is up to 1.4 times more expensive.

3 RESULTS

We study what happens when a political campaign places ads on Facebook to an *audience* (the set of targeted users) containing both users who likely agree with the campaign's views (e.g., to solicit donations, or to increase engagement) as well as users who likely disagree with the campaign's views (e.g., to try and change their minds). We do so first by exploring the impact of the ad platform's relevance estimates on generic ads, and then demonstrate the impact on real-world ads by running ads similar to those of actual campaigns.

3.1 Generic ads

We first aim to isolate the impact of the ad platform's relevance estimates, and to avoid any interference from the reactions of the users themselves, which can be different across groups with different political leanings. To do so, we run several copies of ads that appear identical to users (Figure 1e) but differ in the political leaning of the *landing page*, i.e. the webpage that the ad links to. To achieve that, we configure our servers so that when viewed by a real user, the landing page of all ads is a U.S. government website with instructions on how to register to vote. However, when visited by Facebook's web crawler[5], each ad's landing page shows different content: one serves Trump's campaign content, another serves Sanders'. Since all ads look entirely identical to users, any skew in delivery can only be attributed to Facebook's relevance estimates based on the content of the linked website.

We set a budget of $40 per day for each of our two ad campaigns and run them simultaneously from the same advertising account for two days. We run two identical copies of each ad: one targeting a liberal-leaning audience of 3,000 users in a single city, and the other – a conservative-leaning audience of 3,600 users in the same city.

The two audiences are disjoint, so for each ad, we divide the number of users reached in the ad targeting the liberal-leaning audience by the total number of users reached in both copies; the results are presented in Figure 2A. Even though the users see the same ad in both cases—meaning users' explicit or implicit reactions are no more different than chance—we observe that delivery is skewed according to the political leaning of the landing page (with the ad with Sanders' landing page being delivered to the highest fraction of liberal-leaning users, and Trump's, the lowest). We also calculate the *cost penalty* by comparing prices of reaching the first 1,000 users in each audience for the non-aligned ad versus the aligned ad; the results are presented in Figure 2B. We observe that it costs 1.4 times more for the ad with Sanders' landing page (as perceived by Facebook) to reach the same number of users in a conservative audience than for the ad with Trump's landing page. Conversely, it costs 1.4× more for the ad with Trump's landing page to reach the same number of users in a liberal audience than for the ad with Sanders' landing page.

These results show that the content of the landing page—and not only users' reaction or engagement with the ad, or the competition from other advertisers—plays a significant role in Facebook's ad delivery optimization decisions, and can result in both skewed delivery and differential pricing, despite inclusive targeting by the campaign. As a result, two political campaigns running ads concerning the same issue to the same target audience may reach different sub-populations of that audience and at different prices, only because their landing pages are different.

3.2 Real ads

We now explore the implications of ad delivery optimization for real-world ads that differ both in the ad content and landing page. In this experiment, we run two ads (again, one for Trump, and one for Sanders, using ad creatives *a* and *c* from Figure 1 respectively), each to two audiences of over 30,000 liberal- and conservative-leaning users[6] over a period of seven days in August 2019 and with a daily budget of $100 for each ad and audience combination. For these ads, we specify that we only want to show the ad at most once to each user each week, thus preventing the delivery mechanism from showing the ad to the same subset of users repeatedly. Given our audience size and budget, we expected to reach *almost everybody* in the audience by the end of the run.

The results of this experiment are presented in Figure 3. We first focus on panel A, which shows the cumulative number of users reached over seven days. We can observe two notable effects: the smaller reach of the Trump ad targeting the liberal-leaning audience and the Sanders ad targeting the conservative-leaning audience. Both of these non-aligned ads end up delivering to over 20% *fewer* users than their aligned counterparts. Specifically, when the Trump ad is targeted to the conservative audience, it delivers to a total of 21,792 users; when the Sanders ad is run at the same time and targeted to the same conservative audience, it delivers to only 17,964 users. This difference in reach cannot be attributed to an underlying difference in users' likelihood to click on the ads, as we configured the campaigns to pay per ad impression and optimize for reach, not clicks.

Figure 3B shows that despite equal budgets for all ads, the Sanders campaign targeting liberal users slowed down the spending after the second day and, as an effect, spent less than $450 of the allocated $700. In fact, after this point, the campaign did not reach many more users, as seen in panel A.

We turn to panel C, which shows the cumulative cost per 1,000 unique users to help explain why this effect is occurring. We can immediately notice an increasing cost trend for all ads: as the ads run longer, their cost per 1,000 reached users increases substantially. Presumably, this is because Facebook first delivers the ad to the "cheaper" users in the target audience before deciding to spend our budget on the more "expensive" users (recall, we prevented Facebook from delivering ads to users more than once). However, we can observe that the non-aligned ads are again outliers: both show a substantially *higher* cost per 1,000 users, a difference noticeable from the start of the experiment. By the end, when the liberal ad is delivered to the liberal-leaning audience, it is charged $21 per 1,000 users; when the conservative ad is delivered to the same audience, it is charged over $40 per 1,000 users.

Because the delivery rates slow down after the first day, Figure 3C makes the growth of cost per 1,000 also appear to slow down. Therefore, we turn to Figure 3D which shows this growth as a function of the size of reached audience, rather than time. We observe that the growth is rapid and accelerating, especially for non-aligned ads. Finally, in Figure 3E we show that the ratio between the cost a political campaign pays to show their ad to the non-aligned audience and the cost of their competitor showing to the same audience is relatively stable, between 2:1 and 3.5:1.

Overall, Figure 3 emphasizes three findings relevant to the questions we posed: *First*, by the end of the experiment, the two aligned partisan ads reached over 20,000 users, while the non-aligned ads reached many fewer–Facebook limited the delivery of ads whose content did not agree with the audience's inferred political leaning (Figure 3A). *Second*, among two campaigns trying to reach the same audience, the one that Facebook deems non-aligned will pay a significant cost penalty (see Figure 3E). *Third*, while the cost per 1,000 reached users grows with time in a sub-linear fashion (Figure 3C), it grows super-linearly with the number of users already reached (Figure 3D).

[5]We determined Facebook IP addresses by using the IP address blocks advertised by Autonomous Systems numbers owned by Facebook.

[6]Using Facebook's interest based targeting "Likely engagement with US political content (Conservative)" and "Likely engagement with US political content (Liberal)"

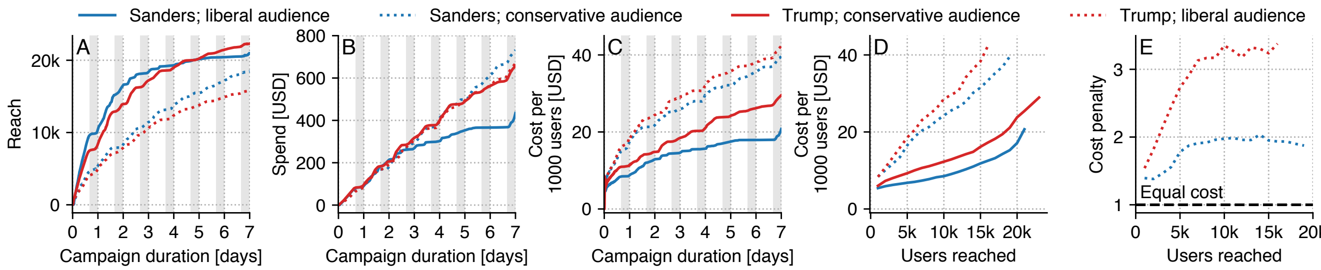

Figure 3: Ads for a political campaign deliver to more users and for a lower cost if the targeted users have the same inferred partisanship. A - the delivery rates are the highest in the beginning of the ad runtime and for aligned audiences. B - initially all ads spend \$100 per day, but the Sanders ad targeting liberal audiences does not spend its full budget. C - the cost of reaching non-aligned audiences is higher. D - the more people have already seen the ad, the more expensive it becomes to show it to even more people and the effect is even stronger for non-aligned audiences. E - the ratio between the cost of a political campaign advertising to a non-aligned audience and their competitor advertising to the same audience.

Taken together, our findings demonstrate the core phenomenon: there's a reach and cost penalty on those campaigns whose political affiliation (or even merely the landing page) is (as inferred by Facebook) not politically aligned with the target audience, as compared to the campaigns targeting the same audience whose political affiliation is inferred to be aligned.

3.3 Robustness of results

In this section we describe our efforts to corroborate our findings and show the robustness of the presented effects to changes in a range of variables.

First, we replicate the method that Ali et al. [4] used to measure racial skew in the delivery of job and housing ads. To that end, we use the Custom Audiences CA_A, CA_B, CA_C, and CA_D described in Table 1 that are based on publicly available voter records from North Carolina. These audiences are designed so that asking Facebook to report delivery statistics by DMA serves as a proxy for obtaining delivery statistics by political affiliation.

We re-use ad creatives from the official Trump and Sanders Facebook pages (similar to Figure 1) and link to the respective campaign's web site. We then run one copy of each ad targeting each of the four Custom Audiences, for a total of 8 individual ads. Our ads are run with a daily budget of \$20 per ad set and use the objective "Traffic" and optimization "Link Clicks" as in prior work [4].

Figure 4 (top row) presents the overall delivery statistics for these two ad creatives, with the delivery statistics of all four instances of each ad aggregated together. We can immediately observe significant differences in delivery: the Trump ad delivers to less than 40% Democrats, while the Sanders ad delivers to almost 70% Democrats. Note that this difference in delivery is despite the fact that all ads are run from the same ad account, at the same time, targeting the same audiences, and using the same goal, bidding strategy, and budget; *the only difference between them is the content and destination link of the ad.* This finding shows that the skewing effect persists even if the advertiser does not explicitly target using the targeting tools offered by Facebook about the user's political leaning. Interestingly, the effect does not persist when we use the FEC donor records instead of voter records to construct custom audiences in the same way (see the bottom row of Figure 4). The absence of a skew for the donor record audiences might suggest that Facebook does not have sufficient information about these users.

Next, we perform a series of additional experiments to verify the robustness of the results to the specific ad copy, audience size, audience geographical location, and the level of detailed targeting. The results are presented in Figure 5, with each experiment in a separate row. We vary three aspects of our experiments:

(1) The size of the audience, as reported by Facebook's Estimated Daily Reach,
(2) The "specificity" of the audience (narrowing the detailed targeting further by attributes such as users' inferred interest in "Donald Trump for President" or "Bernie Sanders" according to Facebook), and
(3) The specific topic of the ad (adding ads that advertise small campaign-branded merchandise that users can purchase, as shown in Figure 1).

We make a number of observations from this experiment. *First*, we observe statistically significant skews in ad delivery along political lines for *all* of our ad configurations. This suggests that such skew is a pervasive property of Facebook's ad delivery system. *Second*, we observe that the skews tend to be less pronounced when the ads are targeting larger audiences (more than 10,000 daily active users). While we do not know the underlying cause of this phenomenon, we hypothesize that the larger audiences provide the platform with a big enough pool of users to afford "relevant" users regardless of their inferred political leaning. On the other hand, we suspect that when running our ads with smaller audiences, Facebook "exhausts" the (small) subset of users in the non-aligned audience (e.g., Sanders advertising to a conservative audience) for whom Facebook believes the ad is, in fact, relevant, and thus pauses or raises the price for delivery, but continues the delivery among the aligned audience.

3.4 Isolating the role of delivery optimization

Ad delivery is a complex process where multiple aspects can influence the makeup of the audience that ultimately sees the ad. Here

DMA(s) [25]	CA_A		CA_B		CA_C		CA_D	
	Dem	**Rep**	**Dem**	**Rep**	**Dem**	**Rep**	**Dem**	**Rep**
Greensboro, Charlotte	70,000	0	0	70,000	70,000	0	0	70,000
Wilmington, Raleigh-Durham, Greenville-(New Bern and Spartanburg)	0	63,137	70,000	0	0	54,000	64,166	0

Table 1: Number of uploaded records for Custom Audiences created using publicly available voter records. We divide the DMAs in the state into two sets, and create two audiences, each with voters registered with one party per DMA set (CA_A and CA_B). We repeat this process with separate voter records (creating CA_C and CA_D). The number of uploaded records does not match, as we uploaded records so as to achieve a match on the Estimated Daily Reach.

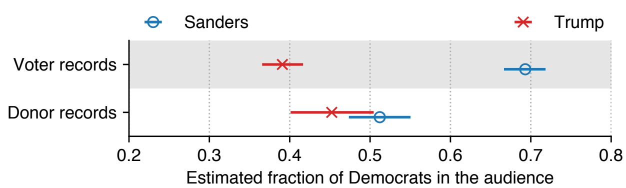

Figure 4: The estimated fraction of Democrats who were shown our ads, targeting both registered voters in North Carolina and political donors. In the case of voter records, the ad delivery to Democrats ranges from 69% for Sanders' ad to only 39% for the Trump's ad. In the case of donor records, we do not see statistically significant differences in ad delivery.

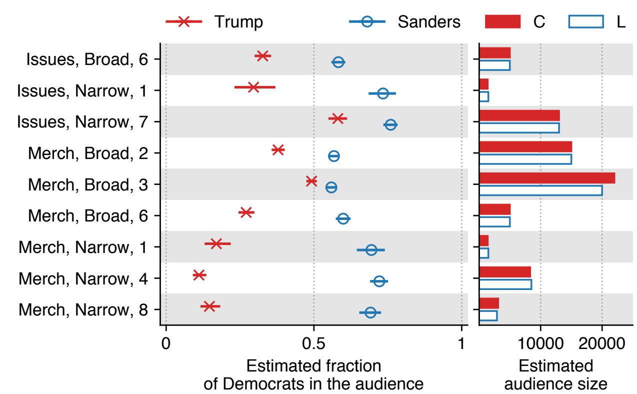

Figure 5: We ran merchandise and issue ads with two levels of targeting specificity (Broad: users with "Likely engagement with US political content (Conservative)" or "... (Liberal)"; Narrow: additional detailed targeting for inferred interest in Donald Trump or Bernie Sanders), and targeting different regions (1: Celina, OH; 2: Dutchess, NY; 3: Lorain, OH; 4: Macclenny, FL; 5: McCormick, SC; 6: Richlands, VA; 7: Saginaw, MI; 8: Slinger, WI). In all cases, Sanders' ads deliver to a larger fraction of Democrats than Trump ads even though they are targeting the same audiences at the same time using the same budgets. The effect is more pronounced for smaller audiences (compare, for example, Merch, Broad, 3 and Merch, Broad, 6).

we discuss how we designed the experiments and analyzed the data in a way that limits the influence of factors other than the delivery optimization itself.

3.4.1 Role of competition. We ran each pair of campaigns targeting a particular audience representing two different political campaigns at the same time and with the same budget. Such a setup is designed to ensure that both campaigns have the same users available for delivery (i.e., if run at different times, the skews could be attributed to different Facebook use patterns by liberals or conservatives) and both are experiencing the same competition from other advertisers (i.e., that it would not be the case that one campaign is under-performing because it happened to run at the same time that another large and wealthy advertiser was targeting those users, whereas another campaign avoided such a collision). Importantly, by selecting "reach" as the campaign optimization goal instead of "traffic", we ensure that Facebook does not have the *direct* monetary incentive to show our ads to users who are more likely to click on them, especially in preference over other ads who might offer such incentive. Thus, running campaigns simultaneously is an effective strategy to isolate the effects of delivery optimization from other extraneous factors. Furthermore, to verify that the skews are not merely the effect of our ads competing with each other, we also re-ran a subset of campaigns separately. The qualitative and quantitative skew effects for those campaigns were similar.

3.4.2 User engagement with our ads. There are a number of ways users can engage with ads, each of which potentially influences future delivery and pricing: reactions (e.g, 'like', 'love') , commenting, and sharing. Facebook's advertising interface reports all such engagements. Additionally, Facebook might be collecting and using *telemetric information*; for example, how long each user spent looking at the ad. This telemetric information is not available to the advertisers (and thus, to us), but might still play a role in ad delivery optimization algorithms.

Some of our ads received reactions, comments, and re-shares from users. We note three important, related observations, that emphasize that our findings about skew in delivery and differential pricing are not merely a function of the ad delivery algorithm's use of user engagement. *First*, we observe consistent skew and price differences in ads that look identical to users, yet trick Facebook into classifying them as partisan (Figure 2). Users do not react differently to ads that appear identical, and, therefore, the entire observed difference can be attributed to Facebook's relevance optimization (and some random effects). *Second*, we observe consistent skew in delivery of ads that had virtually no engagement since they were run on small budgets and only for a few hours (Figure 5). *Third*, we do find a negative correlation between the fraction of positive reactions ('like' and 'love') among all reactions and the price in the longitudinal ads with $\rho = -0.91$, $p_{val} = 0.01$. Taken together, our work shows that although the skew in delivery and differential pricing can be further amplified by users' reactions, their primary source is Facebook's ad delivery optimization's use of classification of the ad and its landing page content.

3.5 Limitations

We note that we can only report on delivery skew that we observed for our own ads; we *cannot* draw any conclusions about how political ads in general (or all ads run by a particular campaign) are delivered. Nonetheless, the fact that we observe strong and statistically significant effects in our small set of ads suggests that the potential negative outcomes for individuals, political campaigns, and society in the context of ad delivery optimization of political advertising are not mere hypotheticals and warrant further scrutiny.

Choice of partisan ads. In this work we chose to advertise two candidates of two opposing American parties in order to clearly present the optimization of delivery to an audience also divided in a binary way (liberal/conservative). However, even within the two parties there can be differing views, as exemplified by the process of primary elections, i.e., choosing one candidate among many to represent the party in general elections. Our current results do not allow us to make strong statements about the potential differential pricing among different candidates of the same party. We leave this investigation, as well as auditing the differential pricing in systems with more than two prominent political parties, to future work.

Role of advertiser's identity. We have repeated a subset of our experiments using another advertising account registered as an advertiser in the area of "Social Issues, Elections, and Politics" and linked to a Facebook page unrelated to the first. Our results were quantitatively and qualitatively similar. This suggests that the effects we observed were not tied to our particular advertising account. Nevertheless, we do not make any statements about the extent to which the observed effects hold when run by real political campaigns with a more established history and larger overall spending than ours.

Audience sizes. We aimed to match our constructed liberal and conservative audiences in size as closely as possible, but the matches are inevitably imprecise as Facebook only provides *estimates* of daily reach[7] rather than audience sizes. Regardless, we always ran both liberal and conservative ads to the same audiences at the same time, so any imbalance in the audience size would affect both ads equally.

4 DISCUSSION

Our findings suggest that Facebook is wielding significant power over political discourse through its ad delivery algorithms without public accountability or scrutiny, and raise open questions in the domains of fairness and accountability in advertising.

4.1 Implications

First, Facebook limits political advertisers' ability to reach audiences that do not share those advertisers' political views in ways that are significantly different from traditional broadcast media. The existence and extent of this skew may not be apparent to advertisers and varies based on their ad's message as well as the destination link used by the campaign. Furthermore, the strength of delivery skews vary for campaigns of different political leanings and targeting different populations, making digital advertising inequitable for political campaigns with identical budgets.

Second, recent moves to restrict political advertisers' targeting options [7, 15, 16]—although valuable from a user privacy perspective [12, 20, 29]—might be undermined by the ad delivery algorithms, and even give companies like Facebook *more* control over selecting which users see which political messages. This selection may be occurring without the users' or political advertisers' knowledge or control. Moreover, it is likely aligned with Facebook's business interests, but not necessarily with societal goals.

Third, today, researchers, regulators, and campaigns lack access to algorithms and data required for a more thorough study of ad delivery implications. In fact, Facebook has actively sought to thwart a recent initiative, NYU Ad Observer,[8] whose goal was to collect such data [21]. Much has already been said about the inadequacy of current ad transparency tools provided by ad platforms for studying ad targeting [24, 31]. Our work draws attention to the need to further expand these efforts to enable scrutiny of ad delivery. It is an interesting open question as to what algorithmic and data sharing advances are needed to enable such auditing while preserving user privacy and ad platform's and advertisers' competitive interests.

4.2 Policy analysis

Today, U.S. law cannot do much, if anything, to *directly* change how platforms deliver political ads. For now, it is likely that the primary regulator of online political ads will not be the government, but rather ad platforms themselves.

The U.S. Congress has addressed conceptually similar "ad delivery issues" in the past, albeit in a different domain. For example, the Federal Communications Commission (FCC) enforces the so-called Equal-Time Rule [1], which originated in 1927 in response to worries that broadcast licensees could unduly influence the outcome of elections. The rule requires that licensees make air time available to all candidates for the same office on equivalent terms. However, the rule only applies to broadcast licensees, and has only narrowly survived constitutional scrutiny in part because it implicates government interests in managing limited broadcast spectrum [6].

Prevailing interpretations of the First Amendment are likely to block efforts to extend the logic of the Equal-Time Rule to digital advertising platforms, which are not regulated like broadcast licensees. As an initial matter, the First Amendment strongly protects political speech, and generally tolerates only narrowly-tailored government regulations [34]. Moreover, the Supreme Court recently declared that "the creation and dissemination of information" constitutes

[7] https://www.facebook.com/business/help/1691983057707189?helpref=faq_content

[8] https://adobserver.org/

speech under the First Amendment [28]. This reasoning, which might expand the "commercial free speech" rights of companies, creates some uncertainty about the government's ability to restrict corporations' use of data in digital advertising.

Looking ahead, it is clear that government regulation of digital political advertising is on firmest legal footing when it requires disclosure about who is speaking to whom, when, and about what [34]. Accordingly, Congress and the FEC can consider transparency requirements that will enable detailed auditing of ad targeting and the delivery optimization as applied to political ads.

4.3 Mitigations

The public, policy makers, researchers, and the campaign managers need more information about the operation of ad delivery algorithms and their real-world effects. Ad platforms could increase transparency around political ads (including key metrics such as targeting criteria, detailed ad metadata, ad budgets, and campaign objectives) to enable further study of the effects of ad targeting and delivery. And they could provide access to and insight into the ad delivery algorithms themselves (including those involved in running the auction, relevance measurement and estimation, and bid and budget allocation on advertisers' behalf), allowing third parties greater ability to study and audit their performance and effect on political discourse. Ad platforms could also disable delivery optimization for political content, or at least allow advertisers to do so. They could also introduce more nuanced user-facing controls for political content delivery. Beyond these mitigations, our work highlights the need for advances that could help set the goals of accountability, fairness, and interpretability in advertising delivery on firm scientific ground. Finally, we call on ad platforms to acknowledge the central role they play in the delivery of political ads, and to collaborate with other key stakeholders—including researchers, political campaigns, journalists, law, and policy scholars—to address that role when it is not aligned with public interests.

ACKNOWLEDGEMENTS

We would like to thank Dean Eckles and David Lazer for their invaluable insights. We are also grateful to the participants and organizers of the REAL ML workshop for their encouragement and constructive feedback. This work was done, in part, while Aleksandra Korolova was visiting the Simons Institute for the Theory of Computing, where she benefited from feedback of participants in the Privacy and Fairness programs, and particularly, from the suggestions of Amos Beimel and Kobbi Nissim. This work was funded in part by a grant from the Data Transparency Lab, NSF grants CNS-1616234, CNS-1916020, CNS-1916153, CNS-1956435, CNS-1943584, and Mozilla Research Grant 2019H1.

REFERENCES

[1] 47 Usc §315 – Candidates For Public Office. https://www.law.cornell.edu/uscode/text/47/315, (Accessed on 2021-01-20).

[2] Ads About Social Issues, Elections Or Politics. https://www.facebook.com/business/help/208949576550051, (Accessed on 2021-01-20).

[3] Agresti, A., and Coull, B. A. Approximate Is Better Than "exact" For Interval Estimation Of Binomial Proportions. *The American Statistician 52*, 2 (1998), 119–126.

[4] Ali, M., Sapiezynski, P., Bogen, M., Korolova, A., Mislove, A., and Rieke, A. Discrimination Through Optimization: How Facebook's Ad Delivery Can Lead To Biased Outcomes. In *ACM Conference on Computer Supported Cooperative Work* (2019).

[5] Angwin, J., and Parris Jr, T. Facebook Lets Advertisers Exclude Users By Race, 2016. https://www.propublica.org/article/facebook-lets-advertisers-exclude-users-by-race/.

[6] Cbs, Inc. V. Fcc. https://supreme.justia.com/cases/federal/us/453/367/, (Accessed on 2021-01-20).

[7] Conger, K. Twitter Will Ban All Political Ads, C.E.O. Jack Dorsey Says, 2019. https://www.nytimes.com/2019/10/30/technology/twitter-political-ads-ban.html.

[8] Creating customer lists with mobile advertiser ids. https://developers.facebook.com/docs/app-ads/targeting/mobile-advertiser-ids/, (Accessed on 2021-01-20).

[9] Eckles, D., Gordon, B. R., and Johnson, G. A. Field Studies Of Psychologically Targeted Ads Face Threats To Internal Validity. *PNAS: Proceedings of the National Academy of Sciences of the United States of America 115*, 23 (2018), 5254–5255.

[10] Facebook: About Facebook Pixel. https://www.facebook.com/business/help/742478679120153, (Accessed on 2021-01-20).

[11] Facebook Ad Library. https://www.facebook.com/ads/library/, (Accessed on 2021-01-20).

[12] Faizullabhoy, I., and Korolova, A. Facebook's Advertising Platform: New Attack Vectors And The Need For Interventions. https://arxiv.org/abs/1803.10099, Workshop on Technology and Consumer Protection (ConPro 2018).

[13] Fowler, E. F., Franz, M. M., Martin, G. J., Peskowitz, Z., and Ridout, T. N. Political advertising online and offline. *American Political Science Review* (2020).

[14] Gelauff, L., Goel, A., Munagala, K., and Yandamuri, S. Advertising for demographically fair outcomes, 2020. https://arxiv.org/abs/2006.03983.

[15] Glazer, E. Facebook Weighs Steps To Curb Narrowly Targeted Political Ads, November 2019. https://www.wsj.com/articles/facebook-discussing-potential-changes-to-political-ad-policy-11574352887.

[16] Google: An Update On Our Political Ads Policy. https://www.blog.google/technology/ads/update-our-political-ads-policy, (Accessed on 2021-01-20).

[17] Green, J., and Issenberg, S. Inside The Trump Bunker, With Days To Go, 2016. https://www.bloomberg.com/news/articles/2016-10-27/inside-the-trump-bunker-with-12-days-to-go.

[18] Imana, B., Korolova, A., and Heidemann, J. Auditing for bias in algorithms delivering job ads. In *Proceedings of The Web Conference* (2021).

[19] Kayser-Bril, N. Automated discrimination: Facebook uses gross stereotypes to optimize ad delivery. *AlgorithmWatch* (2020).

[20] Korolova, A. Privacy Violations Using Microtargeted Ads: A Case Study. *Journal of Privacy and Confidentiality 3*, 1 (2011), 27–49.

[21] Lyons, K. Facebook wants the NYU Ad Observer to quit collecting data about its ad targeting. *The Verge* (2020).

[22] Matz, S. C., Kosinski, M., Nave, G., and Stillwell, D. J. Psychological Targeting As An Effective Approach To Digital Mass Persuasion. *PNAS: Proceedings of the National Academy of Sciences of the United States of America 114*, 48 (2017), 12714–12719.

[23] Merrill, J. Facebook Charged Biden a Higher Price Than Trump for Campaign Ads, October 2020. https://themarkup.org/election-2020/2020/10/29/facebook-political-ad-targeting-algorithm-prices-trump-biden, (Accessed on 2021-01-20).

[24] Mozilla. Facebook's Ad Archive Api Is Inadequate, April 2019. https://blog.mozilla.org/blog/2019/04/29/facebooks-ad-archive-api-is-inadequate/.

[25] Nielsen dma® regions. https://www.nielsen.com/intl-campaigns/us/dma-maps.html, (Accessed on 2021-01-20).

[26] Reach Objective. https://www.facebook.com/business/help/218841515201583, (Accessed on 2021-01-20).

[27] Ribeiro, F. N., Saha, K., Babaei, M., Henrique, L., Messias, J., Oana Goga, F. B., Gummadi, K. P., and Redmiles, E. M. On Microtargeting Socially Divisive Ads: A Case Study Of Russia-linked Ad Campaigns On Facebook. In *Conference on Fairness, Accountability, and Transparency* (Atlanta, Georgia, USA, January 2019), ACM, pp. 140–149.

[28] Sorrell V. Ims Health Inc. https://supreme.justia.com/cases/federal/us/564/552/, (Accessed on 2021-01-20).

[29] Speicher, T., Ali, M., Venkatadri, G., Ribeiro, F. N., Arvanitakis, G., Benevenuto, F., Gummadi, K. P., Loiseau, P., and Mislove, A. On The Potential For Discrimination In Online Targeted Advertising. In *Conference on Fairness, Accountability, and Transparency* (New York, New York, USA, February 2018).

[30] United States V. Internet Research Agency. https://www.justice.gov/file/1035477/download, (Accessed on 2021-01-20).

[31] Upturn. Leveling The Platform: Real Transparency For Paid Messages On Facebook, 2018. https://www.upturn.org/reports/2018/facebook-ads/.

[32] Venkatadri, G., Sapiezyński, P., Redmiles, E. M., Mislove, A., Goga, O., Mazurek, M., and Gummadi, K. P. Auditing Offline Data Brokers Via Facebook's Advertising Platform. In *International World Wide Web Conference* (2019).

[33] Weintraub, E. L. Don't Abolish Political Ads On Social Media. Stop Microtargeting., November 2019. https://www.washingtonpost.com/opinions/2019/11/01/dont-abolish-political-ads-social-media-stop-microtargeting/.

[34] Wood, A. K., and Ravel, A. M. Fool Me Once: Regulating Fake News And Other Online Advertising. *Southern California Law Review 91* (2017), 1223.