

Liberini, Federica; Redoano, Michela; Russo, Antonio; Cuevas, Ángel; Cuevas, Rubén

Working Paper

Politics in the Facebook Era - Evidence from the 2016 US Presidential Elections

CESifo Working Paper, No. 8235

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Liberini, Federica; Redoano, Michela; Russo, Antonio; Cuevas, Ángel; Cuevas, Rubén (2020) : Politics in the Facebook Era - Evidence from the 2016 US Presidential Elections, CESifo Working Paper, No. 8235, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/216631>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Politics in the Facebook Era Evidence from the 2016 US Presidential Elections

*Federica Liberini, Michela Redoano, Antonio Russo, Angel Cuevas,
Ruben Cuevas*

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Politics in the Facebook Era

Evidence from the 2016 US Presidential Elections

Abstract

Through social media, politicians can personalize their campaigns and target specific groups of voters with an unprecedented precision. We assess the effects of such political micro-targeting by exploiting daily advertising prices on Facebook during the 2016 US presidential campaign. We measure the intensity of online campaigns using variation in ad prices charged to reach certain audiences, defined by political orientation, location, and demographic characteristics. We address two fundamental questions: How intensively did social media political campaigns target each audience? How large were any effects on voters? We find that micro-targeted political ads on social media had significant effects when based on geographical location, ideology, ethnicity, and gender. Exposure to these ads made individuals less likely to change their initial voting intentions, particularly among those who had expressed an intention to vote for Donald Trump. We also find that micro-targeted ads reduced turnout among targeted liberals, whereas they increased turnout and support for Trump among targeted moderates.

JEL-Codes: D720, M370, D910.

Keywords: social media, political micro-targeting, elections, advertising, populism, polarization.

Federica Liberini
Department of Economics
University of Bath / United Kingdom
fl571@bath.ac.uk

Michela Redoano
Department of Economics
University of Warwick / United Kingdom
michela.redoano@warwick.ac.uk

Antonio Russo
School of Business and Economics
Loughborough University / United Kingdom
A.Russo@lboro.ac.uk

Angel Cuevas
Department of Telematic Engineering
University Carlos III / Madrid / Spain
acrumin@it.uc3m.es

Ruben Cuevas
Department of Telematic Engineering
University Carlos III / Madrid / Spain
rcuevas@it.uc3m.es

April 17, 2020

We thank Dan Bernhardt, Greg Crawford, James Fenske, Raymond Duch, Marko Koethenbuerger, Massimo Morelli, Jens Prüfer, Francesco Sobbrío, Pilar Sorribas-Navarro and participants to presentations at ASSA Meeting (San Diego), Social Media Workshop (Nuffield College, Oxford), Tilec Conference on Governance of Big Data and Artificial Intelligence (Tilburg), GSE Summer Forum (Barcelona), Warwick, ETH Zurich, EUI, Bath, Brescia, Dr@w Forum (Warwick), Annual Conference NICEP (Nottingham), Public Economic Theory Conference (Hue), Workshop on Political Economy (University of Bolzano), Political Economy Workshop (University of Warwick), IV Media Bias Workshop (WZB Berlin), Workshop on Political Economy (University of Sheffield), X Workshop on Institution, Individual Behavior and Economic Outcomes (Alghero).

“I doubt I would be here if it weren’t for social media, to be honest with you”¹
(President Donald Trump, October 20, 2017)

1 Introduction

Social media, such as Facebook, Twitter and YouTube, have recently become essential tools for political campaigning. According to several reports, these platforms may have played a decisive role in the outcomes of two pivotal 2016 votes: the UK referendum on EU membership, and the US presidential election.² The growing importance of social media raises questions about the impact that this new channel of political communication might have on democratic institutions.³ However, direct evidence of the effect on voters from exposure to political campaigns on social media is still limited. In this paper, we focus on a specific, and previously unexplored, feature of these digital platforms: the fact that they enable campaigns to reach users with *micro-targeted* political advertisements.

Political candidates began using of social media for campaigning to a significant degree during the 2008 US presidential election. Since then the role of social media in the worldwide political arena has grown considerably. As of today, this role is arguably as relevant as – if not more than – that of traditional media.⁴ One of the reasons for the growing importance of social media in the context of political advertising is the rapid growth of the number of individuals who use these platforms, and rely on them for political information.⁵ According to the Pew Research Center ([2016a], [2016b] and [2016c]), in 2016 about 70 percent of Americans used Facebook, which constituted the third most frequently used source of political information during the presidential campaign.⁶ Furthermore, thanks to technolog-

¹Quote retrieved from <https://www.politico.com/story/2017/10/20/trump-social-media-election-244009>.

²For example, the US Senate Select Committee on Intelligence released Volume 2 of its report on Russian Active Measures Campaigns and Interference in the 2016 US Election available at <https://www.intelligence.senate.gov/sites/default/files/documents/ReportVolume2.pdf>

³According to a report published in October 2018 by the UK Parliament’s Digital, Culture, Media and Sport Select Committee investigating the manipulation of social media in elections, an unknown organization spent more than £250,000 in 2018 on Facebook ads that reached over 10 million people in the UK, and pushed for a far harder “Brexit” than (then-)Prime Minister Theresa May had sought.

⁴PEW reported that between 2016 and 2017 the gap between television and online news consumption narrowed from 19 percentage points to 7 percentage points.

⁵Some 20 percent of Americans have a Twitter account. About one-quarter of such individuals report that “a lot” of what they see on social media relates to politics. Figures are somewhat similar for European countries, with a 2016 Eurobarometer survey reporting that 40 percent of Europeans use social media daily, and that about 33 percent (16 percent) of Europeans indicate that the internet (social media) is the major source of “most of their news on national political matters.”

⁶At the end of the 2012 US Presidential campaign, the Pew Research Center reported that 12 percent of Americans had regularly received their campaign news from Facebook. This put the social media site on par with national newspapers. By 2016, these figures had grown substantially. At that time, an estimated

ical advances and the extensive availability of user-generated data (including information on individual interests, networks, and behavioral patterns), platforms like Facebook enable politicians to reach very specific audiences with very specific messages, and to address particular groups with a degree of precision that arguably was not previously possible. This capability is generally referred to as *micro-targeting*.

The aim of this paper is to investigate how micro-targeted political ads on social media affect voters. To address this question, we first collect a novel database of Facebook’s advertising prices during the 2016 US presidential campaign, and we show that they serve as a good proxy for political campaign intensity. To address the effect of the Facebook political campaign on voters, we then combine these data with individual-level information on voter behavior from the American National Election Studies (ANES) electoral survey. As far as we are aware, our paper offers the first attempt in the literature to measure (i) to what extent political campaigns used micro-targeted ads on social media to reach specific groups of users, audiences characterized by common ideology, location and demographics, (ii) the effect that these ads had on voters, in terms of electoral turnout, choice of candidates, and likelihood of changing their minds about which candidate they preferred.

Targeting voters with political ads is, of course, not a new concept. On television, for example, political campaigns target potential voters by placing ads during shows that attract viewers who are likely to have a certain political ideology, or those who belong to a certain demographic group. The key novelty of social media is that it enables campaigns to pinpoint their efforts far more precisely. New technologies allow campaigns (and all advertisers) to exploit extensive quantities of user-generated data, and to trace dynamic behavioral patterns, interests, and networks.⁷ For example, to facilitate the identification of different audiences, in 2016 Facebook began classifying its US users in terms of political orientation (conservative, liberal, and moderate) and interests (on specific candidates, issues, or initiatives).

The first challenge we face is measuring the intensity of political ads campaigns on different Facebook audiences. Unlike traditional media, digital platforms are not required to disclose the source, volume and content of ads, or the amounts spent by campaigns or other political actors (regulated in the US by the Federal Election Campaign Act of 1971).⁸

79 percent of adults who said that they accessed online information (roughly 68 percent of all Americans) used Facebook, and an estimated three-quarters of these Facebook users accessed the site daily.

⁷Social media offer also other advantages, with respect to traditional media, which we do not explore in this paper. For example, social media allow candidates to gain real-time feedback during the course of their political campaigns, and to adjust strategies accordingly. Social media also offer a way to spread campaign messages quickly and at low cost, by exploiting users’ networks.

⁸Facebook recently announced that it will implement measures to increase transparency. This follows reports that more than 3,000 ads linked to Russia, addressing divisive social issues and allegedly favoring Donald Trump, circulated on its platform during the 2016 presidential campaign. These ads were reportedly seen by an estimated 10 million people. Facebook has vowed to enable users to see the content of political ads

To circumvent this lack of information, we exploit variation in Facebook advertising prices to different audiences, defined by location, political orientation and several demographics, for a period encompassing the 2016 US presidential election. Our approach is based on the following idea: the extent to which political campaigns targeted specific audiences on Facebook should affect the demand for ads and thus be reflected in the prices charged to reach such audiences. Specifically, campaigns targeting a given audience should produce a positive demand shift and, for a given supply (which is largely stable during our observation period), push up the price of ads to that audience. By analyzing the variation in prices across time and across audiences, therefore, one can gather information about politicians’ strategies, revealing, for example, the value that politicians place on different types of voters in different states (e.g., partisan or swing states) at different points in time (e.g., after an election poll). Following this logic, we use variation in advertising prices as a measure of the *intensity* of Facebook political campaigning on narrowly defined audiences.

We construct a novel database, gathering daily unit prices for online advertising on Facebook, by scraping the platform’s marketing API (i.e. Facebook’s Ad Manager website). Prices are specific to audiences characterized by political ideology, political interest (in a given candidate), gender, age, ethnicity, education level, and location (within the United States). The database covers a period of approximately 12 months, from April 2016 until March 2017. Based on these data, we construct audience-level measures of intensity of political advertising on social media during the electoral period.

Our second objective is to study the effect of Facebook ads on the behavior of voters. To address this point, we combine the information described above with data from the 2016 American National Election Study (ANES, [2017]). These data contain individual responses regarding behavior at the polls, political preferences, demographic and economic characteristics of a sample of 2,414 American voters interviewed before and after the 2016 US presidential elections. It is also the only countrywide survey that reports information on social media and Facebook usage. We can therefore compare the behavior of respondents belonging to different audiences, conditional on (i) the respondent’s use of Facebook as a source of political information and (ii) the intensity of advertising to the respondent’s audience (as discussed above). To identify the effect of such ads, we rely on the fact that users exposure to Facebook (i.e. how often they access the platform to gather political information) is a pre-determined time invariant characteristic, whereas intensity of the campaign during the period going from the interview through to the elections depends on the *random* assignment

posted on its site, making them visible to any Facebook user. In October 2017, US senators Amy Klobuchar, Mark Warner and John McCain introduced the “Honest Ads Act” bill as a response to the scandal, with the intent of regulating online campaign advertising.

of respondents to different interview dates.

To our knowledge this is the first contribution in the literature comparing different micro-targeting dimensions, both in term of advertisers’ strategies as well as on their effect on (voting) behavior. Our analysis indicates that micro-targeted political ads on Facebook matter and have a significant effect on voter behaviors, particularly when targeting users based on location, political orientation and ethnicity, or gender.

Our findings contribute to the important current debate on the effect of social media on political polarization. As an indirect test for the “echo chamber” theory, according to which social media amplifies polarization by selectively exposing individuals almost solely to content and networks espousing like-minded preferences and ideologies (Allcott et al., [2020]), we show that intensely targeted Facebook ads make users less likely to change their minds. This effect, which is robust to different specifications, appears to be heterogeneous across different groups of individuals. Those who are most likely to stick to their initial positions when highly targeted on the social media platform are men, conservatives and non-white voters. An increase of 10 percent in our measure for the intensity of political campaign exposure reduces the likelihood of changing one’s voting intention by 3 percent overall. This effect amounts to 4.5 percent for men, 0.3 percent for conservatives and 2.3 percent for non-white voters. Furthermore, exposure to social media reduces the probability of changing one’s initial political position among those who hold conservative views and belong to a highly targeted audience, as compared to those who also hold conservative views but are not exposed.

Our data do not provide information on either the identity of the political advertisers or the content of the ads. In spite of this, our results indicate that advertising on Facebook had contrasting effects for the two campaigns. Facebook advertising was effective at convincing Trump’s core supporters to turn out on election day. By contrast, Facebook ads had a negative effect on turnout among liberals. Campaigning on Facebook swayed moderate and less-informed voters toward voting for Trump. By contrast, Facebook advertising reduced turnout among Clinton supporters. These findings are consistent with existing evidence that Clinton’s campaign relied largely on traditional media, while Trump’s primary communication channels were social media, such as Facebook and Twitter.⁹

Our analysis then incorporates a simulation exercise to explore the effect on both voter mobilization and persuasion that would have been triggered by a reduction in the intensity

⁹Using data from the Federal Election Commission, Williams et al. [2017] calculated that during the period from July 1 to November 30, 2016, 8 percent of Clinton’s overall media expenditures went to digital media; the comparison figure for Trump was 47 percent. Figures from Bloomberg [2018] for the same period show that Clinton spent \$28 million on Facebook ads, while Trump spent \$44 million. Overall, Clinton spent more: \$252 million for all advertising, compared Trump’s overall spend of \$176 million.

of political campaigning on Facebook should all else have remained the same. The results from this exercise indicate that a reduction in the intensity of the Facebook campaign would have increased turnout in blue states, and reduced in red states. Under this scenario, Trump would have experienced a sizeable loss of support in those states that were pivotal in the election.

The remainder of the paper is organized as follows: Section 2 discusses the related literature. Section 3 introduces our measure for political campaign intensity using online ads prices. Section 4 analyzes the effect of political ads on voter behavior. Section 5 presents the result of a simulation exercise on the potential impact of political advertising on actual voting behavior, a discussion of the main results, and concluding remarks.

2 Related Literature

This study contributes to a number of literature streams. First, it relates to the literature on the effect on voters from exposure to the media. Most of the literature focuses on traditional media. Several studies find that television campaigns are effective in persuading voters, but less so in mobilizing voter turnout (see, e.g., Gordon and Hartmann [2016]; Gerber et al. [2011]; Krasno and Green [2008]; Della Vigna and Kaplan [2007]). Strömberg [2008] shows that the introduction of radio broadcasting caused an increase in US voter turnout. However, other studies have found that the availability of new media reduces political participation (e.g., Gentzkow [2006] on the effect of television, and Falck et al. [2014] on the effect of broadband internet).

More recently, the above literature has turned its attention to the effects of social media.¹⁰ Bond et al. [2012] find a strong effect on electoral turnout from a mobilization campaign ran on Facebook during the 2010 US congressional elections. Allcott and Gentzkow [2017] show that social media was an important source of information during the 2016 US presidential campaign, and that most American adults were exposed to at least one piece of fake information on the internet. Allcott et al. [2020] study the relationship between exposure to echo chambers on social media and voter polarization, while Enriquez et al. [2019] study voter responses to incumbent performance information disseminated through social media. Aside from the effects on elections, studies have also provided evidence of the role of social media in spreading information that affects individuals participation in street protests (Enikolopov et al., [2019]) and violence against minorities (Müller and Schwartz,

¹⁰A related literature studies how political candidates' use social media. Petrova et al. [2017] track candidates' entry on Twitter during the latest US congressional elections, and find that opening a Twitter account helps new candidates increase donations by up to 3 percent. See also Gainus and Wagner, [2014], and Bright et al. [2017].

[2019]).

We contribute to the above literature in several ways. First, we provide direct evidence of the effect of micro-targeted political advertising campaigns on multiple categories of voters, distinguished by demographics and political orientation. Furthermore, we link exposure to micro-targeted ads on Facebook to individual voting behaviors. We measure the effect on election turnout, choice of candidates, and the likelihood that individuals change their minds about which candidate they prefer during the course of the campaign.

Several authors have hypothesized a link between the polarization of political opinions and the diffusion of social media. Sunstein ([2009], [2017]) and Pariser [2011] argue that social media favors the creation of “echo chambers” in which individuals are exposed only to like-minded sources of information. Furthermore, by giving candidates the opportunity to tailor their messages to the specific views of each recipient, digital media may induce campaigns to take a more extremist tone (Hillygus et al., [2008]). However, relying on evidence from demographics, Boxell et al. [2017] cast doubt on the hypothesis that the Internet and social media cause polarization. Barbera [2015] uses data from Twitter to argue that social media actually expose users to information from individuals with whom they have weak ties and, thus, with whom they are also likely to have contrasting political views. Although polarization is not the central issue in our study, we find that exposure to political ad campaigns on social media reduces the likelihood that individuals change their mind about a political candidate. This evidence of hardening of views lends some support to the hypothesis that social media exposure increases polarization.

Finally, our study belongs to a novel multidisciplinary literature that leverages Facebook Marketing API data to gain insights into socioeconomic issues. González Cabañas et al. [2018] quantify how many Facebook users in the European Union (EU) may be defined as having “sensitive” interests (according to the definition provided by the EU General Data Protection Regulation). They show that 73 percent of EU Facebook users may be defined as having such sensitive interests.¹¹ Zagheni et al. [2017] and Dubois et al. [2018] have used Facebook marketing API data to study the nationalities of immigrants and their assimilation in various countries. Garcia et al. [2018] use Facebook marketing data to perform a comprehensive analysis of the gender divide worldwide.

¹¹In earlier studies, Liu et al. [2014] perform a quantitative analysis of the auction system used by Facebook. The findings demonstrate the huge difference of advertising prices (i.e., CPM and CPC) over different demographic parameters. Saez-Trumper et al. [2014] generate a model to reflect how Facebook users’ activity (e.g., likes, shares) propagates among friends; the authors use a second simplistic model to estimate the number of ad impressions received per user.

3 Online Advertising Prices and Political Campaign Intensity

During electoral periods, the demand for political advertising surges as campaigns strive to reach voters. Audiences are thus exposed to political ads, which partly substitute nonpolitical ones. Given a limited supply of ads, the surge in advertising demand results in higher prices, particularly for audiences that are highly valuable to politicians. This phenomenon has been well documented in the US (Moshary, [2017]).¹² As Figure D.1 suggests, campaigns spend huge amounts to buy advertising space in the media.

The process described above characterizes all media, but social media have a unique feature in that they allow political campaigns to finely target their messages to specific audiences at a very granular level. Based on the above considerations, we argue that observing fluctuations in the prices of ads on social media provides a proxy for the intensity of micro-targeted political campaigning. In the following, we provide a simple theoretical model to formalize the relationship between prices and campaign intensities, and we show preliminary evidence.

3.1 Theoretical Background

Consider a social media audience, i.e. a group of users defined by a set of characteristics (e.g., income, age, gender, and political preferences). We assume that users are willing to tolerate a limited number of ads while browsing the platform. We let this reservation level be \bar{q} . The marginal cost of an ad to the platform is zero and we assume the demand by advertisers is large enough to saturate the reservation quantity of ads. Given these conditions, the total quantity of ads the platform exposes each consumer to is inelastic and equals \bar{q} .

We consider two periods, indexed by $i = \{N, E\}$. Period N is a non-electoral period, so there are no political candidates (C) that try to reach voters. In this period, there is only a generic (non-political) advertiser G , whose inverse demand for ads (per user) on the platform is $p_G(q)$. We assume that this function is decreasing, because the marginal value of an impression on a consumer is decreasing.¹³ We represent this demand in the center panel of Figure I. Given our assumptions, in the non-electoral period the equilibrium (denoted

¹²Unlike the US, other countries tightly regulate political advertising on traditional media. For example, paid television and radio political advertising is generally not allowed in EU member countries. The United Kingdom also forbids paid television and radio from airing political advertisements, but allows political parties a limited and regulated number of political broadcasts in the period immediately preceding the elections. Canada allows political broadcasts, but strictly regulates airway access.

¹³For example, assuming advertising is informative, the more an ad is shown to a consumer, the higher the probability the consumer is already aware of its content. See, e.g., Ambrus et al. [2016].

with subscript N) is such that $q_G^0 = \bar{q}$. Period E is the electoral period. In addition to the demand from G (that we assume invariant), there is a demand for ads by political campaigns, denoted $p_C(q)$. We represent this additional demand in the right panel of Figure I, whereas the left panel depicts the aggregate demand for ads. The equilibrium price of ads, p_E , is higher than in period N (we assume the users’ reservation level of ads does not change). Furthermore, political ads substitute generic ones. As Figure I suggests, the more intense the demand increase, the higher the equilibrium price and the quantity of ads allocated to political campaigns (superscript $+$). Therefore, a higher price for advertising in the electoral period is also indicative of higher exposure to political ads.

We have assumed the supply of ads to a given audience is inelastic, but the correlation between prices and exposure to political ads should exist even if supply is elastic. Nevertheless, descriptive evidence shown in Arrate et al. [2018], corroborates the assumption of inelastic supply, showing that the share of ads over total posts (i.e. content) and ads on individual Facebook accounts remained essentially constant in the period from October 2016 to May 2018 (see Figure D.2; in the figure, the electoral weeks are circled in yellow). The number of active users of the main political orientations has remained very stable during our period of observation. See Figure D.3.

Insert Figure I in about here

3.2 Ad prices as a proxy for the intensity of political advertising

We cannot observe the total volume of political ads to a given audience or to a specific user. Hence, we cannot directly measure the intensity of political ad targeting. However, as argued in the previous section, ad prices should respond to increased demand by political advertisers. We therefore collect ad price data on Facebook and use them to proxy for the intensity whereby political advertisers target different audiences on the platform.

Facebook classifies users based on a set of observable characteristics (e.g., demographics, interests, etc.). A user connecting to Facebook generates ad impression opportunities that the platform sells to advertisers via online auctions, virtually in real time.¹⁴ When intending to bid for an audience on Facebook, an advertiser connects to Facebook’s Ad Manager website and selects an audience, defined along several dimensions. At the time of our data collection, the website showed the range of prices that advertisers pay for ads to the selected audience (Facebook has recently modified the information provided by the Ad Manager, see Appendix B). We collected this range of prices for a large number of audiences, at a daily level, from April 2016 until March 2017, focusing on CPM (Cost-Per-Mille, i.e., the price of

¹⁴See Appendix A for a more detailed description of Facebook’s auction system.

one thousand impressions) as the price metric.¹⁵ We define audiences based on the following dimensions: location (by US state), ideology (conservative, liberal or moderate), age bracket, gender, ethnicity, and education.¹⁶ In the empirical analysis, we focus on the median of the range of prices for each audience.¹⁷ We emphasize that our ad price data are aggregated at the audience level, and do not include any personal data from individual users (see Appendix B.1).

We construct our price indexes to account for a number of issues. First, the relevance of a given political audience may vary according to the context (e.g., in which US state members of the audience reside) and over time. For instance, a “moderate” user in a swing state is more likely to be pivotal than a similar user in a “safe,” partisan state. Secondly, we need to account for the fact that political orientation and preferences for goods and services may be correlated, and that non-political advertisers might bid for the audience with that profile for reasons unrelated to the elections. Notably, recent research shows that digital footprints, such as Facebook “likes,” are strongly correlated with personal attributes and preferences (Kosinski et al., [2013]); as a result, political orientation can serve to identify individuals’ preferences for certain goods and services. This aspect limits the comparability of price fluctuations across different audiences. Finally, Facebook ad prices may be affected by an audience’s size and specificity. For instance, advertisers may be less willing to pay for a broadly defined audience if they seek to reach very specific targets. On the other hand, few advertisers may compete to reach an exceedingly narrow audience (Levin and Milgrom, [2010]). Because the size of the audience matters for the equilibrium price, ensuring that the number of users within a certain audience does not vary over our observational period is crucial.

We account for these issues by building two different price indexes. The first is a relative price index: the relative CPM (p_{sci}). This index is obtained by normalizing the CPM of an audience located in a U.S. state (s), with a specific demographic characteristic (d) and political ideology (i), by the CPM of the general audience located in the same U.S. state and with the same demographic characteristic.

I) Relative CPM for audience defined by State (s) demographics (d) Political Ideology

¹⁵The “M” in CPM represents the Roman numeral for 1,000. The Ad Manager provides two additional price metrics for ads: cost-per-click (CPC) and cost-per-action (CPA). CPC is the rate that websites charge advertisers every time someone clicks on an ad. CPA is the rate that websites charge advertisers every time a user clicks and takes a specific action (such as page likes, video views, mobile app installations, etc.). We focus on CPM because it is the most commonly used measure.

¹⁶This means that our base audience will be, for example, Facebook users in Texas, who have been identified to be conservative. Further narrowing of this audience would mean, for example, distinguishing between female conservative users in Texas, or liberal users who live in California and have completed college.

¹⁷See Appendix B for more details on our data collection procedure.

(i):

$$p_{sdi} = \frac{P_{sdi}}{P_{sd}} \quad (1)$$

The intuition behind this index is that geographic location and demographic characteristics should determine the attractiveness of an audience to generic advertisers. That is, these characteristics serve as proxies for users’ preferences and service. However, ideology should matter mostly (though, possibly, not exclusively) to political advertisers. In other words, conditional on a geographic location and a given set of demographic characteristics (such as gender, age and education), this index can proxy for the intensity of political campaigns under the assumption that the price variation explained by political ideology is driven by demand from political campaigns.

Because political ideology is identified by Facebook through the user’s revealed preferences (i.e., “likes” of posts or articles; and information provided on education level, profession, relationship status, and favorite movies, books, sports, etc.) and network, the political subsample composition is rather time invariant.¹⁸ Thus, the observed variation in our relative price measure is not driven by variation in the number of users with a given political ideology (see Figure D.3 in the Appendix). This rules out endogeneity of political ideology to the electoral campaign (at least in the short term). Figure D.3 shows that moderates tend to be the majority among users in swing states, whereas liberals and conservatives tend to be the majority in safe Democrat and Republican states, respectively. We find almost no variation in the number of users in each category over time.

We also build an additional index, which captures the run-up of the price for a given audience, defined by demographics and political ideology, over a certain time interval, specified by the time between the week the individual was interviewed by the ANES survey and the week of the election.

II) CPM runup for audience defined by State (s) Demographics (d) Political Ideology (i), between Election Week (e) and Interview Week (w):

$$r_{sdp,w} = \frac{p_{sdp,e}}{p_{sdp,w}} \quad (2)$$

A first advantage of this index is that it should not be subject to the influence of generic advertisers that use political ideology as a proxy for users’ preferences for goods and services. The reason is that the CPM run-up exploits price variation over time *within*

¹⁸The Facebook audience defined on the basis of the users “behavior” displays much more variation over time. That is because the “interest” of Facebook users on a specific topic varies, according to their actions, such as clicking on, commenting on, liking, reading and watching posts.

a given politically defined audience. The relevance of this audience to generic advertisers should not vary over the electoral period.¹⁹ Hence, changes in the index should be driven only by changes in the relevance of the audience to political advertisers.

An additional advantage of the CPM run-up is that, when assigned to a respondent of the ANES political survey, it exploits the random assignment of the respondent to the date of the interview, and the subsequent variation in exposure to the social media political campaign. Therefore, despite being unable to observe how many ads each respondent has visualized during the electoral period, we can, by using this index, exploit the randomness of the duration between the two dates used to compare differences in the responses to voting questions.

3.2.1 Discussion and Some Illustrative Examples

Presumably, political campaigns do not attach the same importance to the same type of audiences in different states. By analyzing the variation of ad prices across politically oriented audiences in different states we can obtain a first picture of how ad prices reflect the strategies of presidential campaigns on social media. Figure II shows weekly relative CPM ad prices (as defined above), to audiences defined by geographic location and political orientation, in the weeks preceding the 2016 election. We concentrate on three states: New York, a consistently Democratic state; Florida, a swing state; and Texas, a consistently Republican state.²⁰

The figure indicates that campaigns heavily targeted moderate voters in Texas and even more in Florida. Liberal voters are the most expensive in New York, and least expensive in Texas. By contrast, in the latter state conservatives are significantly more expensive than liberals and virtually as expensive as moderates. The story this picture seems to tell is quite intuitive: moderate voters tend to be highly valuable to competing campaigns, especially in swing states. On the other hand, the high relative prices for voters belonging to the dominant orientation in partisan states is consistent with dominant party candidates focusing on getting out the vote of their own supporters there, as well as with opposing candidates trying to discourage such supporters from voting.

¹⁹The relevance of ideologically defined audience to generic advertisers would only change if the preferences for their generic product or service were to change within the short period of times spanned by our data, which is plausibly unlikely.

²⁰See State Historical Presidential Election Information from “270 to Win”, <https://www.270towin.com/states/>.

3.3 Evidence on the link between ad prices and electoral competition

To show that ad prices effectively capture the intensity of political advertising on social media, we provide two further pieces of preliminary evidence.

Ad prices and electoral cycles. First, we show systematic price increases around the date of elections. We build a worldwide dataset of Facebook ad prices, during an interval of one-and-a-half years (from July 2015 to February 2017). We focus only on democratic countries, as determined by the Freedom House’s Electoral Democracy Rating, which leaves us with 45 elections during the observation period. We collect the average price of ads in each country at a daily level for 200 days before and after each election. As illustrated by Figure IV, average prices sharply rise prior to elections (by about 25 percent on average), and steadily drop afterwards, eventually returning to their initial level. This finding is in line with the assumption that political campaigns enter the advertising market immediately before the election, and leave in the subsequent period. See Appendix C.1 for further details.

Insert Figure IV in about here

Winning probabilities and ad prices. We examine fluctuations in US state-level Facebook ad prices during the 2016 Presidential election campaign, with daily frequency. showing that these prices are correlated with changes in the candidates’ winning probabilities estimated by the FiveThirtyEight.com 2016 Election Forecast. We estimate a dynamic panel model, that relates the daily ad prices (CPM) in each state to the absolute value of the lagged difference among daily winning probabilities of the two candidates in that state. See Appendix C.2 for details. Table I shows that ad prices tend to increase in response to decreases in the relative probability of Trump winning. For instance, in Republican states (where the Republican candidate typically leads the polls), the price appears to increase when the difference between the two candidates narrows (i.e., Clinton is catching up). The opposite applies to Democrat states. This evidence is consistent with the notion that the two candidates devoted a quite different degree of attention to campaigning on social media. Note, however, that this evidence does not imply that our measures of campaign intensities are endogenous to *individual* voting behavior. The dynamic model used in this exercise focuses on the correlation between ad prices at the audience (i.e. aggregate) level and candidate performance at the state-level.

Insert Table I in about here

4 The effect of Facebook ads on voters

We have argued that Facebook ad prices can be used as proxies for the intensity with which political advertisers target audiences on social media. However, what we are ultimately interested in is how exposure to political campaigns on social media, at different levels of intensity, affects voter behavior. To investigate this important question we combine our measures of advertising intensity with individual-level information on voting and “potential” exposure to the political campaign on the social media platform.

It is important to bear in mind that, because we do not observe the source or content of political campaigns, we can only establish whether exposure to *any* sort of political advertising on Facebook has an effect on voters. That is, we cannot distinguish between the effects of different types of campaigns (e.g. favorable or unfavorable message about a given candidate); we can only observe the outcomes determined by a net effect. This aspect also implies that observing weak advertising effects on voters of a given audience could be due either to such voters being unresponsive to ads, or to the fact that opposing campaigns offset each other. Given that we do not observe which ads a given voter sees, we cannot fully disentangle these two possible explanations.

Intuitively, the effect of social media ads on an individual depends on the “volume” of messages directed to the specific audience to which the individual belongs, and the likely extent of exposure to such messages. To analyze the effect of social media ads on voter behavior we have to integrate our measures of advertising *intensity* with information on voters’ *exposure* to social media. To find a measure of *exposure* to advertising on social media, we collect information on individuals’ media consumption, available from the 2016 American National Election Survey (ANES). Based on information on the ANES respondents’ ideological preferences, demographic characteristics and geographical location, we match each respondent to a Facebook audience. Furthermore, based on the individual’s responses regarding how frequently the individual accesses political news on the platform, we build a straightforward indicator of exposure (as detailed in Section 4.1). As a result, we obtain a composite individual-level measure of *treatment* to political advertising on social media, by interacting our audience-specific proxy for *intensity* with individual-specific measure of Facebook *exposure*.

Our identification strategy for the effect of social media campaigning on voting behavior is based on comparing Facebook users and non-users, conditional on the combined level of exposure and intensity of the campaign. Sections 4.1 and 4.2 describe our methodology. Observe that exposure is an *individual-level* characteristic, possibly dependent on the respondent exogenous preferences for media consumption. Intensity, on the other hand, is an

audience-level characteristic, which is exogenous because it depends on the random assignment of survey respondents to their respective interview dates (more on this in the Section 4.1).

4.1 Matching Facebook Audiences to Survey Respondents

The ANES survey covers a sample of over 3,000 American voters interviewed before and after the 2016 US Presidential elections.²¹ Pre-election interviews are randomly conducted during the two months prior to the election and are followed by post-election re-interviewing beginning November 9, 2016. Information on respondents’ political ideology, demographic characteristics (e.g., education, gender and age) and geographical location allows us to match them to a specific audience, as defined in our Facebook ad-price database. Specifically, in 2016 Facebook started classifying American users according to their ideology (liberal, moderate or conservative), based on an algorithm exploiting activity on the site.²² The ANES survey allows us to match respondents according to the same categories thanks to a pre-election liberal-conservative self-placement question.²⁴ Using this information we aggregate subgroups of respondents according to three main categories: (i) we will refer to those considering themselves as slightly conservatives, conservatives and extremely conservatives as the *Conservatives*, (ii) we will refer to those considering themselves as extremely liberal, liberal and slightly liberal as the *Liberals*, and, finally, (iii) we will refer to those who replied they “feel in the middle of the road” and those “who haven’t thought very much about this” as the *Moderates*.²⁵ In what follows we will also refer to *Conservatives* and *Liberals* as partisan voters, and to *Moderates* as swing voters. In addition to these political attributes we allocate each respondent to a Facebook audience on the basis of further demographic characteristics: gender, age, and ethnic origins.

Finally, we use the date of the interview to assign to each respondent the corresponding ad price a political advertiser pays to target the audience that incorporates that type of

²¹Interviews take place either on line or face to face. ANES conducted overall 4,271 pre-election interviews consisting of 1,181 face-to-face and 3,090 online sessions; and 3,649 post-election re-interviews consisting of 1,059 face-to-face and 2,590 online sessions). For more information see online documentation available at <https://electionstudies.org/wp-content/uploads/2016/02/anes-timeseries-2016-methodology-report.pdf>

²²To our knowledge, the only alternative source of data on voter behavior to the ANES survey is the CCES survey²³. Although the latter covers a substantially larger sample of interviewees, it is arguably less suitable for our study. The CCES does not elicit information about whether respondents have a Facebook account or how often they use it. Therefore, using the CCES survey would not help us to properly assess individuals’ exposure to Facebook.

²⁴Specifically, question V161126. The choices were extremely liberal (4.15 percent, liberal, slightly liberal, moderate /middle of the road, slightly conservative, conservative, extremely conservative and, “haven’t thought much about this”

²⁵We keep uninterested/uninformed voters and politically moderate voters together because of their relative small number.

individual. Important for our identification strategy is that (pre-election) interview dates are independent of any political campaign on Facebook. These dates are spread across states over the full pre-interview period (see Figure D.4).

Table II provides an example of the difference in average intensity (relative CPM) of advertising targeted to audiences specified by geographic location, ideology and gender. Figures show that the most intensive campaign is targeted toward Moderate users. We also observe that intensity was generally higher for Republicans than for Democrat voters. Finally, among liberal respondents, women were more intensively targeted than men, the opposite being true among conservatives. That is, among conservatives men were more intensively targeted than women.

Insert Table II in about here

Table III shows that about 30 percent of the ANES respondents does not use Facebook. Among those who have a Facebook account about 60 percent check political news on Facebook on a daily basis. Facebook users are on average younger, more educated, and, in higher proportion, male; politically, Facebook users appear to lean slightly more toward the left than non-users. It is also interesting to note that over 17 percent of respondents admit during the first interview that they are not very interested in politics because “they haven’t thought very much about politics”. These figures are consistent with Facebook’s own estimates of its penetration rate among the US population (see Table IV).

Furthermore, Figure D.3, in the Appendix, shows that daily active users of Facebook were stable between the September 1 and December 1, 2016, and that this trend was constant across gender and ideology. These observations justify our working assumption that the choice of having a Facebook account is exogenous to the choice of reading political news on its social media platform.

4.2 Empirical Strategy

Our goal is to estimate whether political ads on social media affect the voting behavior of individuals. Our identification strategy is based on the assumptions that (i) Facebook usage does not vary over time in the medium run (i.e., over the electoral campaign period) (shown in Figure D.3), and (ii) ANES survey respondents are randomly allocated to pre-election interview dates (as discussed in the Section 4.1).

Having predetermined and time-invariant Facebook usage allows us to build an exogenous measure of *exposure* to political ads on the platform; it also allows us to group ANES survey respondents on the basis of the *potential* frequency with which they can be reached by such ads. This allows us to rule out endogeneity due to medium-run changes in individuals’

Facebook usage, which could invalidate our analysis if driven by political information. For example, respondents may spend more time on Facebook because they become increasingly interested in political news during the electoral campaign, or they may devote less time to Facebook if they spend longer time reading newspapers or watching TV instead.²⁶ Our proxy for exposure, EXP_i , combines the answers from two questions in the ANES survey: (i) whether or not the respondent has a Facebook account (“Do you have a Facebook account used recently?”) and (ii) how many days a week the respondent uses social media to learn about Presidential elections (from zero to seven). We set EXP_i equal to zero when respondent i does not have a Facebook account or uses it only sporadically -up to twice a week- (we will refer to this individual as Facebook non-user), and equal to one if she heavily uses social media to get political information, (we will refer to this individual as Facebook user).²⁷

In addition, the random allocation of respondents to the ANES interview dates provides us with individual-specific windows during which the intensity of advertising campaigns varies. For example, consider two respondents with identical location, political orientation, and demographic characteristics (e.g., two conservative women in Texas, both in their mid-30s). These users belong to the same Facebook audience, but may be interviewed at different times before the election. By keeping track of ad prices for the given audience between the pre-election interview and the election, we can account for how advertising intensities affect the likelihood that each of these users deviates from the initially stated voting intentions, given the different duration of the *treatment*. By estimating the effect of advertising intensity within audience, we can remove confounding factors due to state-level, audience-level or demographic characteristics. Following the discussion of Section 3.2.1, we focus on the relative CPM runup ($r_{a,t}$), as a measure for intensity, INT_{at} . This metric captures the evolution of Facebook relative ads prices from the week prior to the interview up to the week prior to the election ($r_{a,t}$), and varies over audience a (characterized by state, political ideology and demographics) and interview time t .²⁸ The random allocation of respondents to interview dates also avoids the potential endogeneity of ad prices to the expected behavior of users in a given audience. For example, ad prices may react to changes in electoral polls. Figure III provides a graphic illustration of how our identification strategy compares the

²⁶Figures D.2 and D.3 corroborate our choice by showing that the shares of ads over total content and the number of active users, respectively, do not vary over time (particularly during the 2016 electoral campaign).

²⁷We also used a different variation for EXP , with the variable taking value of zero if the respondent does not have or does not use a Facebook account, and taking the values from one to seven depending on the number of days a week the platform is used to collect political information. The results are very similar to those presented in the paper, and they are available upon request.

²⁸Using a period of one week is motivated by studies suggesting that individuals typically forget messages within a matter of days (see, e.g., Gerber et al. [2011]).

behavior of individuals belonging to the same Facebook audience.

The interaction between EXP_i and INT_{at} provides an individual-level measure of the intensity of the political advertising campaign targeting individuals who heavily rely on Facebook to gather political news. This allows us to compare the effect of the *treatment* (reading political ads on Facebook) of individuals potentially belonging to the same targeted audience but differing in their exposure to the social media platform.²⁹

Because voters receive political information from a variety of media outlets, our identification strategy crucially requires controlling for political campaigns conducted outside Facebook. The ANES survey provides a rich set of additional information relative to each respondent’s declared source of political news, such as the internet at large, TV and radio programs, and newspaper readership. To exploit this source of information, we construct a number of explanatory variables that combine detailed information on ads conducted on TV. Specifically, we depart from the TV shows that each individual watches regularly, out of the list of 64 popular TV shows the ANES asks about. We then use data from Nielsen to build a weighted average of the CPM for TV ads targeting Republican and Democrat audiences, respectively, where weights are attributed on the basis of the size of the audience of the show at the time the ad aired.³⁰ The resulting indicator $CPMTV_{it}$ is matched in terms of timing of the interview date of the respondent and allows us to control for the intensity and exposure to political ads displayed on TV.

The baseline model we estimate for respondent i matching audience a , who received pre-election interviews at time t , is as follows:

$$Y_i^{at} = \alpha EXP_i + \beta INT^{at} + \gamma EXP_i \times INT^{at} + \delta CPMTV_i^t + \rho X_i + D_a + T + \epsilon_i^{at} \quad (3)$$

Where X_i is a vector representing individual-specific media consumption from additional information channels dummies (internet, TV, newspapers, talk shows and radio) taking the value of one if the respondent used any of these media to learn about the elections. It also includes additional demographic controls, such as gender, religion, marital status, level of education, ethnic origin, number of children income, home and gun ownership, employment status, voting choice and participation in previous elections. We also add fixed effects for income class, state and political ideology. This allows us to exclude a number of confounding

²⁹A caveat here is that our measure does not provide information on the content of the ad or the identity of the advertiser. Liberini et al. [2020] look at this issue with an online experiment.

³⁰Nielsen is a global market research and data analytic company that generates revenue from its two major business units, "Watch" and "Buy". The "Watch" side of the Nielsen data business tracks the number of people, exposed to certain content, and what actions they take thereafter. The Broadcast TV section of their dataset provides ad-specific information, such as the show during which the ad aired, the time and channel of the ad, as well as the estimated size and demographic characteristics of the audience exposed to the ad.

factors due to audience characteristics. To factor out any media seasonality in the ads distribution and common trend in the campaign development all our specifications account for “day of the week” and “month of the year” dummies. Finally, ϵ_{iat} is the error term. Table V provides summary statistics and definitions of the variables employed in the regressions.

Insert Table V in about here

We estimate Equation (3) for a number of different outcome variables, Y , capturing voting choices and turnout. The principal analysis is conducted on the probability of an individual deviating from her initially stated voting intentions (both in term of turnout and candidate choice). Because vote change is relatively rare, we pool together all different types of deviations, and build an indicator that equals one any time a respondent declares that her likelihood of voting (effective turnout), or her choice of candidate has changed from what she had initially stated in the pre-election interview. We then narrow our attention to focus on the ways in which individuals deviate from intentions regarding support for specific candidates. We estimate our model using outcome variables that are one for the cases in which the turnout or the effective vote for a Republican candidate or a Democrat candidate changes from originally expressed intention. Finally, to ease the interpretation of the results, we consider precise voting decisions, such as voter turnout and candidate choice. We estimate our models using LPM³¹. Standard errors are clustered at state level.

5 Results

This section presents and discusses the main results. We proceed as follows: In Section 5.1 we examine which targeting dimensions are mostly successfully exploited by political advertisers; we use alternative CPM indexes for the estimation of Equation (3). In Section 5.2 we estimate two variations of Equation (3) in which the outcome variable, Y , captures changes between voting intentions (in term of candidate choice and election turnout recorded at the pre-election interview), and actual voting (as recorded at the post-election interview). In Section 5.3, we focus on precise voting outcomes. Here the dependent variable Y is the “actual” voting (again in terms of candidate choice and election turnout) measured at the post-election interview.

³¹Probit regressions produce results nearly identical to OLS average marginal effect. Outputs are available upon request.

5.1 Alternative targeting dimensions

We believe that an important contribution of our empirical analysis is the ability to capture different targeting dimensions of the campaign launched on social media. As discussed in previous sections, sophisticated micro-targeting represents the true advantage of using online platforms rather than traditional media to conduct political campaigns. Thanks to the richness of our data, when estimating Equation (3), we can choose among intensity measures based on targeting of different types. Specifically, we always condition an audience on its geographic location (the state) and political ideology. In addition, we allow the audience to also vary for one of these demographic characteristics: gender, age or ethnicity.³² By comparing the performances of the same model across intensity measures, we are able to examine the relative effectiveness of targeting along these alternative dimensions.

Insert Table VI in about here

Table VI displays the results for this exercise. Each column in the table represents a different targeting dimension. Each of the six panels presents regression results with the outcome variable, Y in Equation (3), being one for, in turn, (i) changes in voting decision, (ii) changes in Trump vote, (iii) changes in Clinton vote, (iv) voter turnout, (v) Trump vote and (vi) Clinton vote. Each of these outcomes refer to the whole sample of the ANES survey respondents.

To assess the differential ability of affecting voters behavior through alternative targeting channels, in this first exercise, we limit ourselves to compare our estimates across columns of the table. In fact, inspection of these initial results suggests that campaign intensity based on age targeting does not have significant effects on any of our outcome variables. By contrast, campaign intensity based on gender or ethnicity seem to be effective, to different extents, on all of our outcome variables. Arguably, our results for the irrelevance of some targeting dimensions could be driven by one of three facts: (i) candidates may not have chosen to target the audience based on those particular channels; (ii) candidates may have chosen to target the audience, but the ads were ineffective; (iii) a candidate may have targeted the audience but the effects were offset by the opposing campaigns. The ANES estimates only provide the net average effect on the targeted audience. So, if the situation arose that, the effect on some respondents in a given targeted audience were opposite in sign and similar in magnitude to the effects on other respondents in the same audience, then we could potentially misinterpret an insignificant effect as an absence of targeting.

³²The education targeting dimension can only be partially analysed in our study, due to limited data availability. The sample of Facebook prices was only collected starting in October 2016, and several ANES respondents do not disclose their highest education level. For this reason, the sample for the targeting based on education is quite restricted, and only delivers inconclusive results.

Although we cannot directly test which of these three scenarios apply to the non significant targeting channels, we can use the full distribution of relative Facebook ad prices to discuss their plausibility. Table VII and Figure D.7 help us to do so. The table lists alternative audiences, defined by ideology and macro-locations (US states classified in terms of expected electoral outcomes). In the columns, we report the daily relative CPM for Facebook ads targeted on the basis of ideology, location and one additional dimension among Gender, Ethnicity and Age.³³ As reported, the relative CPM based on targeting by ideology-location-age and ideology-location-ethnicity dimensions is consistently higher than the relative CPM based on targeting by ideology-location or by ideology-location-age dimensions. Targeting on gender and ethnicity also displays higher variation across subgroups of the populations. These observations hold for Figure D.7, corroborating the hypothesis that candidates plausibly chose the channels of gender and ethnicity over that of age for their Facebook political campaigns. For these reasons, we will present our next set of results employing interchangeably measures of intensity based on audiences defined by location, political ideology combined with either gender or ethnicity.

5.2 Changes in candidate preferences during the political campaign

We start by asking whether being highly exposed to political messages on Facebook has an effect on the likelihood to “change one’s mind”. Broadly speaking, there are two theoretical views on the mechanism by which social media can affect political opinions and behaviors. The “echo chamber” theory hypothesizes that political polarization is amplified on social media where individuals are selectively exposed almost solely to like-minded content and networks in terms of preferences and ideologies. In this context, political advertisers micro-target voters based on similarity of views. In this way social media are viewed as platforms that reinforce the mechanism (see Allcott et al. [2020] for example, and the discussion in Section 2).³⁴ The alternative view is that of Allport’s “contact hypothesis”, arguing that social media facilitate connections between different social groups that would not meet in “real” life. In this way, social media are viewed as platforms that reduce inter-group prejudice, and lower political polarization.

³³CPM indexes are relative with respect to the state-level prices, averaged over the period from September 1 to November 8, 2016.

³⁴Note also that, in contrast to traditional media outlets, at moment of writing, Facebook is not subject to any regulations on political advertising. Political advertisers may send extreme messages targeting voters based on their specific preferences. Hence, exposure to social media could lead these voters’ political views to become more entrenched; thus, they may be less likely to change their initially expressed candidate preferences.

We contribute to the debate on social media and polarization by assessing the likelihood that users exposed to highly intensive campaigns deviate from their initially stated voting intentions. For the hypothesis that social media content increases polarization to hold, highly targeted Facebook users would need to be less likely to change their initial political views compared to similar, less-exposed, voters.

We calculate the probability of changing voting intentions by estimating equation (3) with the dependent variable set as a dummy taking the value of one when the respondent's voting intentions expressed during the pre-election interview differ from what effectively transpired on election day. To carry out this exercise we pool together all the possible voting outcomes (preferences for a candidate, and intention to vote); so Y takes values of one if the answers to the same questions in the two interviews differ and zero if they stay the same.³⁵ Table VIII presents the result for this exercise. The table is divided into six parts. In the first column we estimate equation (3) for the whole ANES survey sample. But, because we are interested in the effect of micro-targeting on specific audiences, we then estimate separate regressions for different potentially targeted subgroups, based on some observable characteristics. For example, if both candidates compete for a state whose winning probability depends on the male electorate, we expect to find a significant effect of intensity on men, but not on women. The second part looks at men and women separately. The third part divides respondents by self-declared political ideology (conservatives, moderates and, liberals) as discussed in Section 4.1. The fourth part divides by past voting records; the fifth by intention to vote as expressed at the pre-election interview; and, the sixth divides voters by ethnic origins (white and non-white).

Insert Table VIII in about here

Table VI shows that, as a whole, targeting on the basis of ethnicity, jointly with ideology, effectively reduces the probability that voters change their initial candidate choice. Column 1 of Table VIII replicates this exercise. Note that the regression results indicate that for the whole sample, using Facebook to gather political information *per se*, as well as being a political target *per se* both marginally affect vote change in a positive way. However, among Facebook users highly targeted voters are less likely to change their minds compared to less-targeted voters. Looking at the results displayed in the subsequent columns we can clearly see that the extent to which political advertising on Facebook affects different groups is heterogeneous. Those who are less likely to change their initial positions when highly targeted on the social media platform are the men, the conservatives, and the non-white

³⁵To be more precise we code Y to equal one if either the respondent expressed an intention to vote and she did not vote, (or the reverse, if she initially said she did not intend to vote, but subsequently did), or if she declared an intention to vote for a given candidate but subsequently either did not vote at all, or voted for a different candidate.

voters. Specifically, taking the whole sample, an increase of 10 percent in the CPM run-up between the interview and the election week reduces the likelihood of changing one’s vote from stated intentions, by 3 percent. This effect amounts to 4.5 percent for Men, 0.3 percent for Conservative and 2.3 percent for non-white voters. There is no significant effect for the remaining subgroups of respondents.

Figure V helps with the interpretation of the results. It represents a graphical visualization of columns (4) to (6) of Table VIII.³⁶ The figure is divided in Panel A and Panel B. Panel A plots the effect of gender-based micro-targeting on the probability of changing candidate choice among Facebook users (the red dotted line) and non-users (the black dotted line) by level of campaign intensity. The vertical lines are the 5 percent confidence intervals for these predicted probabilities.

Insert Figure (V) in about here

For those who declared themselves as holding conservatives political views, the figure indicates the existence of a clear difference between the two groups (Facebook users and non-users) conditional on a high level of campaign intensity. An increase in exposure is associated with a moderate decrease in the probability of changing one’s initial political view among those who use Facebook, and a sharp increase in the likelihood of changing one’s political views for those who do not use the platform. The net effect on Facebook users and non-users is displayed in Panel B. For each ideology profile, the solid lines represent the difference in the predicted probabilities of changing positions between Facebook users and non-users by level of Facebook campaign intensity. The shaded areas surrounding the lines are the 5 percent confidence interval for these differences. The bar charts at the bottom of the tables show the distribution of intensity by audience type. Inspection of the three plots in Panel B confirms that the difference between Facebook users and non-users become statistically significant only for the conservatives when exposure to Facebook political campaign is of high intensity. This difference is consistent with the ”echo chamber” hypothesis. We then conduct a similar exercise by looking specifically at the changes in voting decisions regarding which candidate to support. The new outcome variables, Y , are set equal to one if the respondent supported Trump (Clinton) in the pre-election interview and, on election day, either did not vote or voted for a different candidate. Tables IX and X display the results for these exercises. The tables have the same format as Table VIII. Table IX shows that for Facebook users an increase of 10 percent in the CPM run-up between the interview and the election week reduces the likelihood of changing one’s vote for Trump from originally stated intentions, by 1.10 percent for the whole sample (column 1) and by 4.27 percent for moderate voters (column 2). This effect is positive and equal to 2.31 percent for conservative voters. This is

³⁶All figures are based on the Probit estimates of the equation 3.

also shown in Figure VI (Panel A). However what matters for our purposes is the comparison of Facebook users and non-users, Figure VI (Panel B) helps with the visualization of this comparison by ideology.

Insert Tables (IX) (X) , in about here

From the inspection of Panel B those Facebook users who are classified as conservatives are less likely to change their initial position with respect to Mr Trump compared to those conservatives who do not use Facebook. A similar pattern, albeit less strong and significant, is found also among those who define themselves politically moderate and/or not very interested in politics.

Insert Figure VI, in about here

It is also interesting to compare the distribution of intensity by political ideology by looking at the bar charts. In line with the fact that we find stronger results for conservative voters, conservatives look also to be as the most heavily targeted group.

Turning attention toward changes in preferences for Clinton (Table X), a similar pattern emerges. However, the results are much less significant than those for Trump. Figure VII confirm these results. These results are consistent with the "echo chamber" theory that social media increases political polarization. Our results show that this is the case at least for conservative voters and to lesser extent for moderates if they are heavily targeted.

However, it is possible that those who are constantly keeping up with the political news on social media may become better informed compared to those who do not. Thus, these people who devote time and energy into keeping abreast of political news may be more difficult to persuade to change their political positions. This is an hypothesis that we can test by exploiting questions from the pre- and post- election surveys by assessing respondents' political knowledge. During the first and second interviews, respondents answer a number of questions about current policy debates.³⁷ We use the ratio of the number of correct post-election answers over the number of correct pre-election answers as a crude indicator of the improvement of political knowledge, *PolKnow*. We re-estimate Equation (3) where the outcome variable *Y* is now *PolKnow*.

Insert Table (XI), in about here

The results of this exercise are displayed in Table XI). In the first column, we ignore the interaction between exposure and intensity and we only include *EXP* as a regressor to control for the effect of Facebook. We then augment this basic specification with various measures

³⁷The questions are: "For how many years is a senator elected?", "Which is the program for which the Fed Government spend most?", "Which is the party with most members in the House before elections?", "Which is the party with most members in the Senate before elections?" These are questions V161513, V161514, V161515, and V161515 in the ANES survey(2017). During the post-election survey respondents are asked to "... recall the office of: Joe Biden, Paul Ryan, Angela Merkel, Vladimir Putin, Justice Roberts."

for *INT* and its interaction with *EXP*. The results show that reading about the political campaign on Facebook does not improve respondents' knowledge about politics. The same is true for the results of listening to radio or watching television and talk shows or even reading newspapers instead. A positive effect on political knowledge comes from using internet. As shown on the table, the results show, quite clearly that becoming more knowledgeable is not the key to understanding why Facebook consumers are less likely than others to change their minds, and why they remain wedded to the views they initially expressed. Political polarization may be the more obvious answer, but a deeper understanding of what is happening on on this front goes beyond the scope of this paper and is left to future research.

5.3 The effect of Facebook advertising on voting outcomes

In the previous section we established that respondents who receive intense political targeting, and who rely on Facebook as a source of political information, are less likely to revise their initial positions compared to those who do not. The analysis so far does not allow us to gain insights on the effects of Facebook political ads on specific voter decisions. In this section we attempt to contribute to this point by focusing on estimating the effect of political advertising on Facebook on the two main goals of a political campaign (see Section 2), namely mobilization and persuasion. A caveat here, by estimating voting decisions rather changes in voting decisions we gain ease in interpretability, at a loss of identification precision. This is because any specific outcome decision (e.g., whether or not turnout to vote in and election or the choice of a specific candidate) is influenced by the Facebook campaign run during the window in the wake of the pre-election interview and by any campaigns to which they were exposed prior to the interview. We can only control for this situation by using additional explanatory variables.

5.3.1 Voter turnout

Whether advertising helps to mobilize voters and increase turnout in elections is a fundamental question in political economy. We contribute to this inquiry by providing the first evidence on the effect on turnout from micro-targeted political advertising conducted on social media. We estimate variations of equation (3), in which the dependent variable is a dummy taking the value of one if the respondent has voted in the 2016 presidential election and zero otherwise. The full set of controls listed in the previous section is included in all our specifications.³⁸ We report the results in Table XII.

³⁸These are not reported in the tables. The estimations are available upon request

Insert Table XII in about here

The table shows that there is a negative effect on turnout from both exposure and intensity, when each is taken separately; however, there is a positive cumulative effect of the two when interacted with each other. This suggests that exposure to social media has a negative effect on turnout *per se*, which is in line with findings by Falck et al. [2014]. Yet, this negative exposure effect decreases with the intensity of the political campaign. Our main coefficient of interest, the interaction between Facebook *Exposure* and *Intensity*, shows that overall an increase in the level of intensity by 10 percent is associated with an increase in Facebook users' turnout of about 0.9 percent. However, if we look into the effect on specific subgroups of voters, in some cases, this combined effect is much larger and more significant. This is true for the moderates (i.e., those who do not have a strong political ideology, and/or those who are not very engaged in politics), for those who declared a preference for Trump in the pre-election interview and for those who define themselves as white. For all of these subgroups, reading political news on Facebook is associated with higher voting turnout. Only those who previously supported Clinton show a fall in the likelihood that they will cast ballots on election day.

It is also interesting to examine and compare the effects of other media outlets on voting. Note that any effect can only be interpreted in terms of correlation, because our identification strategy only applies to the effects of political ads on Facebook. Overall, surfing the internet is associated with higher turnout; this effect is positive and significant for the whole sample and for most subgroups. The exceptions are liberal and conservative voters, those who planned to vote for Trump, and the whites. Other media outlets appear to have only a weak correlation with voting. Noticeably, reading newspapers is associated with a lower likelihood of voter turnout among women and Clinton supporters.

To give a more precise indication of the combined magnitude of the effects of intensity and exposure to political campaigns on Facebook, we again use graphical illustrations. Following Table XII, each figure depicts regression results for a different subgroup of respondents. As presented in Figure VIII, the top panel reports predicted probabilities, and the bottom panel reports differences in these probabilities as well as the distribution of users over the intensity spectrum. The results on the effect of Facebook on turnout by ideology show different effects on the three groups of voters who are identified to be Conservative, Moderate or Liberal. Increasing intensity of the campaigning has a striking negative effect on turnout among liberal Facebook users, a positive effect on moderates; and an almost null effect on conservatives. It is also worth noting that increasing intensity of campaigning on Facebook appears to have a strong positive effect of turnout among liberal voters, but almost no effect (an almost flat line as indicated on the graph in Panel A) for the other two

voter groups (conservatives and moderates). One could legitimately wonder why our results would show that also non-Facebook users differ in predicted voting behavior according to the intensity of micro-targeted political campaigns that they likely did not see, since they were not in the platform. We believe that there is an obvious explanation for this. Because political campaigns advertise through multiple media channels during the electoral period, there is a structural correlation between how valuable an audience is on social media and how valuable it is on all traditional media. Figure D.6 validates that there is, indeed, evidence for this interdependence, as the relative CPM of ads shown on Facebook during the pre-election interview and the CPM of ads shown at the same time on television display a clear positive correlation. Following this argument, we can interpret the non-users' predicted probability of turnout as being affected by similar campaigns running on traditional media.

Turning to Panel B in Figure VIII, we can see that the most significant effect is the negative difference between the predicted turnout probabilities for liberal Facebook users and non-users. What convinced liberal Facebook users to stay away from the polls? Although our data do not provide information on the content of the campaign, or on the identity of the advertisers, a plausible conjecture is that the campaigns were on average not in favour of Clinton's election. By looking at the distribution of intensity at the bottom of Panel B, we can clearly see that the large majority of liberals received intense campaigns on Facebook. By contrast, increasing political campaign intensity has virtually no net effect on turnout among conservatives, but a positive effect among moderates on Facebook.

The same Figure VIII also looks at whether ideological targets might be based on past electoral preferences (i.e. voting in 2012). Panel A shows that, for high levels of campaign intensity, a big drop in turnout occurs among those who supported Trump but did not accessed Facebook, while a flat trend for those Trump supporters who were exposed to the campaign on Facebook. This points at a counteractive effect of Facebook campaign against aggressive campaigns on alternative media outlets.

Figures D.9, in the Appendix, turn to ethnicity and gender differences. We find no particular difference in predicted turnout among white and non-white voters. We find that exposure to Facebook campaigns has the opposite effect on men and women. Men are less likely to vote, and women are more likely to vote in the wake of exposure to the online campaigning. Panel B indicates that the net effect is negative and significant for men and positive, but not significant, for women.

5.3.2 Choice of candidates

The next question we address is whether exposure to micro-targeted political ads on Facebook is effective in persuading voters toward a specific candidate. Table XIII displays the

results for the estimation of equation (3), when the dependent variable is equal to one if the respondent voted for Trump and zero otherwise. Table XIV repeats the analysis for Clinton.

Insert Tables XIII and XIV in about here

Table XIII suggests that the most significant and positive effects of the Facebook campaigns on the Trump vote are among moderates and those who were planning all along to vote for Trump (Trump intention). We analyze these effects with the help of figures IX, which report the effects of *exposure*, *intensity*, along with their interaction.

Insert Figure IX in about here

The results for the subgroups based on ideology displayed in Figure IX are the most interesting. There is a strong and positive effect on the Trump vote among moderates, this group includes those who are not very interested in politics. Moderates reading political news on Facebook regularly are more likely to vote for Trump as campaign intensity increases. The opposite effect is registered for those who do not use Facebook. There is virtually no difference for liberal and conservative Facebook users compared with their counterpart non-users. The bottom graph shows that for high levels of campaign intensity (1.2); among Moderates, Facebook users are up to 15 percentage points more likely to cast a vote for of Trump compared to non-users. A possible explanation could be that individuals who rely on Facebook to learn about politics are exposed to a high volume of political messages aligned with Trump. We find very little effect of social media ads on the choice of candidate of partisan voters. This is rather intuitive. From the top panel of Figure IX, we see that while partisan voters are hardly affected by exposure to the social media, moderate voters who access Facebook regularly are up to twice as likely to support the Republican candidate than neutral voters who do not use Facebook.

Figure D.10, in the Appendix, turns to differences in the effect of Facebook campaigns across ethnicity and gender. The top panel shows that there is virtually no difference in the predicted probabilities of voting for Trump of Facebook users and non-users among white respondents targeted at different levels of intensity. Instead, for non-white respondents we find that the probability of supporting Trump decreases as the campaign intensifies, but the effect is less pronounced for individuals who are exposed to the social media campaigns. This difference in the predicted probabilities among Facebook users and non-users is particularly significant around the average of the intensity distribution (1 to 1.1), as panel B indicates.

The Figure also shows that there effect on Trump support for Facebook users is not significantly different from that on non-users among both men and women. These results are consistent with the idea that Trump concentrated his targeting effort on voters who are moderate or uninterested in politics. The potential "swing" voters were key to the candidate's support base. This is also consistent with previous findings showing that the effects of

advertising have differential impacts on voters depending on their level of political knowledge and sophistication (Huber and Arceneaux, [2007]; Ridout and Franz, [2011]; Zaller, [1992]).

What is particularly interesting is that all these results are exclusive to the vote cast in favor of the Republican candidate. We repeat the whole analysis for the Clinton vote, and, surprisingly, we find that micro-targeting had no significant effect. As shown in Table XIV, this is true for all types of micro-targeting, and holds for all subgroups of voters. Figures D.11 and D.12, in the Appendix, graphically confirms these findings.

6 Discussion

The overall picture that emerges from these results is that social media campaigns lower the likelihood that individuals will deviate from initially stated voting intentions. Furthermore, exposure to intensive ad campaigns on Facebook decreases turnout among liberals, and increases turnout among moderates and conservatives. Such exposure increases moderates' support for the Republican candidate.

To make our results more tangible, we now attempt to translate them in terms of gain or loss of actual votes for the two main candidates. We start by asking what turnout would have been had the Facebook campaign been less intense. Figure X maps this effect over the 50 US states. On the left, we show how many fewer votes each state would have recorded from Facebook users who supported Trump as percentage of total turnout, had the campaign been 1 percent less intense. On the right, we show how many more votes each state would have recorded from Facebook users who supported Clinton, had this same scenario (reduced Facebook campaign intensity of 1 percent) taken place. We compute these changes by multiplying the overall turnout in each state by the share of ANES survey respondents who voted, used Facebook to gather political news, and were either Trump or Clinton supporters (for the left and right maps, respectively). We then multiply these turnout figures by the marginal effects estimated in columns (10) and (9) of Table XII.

The results show that votes coming from Facebook users who were initially supporting Trump would have been lost prominently in states that are Republican strongholds, such as Alabama, Mississippi, Louisiana, Arizona, North Dakota, Montana, and Tennessee. By contrast, votes coming from Facebook-users who were initially supporting Clinton would have been gained in states that are Democrat strongholds, such as California, Oregon, New York, Massachusetts, Rhode Island, Washington state, and the District of Columbia. However, the increase in turnout from Clinton supporters would have been almost twice as large as the reduction in turnout from Trump supporters.

We then move on to the question of how many moderate voters Trump would have lost

in each state, had the intensity of the Facebook campaign been 1 percentage point lower. We calculate the number of moderate voters for Trump using the same provisions used to generate the previous maps. We then multiply these votes by the marginal effect of intensity estimated in Table XIII, column (5). The results (mapped by state in Figure XI) show that the Republican candidate would have suffered the biggest losses in some of the swing states (Florida, Pennsylvania, New Hampshire) but, interestingly, also in those states that turned out to have had a big impact on his overall win, such as Michigan, Wisconsin, and Minnesota.

In conclusion, the curious and partly unexpected result of our study is that the overall beneficiary of the Facebook campaigns was the Republican candidate. We cannot make any inferences regarding the content or the framing and alignment of the campaigns. However, it seems that the campaigns on depressed turnout in blue states, and enhanced turnout in red states. Moderate voters were persuaded to vote for Trump in states that proved to be crucial in securing the Electoral College votes needed to win the election. Thus, while the effects we find are of a rather small magnitude, they appear to have been significant in terms of swaying a relatively small group of individuals who might have had a disproportionately large impact on the overall election results.

7 Conclusions

This paper has examined the effect of Facebook’s micro-targeted political ads on the behavior of voters in the 2016 US presidential election. We rely on the fact that social media exploit extensive quantities of user-generated data that can be used to help politicians personalize their election campaigns, and to target with increasing precision those voters who may be decisive in determining election outcomes. We exploit a unique, novel dataset which allows us to assess the effects and the power of political “micro-targeting” by using daily Facebook advertising prices as a new measure to proxy for the intensity of online political campaigns.

In addition to making this methodological contribution, this paper is the first, so far as we are aware, to compare micro-targeting dimensions, both in terms of advertising strategies and effect of these strategies on voter behavior. We have shown that micro-targeted political ads on Facebook matter, and that they have a significant effect on voter behavior, particularly among users who are targeted based on ethnicity, gender, location, and political orientation.

An important result is that the findings provide direct evidence that political advertising on Facebook exacerbates political polarization. Highly targeted users of the platform are less likely to change their minds compared to less-targeted users. This effect is particularly striking among men, conservatives, and non-white voters. Our results also indicate that

learning about politics on Facebook does not make voters more politically informed, but rather makes them more likely to stick with the voting choices and intentions they already hold. These findings are also consistent with evidence that Facebook advertising intensifies political polarization.

Our analysis estimates the net effect of the campaigns on political outcomes. Our results indicate that advertising on Facebook was effective at convincing Trump’s core supporters to turn out on election day, and that such advertising had a negative effect on turnout among liberals, Clinton’s core supporters. Moreover, campaigning on Facebook swayed moderate and less-informed voters to cast votes for Trump. These findings are consistent with financial campaign reports showing that Trump spent more heavily than Clinton on using social media as campaign vehicles. It could also suggest, of course, that his messages may have been better designed.

As a first attempt to shed light on how social media can affect politics, we believe that our measures of intensity represent a step forward in advancing our understanding of how these important and increasingly commonplace channels of political communication have the potential to rewrite the political campaign textbook. By resorting to a revealed preference strategy, we are able to describe the political strategy that might explain the patterns found in the variation of intensity of political campaigns on the Facebook platform. Given the outsized role that some audiences in specific states (e.g. moderate voters in swing states) play in the conduct of the US national presidential campaigns, we thus can infer which political candidates would be willing to compete to capture the votes of the voters in these audiences. This is a first step in a much broader research agenda needed to better understand how growing online communication platforms affect the conduct of political campaigns and the outcomes of democratic elections.

References

- [1] Allcott, H., Braghieri L., Eichmeyer S. and M. Gentzkow. "The Welfare Effects of Social Media," *American Economic Review*, 110 (2020), 629-76.
- [2] Allcott, H. and M. Gentzkow. "Social Media and Fake News in the 2016 Election." *Journal of Economic Perspectives*, 31 (2017), 211-236.
- [3] Allport, G. W.. "The nature of prejudice. Cambridge," MA: Perseus Books (1954).
- [4] Ambrus, A., Calvano E. and M. Reisinger. "Either or Both Competition: A 'Two-Sided' Theory of Advertising with Overlapping Viewerships," *American Economic Journal: Microeconomics* 8 (2016), 189-222.
- [5] American National Election Studies, University of Michigan, and Stanford University, "ANES 2016 Time Series Study. ICPSR36824-v2." Ann Arbor, MI: Inter-university Consortium for Political and Social Research (2017).
- [6] Barberá, P.. "How Social Media Reduces Mass Political Polarization. Evidence from Germany, Spain, and the U.S," Working Paper, NYU (2015).
- [7] Bloomberg (2018) <https://www.bloomberg.com/news/articles/2018-09-20/facebook-to-give-less-direct-support-to-trump-in-2020-campaign>.
- [8] Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D., Marlow, C., Settle, J. E., and Fowler, J. H.. "A 61-million-person experiment in social influence and political mobilization," *Nature*, 489 (2012), 295.
- [9] Boxell, L., Gentzkow, M. and J.M. Shapiro. "Is the internet causing political polarization? Evidence from demographics," NBER Working Paper No. 23258 (2017).
- [10] Bright, J., Hale, S. A., Ganesh, B., Bulovsky, A., Margetts, H., and Howard, P.. "Does Campaigning on Social Media Make a Difference? Evidence from candidate use of Twitter during the 2015 and 2017 UK Elections," arXiv preprint arXiv (2017), 1710.07087. <https://www.technologyreview.com/s/609229/how-tweets-translate-into-votes/>
- [11] DellaVigna, S. and E. Kaplan. "The Fox News effect: Media bias and voting," *Quarterly Journal of Economics*, 122 (2007), 1187-1234.
- [12] Dubois, A., Zagheni, E., Garimella, K. and I. Weber. "Studying Migrant Assimilation Through Facebook Interests," arXiv (2018), 1801.09430v2
- [13] Enikolopov, R., Makarin, A. and M. Petrova. "Social Media and Protest Participation: Evidence from Russia," Forthcoming in *Econometrica*.
- [14] Enriquez, J.R., Larreguy, H., Marshall, J. and A. Simpser "Online Political Information: Facebook Ads, Electoral Saturation, and Electoral Accountability in Mexico," Working paper, Harvard University, (2019).
- [15] Eurobarometer. "Media use in the European Union report," European Commission, 78 (2012), 8.
- [16] Falck, O., Gold, R., and S. Heblich. "E-lections: Voting Behavior and the Internet," *American Economic Review*, 104 (2014), 2238-65.
- [17] Gabler, N.. "The internet and social media are increasingly divisive and undermining of democracy," Accessed at <http://www.alternet.org/culture/digital-divide-american-politics> on February 19, 2017.
- [18] Gainous, J., and Wagner, K. M.. "Tweeting to power: The social media revolution in American politics," Oxford University Press (2014).

- [19] Garcia, D., Mitike Kassa, Y., Cuevas, A., Cebrian, M., Moro, E., Rawan, I., and R. Cuevas, “Analyzing gender inequality through large-scale Facebook advertising data,” *Proceedings of the National Academy of Sciences*, 115 (2018), 6958-6963.
- [20] Gentzkow M.. “Television and Voter Turnout,” *Quarterly Journal of Economics*, 121 (2006), 931-972.
- [21] Gerber, A. S. , J. G. Gimpel, D. P. Green, and D. R. Shaw, “How Large and Long-lasting Are the Persuasive Effects of Televised Campaign Ads? Results from a Randomized Field Experiment,” *American Political Science Review*, 105 (2011), 135-150.
- [22] Gold, M., and E. Dwoskin “Trump Campaign’s Embethnicity of Facebook Shows Company’s Growing Reach in Elections.” *The Washington Post*, (2017).
- [23] González Cabañas, J., Cuevas, A. and R. Cuevas. “Unveiling and Quantifying Facebook Exploitation of Sensitive Personal Data for Advertising Purposes,” *USENIX Security Symposium* (2018), Baltimore.
- [24] Gordon, B.R., Hartmann, W.R.. ”Advertising competition in presidential elections,” *Quantitative Marketing and Economics*, 14 (2016), 1–40.
- [25] Hillygus, D., Sunshine and T. Shields. “The Persuadable Voter: Wedge Issues in Presidential Campaigns,” Princeton, NJ: Princeton University Press. (2008)
- [26] Huber, G. A., Arceneaux, K.. “Identifying the persuasive effects of presidential advertising,” *American Journal of Political Science*, 51 (2007), 957–977.
- [27] Kosinski M., Stillwell D. and T. Graepel. “Digital records of behavior expose personal traits,” *Proceedings of the National Academy of Sciences*, 110 (2013), (15)
- [28] Krasno, J. S. and D. P. Green. “Do Televised Presidential Ads Increase Voter Turnout? Evidence from a Natural Experiment,” *Journal of Politics*, 70 (2008), 245–261.
- [29] Johnston, R., Hagen, M. G., and K.H. Jamieson. “The 2000 presidential election and the foundations of party politics.” Cambridge University Press (2004).
- [30] Levin, J., and P. Milgrom.. ”Online Advertising: Heterogeneity and Conflation in Market Design,” *American Economic Association P&P*, 100 (2) (2010): 603-07.
- [31] Liberini F., Redoano M., Russo A., Cuevas A. and R. Cuevas. Good Ad, Bad Ad. Effect of Positive and Negative Political Content of Online Advertising. Mimeo, University of Warwick (2020).
- [32] Liu, Y., Kliman-Silver, C., Bell, R., Krishnamurthy, B. and A. Mislove. “Measurement and Analysis of OSN Ad Auctions,” *Proceedings of the Second ACM Conference on Online Social Networks*, (2014) Dublin, Ireland.
- [33] Moshary, S.. “Price Discrimination in Political Advertising: Evidence from the 2012 US Presidential Elections,” *Working Paper* (2017), University of Chicago Booth School of Business.
- [34] Müller, K. and Schwarz, C., “Fanning the Flames of Hate: Social Media and Hate Crime,” *Working Paper* (2019), Available at SSRN: <https://ssrn.com/abstract=3082972>
- [35] Pariser, E.. “The filter bubble: What the internet is hiding from you,” New York, NY: Penguin Press. (2011)
- [36] Petrova, M., Sen, A. and P. Yildirim, “Social Media and Political Donations: New Technology and Incumbency Advantage in the United States,” *CEPR Discussion Paper No. DP11808* (2017).
- [37] Pew Research Center, (May 2016), “News Use Across Social Me-

- dia Platforms 2016” <http://www.journalism.org/2016/05/26/news-use-across-social-media-platforms-2016/>.
- [38] Pew Research Center, (October 2016), “The Political Environment on Social Media”, <http://www.pewinternet.org/2016/10/25/the-political-environment-on-social-media/>.
 - [39] Pew Research Center, (November 2016)2, “Social Media Update 2016”, <http://www.pewinternet.org/2016/11/11/social-media-update-2016/>
 - [40] Ridout, T. N. and M. Franz. “The persuasive power of campaign advertising,” Philadelphia, PA: Temple University Press (2011).
 - [41] Saez-Trumper, D., Liu, Y., Baeza-Yates, R., Krishnamurthy, B. and A. Mislove. “Beyond CPM and CPC: Determining the Value of Users on OSNs,” *Proceedings of the Second ACM Conference on Online Social Networks*. (2014) Dublin, Ireland.
 - [42] Santa-Clara, P. and Valkanov, R., “The Presidential Puzzle: Political Cycles and the Stock Market,” *The Journal of Finance*, 58 (2003), 1841-1872.
 - [43] Sattler, T., “Do Markets Punish Left Governments?” *Journal of Politics*, 75 (2013), 343–356.
 - [44] Strömberg, D., “How the Electoral College Influences Campaigns and Policy: The Probability of Being Florida,” *American Economic Review*, 98 (2008), 769-807.
 - [45] Sunstein, C. R.. “Republic.com 2.0,” Princeton, NJ: Princeton University Press (2009).
 - [46] Sunstein, C. R.. “Republic: Divided democracy in the age of social media,” Princeton, NJ: Princeton University Press. (2017).
 - [47] Williams C. B., G. J. Gulati, “Digital Advertising Expenditures in the 2016 Presidential Election,” *Social Science Computer Review* , 36 (2017), 406 - 421.
 - [48] Zagheni E., Weber I. and G. Krishna. “Leveraging Facebook’s Advertising Platform to Monitor Stocks of Migrants,” *Population and Development Review*, 43 (2017), 721-734
 - [49] Zaller, J.. “The nature and origins of mass opinion,” New York, NY: Cambridge University Press (1992).

Tables

Table I: Candidates Winning Probabilities and CPM Prices

Dep. Variable: Relative CPM Median	Republican Partisan States	Democrat Partisan States	Swing States
Lag of CPM median	0.533** (0.031)	0.545** (0.026)	0.502** (0.042)
Abs. Value of Difference in Winning Prob	-2.887** (0.640)	5.814** (0.907)	1.687** (0.508)
Abs. Value of Difference in Winning Prob $\times 1[T > C]$			-2.390** (0.921)

Note: results are based on the panel of state daily CPM prices, for the three months preceding the 2016 Presidential elections. The estimated equation is $cpm_{s,t} = \alpha cpm_{s,t-1} + \beta_1 PDIFF_{s,t-1} + \beta_2 (PDIFF_{s,t-1} \times TADV_{s,t-1}) + e_{s,t}$ where $PDIFF_{s,t} = |(P_{s,t}^T - P_{s,t}^C)|$, and $P_{s,t}^T$ and $P_{s,t}^C$ indicate the winning probability of Trump and Clinton, respectively. $TADV_{s,t}$ is an indicator variable with value 1 for all cases where Trump has the lead on the election forecast. A GMM one-step estimator was used, with a restricted number of 2 lags for both the dependent variable, and the predetermined and endogenous variables used as instruments. Number of Observations is 1,977 for Republican Partisan States, 1,376 for Democrat and 859 for Swing States. State grouping (Republican, Democrat and Swing) is based on historical election results.

Table II: Price Ratio between Interview and Election

Political Ideology	Male	Female
Liberal - very strong	0.988	0.990
Liberal - strong	0.978	0.986
Liberal - mild	0.977	0.978
Moderate	1.041	1.004
Conservative - mild	1.024	1.001
Conservative - strong	1.029	1.004
Conservative - very strong	1.013	1.003

Note: average relative CPM for audiences categorized according to gender and ideology. Figures span the period from the 1st of September until the 8th of November 2016.

Table III: Facebook users and non-users

Variable	non-Users		Users	
	Freq.	Mean	Freq.	Mean
Turnout Intention	989	0.934	1,425	0.940
Turnout (Effective)	871	0.873	1,233	0.862
Clinton Vote Intention	997	0.450	1,430	0.469
Clinton Vote (Effective)	994	0.427	1,428	0.438
Trump Vote Intention	997	0.384	1,430	0.341
Trump Vote (Effective)	994	0.383	1,428	0.349
Liberals	997	0.472	1,430	0.504
Moderate	997	0.109	1,430	0.095
Conservative	997	0.418	1,430	0.401

Note: mean voting behaviour and ideology distribution among Facebook users and non-users interviewed for the ANES 2016.

Table IV: Facebook Penetration Rates among Online Adults

	PEW	ANES
All Online Adults	79	67
Men	75	60
Women	83	74
18-29	88	77
30-49	84	74
50-64	72	65
65+	62	51
High School Degree or Less	77	59
Some College	82	68
College+	79	70
Less than \$30K/year	84	66
\$30K - \$49,999K	80	70
\$50K - \$74,999K	75	68
\$75K+	77	67

Note: Facebook penetration rates (% of population with a Facebook account) for the population of American adults with access to Internet, as reported by two different sources. The first column is based on the PEW, Social Media update survey conducted on March-April 2016. The second column is based on the ANES 2016.

Table V: The sample of Facebook users and non-users

Variable	Obs	Mean	Std. Dev.	Min	Max
Any Change of Voting Behaviour	2,964	0.182	0.386	0	1
Change of Trump Vote	2,964	0.093	0.291	0	1
Change of Clinton Vote	2,964	0.104	0.305	0	1
Effective Turnout	2,963	0.877	0.328	0	1
Effective Trump Vote	2,943	0.383	0.486	0	1
Effective Clinton Vote	2,943	0.419	0.493	0	1
Voted in Primary	2,960	0.504	0.500	0	1
Turnout in 2012	2,954	0.828	0.377	0	1
Voted Democrat in 2012	2,928	0.466	0.499	0	1
Party Registered in 2012	2,964	0.575	0.494	0	1
Internet Access	2,964	0.951	0.215	0	1
Facebook Account	2,964	0.563	0.496	0	1
Facebook CPM Ratio (generic)	2,964	1.082	0.086	0.847	1.525
Facebook CPM Ratio (ethnicity Targeting)	2,950	1.076	0.083	0.847	1.525
Facebook CPM Ratio (Gender Targeting)	2,940	1.071	0.080	0.837	1.422
Facebook CPM Ratio (Age Targeting)	2,964	1.087	0.105	0.672	1.628
Facebook CPM Ratio (Education Targeting)	1,286	1.079	0.069	0.881	1.359
Media Used: Internet	2,936	0.653	0.476	0	1
Media Used: Newspapers	2,936	0.538	0.499	0	1
Media Used: Radio	2,936	0.584	0.493	0	1
Media Used: TalkShows	2,936	0.657	0.475	0	1
Media Used: TV	2,936	0.870	0.336	0	1
CPM TV	2,316	0.186	0.070	0.055	0.708
Gender (female = 1)	2,940	0.532	0.499	0	1
Age	2,888	50.694	17.272	18	90
Any Children	2,961	0.326	0.469	0	1
Married	2,951	0.524	0.500	0	1
Has a Degree	2,942	0.580	0.494	0	1
Unemployed	2,959	0.052	0.223	0	1
Importance of Religion (high = 1)	2,952	0.654	0.476	0	1
Owns a Gun	2,798	0.335	0.472	0	1
Owns a House	2,953	0.660	0.474	0	1

Note: descriptive statistics for the variables used in the regression analysis. We exclude the sample of individuals who reported to have voted before the 8th of November 2016 (“early-vote”). Other factors (ideology, income and US state) are controlled for in our model via a set of dummy variables.

Table VI: Effect of Facebook campaigns along Alternative Targeting Channels

	Only Ideology & Location (1)	also ethnicity (2)	also Gender (3)	also Age (4)
Dep. Var.: Change in Any Voting Decision				
News Access on FCBK (Exposure)	0.273 (0.206)	0.417* (0.213)	0.214 (0.229)	-0.0787 (0.164)
CPM runup (Intensity)	-0.00910 (0.154)	0.0987 (0.161)	0.255 (0.169)	-0.0480 (0.0922)
FCBK Exposure x Intensity	-0.255 (0.189)	-0.390** (0.197)	-0.203 (0.213)	0.0695 (0.149)
Dep. Var.: Change in Trump Vote				
News Access on FCBK (Exposure)	0.203 (0.158)	0.235 (0.169)	0.495*** (0.184)	0.0377 (0.130)
CPM runup (Intensity)	-0.0295 (0.117)	-0.0307 (0.127)	0.338** (0.133)	-0.0704 (0.0689)
FCBK Exposure x Intensity	-0.171 (0.144)	-0.202 (0.156)	-0.447*** (0.173)	-0.0188 (0.118)
Dep. Var.: Change in Clinton Vote				
News Access on FCBK (Exposure)	0.147 (0.163)	0.289* (0.168)	-0.250 (0.161)	-0.0604 (0.132)
CPM runup (Intensity)	-0.009 (0.122)	0.120 (0.127)	-0.110 (0.123)	0.0230 (0.0763)
FCBK Exposure x Intensity	-0.151 (0.149)	-0.284* (0.154)	0.217 (0.149)	0.0397 (0.120)
Dep. Var.: Turnout				
News Access on FCBK (Exposure)	-0.0762 (0.181)	-0.168 (0.189)	-0.381* (0.197)	0.139 (0.138)
CPM runup (Intensity)	0.200 (0.136)	0.117 (0.138)	-0.238 (0.148)	-0.0291 (0.0743)
FCBK Exposure x Intensity	0.0492 (0.165)	0.134 (0.174)	0.334* (0.183)	-0.149 (0.125)
Dep. Var.: Trump Vote				
News Access on FCBK (Exposure)	-0.0933 (0.192)	-0.116 (0.208)	-0.395* (0.216)	0.0685 (0.157)
CPM runup (Intensity)	-0.0992 (0.147)	-0.123 (0.154)	-0.319* (0.154)	0.00855 (0.0950)
FCBK Exposure x Intensity	0.106 (0.176)	0.127 (0.192)	0.389* (0.202)	-0.0434 (0.143)
Dep. Var.: Clinton Vote				
News Access on FCBK (Exposure)	-0.212 (0.193)	-0.298 (0.203)	-0.0623 (0.202)	0.210 (0.161)
CPM runup (Intensity)	0.171 (0.145)	0.126 (0.147)	-0.0489 (0.150)	-0.0628 (0.0950)
FCBK Exposure x Intensity	0.152 (0.177)	0.232 (0.186)	0.0133 (0.187)	-0.236 (0.146)

Note: reported coefficients are estimated with a LPM on a number of alternative outcome variables (change in voting outcome, turnout, Trump vote and Clinton Vote). The different columns differ in the indicator used for the intensity of the Facebook ads. Each column uses an index based on a different set of audience characteristics, with the baseline, Column (1), being users ideology and location. In Column (2) targeting is based on ideology, location and ethnicity. In Column (3) also on Gender and in Column (4) also on Age.

Table VII: Distribution of Facebook CPM prices across different audiences

	Only Ideology & Location	also Gender	also ethnicity	also Age
Users Ideology				
Active Liberal	0.988	1.000	0.991	1.028
Very Liberal	0.951	0.958	0.971	0.961
Liberal	0.953	0.992	1.010	0.975
Moderate	1.019	1.077	1.030	1.016
Conservative	1.032	1.092	1.000	1.033
Very Conservative	1.017	1.052	0.972	1.030
Active Conservative	1.128	1.189	1.023	1.207
States Alignment				
Battleground	1.069	1.004	0.986	1.067
Safe Democrat	1.011	1.070	0.986	1.025
Safe Partisan	1.029	1.052	1.008	1.040
Swing	1.011	1.037	1.004	1.032
Rustbelt	1.000	1.058	0.999	1.029
Small Margin	1.011	1.051	0.992	1.019

Note: mean daily relative CPM prices for Facebook ads displayed during the period 01 September 2016 - 08 November 2017 and targeted to various audiences. The first column reports the price of ads targeted on the basis of user location and ideology. The remaining columns report the price of ads targeted on the basis of location, ideology and an additional demographic characteristic such as Gender (Column (2)), ethnicity (3) and Age (4).

Table VIII: The effect of Facebook on Change in Voting Behavior

Dep. Variable: Change in Voting Preferences	All (1)	Women (2)	Men (3)	Liberal (4)	Moderate (5)	Conservative (6)	Vote Dem 2012 (7)	Vote Rep 2012 (8)	Clinton Intention (9)	Trump Intention (10)	White (11)	Non-white (12)
News Access on FCBK (Exposure)	0.417* (0.213)	0.225 (0.306)	0.648* (0.330)	-0.308 (0.403)	0.679 (0.467)	0.631* (0.320)	0.492 (0.333)	0.0519 (0.317)	0.241 (0.285)	0.259 (0.278)	0.342 (0.231)	1.251* (0.599)
CPM runup (Intensity)	0.0987 (0.161)	0.0496 (0.236)	0.154 (0.238)	-0.131 (0.332)	0.226 (0.345)	0.581** (0.261)	0.219 (0.270)	-0.0876 (0.260)	-0.0491 (0.217)	-0.00479 (0.240)	0.0116 (0.177)	0.945* (0.485)
FCBK Exposure x Intensity	-0.390** (0.197)	-0.228 (0.282)	-0.598* (0.305)	0.255 (0.373)	-0.613 (0.435)	-0.609** (0.293)	-0.454 (0.308)	-0.0227 (0.292)	-0.218 (0.262)	-0.219 (0.251)	-0.316 (0.212)	-1.185** (0.559)
Other sources of Political News												
TV	-0.0293 (0.0475)	-0.0582 (0.0654)	0.0216 (0.0760)	0.0167 (0.0830)	-0.136* (0.0789)	0.0497 (0.0778)	0.00831 (0.0629)	-0.0256 (0.0865)	-0.0820 (0.0651)	-0.0572 (0.0713)	-0.0562 (0.0608)	0.0225 (0.0936)
Newspapers	0.00439 (0.0195)	-0.0121 (0.0271)	0.0178 (0.0302)	-0.0307 (0.0334)	0.0491 (0.0371)	0.0143 (0.0350)	0.00115 (0.0275)	-0.0297 (0.0347)	0.0370 (0.0265)	0.00949 (0.0254)	-0.0169 (0.0224)	0.0498 (0.0470)
Talk Shows	-0.0489** (0.0222)	-0.0586* (0.0309)	-0.0418 (0.0344)	0.00185 (0.0442)	-0.112*** (0.0395)	0.0206 (0.0383)	-0.0545* (0.0321)	-0.0847** (0.0385)	-0.0530* (0.0302)	-0.0336 (0.0297)	-0.0538** (0.0246)	-0.0327 (0.0556)
Internet	-0.0234 (0.0209)	-0.0129 (0.0276)	-0.0325 (0.0335)	-0.0435 (0.0391)	-0.0272 (0.0383)	0.0297 (0.0354)	-0.0481* (0.0290)	0.0298 (0.0346)	-0.0586** (0.0287)	-0.0377 (0.0276)	-0.0378 (0.0241)	-0.00423 (0.0492)
Radio	0.00509 (0.0191)	0.0459* (0.0262)	-0.0301 (0.0286)	-0.0119 (0.0293)	0.00703 (0.0376)	-0.00305 (0.0361)	0.00308 (0.0259)	-0.00636 (0.0348)	-0.0131 (0.0250)	0.00326 (0.0273)	0.0139 (0.0215)	-0.0125 (0.0467)
CPM for TV ads	0.107 (0.127)	0.0576 (0.187)	0.0224 (0.189)	0.267 (0.212)	0.148 (0.255)	-0.0177 (0.204)	0.402** (0.197)	-0.335* (0.200)	0.0655 (0.181)	0.133 (0.186)	0.100 (0.149)	0.131 (0.250)
Observations	2,076	1,147	929	571	818	687	1,031	688	944	742	1,537	539
R-squared	0.134	0.164	0.202	0.289	0.198	0.238	0.172	0.216	0.220	0.256	0.157	0.283

Note: LPM estimates. The dependent variable is 1 if the respondent changed her voting behaviour, compared to the intentions revealed at the pre-election interview. This pertains any change in turnout or presidential vote. The CPM runup is measured as the ratio between the relative CPM price at the election week and the relative CPM price at the week of the pre-election interview. All models include Income Class FE, State FE, Day of the Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, party registration, vote at the Primary, gender, religion, age, marital status, education, ethnicity, number of children, employment status, home ownership and gun ownership status.

Table IX: The effect of Facebook on Change in Trump Support

Dep. Variable: Change in Trump Voting	All (1)	Women (2)	Men (3)	Liberal (4)	Moderate (5)	Conservative (6)	Vote Dem. 2012 (7)	Vote Rep. 2012 (8)	Clinton Intention (9)	Trump Intention (10)	White (11)	Non-white (12)
News Access on FCBK (Exposure)	0.495*** (0.184)	0.539** (0.283)	0.460* (0.276)	-0.0793 (0.137)	0.625* (0.336)	0.806** (0.378)	0.415** (0.209)	0.693* (0.385)	0.0216 (0.136)	0.712** (0.342)	0.489** (0.225)	0.452 (0.305)
CPM runup (Intensity)	0.338** (0.133)	0.388 (0.247)	0.259 (0.197)	0.0749 (0.111)	0.106 (0.218)	0.995*** (0.296)	0.284* (0.152)	0.511 (0.313)	-0.0389 (0.111)	0.661** (0.288)	0.278* (0.158)	0.413 (0.265)
FCBK Exposure x Intensity	-0.447*** (0.173)	-0.498* (0.265)	-0.416 (0.261)	0.0621 (0.128)	-0.533* (0.314)	-0.764** (0.351)	-0.376* (0.196)	-0.620* (0.358)	-0.0122 (0.127)	-0.641** (0.316)	-0.441** (0.209)	-0.420 (0.293)
Other sources of Political News												
TV	-0.0199 (0.0374)	-0.00262 (0.0446)	-0.0283 (0.0712)	0.0302 (0.0252)	-0.129** (0.0667)	0.0559 (0.0671)	0.00756 (0.0466)	0.0393 (0.0769)	-0.0642 (0.0490)	-0.0591 (0.0685)	-0.0199 (0.0488)	-0.0258 (0.0604)
Newspapers	-0.0176 (0.0153)	-0.0369* (0.0205)	0.00146 (0.0257)	-0.0271 (0.0176)	-0.00178 (0.0276)	-0.00630 (0.0326)	-0.0210 (0.0173)	-0.0203 (0.0329)	0.00210 (0.0153)	0.00971 (0.0253)	-0.0309 (0.0193)	0.00945 (0.0233)
Talk Shows	-0.0245 (0.0170)	-0.0414* (0.0227)	-0.0148 (0.0270)	-0.000125 (0.0183)	-0.0475* (0.0282)	0.00950 (0.0366)	-0.00843 (0.0179)	-0.0752* (0.0368)	-0.0181 (0.0154)	-0.0385 (0.0296)	-0.0353* (0.0203)	-0.000723 (0.0332)
Internet	-0.00383 (0.0161)	0.00384 (0.0199)	-0.0112 (0.0278)	0.000679 (0.0167)	-0.00839 (0.0280)	0.0233 (0.0322)	-0.0166 (0.0171)	0.0303 (0.0329)	0.0142 (0.0145)	-0.0315 (0.0275)	-0.0171 (0.0197)	0.00296 (0.0272)
Radio	0.000768 (0.0147)	0.0173 (0.0192)	0.00170 (0.0233)	-0.0168 (0.0127)	0.00336 (0.0271)	0.0102 (0.0342)	0.00351 (0.0153)	-0.0188 (0.0332)	-0.0184 (0.0120)	0.00503 (0.0270)	0.00236 (0.0177)	-0.00240 (0.0248)
CPM for TV ads	-0.00322 (0.0941)	-0.0140 (0.121)	-0.0903 (0.149)	-0.0477 (0.0763)	0.0303 (0.170)	-0.0197 (0.190)	0.125 (0.103)	-0.191 (0.193)	0.0123 (0.0814)	0.152 (0.186)	0.0589 (0.125)	-0.131 (0.128)
Observations	2,076	1,147	929	571	818	687	1,031	688	944	742	1,537	539
R-squared	0.120	0.165	0.177	0.273	0.193	0.232	0.178	0.189	0.150	0.262	0.143	0.339

Note: LPM estimates. The dependent variable is 1 if the respondent changed her Trump vote position, compared to the intentions revealed at the pre-election interview. The CPM runup is measured as the difference between the relative CPM price at the election week and the relative CPM price at the week of the pre-election interview. All models include Income Class FE, State FE, Day of the Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, party registration, vote at the Primary, gender, religion, age, marital status, education, ethnicity, number of children, employment status, home ownership and gun ownership status.

Table X: The effect of Facebook on Change in Clinton Support

Dep. Variable: Change in Clinton Voting	All (1)	Women (2)	Men (3)	Liberal (4)	Moderate (5)	Conservative (6)	Vote Dem 2012 (7)	Vote Rep 2012 (8)	Clinton Intention (9)	Trump Intention (10)	White (11)	Non-white (12)
News Access on FCBK (Exposure)	0.289* (0.168)	0.280 (0.235)	0.224 (0.246)	-0.344 (0.397)	0.614 (0.389)	0.266 (0.173)	0.440 (0.298)	-0.0447 (0.171)	0.241 (0.285)	0.195 (0.120)	0.214 (0.169)	0.903 (0.580)
CPM runup (Intensity)	0.120 (0.127)	0.000876 (0.182)	0.261 (0.183)	-0.0185 (0.329)	0.348 (0.293)	0.343* (0.145)	0.296 (0.235)	-0.0860 (0.159)	-0.0491 (0.217)	0.125 (0.0862)	0.00415 (0.134)	0.800* (0.468)
FCBK Exposure x Intensity	-0.284 (0.154)	-0.271 (0.215)	-0.223 (0.228)	0.298 (0.368)	-0.601* (0.362)	-0.248 (0.159)	-0.412 (0.276)	0.0439 (0.157)	-0.218 (0.262)	-0.173 (0.108)	-0.214 (0.155)	-0.855 (0.545)
Other sources of Political News												
TV	-0.0427 (0.0406)	-0.0604 (0.0576)	-0.0187 (0.0550)	-0.0140 (0.0854)	-0.0585 (0.0696)	-0.0699 (0.0695)	-0.0401 (0.0558)	-0.0416 (0.0605)	-0.0820 (0.0651)	-0.0391 (0.0348)	-0.0569 (0.0490)	-0.0232 (0.0843)
Newspapers	0.0261* (0.0151)	0.0216 (0.0227)	0.0244 (0.0198)	0.00529 (0.0317)	0.0626* (0.0319)	0.0263 (0.0166)	0.0225 (0.0249)	0.00343 (0.0171)	0.0370 (0.0265)	0.0125 (0.00770)	0.0147 (0.0157)	0.0433 (0.0449)
Talk Shows	-0.0306* (0.0171)	-0.0297 (0.0254)	-0.0326 (0.0242)	-0.00966 (0.0424)	-0.0651 (0.0344)	0.0183 (0.0167)	-0.0573* (0.0295)	-0.0259 (0.0173)	-0.0530* (0.0302)	0.00326 (0.00890)	-0.0307* (0.0177)	-0.0274 (0.0495)
Internet	-0.00892 (0.0164)	-0.00499 (0.0234)	-0.0125 (0.0227)	-0.0444 (0.0384)	0.00758 (0.0341)	0.0201 (0.0182)	-0.0278 (0.0266)	0.0250 (0.0161)	-0.0586** (0.0287)	0.00531 (0.00996)	-0.0115 (0.0174)	-0.00187 (0.0466)
Radio	-0.00328 (0.0149)	0.0209 (0.0216)	-0.0331 (0.0203)	-0.00599 (0.0287)	-0.0245 (0.0332)	0.0130 (0.0174)	-0.0159 (0.0235)	0.00456 (0.0164)	-0.0131 (0.0250)	-0.00159 (0.00978)	-0.00100 (0.0158)	0.00650 (0.0448)
CPM for TV ads	0.252 (0.217)	0.369 (0.261)	-0.0726 (0.299)	0.443 (0.412)	0.0760 (0.381)	-0.250 (0.277)	0.0193 (0.315)	0.299 (0.279)	0.487 (0.325)	0.0926 (0.183)	0.210 (0.243)	-0.535 (0.559)
Observations	2,076	1,147	929	571	818	687	1,031	688	944	742	1,537	539
R-squared	0.123	0.136	0.205	0.249	0.174	0.242	0.146	0.232	0.220	0.214	0.140	0.228

Note: the dependent variable is 1 if the respondent changed her Clinton vote position, compared to the intentions revealed at the pre-election interview. The CPM run up is measured as the difference between the relative CPM price at the election week and the relative CPM price at the week of the pre-election interview. All models include Income Class FE, State FE, Day of the Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, party registration, vote at the Primary, gender, religion, age, marital status, education, ethnic origin, number of children, employment status, home ownership and gun ownership status.

Table XI: Effect of Facebook on Information

Dep. Variable:	None	Ideology	also	also	also
Improved Political Knowledge		& Location	Ethnicity	Gender	Age
	(1)	(2)	(3)	(4)	(5)
News Access on FCBK (Exposure)	0.0585 (0.0376)	0.206 (0.458)	-0.285 (0.471)	0.101 (0.446)	0.314 (0.380)
CPM Runup (Intensity)		0.482 (0.317)	0.551* (0.324)	0.248 (0.313)	0.236 (0.203)
FCBK Exposure x Intensity		-0.137 (0.425)	0.320 (0.440)	-0.0399 (0.416)	-0.234 (0.346)
Other Sources of Political News					
Internet Access	0.0226 (0.0908)	0.0139 (0.0904)	0.0121 (0.0901)	0.0209 (0.0908)	0.0209 (0.0908)
TV	-0.0343 (0.0990)	-0.0385 (0.0990)	-0.0441 (0.0987)	-0.0353 (0.0990)	-0.0334 (0.0990)
Newspapers	0.00462 (0.0388)	0.00525 (0.0387)	0.00451 (0.0387)	0.00510 (0.0388)	0.00514 (0.0388)
Talk Shows	-0.0440 (0.0463)	-0.0465 (0.0464)	-0.0502 (0.0464)	-0.0450 (0.0463)	-0.0433 (0.0463)
Internet	0.0751* (0.0445)	0.0796* (0.0446)	0.0814* (0.0445)	0.0775* (0.0447)	0.0739* (0.0446)
Radio	-0.00859 (0.0406)	-0.00697 (0.0408)	-0.00696 (0.0407)	-0.00791 (0.0406)	-0.00985 (0.0407)
CPM for TV ads	-0.307 (0.252)	-0.321 (0.252)	-0.342 (0.252)	-0.316 (0.252)	-0.302 (0.253)
Observations	1,888	1,888	1,888	1,888	1,888
R-squared	0.084	0.085	0.089	0.084	0.085

Note: LPM estimates. The dependent variable is the share of correct answers to political knowledge questions given during the post-election interview over the share of correct answer given during the pre-election interview. The table reports mean daily relative CPM prices for Facebook ads displayed during the period 01 September 2016 - 08 November 2017. The first column reports the price of ads targeted on the basis of user location and ideology. The remaining columns report the price of ads targeted on the basis of location, ideology and an additional demographic characteristic such as Gender, ethnicity or Age.

Table XII: The effect of Facebook on Voter Turnout

Dep. Variable:	All	Women	Men	Liberal	Moderate	Conservative	Vote Dem2012	Vote Rep2012	Clinton Intention	Trump Intention	White	Non-white
Effective Turnout	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
News Access on FCBK (Exposure)	-0.383** (0.187)	-0.299 (0.266)	-0.521* (0.311)	0.0170 (0.340)	-0.864** (0.342)	-0.329 (0.320)	-0.248 (0.249)	-0.506 (0.316)	0.438* (0.261)	-0.612* (0.323)	-0.505** (0.204)	0.0437 (0.468)
CPM runup (Intensity)	-0.239* (0.140)	-0.218 (0.242)	-0.421* (0.217)	0.0352 (0.262)	-0.304 (0.249)	-0.474* (0.262)	-0.0129 (0.186)	-0.318 (0.257)	0.125 (0.176)	-0.576** (0.269)	-0.348** (0.152)	0.0400 (0.345)
FCBK Exposure x Intensity	0.330* (0.173)	0.289 (0.245)	0.429 (0.293)	-0.0209 (0.319)	0.733** (0.319)	0.330 (0.292)	0.210 (0.231)	0.430 (0.289)	-0.416* (0.243)	0.555* (0.298)	0.444* (0.190)	-0.0923 (0.428)
Other sources of Political News												
TV	0.00681 (0.0424)	0.0228 (0.0554)	-0.000185 (0.0678)	0.0361 (0.0797)	0.0243 (0.0695)	-0.0370 (0.0599)	-0.0194 (0.0459)	0.154 (0.0980)	0.0412 (0.0547)	0.0356 (0.0638)	0.0139 (0.0527)	-0.0201 (0.0809)
Newspapers	-0.0174 (0.0153)	-0.0342* (0.0201)	0.00277 (0.0251)	-0.00795 (0.0286)	-0.0201 (0.0295)	-0.0319 (0.0271)	-0.0132 (0.0203)	-0.00894 (0.0233)	-0.0436* (0.0230)	-0.00103 (0.0232)	-0.00626 (0.0172)	-0.0212 (0.0393)
Talk Shows	0.0261 (0.0178)	0.0233 (0.0239)	0.0379 (0.0298)	0.00202 (0.0351)	0.0411 (0.0333)	-0.00333 (0.0293)	0.0195 (0.0243)	0.0432 (0.0267)	0.0290 (0.0268)	0.0318 (0.0261)	0.0425* (0.0194)	-0.0437 (0.0501)
Internet	0.0627** (0.0172)	0.0685*** (0.0223)	0.0482* (0.0282)	0.0485 (0.0345)	0.0966*** (0.0317)	0.0261 (0.0264)	0.0626*** (0.0240)	0.0426* (0.0242)	0.0601** (0.0238)	0.0362 (0.0253)	0.0679*** (0.0192)	0.0571 (0.0451)
Radio	-0.00386 (0.0151)	0.00159 (0.0206)	-0.0203 (0.0237)	0.00681 (0.0257)	-0.0103 (0.0299)	0.00862 (0.0276)	0.0124 (0.0204)	-0.00674 (0.0232)	0.000977 (0.0214)	0.00520 (0.0243)	-0.00843 (0.0167)	0.0190 (0.0402)
CPM for TV ads	-0.0141 (0.111)	0.113 (0.141)	-0.114 (0.178)	-0.0422 (0.173)	0.0955 (0.242)	-0.254 (0.171)	-0.180 (0.167)	0.0293 (0.132)	-0.0666 (0.162)	-0.0887 (0.159)	-0.0418 (0.133)	-0.135 (0.238)
Observations	2,075	1,146	929	570	818	687	1,030	688	943	742	1,537	538
R-squared	0.179	0.227	0.229	0.255	0.274	0.254	0.172	0.250	0.173	0.249	0.186	0.300

Note: LPM estimates. The dependent variable is 1 if the respondent votes and zero otherwise. The CPM runup is measured as the ratio between the relative CPM price at the election week and the relative CPM price at the week of the pre-election interview. All models include Income Class FE, State FE, Day of the Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, party registration, vote at the Primary, gender, religion, age, marital status, education, ethnicity, number of children, employment status, home ownership and gun ownership status.

Table XIII: The effect of Facebook on Trump Vote

Dep. Variable: Effective Trump Vote	All (1)	Women (2)	Men (3)	Liberal (4)	Moderate (5)	Conservative (6)	Vote Dem. 2012 (7)	Vote Rep. 2012 (8)	Clinton Intention (9)	Trump Intention (10)	White (11)	Non-white (12)
News Access on FCBK (Exposure)	-0.274 (0.212)	-0.247 (0.305)	-0.407 (0.325)	0.0418 (0.174)	-0.841** (0.383)	0.249 (0.428)	-0.191 (0.273)	-0.251 (0.419)	-0.00991 (0.115)	-0.698** (0.342)	-0.302 (0.263)	-0.384 (0.364)
CPM runup (Intensity)	-0.308** (0.152)	-0.401 (0.263)	-0.264 (0.247)	-0.00341 (0.130)	-0.618** (0.267)	-0.191 (0.323)	-0.133 (0.194)	-0.597* (0.313)	-0.0733 (0.0922)	-0.645** (0.288)	-0.341* (0.185)	-0.305 (0.288)
FCBK Exposure x Intensity	0.269 (0.199)	0.267 (0.283)	0.374 (0.305)	-0.0318 (0.164)	0.805** (0.357)	-0.209 (0.394)	0.212 (0.257)	0.227 (0.385)	0.0123 (0.106)	0.627** (0.316)	0.278 (0.246)	0.394 (0.339)
Other sources of Political News												
TV	0.0261 (0.0435)	0.0655 (0.0531)	-0.0114 (0.0766)	-0.0273 (0.0690)	0.0949 (0.0697)	-0.0168 (0.0902)	-0.0861 (0.0585)	0.111 (0.0871)	-0.0677 (0.0484)	0.0566 (0.0686)	0.0596 (0.0562)	-0.0456 (0.0743)
Newspapers	-0.0408** (0.0179)	-0.0317 (0.0240)	-0.0221 (0.0293)	0.0210 (0.0219)	-0.0509 (0.0346)	-0.0660* (0.0347)	-0.00602 (0.0233)	-0.0293 (0.0333)	0.000605 (0.0128)	-0.00850 (0.0254)	-0.0504** (0.0223)	-0.0144 (0.0288)
Talk Shows	0.0236 (0.0201)	0.0462* (0.0265)	-0.0142 (0.0325)	-0.00101 (0.0217)	0.0497 (0.0363)	0.0112 (0.0421)	-0.0183 (0.0246)	0.0369 (0.0403)	-0.0204 (0.0140)	0.0396 (0.0297)	0.0375 (0.0240)	-0.0412 (0.0362)
Internet	0.00898 (0.0196)	0.00976 (0.0265)	0.00971 (0.0309)	-0.000882 (0.0277)	0.0173 (0.0369)	0.0150 (0.0377)	-0.0167 (0.0256)	0.0156 (0.0379)	0.0158 (0.0142)	0.0325 (0.0275)	-0.00194 (0.0241)	0.00853 (0.0317)
Radio	-0.0205 (0.0181)	-0.0327 (0.0243)	-0.0119 (0.0287)	-0.0222 (0.0214)	-0.00599 (0.0345)	-0.0353 (0.0387)	0.00222 (0.0218)	-0.0153 (0.0364)	-0.0184* (0.0111)	-0.00441 (0.0270)	-0.0328 (0.0218)	0.0103 (0.0321)
CPM for TV ads	0.0555 (0.111)	0.194 (0.154)	-0.180 (0.169)	-0.0238 (0.0908)	0.307 (0.231)	-0.174 (0.226)	-0.114 (0.120)	0.189 (0.223)	0.0450 (0.0762)	-0.147 (0.186)	0.0517 (0.154)	-0.0408 (0.160)
Observations	2,066	1,143	923	571	810	685	1,025	686	942	740	1,533	533
R-squared	0.532	0.575	0.539	0.310	0.429	0.404	0.313	0.294	0.132	0.262	0.518	0.537

Note: LPM estimates. The dependent variable is 1 if the respondent voted for Trump and 0 otherwise. The CPM runup is measured as the difference between the relative CPM price at the election week and the relative CPM price at the week of the pre-electoral interview. All models include Income Class FE, State FE, Day of the Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, party registration, vote at the Primary, gender, religion, age, marital status, education, ethnicity, number of children, employment status, home ownership and gun ownership status.

Table XIV: The effect of Facebook on Clinton Vote

Dep. Variable: Effective Clinton Vote	All (1)	Women (2)	Men (3)	Liberal (4)	Moderate (5)	Conservative (6)	Vote Dem. 2012 (7)	Vote Rep. 2012 (8)	Clinton Intention (9)	Trump Intention (10)	White (11)	Non-white (12)
News Access on FCBK (Exposure)	-0.117 (0.201)	-0.00742 (0.294)	-0.142 (0.292)	-0.348 (0.418)	-0.0650 (0.403)	-0.189 (0.257)	-0.176 (0.350)	0.0253 (0.258)	0.457 (0.293)	-0.00723 (0.0742)	-0.137 (0.217)	-0.0588 (0.510)
CPM runup (Intensity)	-0.0359 (0.148)	0.159 (0.263)	-0.288 (0.231)	-0.00585 (0.304)	-0.102 (0.311)	-0.0522 (0.205)	-0.0532 (0.255)	0.285 (0.206)	0.0867 (0.205)	0.0856 (0.0701)	-0.0674 (0.158)	0.0193 (0.403)
FCBK Exposure x Intensity	0.0599 (0.186)	-0.0358 (0.271)	0.0765 (0.271)	0.300 (0.390)	-0.0319 (0.375)	0.168 (0.233)	0.0986 (0.328)	-0.0509 (0.235)	-0.439 (0.275)	0.0117 (0.0700)	0.0830 (0.202)	0.00667 (0.466)
Other sources of Political News												
TV	-0.0139 (0.0461)	-0.0144 (0.0617)	0.00417 (0.0718)	0.0639 (0.115)	-0.0567 (0.0722)	-0.0853 (0.0855)	0.0735 (0.0737)	0.0481 (0.0573)	0.0904 (0.0646)	-0.0379 (0.0337)	-0.0443 (0.0609)	0.0412 (0.0877)
Newspapers	0.0178 (0.0176)	0.000322 (0.0248)	0.0213 (0.0267)	-0.0113 (0.0385)	0.0399 (0.0349)	-0.00562 (0.0225)	0.00494 (0.0294)	0.00725 (0.0208)	-0.0384 (0.0265)	0.00935 (0.00671)	0.0471** (0.0200)	-0.0401 (0.0436)
Talk Shows	0.00583 (0.0201)	-0.00959 (0.0278)	0.0394 (0.0316)	0.0253 (0.0459)	-0.0171 (0.0390)	-0.00400 (0.0275)	0.0397 (0.0331)	0.00159 (0.0249)	0.0573* (0.0301)	-0.00125 (0.00798)	0.0141 (0.0220)	-0.0246 (0.0520)
Internet	0.0375* (0.0195)	0.0320 (0.0267)	0.0332 (0.0300)	0.0196 (0.0433)	0.0378 (0.0380)	0.0212 (0.0253)	0.0553* (0.0322)	0.0196 (0.0207)	0.0566** (0.0286)	0.00395 (0.00894)	0.0465** (0.0219)	0.0469 (0.0507)
Radio	0.0224 (0.0182)	0.0370 (0.0248)	0.00445 (0.0285)	0.0346 (0.0368)	0.0143 (0.0384)	0.0457* (0.0252)	0.00984 (0.0291)	0.00875 (0.0214)	0.0141 (0.0250)	0.00746 (0.00707)	0.0263 (0.0207)	0.0287 (0.0453)
CPM for TV ads	0.0141 (0.120)	0.126 (0.165)	-0.0675 (0.176)	-0.0550 (0.244)	-0.263 (0.268)	0.0467 (0.165)	0.00987 (0.200)	-0.0636 (0.144)	-0.0636 (0.183)	0.0691 (0.0980)	0.0378 (0.128)	-0.0787 (0.284)
Observations	2,066	1,143	923	568	810	685	1,025	686	942	740	1,533	533
R-squared	0.556	0.565	0.598	0.386	0.456	0.465	0.341	0.272	0.223	0.212	0.585	0.499

Note: LPM estimates. The dependent variable is 1 if the respondent voted for Clinton and 0 otherwise. The CPM runup is measured as the difference between the relative CPM price at the election week and the relative CPM price at the week of the pre-electoral interview. All models include Income Class FE, State FE, Day of the Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, party registration, vote at the Primary, gender, religion, age, marital status, education, ethnicity, number of children, employment status, home ownership and gun ownership status.

Figures

Figure I: *Effect of targeting on audience-specific ad prices*





Figure II: *Facebook CPM Prices across differently aligned states*



Notes. Relative CPM (weekly average of daily median) for the states of New York, Florida and Texas. The period is from July the 4th 2016 (2016w27) to the week after the elections (2016w46). The horizontal line (at 1) represents the state baseline CPM. The vertical line indicate election day. Audiences are divided between Conservative, Liberal and Moderate users.

Figure III: *Identification Strategy*

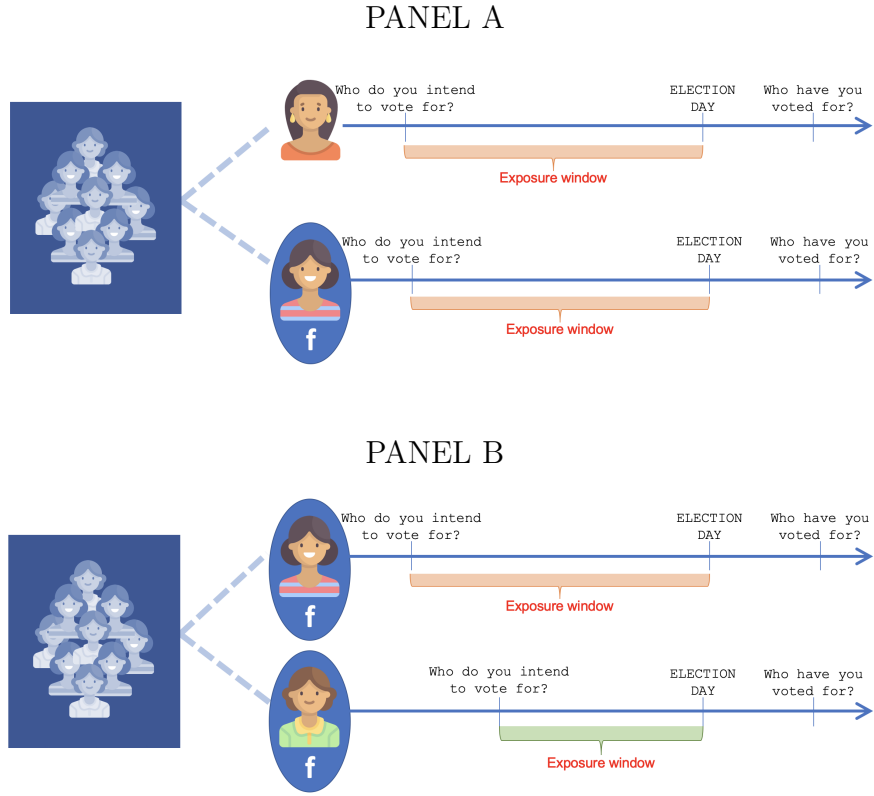

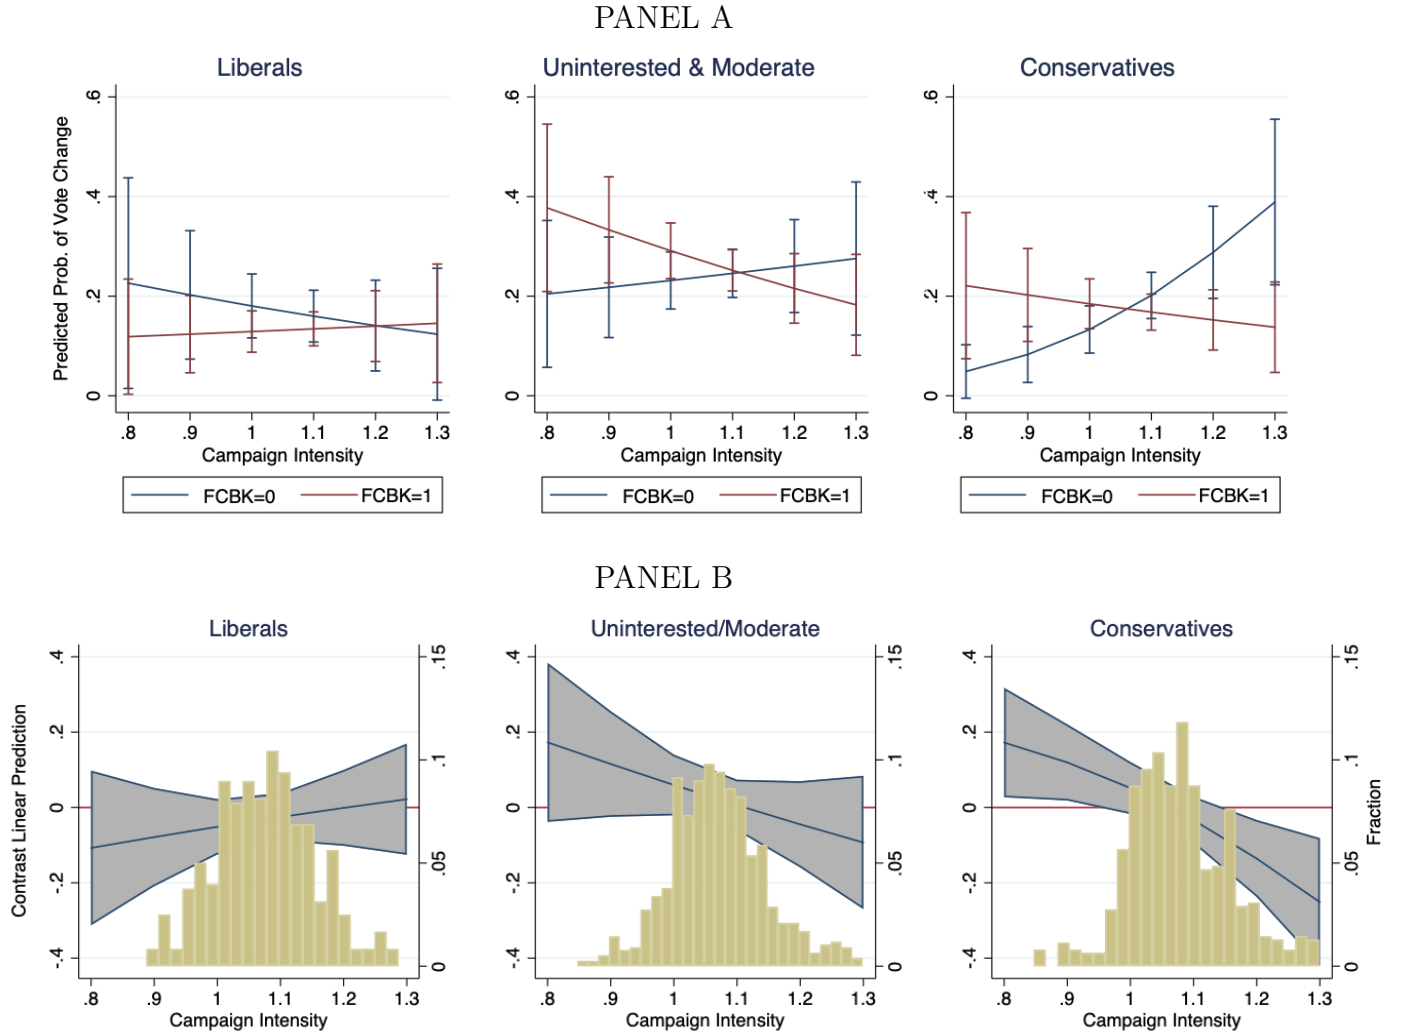


Figure IV: *Facebook CPM Prices in Democratic Countries with Fair Elections*



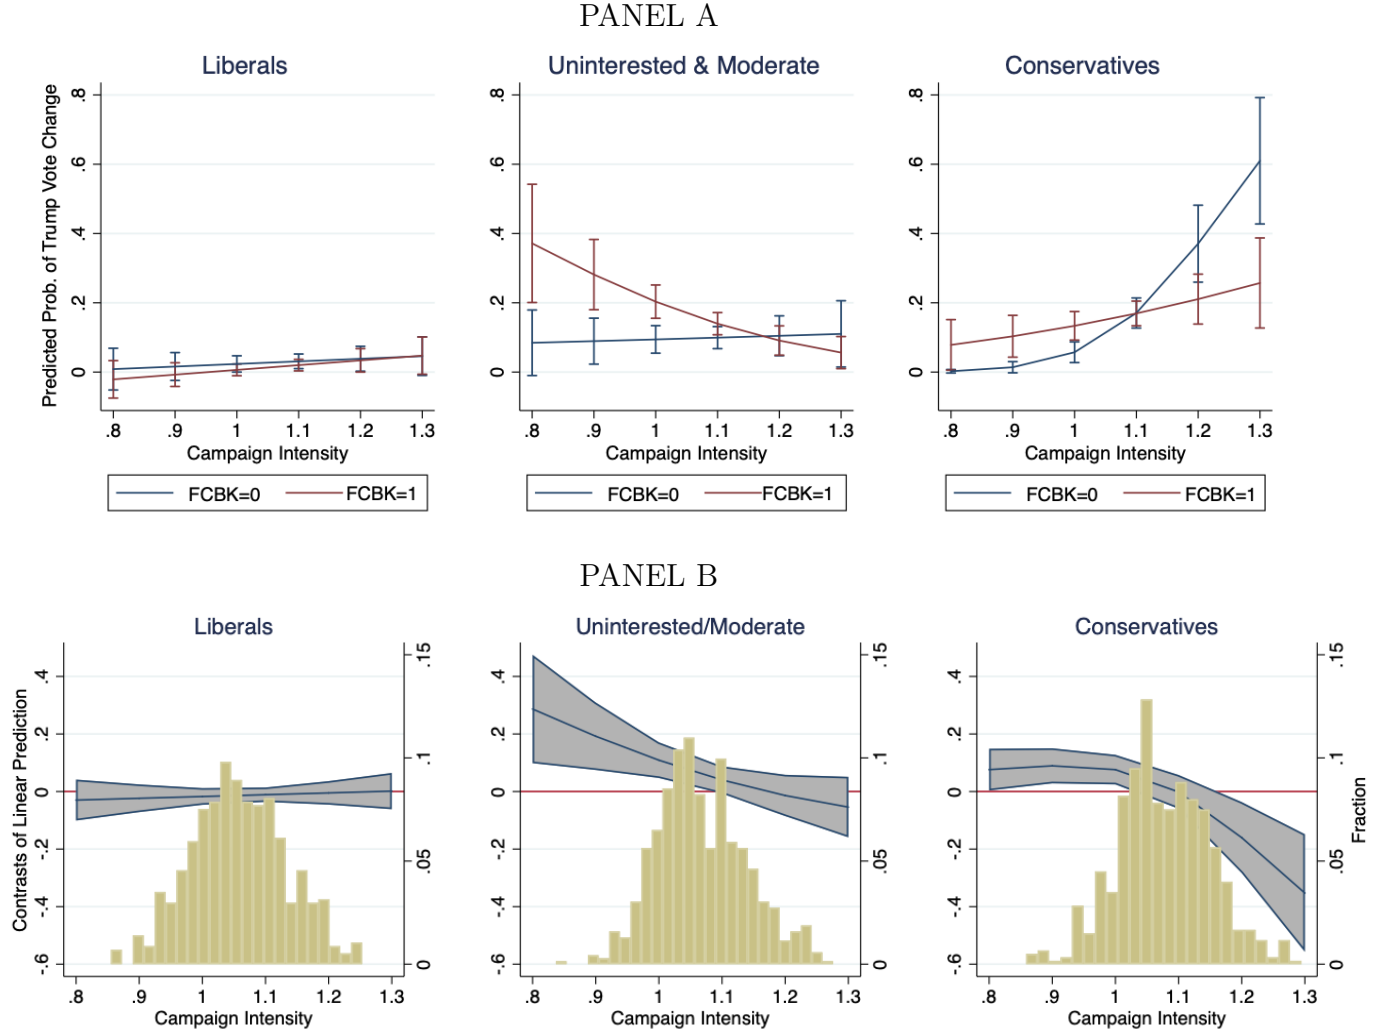
Notes: Facebook daily CPM median prices for Country-specific audiences. Prices are measured as difference from global trend, and conditioned on country and time fixed effects. Sample covers all 45 Fair Elections that took place between July 2015 and February 2017. Free Elections defined on the basis of the Freedom House Electoral Democracy rating.

Figure V: *Effect of Facebook micro-targeted campaigns on vote change: ideology*



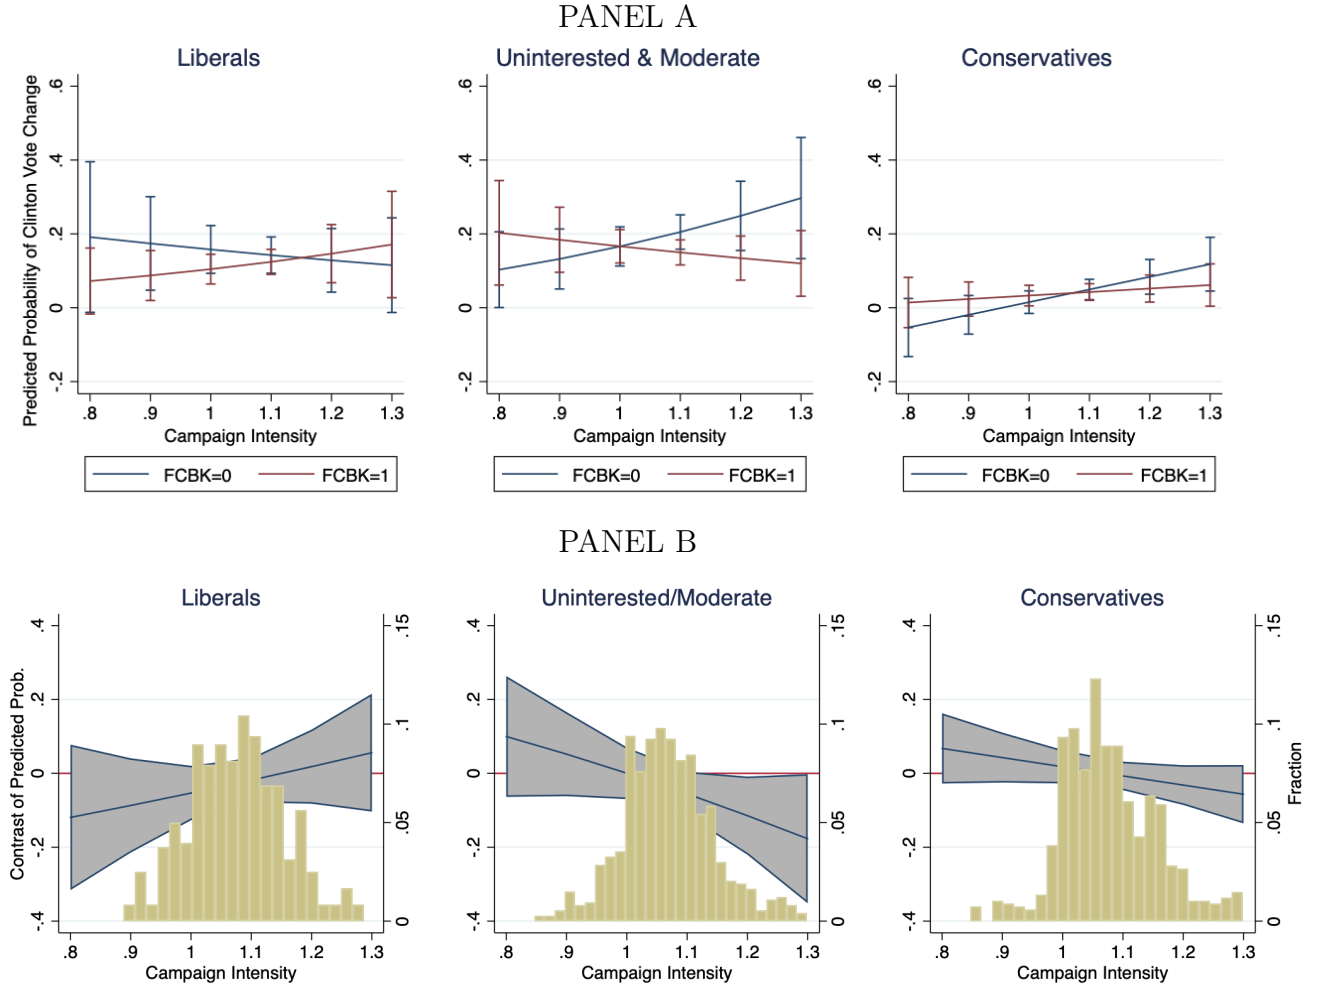
Notes: Panel A reports the predicted probability of changing vote for Facebook users and non-users at different campaign intensity levels. Plots are derived from Probit estimates of model (3) as in columns (4), (5) and (6) of Table VIII. Panel B displays the difference in vote change between Facebook users and non users, with 95 percent confidence intervals. The bar charts at the bottom of the plots display the distribution of campaign intensity for the respective group of voters. For both panels campaign intensity is based on audiences defined by political ideology and ethnicity.

Figure VI: *Effect of Facebook micro-targeted campaigns on Trump vote change: ideology*



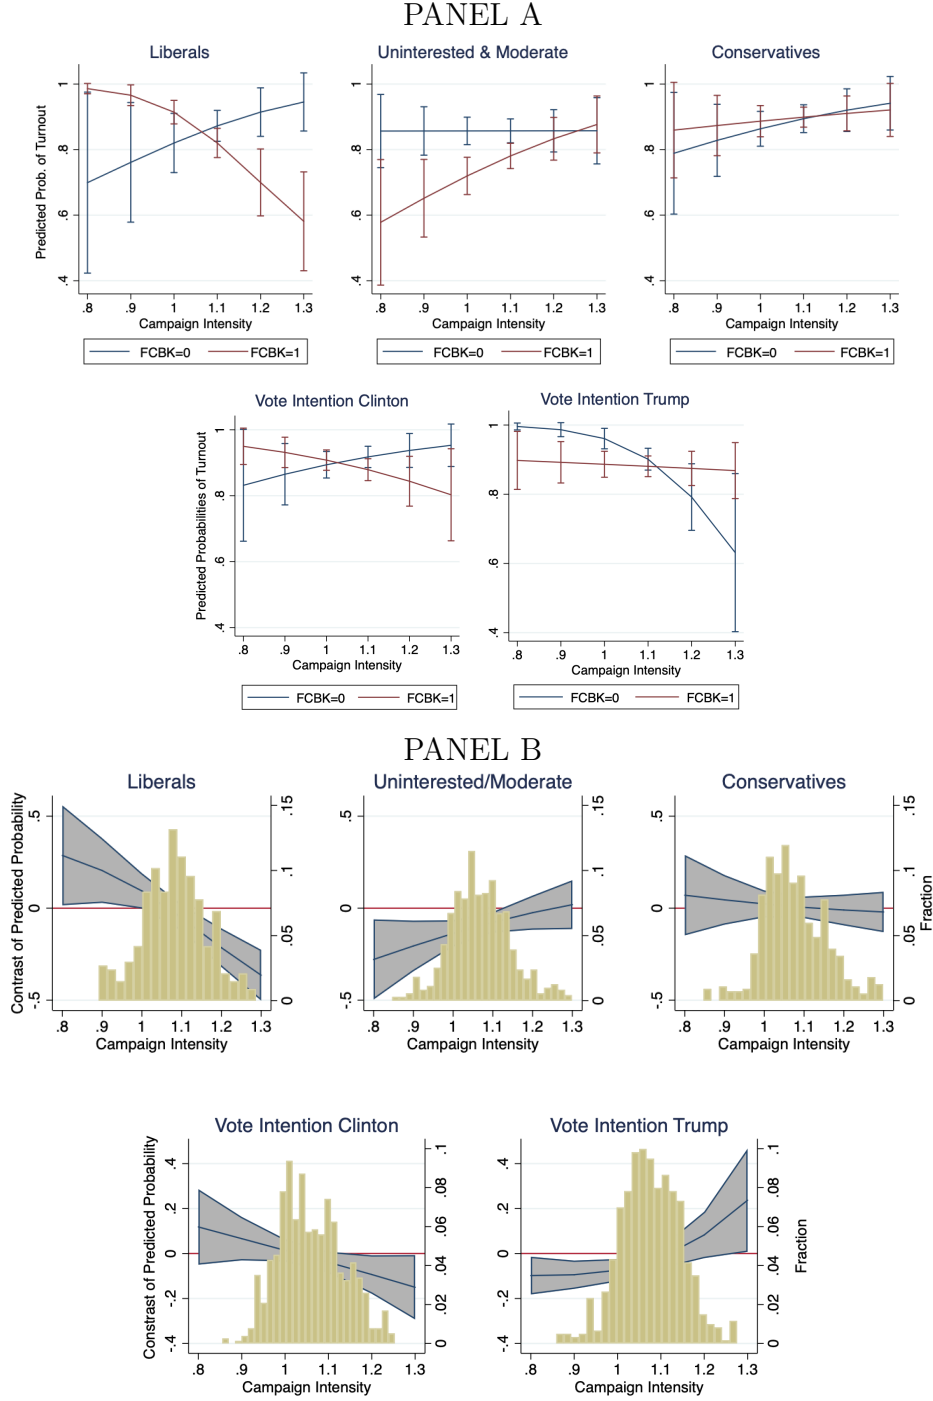
Notes: Panel A reports the predicted probability of changing vote for Facebook users and non-users at different campaign intensity levels. Plots are derived from Probit estimates of model (3) as in columns (4), (5) and (6) of Table IX. Panel B displays the difference in vote change between Facebook users and non users, with 95 percent confidence intervals. The bar charts at the bottom of the plots display the distribution of campaign intensity for the respective group of voters. For both panels campaign intensity is based on audiences defined by political ideology and ethnicity.

Figure VII: *Effect of Facebook micro-targeted campaigns on Clinton vote change: ideology*



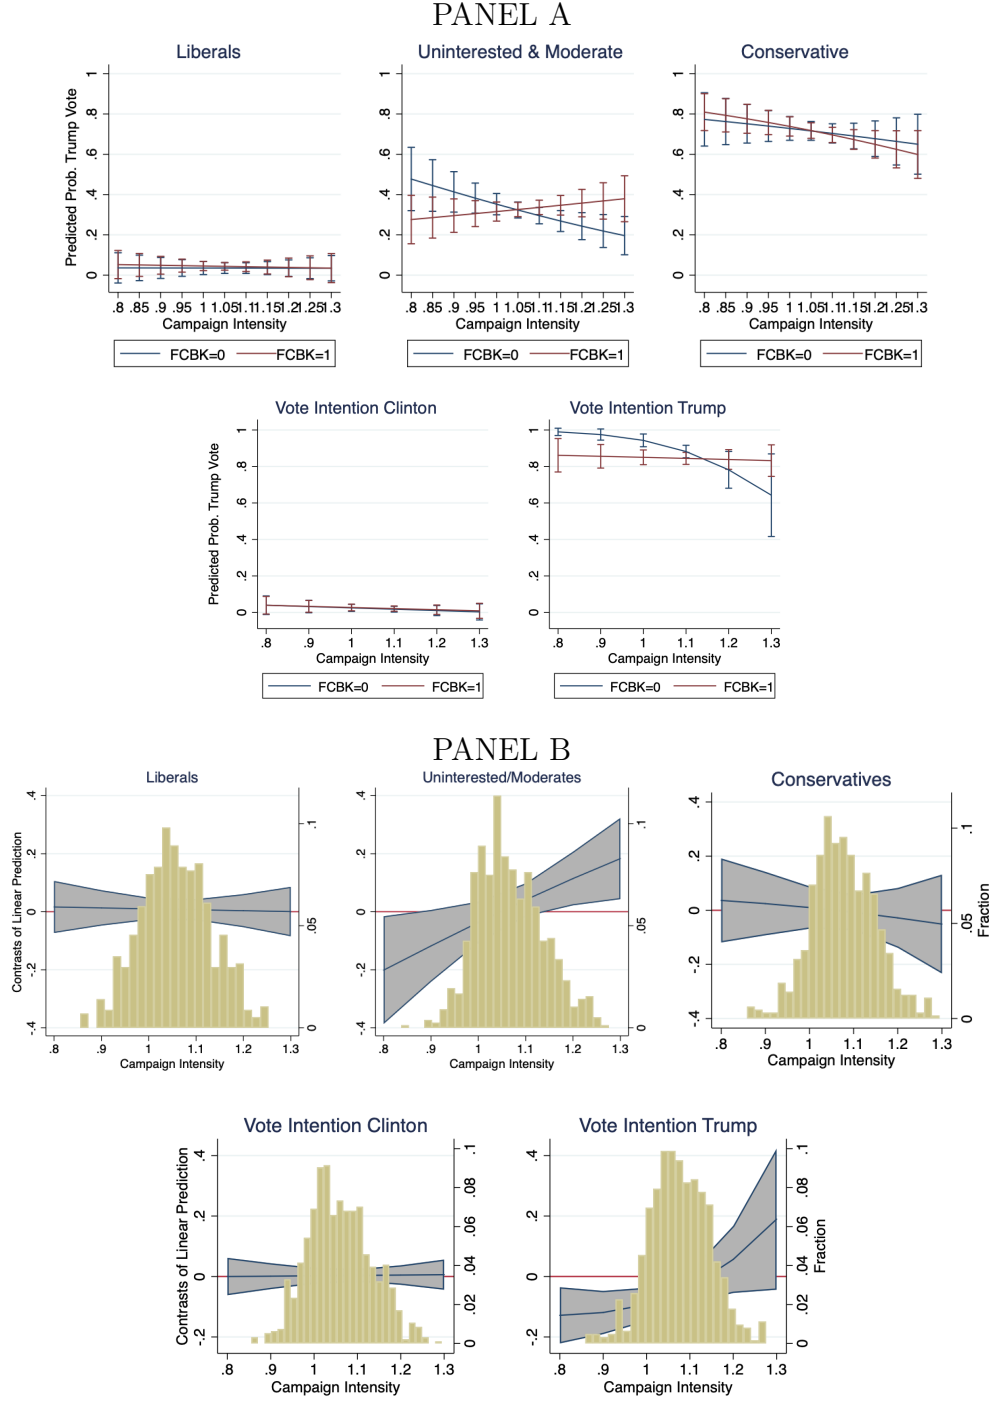
Notes: Panel A reports the predicted probability of changing vote for Facebook users and non-users at different campaign intensity levels. Plots are derived from Probit estimates of model (3) as in columns (4), (5) and (6) of Table X. Panel B displays the difference in vote change between Facebook users and non users, with 95 percent confidence intervals. The bar charts at the bottom of the plots display the distribution of campaign intensity for the respective group of voters. For both panels campaign intensity is based on audiences defined by political ideology and ethnicity.

Figure VIII: *Effect of Facebook micro-targeted campaigns on Turnout: Ideology*



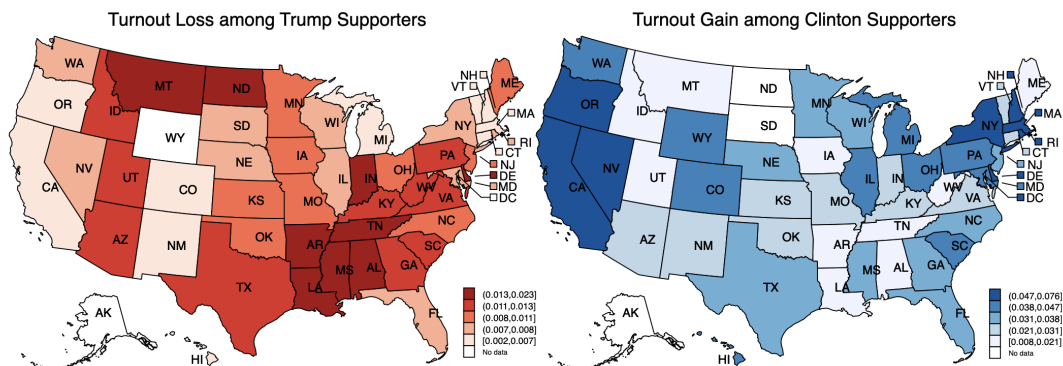
Notes: Panel A reports the predicted probability of going to vote for Facebook users and non-users at different campaign intensity levels. Plots are derived from Probit estimates of model (3) as in columns (4), (5), (6), (9) and (10) of Table XII. Panel B displays the difference in turnout between Facebook users and non users, with 95 percent confidence intervals. The bar charts at the bottom of the plots display the distribution of campaign intensity for the respective group of voters. For both panels campaign intensity is based on audiences defined by political ideology and ethnicity.

Figure IX: *Effect of Facebook micro-targeted campaigns on Trump vote: ideology*



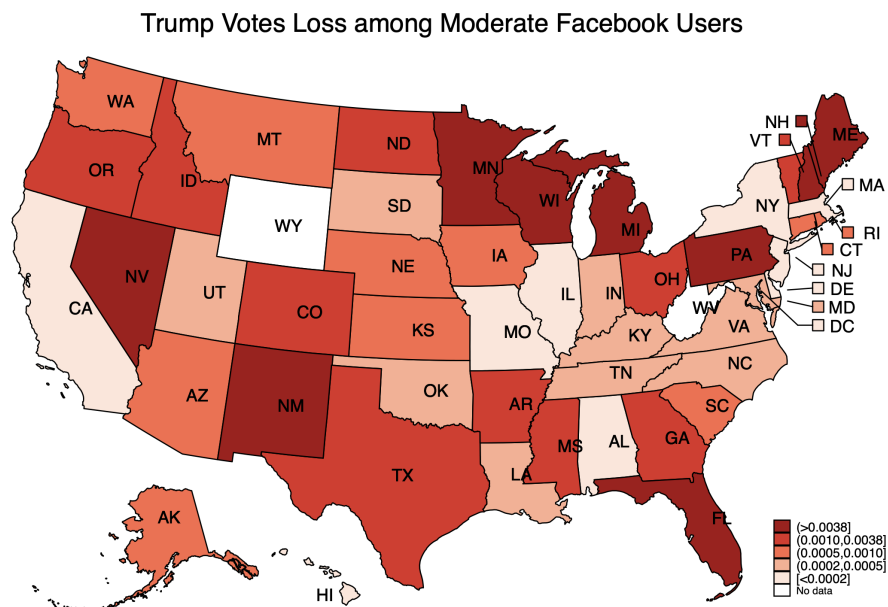
Notes: Panel A reports the predicted probability of voting for Trump for Facebook users and non-users at different campaign intensity levels. Plots are derived from Probit estimates of model (3) as in columns (4), (5), (6), (9) and (10) of Table XIII. Panel B displays the difference in turnout between Facebook users and non users, with 95 percent confidence intervals. The bar charts at the bottom of the plots display the distribution of campaign intensity for the respective group of voters. For both panels campaign intensity is based on audiences defined by political ideology and ethnicity.

Figure X: *Effect on Turnout of Marginal Reduction in Intensity of Facebook Campaign*



Notes: effect of a marginal reduction (by 1%) in the intensity of the Facebook Campaign on turnout. The left map uses the total marginal effect of intensity estimated in column (10) of Table XII, and the right map uses the total marginal effect of intensity estimated in column (9) of the same Table. The effect is computed on the turnout of individuals who use Facebook and intend to vote for Trump, in the left map, and vote for Clinton, in the right map. It is then expressed as a share of the total turnout of individuals who use Facebook.

Figure XI: *Effect on Trump Vote of Marginal Reduction in Intensity of Facebook Campaign*



Notes: effect of a marginal reduction (by 1%) in the intensity of the Facebook Campaign on Trump Votes. It uses the total marginal effect of intensity estimated in column (5) of Table XIII. The effect is computed on the Trump votes cast by Moderate individuals who use Facebook. It is then expressed as a share of the State-specific vote margin.

Appendix

A Advertising Auctions on Facebook

To allocate advertising opportunities on its website, Facebook adopts an auction system based on the Vickrey-Groves-Clark (VCG) mechanism. We present this mechanism in an informal way. Consider an auction where a set of goods is being sold. In the case of Facebook, these goods correspond to advertising space on the “wall” of a set of users defined by certain characteristics (e.g., demographics, interests, etc.), i.e. an *audience*. For each of these goods, bidders (advertisers) announce the maximum price they are willing to pay. Bidders cannot see other people’s bids (sealed-bid auction). The auction closes once all the bids are in.

In the standard VCG auction, the auctioneer calculates for each bid the marginal loss to the other bidders if the bid were successful. For those who would have obtained the goods if the bid in question were ignored, the loss is equal to the bid they placed (i.e., their declared willingness to pay). The loss is instead zero to all bidders who would not have obtained the good, even if the bid considered were ignored. The auctioneer allocates the good to the highest bidders and charges them the marginal loss their bid has caused to others.

To fix ideas, suppose there is only one good to be allocated. In this case, the auction system would allocate the good to the highest bidder, and charge the second-highest bid (i.e., the loss to the bidder who would have received the good otherwise) to the winner. That is, the mechanism reduces to a second-price auction when a single good is for sale.

It can be shown that this mechanism maximizes the aggregate utility of bidders, since all the goods are attributed to the participants with the highest willingness-to-pay. Furthermore, if agents are fully rational and in the absence of collusion, the willingness to pay is reported truthfully. This is because only the marginal harm to other bidders will be charged to each participant, making truthful bidding a (weakly) dominant strategy. However, this type of auction does not maximize the seller’s revenue. According to Facebook, though, this disadvantage is unimportant. The reason is that, although some revenue may be sacrificed in the short run, in the long run the mechanism improves the relevance of the ad to the selected audience and, therefore, the effectiveness of advertising on the platform.

Facebook’s auction system is in fact more complex than the standard VCG auction, because it considers not only the marginal loss to other bidders (advertisers), but also the loss to the users who get exposed to the ads. Although the company does not disclose the details of this procedure, we can summarily describe it as follows. For a given audience, Facebook calculates a *relevance score* associated to each ad proposed by the respective bidders. The higher this score, the higher the imputed loss to the audience when that advertiser’s bid is

unsuccessful. The score captures the cost to the audience of not seeing the ad. Facebook uses the relevance score to determine which bids are successful and how much the winners should be charged (that is, the size of the combined loss on users and other advertisers). This process is controlled by a Facebook automated algorithm, which is not observed by the researcher nor by the advertiser.³⁹

According to Facebook, this system minimizes the probability that ads are shown to uninterested audiences. Furthermore, it increases the price advertisers should expect to pay when attempting to reach an audience that is also targeted by other relevant ads. Furthermore, less relevant ads cause higher "social" harm, so the price the advertiser pays increases as well.

B Collection of Facebook data

Our main source of information is Facebook's Ad Manager website.⁴⁰ This website provides an interface that advertisers can use to place ads on Facebook.

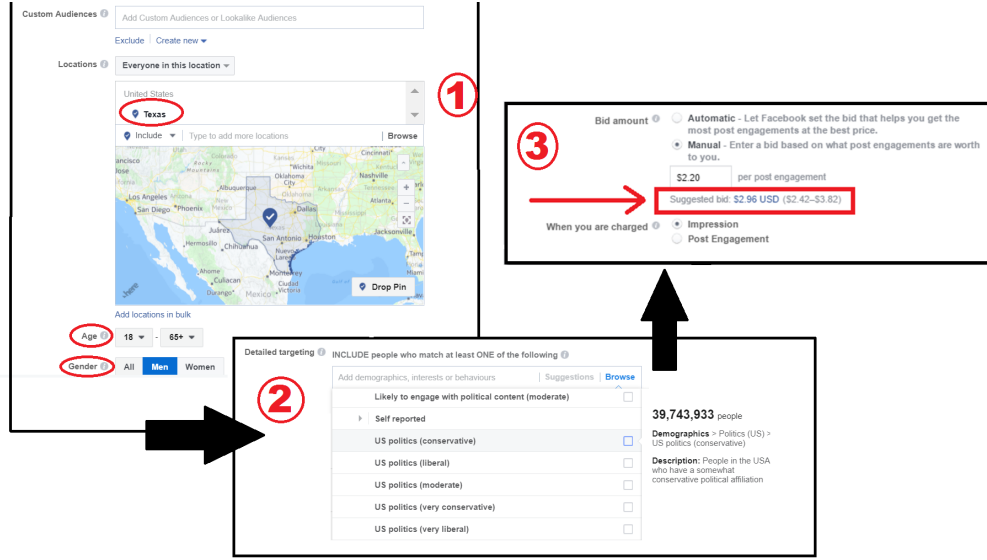
An advertiser using the Ad Manager selects an audience, defined by any combination of geographical location, demographical characteristics, personal interests and behaviors. The advertiser also selects a compensation method: CPM, CPC or CPA. Cost per Mille (CPM), is the unit price paid every time the ad is shown a thousand times to users in the selected audience. Advertisers who care for driving traffic to a specific web-page generally prefer to pay per click by the selected audience. This compensation method is referred to as Cost per Click (CPC). Finally, Cost per Action is a price paid every time a consumer adopts a certain action (e.g., visiting the advertiser's webpage).

The advertiser can then choose a maximum bid for impressions on that audience. At the time of our data collection (2016-2017), Facebook displayed a suggested *bid range*. The range of suggested bids indicated a minimum, a median and a maximum bid. According to Facebook, this is the range of prices that are currently being paid by advertisers winning auctions for the selected audience. See Figure XII for a screenshot of the page reporting the bid range. We queried the Ad Manager to collect the bid range for audiences defined exclusively by country of residence, for all OECD countries, using CPM, CPC and CPA metrics. The queries took place for every 15 minutes during our period of observation. We averaged the data at daily level. Note that, since late 2019, Facebook has partly revised the Ad Manager, which does not report bid ranges anymore. Instead, the website provides a range of projected hits to advertisers conditional on the size of the budget they select.

³⁹see https://www.facebook.com/business/help/430291176997542?helpref=faq_content

⁴⁰<https://www.facebook.com/business/learn/facebook-ads-reporting-ads-manager>

Figure XII: Facebook Ad Manager - Screenshot



The Ad Manager also provides advertisers with information about the expected size of the audience they intend to reach. Specifically, the website reports the expected number of daily active users that belong to the selected audience. See Figure XII. We collect this information as well to keep track of audience size.

B.1 Ethical Considerations

In this section we explicitly explain why the use of the Facebook marketing API data does not create any privacy risk for Facebook users, even though we believe it something that can be clearly inferred from the explanations of the data collection.

The European Union’s GDPR, one of the most restrictive data protection regulations to date, defines what personal data is in its Article 4: “*personal data*’ means any information relating to an identified or identifiable natural person (‘data subject’); an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person”.

The Facebook marketing API only provides aggregate numbers of targeted audiences. For instance if someone queries for the audience “Users in Spain, aged between 20 and 30, male and interested in Science” the only information she will get is the number of Facebook users within that audience. To the best of our knowledge it is unfeasible to retrieve the identity of the actual users included in that audience. Based on the GDPR definition of

personal data the Facebook marketing API data used in our research cannot be considered as personal data. Therefore, there is no privacy/data protection risks for Facebook users derived from our research.

C Preliminary evidence on ad prices

C.1 Online Ads Prices and Electoral Cycles around the World

We build a worldwide dataset of Facebook ad prices, covering 213 countries, during an interval of one-and-a-half years (from July 2015 to February 2017). We focus only on democratic countries, as determined by the Freedom House’s Electoral Democracy Rating. We select countries defined as “Free” or “Partly Free” that held democratic elections during the five years prior the beginning of our time frame. This leaves us with a sample of 45 fair elections held during our study period.⁴¹ We use country-specific CPM as the measure of unit ad prices. We filter country fixed effects, day-of-the-week cyclicalities, and time trends away from the raw-price series, and then normalize time around the election date of each country. We restrict the analysis to a time window of 200 days before and after the election. Averaging by day across countries allows us to assess whether Facebook prices grew worldwide around the time of a political electoral campaign.

As illustrated by Figure IV, average worldwide prices sharply rise during the two months prior to elections (an average increase of around 50 (USD) cents, which corresponds to about a 25 percent increase in prices), and steadily drop over the three months following the election date. After this five-month-long window, prices seem return to their equilibrium level.

The fact that prices drop more slowly than they rise is consistent with the intense media attention typical during the first trimester of any newly elected government. Notably, one would expect a very sharp price drop, following the election date. Instead, figure IV shows a rather evident lag in the movement of the CPM back to its pre-election equilibrium level. We believe that this is a mechanical effect, induced by the design of the social media advertising “dashboard” (so called “ad manager”). As previously mentioned, advertisers are fully informed of the median price that wins the auction for advertising to their audience of interest, up to a moment before placing their own bid. This causes a lag in the absorption of the election effect. To illustrate that this mechanism holds for any event affecting the social media advertising market, we produce a similar plot to that of figure IV, around the

⁴¹The average of a country/territory Political Rights and Civil Liberties ratings is called the Freedom Rating, and it is this figure that determines the status of a country/territory as Free (1.0 to 2.5), Partly Free (3.0 to 5.0), or Not Free (5.5 to 7.0). See <https://freedomhouse.org/report/freedom-world-2018-tablecountry-scores>.

commercial event known as “Black Friday,” the day after the US Thanksgiving holiday. See figure D.5. The figure shows that even for this non-politically relevant effect, the normalized prices return to their initial equilibrium with an evident lag.

C.2 Predicted winning probabilities and ad prices during the 2016 presidential campaign

For this analysis, we focus again on the 2016 US presidential election. The reason for this restriction is that US users are the only ones for whom Facebook keeps track of political orientation.

We examine the correlation between electoral winning probabilities and ad prices, based on the assumption that candidates adjust the intensity of their campaigns in response to changes in the estimated probability of winning the election. We collect daily state-specific winning probabilities for the two main presidential candidates (Trump and Clinton), during the three months preceding the election (from the 8th of August, until the 8th of November 2016). These data come from the FiveThirtyEight.com 2016 Election Forecast, which published daily winning probabilities, estimated using a comprehensive model that accounts for the result of local state-level polls, as well as macroeconomic trends and historical political outcomes.⁴² We then estimate a dynamic panel model, that relates the daily CPM prices to the absolute value of the lagged difference among the winning probabilities of the two candidates. The model is estimated using a generalized method of moments (GMM) approach, and accounts for cyclical fluctuations in internet usage (day-of-the-week dummies) and for a time trend (week-of-the-year dummies).⁴³

We report the results of this exercise in Table I. They show that, at the state level, daily Facebook ad prices are significantly affected by changes in the winning probabilities of the two candidates. Interestingly, it appears that ad prices tend to increase in response to decreases that emerge in the relative probability of Trump winning. Indeed, in Republican states (where the Republican candidate leads the polls rather consistently), the price appears to increase when the difference between the two candidates narrows (i.e., when Clinton is catching up). The opposite applies to Democrat states (i.e., price appears to increase in Democratic states when Trump is catching up to a Clinton lead). Finally, in swing states, the price increases in either situation – when Trump leads and Clinton appears to be closing


⁴²The website also computes winning probabilities simply based on local, state-level polls. When we repeated our test using this series, the results were similar.

⁴³Our methodological approach draws from a literature on the effects of electoral forecasts and outcomes on financial markets and stock prices. See, e.g., Santa-Clara and Valkanov [2003] and Sattler [2013] for the effect of electoral outcomes.

the gap, or when Clinton leads and Trump appears to be gaining.

D Supplementary Figures


Figure D.1: *US Digital vs. Broadcast TV Political Ad Spending*



Source: FEC, The Cook Political Report. Borrell Associates, 2017.

(*) indicates forecasted figures.

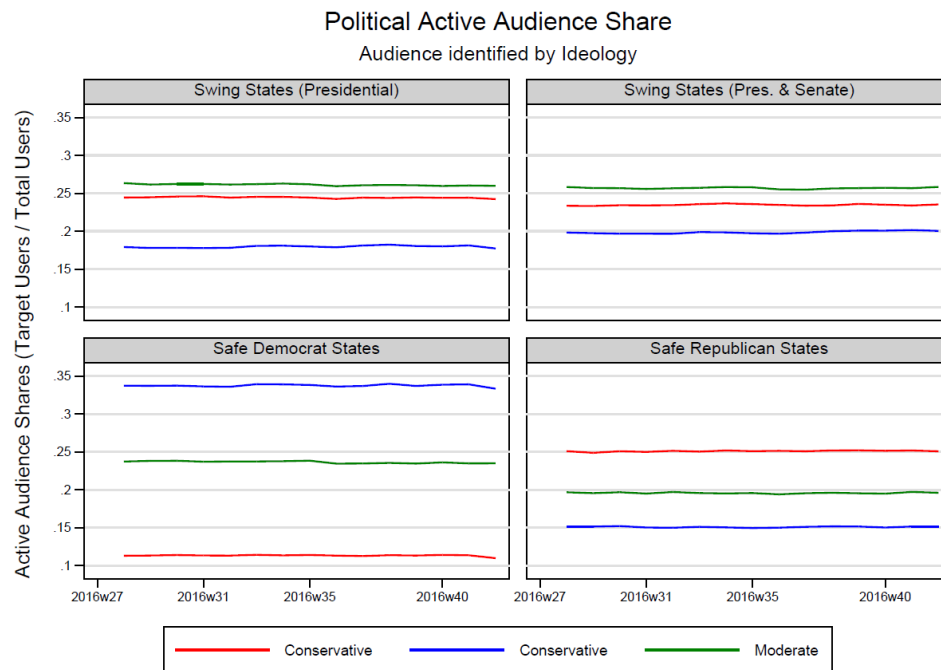
Figure D.2: *Inelastic Ads Supply*



Note: evolution of the metric $AE = \frac{\#ads}{\#posts + \#ads}$ for the period October 2016 - May 2018. Each week presents a boxplot of the metric that includes the distribution of AE across all the users with active sessions. AE is presented as percentage.

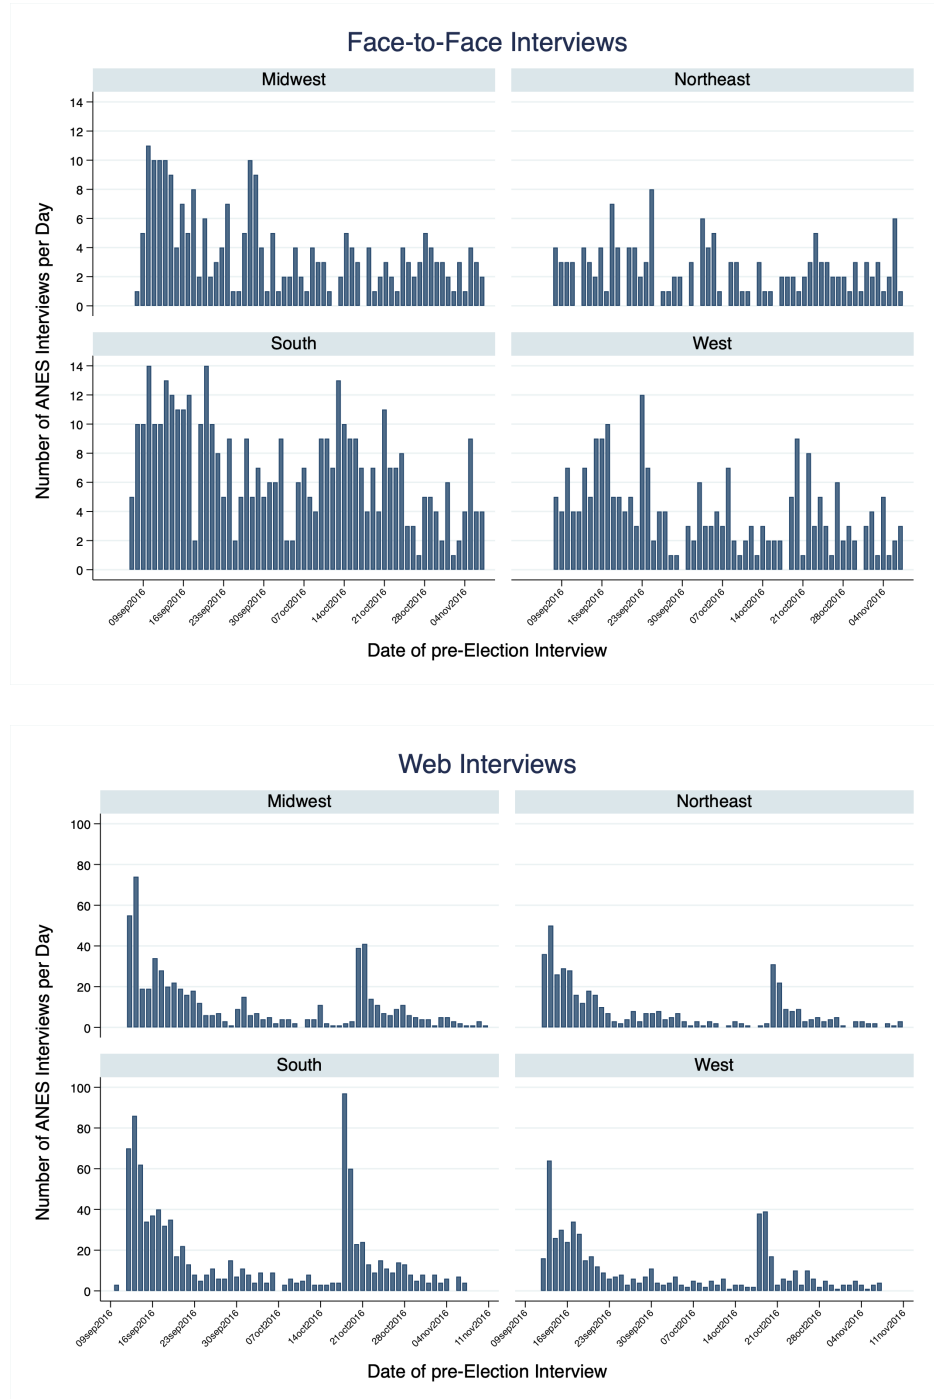
Source: Arrate et al. (2018), Fig. 5.

Figure D.3: *Number of Facebook Active Users over time*



Notes. share of Facebook users over total users for the period from July the 4th 2016 (2016w27) until the week prior to the elections (2016w44).

Figure D.4: *ANES pre-Election Interview Dates*



Notes: distribution of ANES respondents over days of interview, for both the fact-to-face and the web survey mode.

Figure D.5: *Facebook CPM Prices in US across the Black Friday*

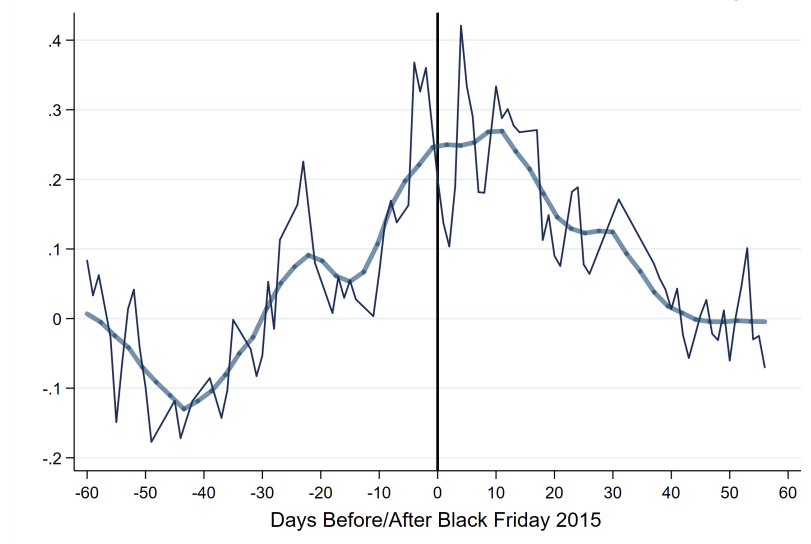
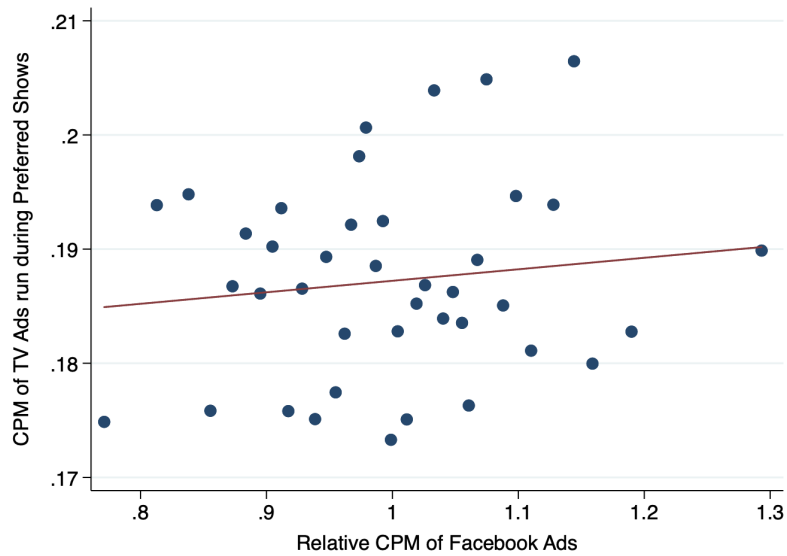
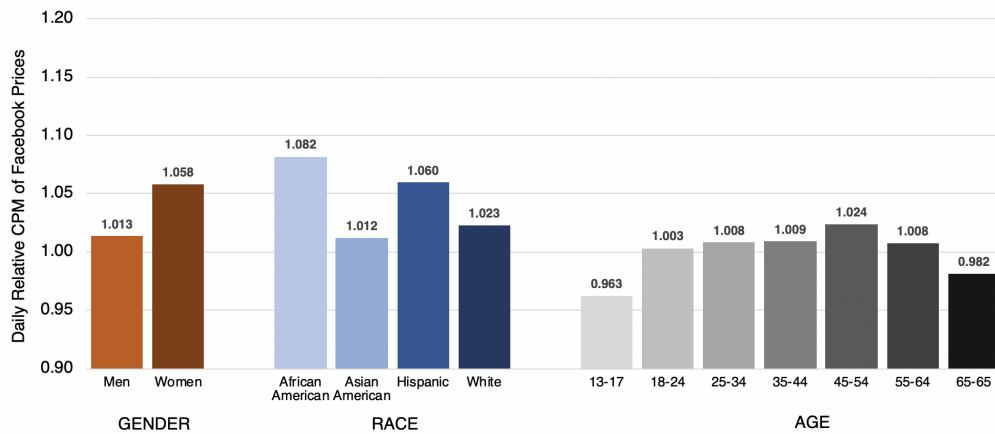


Figure D.6: *Correlation in the Campaign Intensity across Media Outlets*



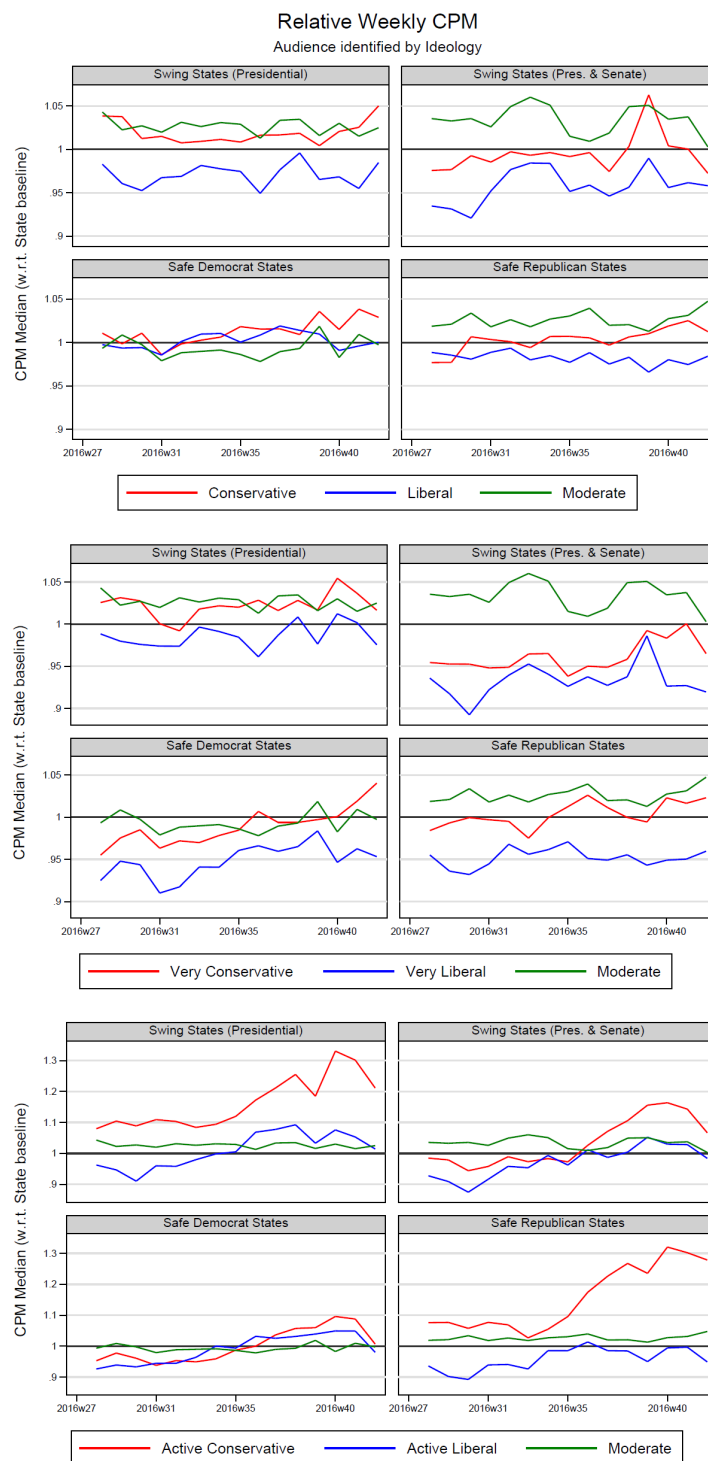
Notes: correlation between the relative CPM of ads shown on Facebook during the pre-election interview and the CPM of ads shown at the same time on the TV shows regularly watched by the ANES respondents. Individuals are binned over the distribution of Facebook relative CPM prices.

Figure D.7: *Mean Facebook CPM prices across Targeted Demographics*



Notes: mean daily relative CPM prices for Facebook ads displayed during the period 01 September 2016 - 08 November 2017. Prices are reported according to the categories of each targeting dimension (Gender, Ethnicity and Age) that is additional to Location and Ideology targeting.


Figure D.8: *Intensity of Political Campaign across States Type - Partisan vs. Moderate Users*



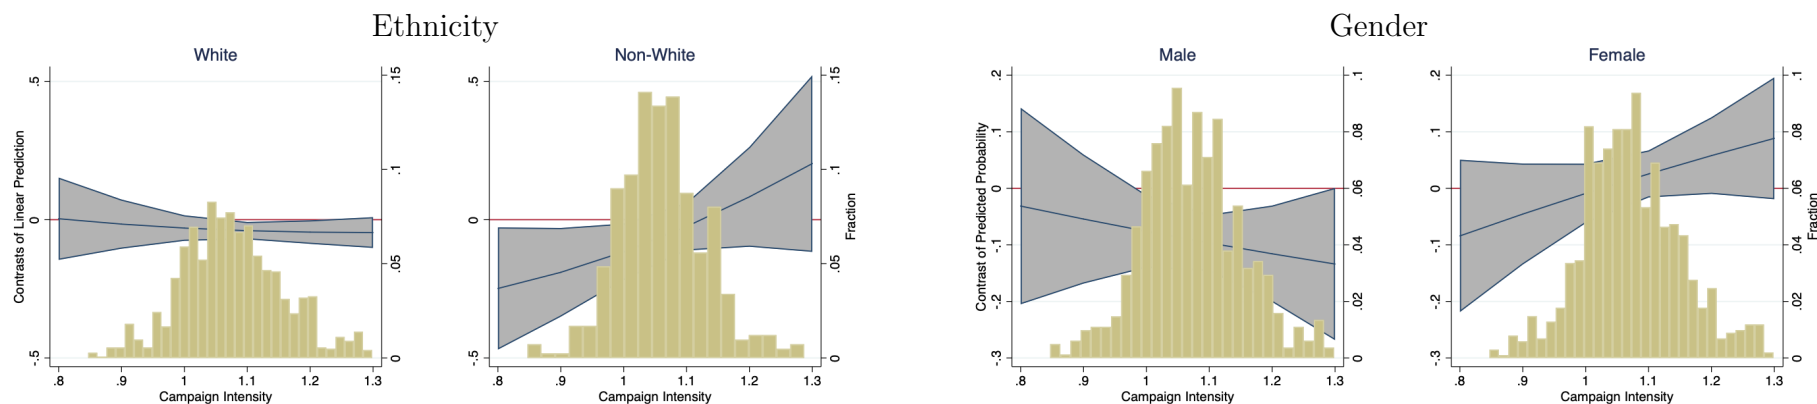
Notes: relative CPM (weekly average of daily median) for different ideology profiles and state types. The period runs from July the 4th 2016 (2016w27) until the week prior the elections (2016w44). The horizontal lines represent the state baseline CPM.

Figure D.9: *Effect of Facebook micro-targeted campaigns on Turnout: Demographics*

PANEL A




PANEL B




Notes: Panel A reports the predicted probability of going to vote for Facebook users and non-users at different campaign intensity levels. Plots are derived from the Probit estimates of model (3) as in columns (11) and (12), for ethnicity, and (2) and (3), for gender, of Table XII. Panel B displays the difference in turnout between Facebook users and non users, with 95 percent confidence intervals. The bar charts at the bottom of the plots display the distribution of campaign intensity for the respective group of voters. For both panels campaign intensity is based on audiences defined by political ideology and ethnicity.

Figure D.10: *Effect of Facebook micro-targeted campaigns on Trump Vote: Demographics*

PANEL A

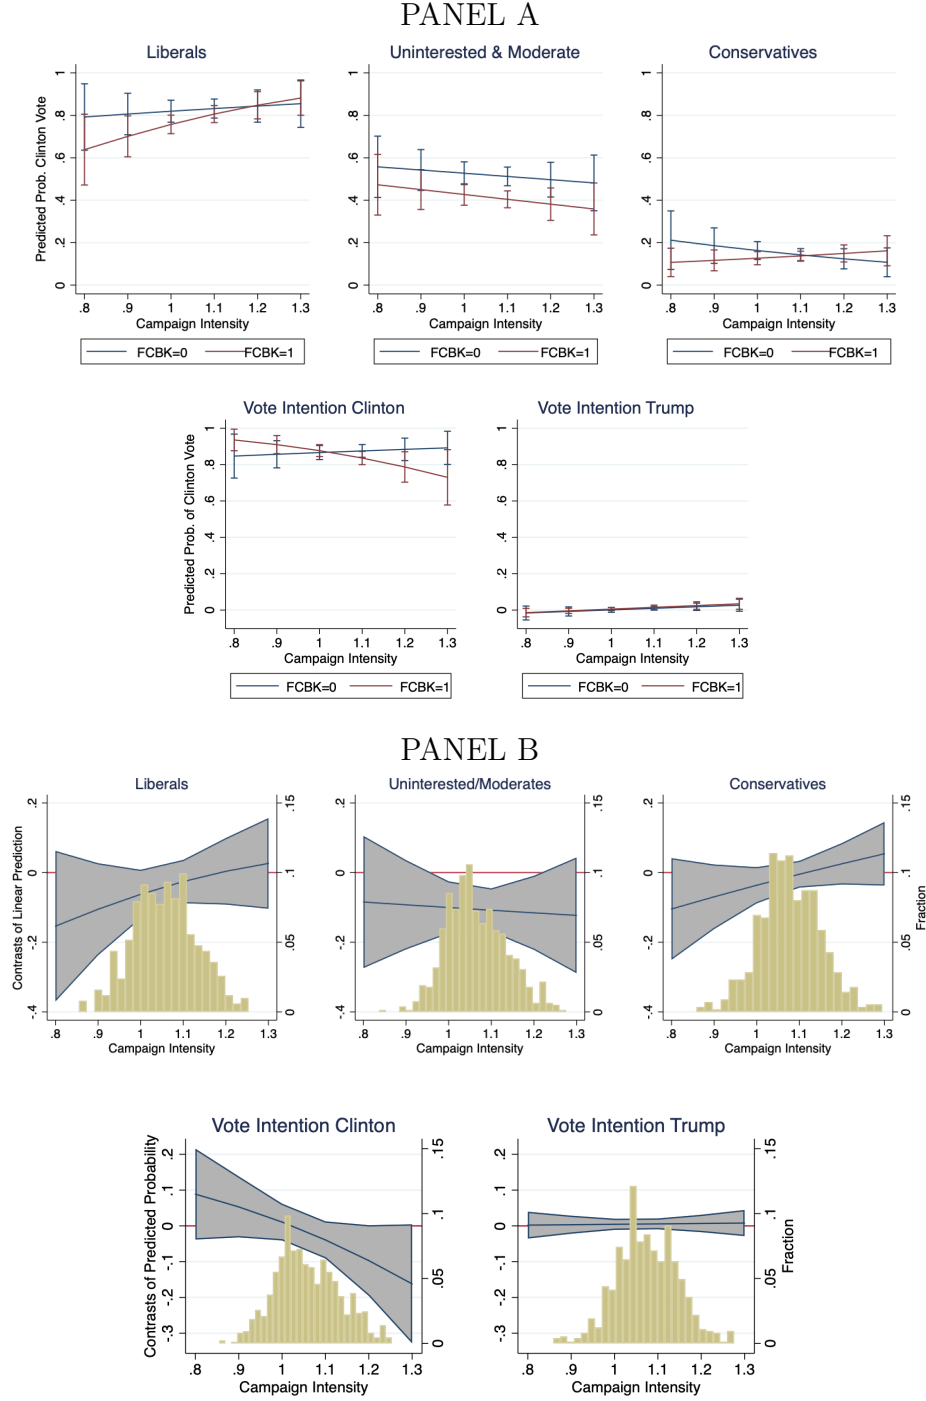


PANEL B



Notes: Panel A reports the predicted probability of voting for Trump for Facebook users and non-users at different campaign intensity levels. Plots are derived from the Probit estimates of model (3) as in columns (11) and (12), for ethnicity, and (2) and (3), for gender, of Table XIII. Panel B displays the difference in Trump Vote between Facebook users and non users, with 95 percent confidence intervals. The bar charts at the bottom of the plots display the distribution of campaign intensity for the respective group of voters. For both panels campaign intensity is based on audiences defined by political ideology and ethnicity.

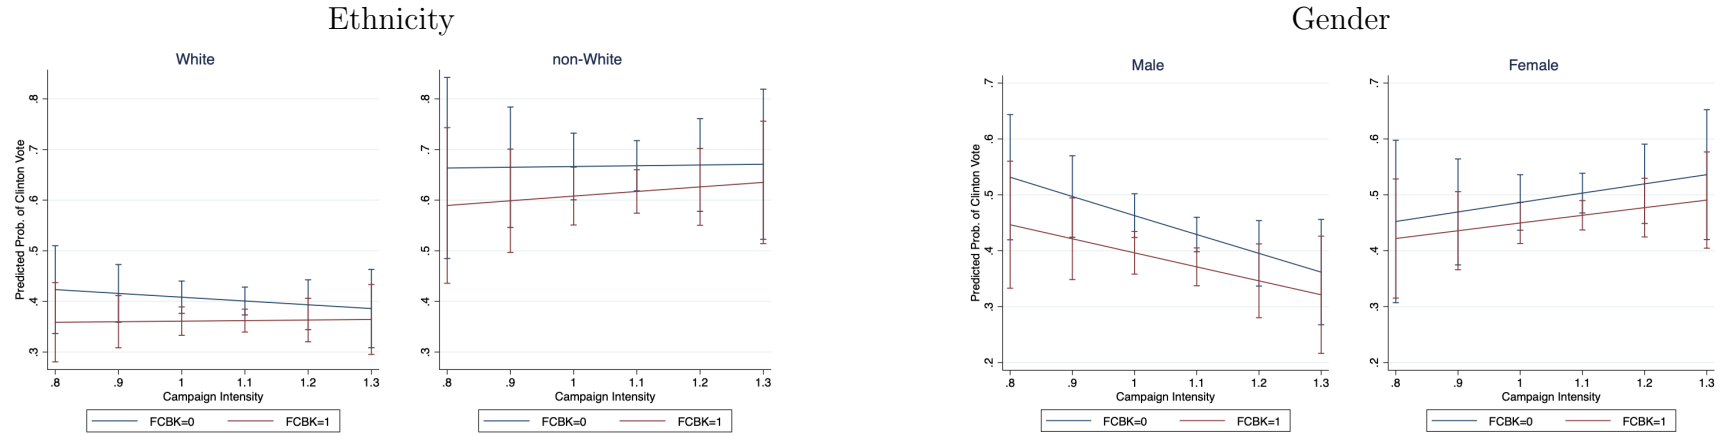
Figure D.11: *Effect of Facebook micro-targeted campaigns on Clinton vote: ideology*



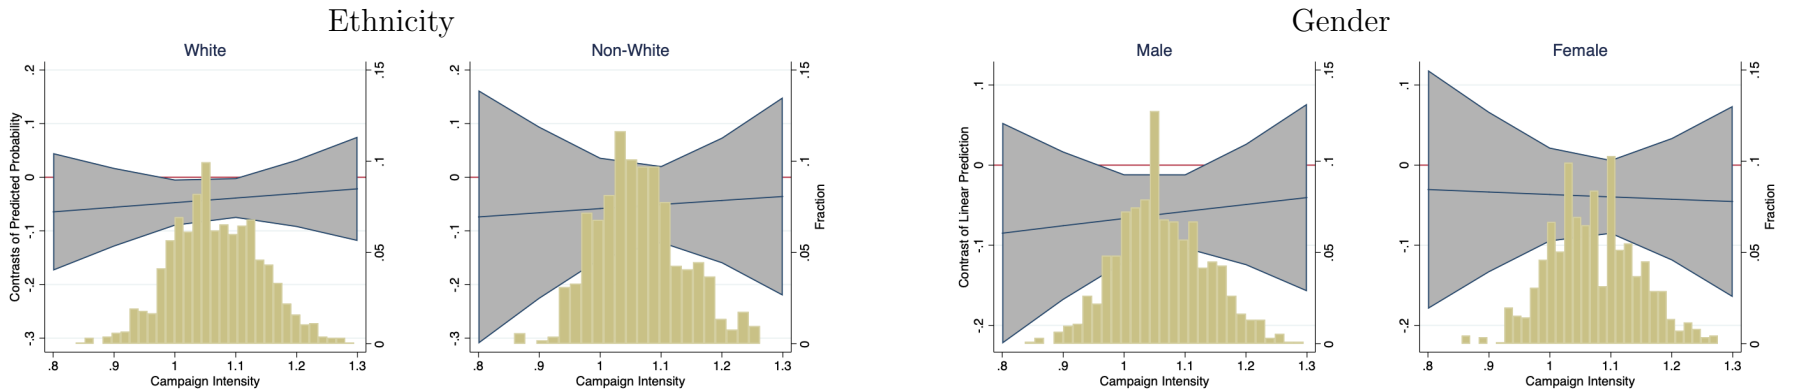
Notes: Panel A reports the predicted probability of voting for Clinton for Facebook users and non-users at different campaign intensity levels. Plots are derived from Probit estimates of model (3) as in columns (4), (5), (6), (9) and (10) of Table XIV. Panel B displays the difference in turnout between Facebook users and non users, with 95 percent confidence intervals. The bar charts at the bottom of the plots display the distribution of campaign intensity for the respective group of voters. For both panels campaign intensity is based on audiences defined by political ideology and ethnicity.

Figure D.12: *Effect of Facebook micro-targeted campaigns on Clinton Vote: Demographics*

PANEL A



PANEL B



Notes: Panel A reports the predicted probability of voting for Clinton for Facebook users and non-users at different campaign intensity levels. Plots are derived from the Probit estimates of model (3) as in columns (11) and (12), for ethnicity, and (2) and (3), for gender, of Table XIV. Panel B displays the difference in Clinton Vote between Facebook users and non users, with 95 percent confidence intervals. The bar charts at the bottom of the plots display the distribution of campaign intensity for the respective group of voters. For both panels campaign intensity is based on audiences defined by political ideology and ethnicity.