



Oct 2018

No.389

Politics in the Facebook Era
Evidence from the 2016 US Presidential Elections
Federica Liberini, Michela Redoano, Antonio Russo,
Angel Cuevas and Ruben Cuevas

WORKING PAPER SERIES

Centre for Competitive Advantage in the Global Economy

Department of Economics



Politics in the Facebook Era

Evidence from the 2016 US Presidential Elections*

Federica Liberini[†] Michela Redoano[‡] Antonio Russo[†]
Angel Cuevas[§] Ruben Cuevas[§]

October 20, 2018

Abstract

Social media enable politicians to personalize their campaigns and target voters who may be decisive for the outcome of elections. We assess the effects of such political “micro-targeting” by exploiting variation in daily advertising prices on Facebook, collected during the course of the 2016 U.S. presidential campaign. We analyze the variation of prices across political ideologies and propose a measure for the intensity of online political campaigns. Combining this measure with information from the ANES electoral survey, we address two fundamental questions: (i) To what extent did political campaigns use social media to micro-target voters? (ii) How large was the effect, if any, on voters who were heavily exposed to campaigning on social media? We find that online political campaigns targeted on users’ gender, geographic location, and political ideology had a significant effect in persuading undecided voters to support Mr Trump, and in persuading Republican supporters to turn out on polling day. Moreover the effect of micro-targeting on Facebook was strongest among users without university or college-level education.

*We thank Greg Crawford, Marko Koethenbueger, James Fenske and seminar participants at Universities of Warwick, ETH Zurich, Bath, Brescia, Dr@w Forum (Warwick), Annual Conference NICEP (Nottingham), PET 18 (Hue), Workshop on Political Economy (University of Bolzano), Political Economy Workshop (University of Warwick).

[†]ETH Zürich, Department of Economics.

[‡]University of Warwick, Department of Economics.

[§]University Carlos III, Department of Telematic Engineering.

“I doubt I would be here if it weren’t for social media, to be honest with you” (President Donald Trump. October 22, 2017)

1 Introduction

This paper investigates the role played by social media in shaping political campaigns and election outcomes. Social media, including Facebook, Google, Twitter and YouTube, have gained prominence as tools for political campaigning in the United States and in several other countries. The 2008 U.S. presidential election was the first in which candidates used social media for direct political campaigning. Since then, the role of social media in the worldwide political scene has grown considerably.

According to the 2016 Pew Research Center (Pew Research Center, 2016a, 2016b and 2016c), about 70 percent of Americans have a Facebook account, and most of them access it on a daily basis. Some 20 percent of Americans have a Twitter account. About one-quarter of these individuals who have social media accounts report that “a lot” of what they see on such platforms is related to politics.¹ Facebook ranked as the third-most-cited “main source” of information for the 2016 U.S. presidential election.²

While Democratic candidate Hillary Clinton’s campaign was run mainly on traditional media, social media, particularly Facebook and Twitter, were among Republican candidate Donald Trump’s primary communication channels. Using data from the the Federal Election Commission, Williams, Girish and Gulati (2017) calculated that during the

¹Figures 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 Internet (social media) as the major source of “most of their news on national political matters”.

²At the end of the 2012 U.S. 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. By then an estimated 79 percent of adults who use online information access (roughly 68 percent of all Americans) used Facebook, and an estimated three-quarters of these Facebook users accessed the site daily.

period from July 1 to November 30 2016, Hillary Clinton’s campaign spent 8 percent of overall media expenditures on digital media, whereas Donald Trump’s campaign spent over 47 percent.³ Moreover, in the same period, according to Bloomberg (2018), Donald Trump’s campaign spent 44 million dollars on Facebook ads compared with 28 million dollars by Hillary Clinton’s campaign. Many practitioners, scholars and journalists have put forward the idea that Facebook and Twitter may have significantly contributed to Donald Trump’s election as US President.

Social media provide politicians with new and more sophisticated channels for reaching voters and targeting their messages across audiences with different political orientations. Targeting on Internet platforms, such as Google and Facebook, is potentially much more precise than on traditional media outlets, thanks to technologies such as behavioral micro-targeting (i.e. the tracing of dynamic behavioral patterns, interests and networks), that exploit extensive quantities of user-generated data. For example, to facilitate the identification of different audiences, in 2016 Facebook began classifying its U.S. users in terms of political orientation (conservative, liberal and moderate) and interests (on specific candidates, issues, or initiatives). As a result, political campaigns are increasingly relying on social media, while comparatively reducing their focus on traditional media outlets.⁴ However, targeting is not a new concept; on television, for example, political campaigns target potential voters by paying for advertisements during shows that attract viewers who lean toward a certain political ideology, or belong to a certain demographic group. The key advantage of social media is that targeting can be far more precise.⁵

As far as we are aware, this paper offers the first attempt in the political and eco-

³According to the authors, Clinton spent 252 million dollars for advertising in total and Trump 176 million dollars.

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

⁵Social media offer also other advantages, with respect to traditional media, which we do not explore in this paper. For example, they allow candidates to observe real time feedback on the course of their political campaign, and adjust strategies accordingly. They also offer to spread campaign messages quickly and at low cost, by exploiting users networks.

conomic literature to measure (i) the extent of the political campaigns conducted via Facebook to micro-target voters, and (ii) the effect that these campaigns have had on the behavior of those voters who relied mainly on social media to gather political information.

To address the first point, we need information on the intensity of political campaigns conducted on Facebook to reach different voters. This is a difficult task, because, unlike traditional media, digital platforms are not required to disclose information regarding the source, amount and content of campaign spending (regulated in the United States by the Federal Election Campaign Act of 1971). At present, Facebook does not disclose information about the content of ads.⁶ Given this lack of information, we conduct our analysis exploiting variations in Facebook advertising prices for different audiences, as defined by locations, political ideology and demographics, during the critical campaign months leading up to the 2016 U.S. presidential election. The idea is that if political micro-targeting was extensively used as a way to reach voters during the campaign, its effect on the online advertising market should be observable through the difference in prices across various political audiences. Specifically, given the size of the audiences (which, as we document below, are largely stable during our observation period), higher competition between the candidates targeting a given audience (a positive demand shift) pushes up the price of advertising to that audience. This, in turn, conveys 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 variations in advertising prices across political audiences as a measure of the *intensity* of Facebook political campaigning.

⁶Facebook recently announced that it will implement measures to increase transparency following 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 posted on its site, making them visible to any Facebook user. In October 2017, U.S. 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.

To this end, we construct a novel database, based on a scraping technology, which gathered daily unit prices for online advertising on the social media platform. Prices are specific to audiences characterized by political ideology, political interest (in a given candidate), gender, age, race, education level, and location (within the United States). The database covers a period of approximately 12 months, from April 2016 until March 2017. The output of this exercise provides audience-level predictions for exposure to political advertising. It is important to stress that the use of the Facebook marketing API (Application Programming Interface) data does not create any privacy risk for Facebook users.⁷

Our second objective is to study the effect of Facebook political advertising on voter behavior. To address this point, we combine information on individual exposure to political campaigns with data from the 2016 American National Election Study (ANES, 2017). This data contains individual responses regarding social media and Facebook usage, as well as voting behaviors, political preferences, demographic and economic characteristics of a sample of 2,414 American voters interviewed before and after the 2016 U.S. presidential elections. Political ideology, demographic characteristics and geographical location allow us to match each user to a specific audience, as defined in our Facebook ad-price data. Intensity to online digital campaigns aimed at the audience to which the user is matched gives us a proxy for the intensity of the online campaign during the period preceding or following the interview date. This makes it possible to compare different respondents in different markets, in terms of their relative exposure to different levels of political campaigning. We focus on

⁷The 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.

the effect of advertising exposure on voter turnout and candidate choice. We are able to identify this effect, due to the fact that users exposure to the social media platform (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 up to the elections depends on the *random* assignment of respondents to different interview dates.

Overall, our results indicate that online political campaigns targeting Facebook users by gender, location and political allegiance had a significant effect on voting behavior. The results suggest that social media campaigns were particularly successful in persuading undecided voters to support Mr Trump and in persuading Republican supporters to turn out on election day. Specifically, we find that targeted Facebook campaigning increased turnout among core Republican voters but not among Democrats or Independent voters. This positive effect on Republican voters' turnout is large in magnitude. Our estimates indicate that exposure to political ads on Facebook adds between five and ten percent to the probability to vote. Furthermore, we find that targeted Facebook campaigning increased the probability that a previously non-aligned voter would vote for Donald Trump, by at least five percent if the voter used Facebook regularly. Similar figures hold for those who do not have a university or college- degree. We also find that this micro-targeting was ineffective for Clinton, failing to boost turnout or to sway voters in her favor. Finally, targeted Facebook campaigning appears to have reduced the probability of a voter changing their mind about which candidate to support among male individuals without college education and supporting the Republican party (although the coefficients are weakly significant). This result provides some support for the hypothesis that exposure to social media strengthens polarization. Finally, a simple test carried out at the end of our analysis suggests that reading political ads on Facebook does not make individuals more politically informed.

The remainder of the paper is organized as follows. The next Session discusses

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. The last part of the paper presents a discussion of the results and concluding remarks.

2 Related Literature

This study contributes to a number of literature streams. First, it relates to the literature that focuses on the effect of different media outlets in channeling the messages of political campaigns to voters. This literature obtains mixed results. Political campaigns on television seem to be effective in persuading voters (for example Lovett (2015); Gerber et al. (2011); Freedman, Franz and Goldstein (2004); Huber and Arceneaux (2007); and Gordon and Hartmann (2016)), but less so in mobilizing voter turnout (Krasno and Green (2008), Huber and Arceneaux (2007), and Gordon and Hartmann (2013)). The introduction of radio broadcasting is associated with an increase in voter turnout in the United States (Strömberg, 2008), while the availability of television (Gentzkow, 2006) and, more recently, high-speed Internet (Falck et al., 2014) is associated with lower voter turnout. DellaVigna and Kaplan (2007) investigate the effect of news bias in media outlets and find that the introduction of Fox News significantly affected the Republican vote shares in the presidential elections held in 1996 and 2000. We contribute to this literature by providing a novel measure of exposure to political targeting on social media and measuring the effect of this targeting on voter turnout and choice of candidate.

There is a literature that focuses on the link between voter targeting and the candidates' key campaign goals. Two main results emerge from this literature: when aiming to persuade voters, politicians should target swing or undecided voters to try to reach those who may switch preferences (Lindbeck and Weibull, 1987); when aiming to mobilize voters to turn out, politicians should target “natural” core supporters to try to reach those who do not

usually vote (Nichter, 2008). There is some empirical evidence for both of these strategies. For example, Hillygus and Shields (2008) find that candidates target weak partisans of the opposing party. Johnston, Hagen and Jamieson (2004) and Shaw (2006) found that in the 2000 and 2004 US presidential elections, both parties concentrated their campaign visits and television ads on the same set of battleground states. Fletcher and Slutsky (2011) show that parties target the media markets that contain the most persuadable voters. Ridout et al. (2012) argue that Democratic and Republican presidential candidates target different genres with their television advertisements. By analyzing the variation in Facebook advertising prices for different audiences and states, we can shed some light on political targeting strategies on social media during the 2016 U.S. presidential campaign. Recently, the literature has devoted increasing attention to how campaigning adapts to the diffusion of the Internet and social media. Bond et al. (2012) estimate that about 340,000 extra people turned out to vote in the 2010 U.S. congressional elections because of a single Facebook political mobilization message. Their results suggest that micro-targeting is an effective way to reach voters. Allcott and Gentzkow (2017) show that social media was an important source of information during the 2016 US presidential elections campaign and that most American adults were exposed to at least one piece of fake information on the Internet. Furthermore, individuals with ideologically segregated social media networks are more likely to believe fake stories.

Several authors, including Sunstein (2001; 2009; 2017), Pariser (2011), and Gabler (2016), claim that the rise of the Internet and social media has contributed to voters' political polarization. They contend 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. Sunshine, Hillygus and Shields (2008) show that candidates are more likely to campaign on wedge issues (such as immigration, abortion, and religion) when the forum is not public. However, relying on evidence from demographics, Boxell et al. (2017) cast some doubt on the hypothesis that

the Internet and social media cause polarization. In a similar vein, Barberá (2015) uses data from Twitter to argue that social media actually reduce polarization, facilitating exposure to information from individuals with whom users have weak ties and are thus more likely to have contrasting political views.

With very few exceptions (cited above), previous literature relies on indirect evidence of the effects of social media and microtargeted ads on voter behavior.⁸ This is most likely due to the lack of data regarding these media outlets. We circumvent this problem, by collecting the data on Facebook advertising prices, and building a metric for the intensity of the political campaign conducted on Facebook. Furthermore, we contribute to literature by linking Facebook advertising data to individual voting behavior and personal characteristics. Comparing Facebook users and non-users’ political behavior, we are able to measure the effect of the online campaign on different audiences.

Finally, our study belongs to a novel multidisciplinary literature that leverages Facebook Marketing API data to address socio-economic issues. González Cabañas et al. (2018) quantify how many Facebook users in the European Union may be labeled with “sensitive” interests (according to the definition provided in the General Data Protection Regulation). They show that 73 percent of EU Facebook users may be labeled with such sensitive interests.⁹ Zagheni et al. (2017) and Dubois et al. (2018) have used Facebook marketing API data to study migration stocks and assimilation, collecting the number of expats of different nationalities living in a particular country. Garcia et al. (2018) use Facebook marketing data to perform a comprehensive analysis of the gender divide world-

⁸A few studies focus on political candidates’ use of Twitter (see Gainus and Wagner, (2014), and Bright et al. (2017). Petrova et al. (2017) track candidates’ entry on Twitter during the latest U.S. congressional elections, and find that opening a Twitter account helps new candidates increase their donations by up to 3 percent.

⁹In earlier studies, Liu et al. (2014) perform a quantitative analysis of the auction system in Facebook demonstrating 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 FB users’ activity (e.g., likes, shares) is propagated to friends together with a second simplistic model that guesses the number of ad impressions received per user.

wide.

3 Online Advertising Prices and Political Campaign Intensity

During electoral campaigns, political advertising takes over available airtime and advertising space in traditional media, particularly in countries with low regulation, such as the United States.¹⁰ During this time, non-political advertisers find the market to be highly saturated, and experience considerable increase in ad prices. Price fluctuations are even stronger in the case of the Internet and social media, where no regulation applies. During campaign season, candidates compete with each other to attract the attention of politically relevant audiences, (see Figure 1). They do so by entering the digital auctions held by Internet platforms (details on the functioning of these auctions are discussed in the Appendix, sections A.1 and A.2). For these reasons, we argue that observing online ads price fluctuations provides a novel tool to proxy for the intensity of political campaigning. In fact, we expect that the willingness to pay for ads targeted at a given audience will be, for any political candidate, proportional to the probability of such ads in convincing members of that audience to vote for that candidate. In turn, the value of convincing a selected audience, depends on the perceived effect on the election by gaining support from voters in that audience will have on winning the election (Moshary, 2017).

¹⁰The United States has a very liberal market for broadcasting political messages on television and radio media. By contrast, many other countries restrict the use of media to broadcast advertising with political content. For example, paid television and radio political advertising is generally not allowed in member countries of the European Union. The United Kingdom and Ireland also forbid paid television and radio from airing political advertisements, but allow political parties a limited and regulated number of political broadcasts in the period immediately preceding the elections. Canada allows political broadcasts, but strictly regulates airways access.

3.1 Theoretical Background

The mechanism mentioned above can be illustrated through a simple model. Consider a given audience, defined by a set of characteristics (e.g., income, age, gender, and political preferences). Assume consumers within the audience are identical. We assume that consumers are willing to tolerate a limited number of ads while browsing the digital platform. Let \bar{q} be the reservation level of advertising that consumers are willing to be exposed to. This level is the highest quantity of ads such that users do not lose interest in the platform. Assume that \bar{q} is known to the platform, and that there is no marginal cost of providing ads. Given these conditions, the quantity of ads the platform exposes each consumer to is inelastic and equal to \bar{q} ¹¹ (see Figure 2).

We consider two states of the world, indexed by $i = \{0, E\}$. State 0 is such that there are no elections and, thus, no political candidates (C) aiming to reach voters on the platform. In this state, there is only a generic, non-political, advertiser G , whose demand for ads (per consumer) is $p_G(q)$.¹² In this state, the equilibrium (denoted with subscript 0) is such that $q_G^0 = \bar{q}$. State E , instead, captures a situation when elections approach and political candidates look to reach voters. Hence, in addition to the demand from the generic advertiser (that we assume is the same as in the no-elections state), there is a demand for ads by political candidates, denoted $p_C(q)$. We represent this additional demand in the right panel of Figure 2, whereas the left panel depicts the aggregate demand for ads. The equilibrium (subscript $i = E$) is such that the price of ads, p_E , is higher than in state 0 (we assume the consumers' tolerance for ads does not change between the two states). Furthermore, although the total quantity of ads (per consumer) remains equal to \bar{q} , ads that would otherwise be allocated to the generic advertiser are allocated to the political

¹¹We assume demand by advertisers is large enough to saturate the reservation quantity of ads.

¹²We assume that $p'_G(q) < 0$, as the marginal value of an impression on a consumer is decreasing. For 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, Calvano and Reisinger (2016).

campaigns. In Figure 2, we also consider a more relevant increase in demand from political candidates. As the figure suggests, the stronger the demand increase, the higher is the price in equilibrium and, importantly, the larger the quantity of ads allocated to political campaigns (superscript +). Therefore, a higher price for advertising in proximity of elections is also indicative of higher exposure to political ads. Note that, in the empirical application, we consider several different audiences that could differ in their tolerance to ads. Although we do not observe this tolerance (nor the total quantity of ads to which consumers are exposed), we shall control for many observable factors (income, age, gender, location, education, etc.) that are reasonably correlated with ad tolerance.

Insert Figure 2 in about here

3.2 The Construction of a New Measure for Intensity of Political Campaigns

Motivated by the above discussion, we collect data on online ad prices for the largest social media network, Facebook. We then construct price indexes that proxy for the intensity of (Facebook) online political campaigns. By campaign intensity, we mean the effort made by candidates to influence an audience; this effort is directly reflected in the strength of the bidding competition to target that audience. To show that our online ad price indexes effectively capture the intensity of political campaigns, we provide graphical evidence of non-negligible cyclical price increases around the time of democratic elections, at a worldwide level. Finally, we show that during the 2016 U.S. presidential elections, changes in daily state-level price fluctuations were significantly correlated to changes in the candidates winning probabilities, as estimated using state-level election polls. We emphasize that, in our empirical application, we can neither observe the total volume of political ads nor track the number of ads visualized by a specific user. Hence, we cannot directly measure campaign in-

tensity. However, for the reasons illustrated above, we know that an increase in competition among bidders will affect prices for targeting a given audience. Thus, we base our measure for intensity of Facebook political campaign on online ad prices.

Advertisers on Facebook compete for ad slots to target an audience. Facebook offers multiple layers of targeting, so that the advertiser can direct their ads to their preferred audiences. (Appendix A.1 gives more detail on the audience-selection procedures.) Once an audience is selected, three metrics determine the price paid to the platform: cost-per-click (CPC), cost-per-action (CPA) and cost-per-mile (CPM).¹³ Advertisers can pick the most appropriate one, depending on their objective and business model. For conversion-oriented objectives, such as subscriptions or retail sales, advertisers usually pay per click (CPC) or per action (CPA), whereas for exposure objectives, such as political campaigns, advertisers typically choose to pay for impressions (CPM). Once the bidding price is placed, Facebook allocates ad spaces according to a mechanism based on the Vickrey-Clarke-Groves auction (described in Appendix A.2).

To build our online advertising dataset, we scrape daily prices for U.S. state- and ideology- specific audiences, during the three months preceding the 2016 U.S. presidential election. We extract both CPM and CPC. We further condition on age bracket, gender, race, or education.¹⁴ We use these data to measure the intensity of the online political campaign during the election campaign, and its effects on voting outcome. There are few groups in the area of Computer Science who have leveraged Facebook Marketing API data (in a completely privacy-preserving manner) to address socio-economic problems, the main contributions to

¹³Cost-per-click (CPC) is the rate that websites charge advertisers every time someone clicks on an ad. Cost-per-action (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 installs, etc.). Cost-per-mile (CPM) is a marketing term used to denote the price of 1,000 advertisement impressions on one webpage. If a website publisher charges \$2.00 CPM, that means an advertiser must pay \$2.00 for every 1,000 impressions of its ad. The “M” in CPM represents the Roman numeral for 1,000.

¹⁴This 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.

the literature have been discussed in the previous session. To use our data, we construct a price index to account for a number of issues. First, online advertising prices are affected by audience size and specificity. Keeping the number of advertisers constant, one can expect the price of targeting an audience to be decreasing in the audience size (because of higher supply), and increasing in its specificity (because the more narrowly defined the audience, the more likely users in the audience are to match the advertiser’s ideal target). Since the size of the audience matters for the bidding price, it is crucial to ensure that the number of users matching a certain audience does not vary over our observational period. Second, we need to account for the fact that a political profile can be correlated with certain consumer preferences, and that non-political advertisers might participate bidding for the audience with that same political profile for reasons unrelated to the electoral campaign. Notably, recent research shows that digital footprints, such as Facebook “likes” are strongly correlated with personal attributes and preferences (Kosinski et al., PNAS 2013); as a result, political ideology can be exploited to identify individuals who are more likely to be interested in certain goods and services. This limits comparability of price fluctuations across different audiences. Finally, there is an issue of relevance of a given political profile in different contexts, and at different points in time. Facebook users who fit a certain political profile could be key, when located in certain geographical area, but not in others; think, for instance, of the relevance of an ideologically “moderate” user in a swing state compared to the same user in a safe partisan state. A user’s relevance might be affected by exit polls, or by exogenous informational shocks (e.g., an event that shifts public opinion at large, such as a scandal).

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

Relative CPM for audience defined by State (s) Characteristics (c) Political Ideology (i) :

$$p_{sci} = \frac{P_{sci}}{P_{sc}} \quad (1)$$

The intuition behind this measure is that geographic location and demographic characteristics can be used by all generic advertisers (G) to proxy for users’ preferences, whereas ideology matters only to political advertisers, the candidates (C), who know what subset of the population they need to convince in order to secure an electoral win. In other words, conditional on a geographic location and a given set of demographic characteristics (e.g., women above 40 years of age, with a college degree, residing in New York), only political candidates should be willing to pay the premium required to target the subset of users also identified as having a certain political ideology.

Because political ideology is identified by Facebook through the user’s revealed preferences (i.e.. profile characteristics, such as likes, education, profession, relationship status, favorite movies, books, sports, places, etc.) and network, the political subsample composition is rather time-invariant.¹⁵ This ensures that observed variation in our relative price measure is not driven by variation in the number of users with a given political ideology (see Figure 3), which rules out endogeneity of political ideology to the electoral campaign (at least in the short term). We naturally observe that swing states have the majority of users classified as moderate, whereas the safe Democrat and Republican states have the majority of users classified as liberal and conservative, respectively. But we find no variation in the number of users in each category over time.

To measure the way the campaign affects changes in respondents’ voting behavior,

¹⁵The 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, commenting, liking, reading, watching posts.

we also build an additional index, which captures the CPM price run-up over a given time interval.

Relative CPM Run-up for audience defined by State (s) Characteristics (c) Political Ideology (i), between Election Week (e) and Interview Week (w):

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

The advantage of using this index is that, when assigned to a respondent of the ANES political survey, it exploits the random assignment of the respondents to the date of the interview, and the subsequent variation in exposure to the social media political campaign. So, despite being unable to observe how many ads each respondent has visualized over the time she spent on Facebook during the campaign, we can exploit the randomness of the duration between two dates used to compare differences in the responses to voting questions.

3.2.1 Discussion and Some Illustrative Examples

Focusing on the U.S. market during the 2016 presidential campaign allows us to collect Facebook ad prices that are specific to ideologically defined audiences. Exploring the variation of these prices across audiences located in states that are more or less strategic for a given candidate allows us to investigate the intensity of social media use for a specific campaign strategy. We illustrate the intuition behind our approach with a visual example. Figure 5 shows average weekly relative CPM prices for Facebook online ads, targeted at audiences identified by a geographic location and a political ideology, in the weeks preceding the 2016

elections. We use relative CPM prices, to filter out state-level effects, which may possibly be driven by confounding economic factors. In the example, the geographic location is selected to be one of three states with rather consistent political alignment: New York, a Democratic state; Florida, a swing state; and Texas, a Republican state.¹⁶ The figure suggests that moderate voters (those with weak or unformed preference for a specific candidate) have been aggressively targeted, and particularly so in swing states, rather than in partisan states. At the same time, liberal voters are the most expensive to target in New York, and the least expensive to target in Texas. This picture seems to tell a story that is very much aligned with political economy theory: undecided voters in swing states are the most valuable voters, because obtaining their support can buy a candidate the electoral win; and, at the same time, candidates focus on getting out the vote of their own supporters in partisan states, to secure the expected results. By contrast, in states where the opposition party dominates, candidates do not invest in attempts to flip support of moderate and partisan voters aligned with the competing candidate, and they do not invest in get-out-the-vote campaigns.

3.3 Supporting Evidence on the link between Advertising Prices and Political Contestability

We conduct two empirical exercises to show that our approach provides robust indicators for the intensity of political campaigning through online advertising. The first exercise shows that online ad prices follow the electoral cycle. The second exercise shows that there is a significant correlation between fluctuations in online ad prices and changes in the winning probabilities of political candidates.

¹⁶The political alignment of these states has remained constant since at least 1988 (see State Historical Presidential Election Information from “270 to Win”, <https://www.270towin.com/states/>). Following California, which has the largest population of any state, Texas, Florida and New York are the most populous states in the country,). These states also lists among the four most populated of the US.

3.3.1 Online Ads Prices and Electoral Cycles around the World

For the first exercise we take a step back from the detailed analysis of the 2016 U.S. presidential election. We build a worldwide dataset, to show that our approach is not circumstantial only to the U.S. context, but rather holds for all electoral cycles. We scrape country-specific daily prices for a large set of 213 countries, during an interval of one-and-a-half years (from July 2015 until February 2017). To identify the subset of democratic countries, where information flow is not controlled or biased by authoritative governments, we refer to the ranking provided by the Electoral Democracy Rating published by Freedom House. We select countries defined as “Free” or “Partly Free” that held democratic elections during the past five years. This leaves us with a sample of 45 fair elections held during the time spanned by our data.¹⁷ 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 in Figure 4, average worldwide prices sharply rose during the last two months before the elections (an average increase of around 50 USD cents, which corresponds to about a 25 percent increase in prices), and steadily dropped over the three months after the election date. After this five-month-long window, prices seem to have returned to their equilibrium level. This is in line with the assumption that new advertisers (the political candidates, C) enter the market with very aggressive campaigns immediately before the elections, and then leave the market in the subsequent period. The fact that

¹⁷The 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).

prices drop more slowly than they rise is also consistent with the media frenzy typically surrounding the first trimester of any newly elected government.

3.3.2 Candidates' Popularity and Ads Prices during the 2016 US Presidential Campaign

Having established that Facebook online ads do respond to the electoral cycle of democratic countries, we proceed with a second exercise: to investigate whether fluctuations observed in the prices of these ads can be used to measure changes in the political campaigns conducted on the social media site. To this purpose, we return to the U.S. market alone. The reason for this restriction is that U.S. users are the only ones for whom Facebook has started, since 2014, to classify political preferences. We look at the correlation between electoral winning probabilities and online advertising prices. We base our test on the assumption that candidates would adjust the intensity of their campaigns in order to respond to changes in their estimated probability of winning the election. We collect daily state-specific winning probabilities for both presidential candidates, during the three months preceding the election (from the 8th of August, until the 8th of November 2016). Our source is the FiveThirtyEight 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.¹⁸ 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-

¹⁸The website also computes winning probabilities simply based on local state-level polls. We repeated our test using also this series, and found similar results.

the-week dummies) and for a time trend (week-of- the-year dummies).¹⁹ The results from this exercise show that, at the state level, daily Facebook ad prices are significantly affected by changes in the winning probabilities of the two candidates. In fact, as shown in Table 1, CPM prices drop, as political competition between the two candidates falls, due to an increase of Trump’s chances of winning the presidential election.

Insert Table 1 in about here

4 The Effect of Facebook on Political Behavior

The previous discussion shows that Facebook ad prices can be used to measure the intensity of political campaigns. In fact, our data provides evidence that the social media platform was largely employed during the 2016 US Presidential election. However, the question remains of whether social media, in general, can be successfully used to affect voting behavior. Our paper proceeds in this direction and combines online political campaigning data with individual level information on turnout and voting outcomes. In studying the effectiveness of online political advertising, we focus on the two recognized purposes of political campaigns: mobilization and persuasion of voters.

A fundamental point is that analyzing the effect of political campaigns requires to integrate our measure of social media campaign *intensity* with information on voters’ *exposure* to this novel media channel. This is because the overall effect of the campaign depends on the “volume” of political messages directed to a specific audience, but also on how much individuals of that audience are open to receive those information. To find a

¹⁹Our methodological approach draws from a literature on the effects of electoral forecasts and outcomes on financial markets and stock prices. See Herron (2000) and Knight (2006) for the effect of winning probabilities on stock returns, and Santa-Clara and Valkanov (2003) and Sattler (2013) for the effect of electoral outcomes. More recent studies are Girardi and Bowles (2018) and Wagner et al. (2017).

measure of *exposure* to Facebook advertising, we collect information on individuals’ media consumption. These are provided by the 2016 American National Election Survey (ANES), among numerous other respondent characteristics. Our strategy consists in exploiting ideological preferences, demographic characteristics and geographical location, to match each ANES respondent to a Facebook online advertising audience. Survey questions on whether the respondent uses social media networks, Facebook specifically, and on how frequently she accesses political news on the platform are used to build a simple indicator of exposure. Doing so, allows us to obtain the composite individual-level measure of *treatment* to the political campaign, by interacting our audience specific proxy for *intensity* with individual specific measure of Facebook *exposure*. Note that exposure is an individual level characteristic, possibly dependent on the respondent exogenous preferences on media consumption. Intensity, on the other hand, is an audience level characteristic, exogenous, because dependent on the random assignment of survey respondents to their respective interview date. The 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 to the campaign. The next section describes in details the methodology we employ.

4.1 Matching Facebook Audiences to Survey Respondents

The 2016 American National Election Study (ANES, 2017) provides individual responses regarding social media and Facebook usage, as well as voting behaviors, political preferences, demographic and economic characteristics for a sample of 2,414 American voters interviewed before and after the 2016 US Presidential elections. Pre-election interviews are randomly conducted during the two months prior to the 2016 elections and are followed by post-election re-interviewing beginning November 9, 2016. Political ideology, demographic characteristics

and geographical location, allows us to match each user to a specific audience, as defined in our Facebook ads-price database. Date of interview, respondent’s state of residence, political ideology, gender, education and age allow us to assign to each respondent the ad price a political advertiser would have paid to target an audience made of subjects with the same profile. The overall effect to online digital campaign for that respondent will later be derived by the *intensity* of the online campaign during the period preceding or following the interview date, conditional on her individual *exposure* to Facebook platform as a mean to access campaign information. This makes it possible to compare different respondents in different markets in terms of their relative exposure to different levels of political campaigning. We focus on the effect of advertising *exposure* on voter turnout and candidate choice.

Table 2 provides an example of difference in intensity of the campaigns targeted to audiences specified by geographic location, ideology and gender. Comparing prices at the interview date with prices at the election date, we see that the most intensive campaign is always targeted at Moderate users. We also observe that campaign intensity was generally higher for republican than for democrat voters. Finally, we find that when it comes to liberal respondents, women were more intensively targeted than men, while the opposite holds for conservative respondents.

Insert Table 2 in about here

From the inspection of Table 3, instead, we learn that of the 2,414 ANES respondents for which we have information, nearly 30 percent of them does not use Facebook (note that the site penetration rate for the US population was equal to 88.5% in 2016); of those who have an 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 they appear to lean slightly more towards 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" . We assume that the choice of having a Facebook account is exogenous to the choice of reading political news on the social media (see Figure3, which shows that Facebook daily active users were stable between the 1st of September and the 1st of December 2016, and that this trend was constant across gender and ideology).

4.2 Empirical Strategy and Main Results

Our key explanatory variables combine measures of *intensity* and of *exposure* to Facebook political campaigns, allowing for their effect to interact for all respondents included in the 2016 ANES sample. We propose two metrics for intensity: the Relative CPM, (p_{sci}), (*INT-1*) and the Relative CPM ratio ($r_{sci,w}$) (*INT-2*), both measures have been discussed in Section 3.3 and are based on average CPM prices recorded in the week prior to the interview (p_{sci}) , and on prices recorded in the week prior to the election ($r_{sci,w}$). Using a period of one week is motivated by studies suggesting that individuals typically forget messages within a matter of days (e.g., Gerber et al. 2011a; Hill et al. 2012; Patterson and McClure 1977; Sears and Kosterman 1994). These measures differ by State, date of the interview and audience type. Both indexes capture the intensity of the campaign but the second measure also takes into account the trend of the campaign up to the election day; for example p_{scp} measures the ratio between the average CPM of Republican Women in Texas and the average CPM of all Women in Texas, while r_{scp} measures the run-up in p_{scp} recorded between the election week and the week before the first interview. So $p_{scp,w}$ captures campaign intensity, and r_{scp} captures both intensity and duration.

Individuals exposure to political information spread through Facebook, instead,

depends on the frequency at which respondents use the social media platform. To construct this measure, EXP , we combine the answers to two questions asked in the 2016 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 equal to zero when the respondent does not have a Facebook account or uses it only sporadically (up to twice a week) and one if heavily relies on the social media for learning about the campaign.²⁰ The interaction between EXP and INT provides a personalized measure of the overall individual level intensity of the political campaign on Facebook. A *caveat* here, our measure does not provide information on the content of the advert or on the identity of the advertiser. Liberini et al (2018) look at this issue with an online experiment. The baseline model we estimate for respondent i matching audience a , who was interviewed at time t , is as follows:

$$Y_{iat} = \alpha EXP_i + \beta INT_{at} + \gamma EXP_i \times INT_{at} + \delta X_i + D_a + \epsilon_{iat} \quad (3)$$

We estimate this model for a number of different outcome variables, Y , measuring voting behavior and participation. We consider, in turn, voter turnout, candidate choice, and between pre-and-post elections interviews' preference switch. Moreover all the regressions include an additional set of controls, X , to take into account factors influencing political behavior and preferences, and a set of ideology dummies D_a , which account for time invariant differences among the ideological audiences. The control variables are: gender, religion, marital status, level of education, race, number of children income, home and gun ownership, employment status, political orientation, intentions and participation in previous elections as well as day of the week and month of the year of the pre-election interview. We also include information on exposure to other media: television, press and Internet. We estimate

²⁰We also experiment with 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 the values from 1 to seven depending on the number of days a week the account is used to collect political information, the results are very similar to those presented in the paper and can be seen upon request.

this model using LPM.²¹

One crucial contribution of our empirical analysis is that we are able to capture different targeting dimensions of the campaign launched on social media. Our previous discussion stressed how sophisticated micro-targeting represents the true advantage of conducting political campaigns on online platforms versus traditional media. When estimating the model in equation (3), we interchangeably use intensity measures based on targeting of different types. In fact, INT_{at} entirely depends on the definition of the audience, a . Specifically, we always condition an audience on its geographic location (the state) and political ideology. In addition, then, we allow the audience to also vary for one of these demographic characteristics: gender, educational level, age and race.²² Comparing the performance of the same model across intensity measures allow us to say something about the effectiveness of targeting along these different dimensions. In fact, our first result is to show that campaign intensity based on race and age targeting does not seem to have significant effect on any of our outcome variables. On the contrary, campaign intensity based on gender or educational level seem to be effective, to different extents, with all our outcome variables.

4.2.1 The Effect on Voter Turnout

One of the main questions the political economic literature has addressed is whether advertising helps persuade potential voters to go to the ballot. The results are mixed and have been discussed in Section 2. As far as we are aware, our methodology contributes to this stream of literature by providing the first evidence on how micro-targeted political advertising on social media affects the likelihood to cast the ballot. We estimate variations of equation(3), where the dependent variable is a dummy taking value of one if the respondent has voted in the general election and zero otherwise. The full set of controls listed in the previous section

²¹Probit regressions produce nearly identical to OLS average marginal effect. Outputs are available upon request.

²²For identification purposes, we only consider one of these dimension at a time.

are included in all our specifications²³. The results for this exercise are displayed in Table 4.

Insert Table 4 in about here

The table is divided in two panels; to proxy for the intensity of the campaign, we employ the relative CPM ($INT - 1$), in the top panel, and the relative CPM run-up ($INT - 2$), in the second panel. In each column, we estimate the model using the CPM prices for different targeting dimensions: by political ideology only, and by political ideology combined, in turn, with gender, age, race and education. The overall result we observe from this table is that turnout is negatively affected by political campaigns targeted at a combination of political ideology and gender. Specifically, from the second column and first panel of the table, we find a negative effect on turnout of exposure and intensity, when taken separately, but a positive cumulative effect of the two, when interacted. This suggests that individuals who access the social media on a daily basis are less likely to go to vote, compared to those who do not have a Facebook account. Yet, the exposure effect is less negative, the higher the intensity of the political campaign targeted at these users.

Insert Table 5 in about here

We now focus on the effect of gender-based micro-targeting on vote turnout (column 2, Table 4). To give a more precise indication of the magnitude of the effects of intensity and exposure to the social media political campaigns, we resort to a number of graphical illustrations. Figure 9 shows the joint effect of accessing Facebook and being exposed to campaigns of different intensity. The figure is a contour plot, with campaign intensity on the horizontal axes and campaign exposure (the number of days a week the respondent reads about politics on Facebook) on the vertical axes. Different colors are associated with

²³These are not reported in the tables and are available upon request

different levels of the predicted probability of turnout, from blue (below 60 percent) to red (close to 100 percent).

Insert Figure 9 in about here

The graph simply shows that for low and medium Facebook exposure, turnout drops as campaign intensity increases. It seems, instead, that there is no significant effect of campaign intensity on turnout for those people who access Facebook on a daily basis. These are average effects, for the overall sample of respondents. Because we are interested in the effect of micro-targeting, however, we expect these effects to be only relevant for the subgroups of respondents that are the aim of the political campaign. If both candidates compete for a State whose winning probability depends on turnout of men, for example, we expect to find a significant effect of intensity on men, but not on women. Because the effective targeting is based on gender and ideology, we re-estimate the model of column 2, table 4, for subsets defined by gender, past vote outcome and current ideology. Results are displayed in table 5, and show that the targeted campaign was effective for male (column 3), but not for female (column 2) respondents, and it was effective for republican (column 8), but not for liberal or neutral voters (columns 6 and 7, respectively). We also look at whether ideological target might be based on past electoral preferences (i.e. voting outcomes in 2012), but find this not to be the case (see columns 4 and 5 of table 5)²⁴. For a quantitative interpretation of these results, we look at the effects of campaign exposure and campaign intensity across gender and ideology separately, in figures 10 and 11 respectively. The top part of each figure shows trends in the predicted probability of turnout, for respondents who rarely use Facebook as a media tool to access political news and for those who, instead, do so on a regular basis. In the bottom part we test whether the difference in these predicted probabilities, conditional on different levels of campaign intensity, is significantly different from zero. Figure 10 tells us

²⁴We find that subgrouping respondents over their past voting preferences never yield significant targeting effects, also for the other outcome variables considered in the analysis. These additional results are omitted from the remainder of the paper, but are available upon request.

that male respondents are less likely to turn out for elections, the more intense the political campaign to get out their vote. However, this effect is much stronger for respondents who rarely use Facebook, compared to those who use it on a regular basis. In fact, the bottom plot tells us that for medium levels of campaign intensity ($INTEN1 = 1.1$), the mere fact of using Facebook on a regular basis increases the predicted probability of turnout by 5 percentage point. These results do not hold for female respondents. Next, we look at the differential effect of micro-targeting across subgroups defined by different political ideologies. Consistently to results in column 6 of Table 5, we see in the top plot of Figure 10 that higher campaign intensity levels correspond to lower turnout rates for all respondents, except for Republican voters who access Facebook on a regular basis. For these, at high levels of campaign intensity, the difference between using Facebook rarely or on a regular basis can result in between 5 and 10 percentage points difference in turnout rates. Note that our results are consistent with the idea that Trump invested more in campaigning on social media and, presumably, targeted his core supporters, i.e. men and conservatives. It is also interesting to find that exposure to social media, *per se*, has a negative effect on turnout, which is in line with findings by Falck et al. (2014).

4.2.2 Choice of the candidates

The next question we address is whether the Facebook campaign is more effective in moving voters toward a specific candidate. Table 6 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 8 repeats the analysis on the dependent variable collecting Clinton votes. These tables have the same format of Table 4. We depart from our results on the vote for the Republican candidate.

Insert Tables (6) and (7) in about here

First, we are interested in understanding whether micro-targeting is equally successful, when based on different personal characteristics. From an inspection of the table, we immediately see that gender and educational level are two dimensions over which targeting seem to be significantly effective, but that the same does not hold for targeting based on race and age. We also find that the effects are significant when using the CPM ratio, which accounts for both intensity and duration of the political campaign (bottom panel of table 6). Comparing columns (2) and (5), bottom panel, we find that the direction of the effect is the same across targeting dimensions, but the magnitude is higher for targeting based on ideology and gender, versus ideology and educational level. Figure 12 allows us to interpret the results from column (2). The figure is, again, a contour plot, showing the joint marginal effect of campaign exposure (on the vertical axis) and intensity (on the horizontal axis). Individuals who have no access to Facebook are decreasingly likely to cast their vote in favor of the Republican candidate, as the campaign intensity targeted to an audience with their political ideology and gender intensifies. Their predicted probability of supporting Trump goes from 0.5 to 0.15. On the other hand, this effect reverts as individuals become more exposed to Facebook content, because they access the social media more often to read about political news. Individuals who access Facebook on a daily basis are, in fact, more likely to vote for the Republican candidate when the campaign intensifies, as their predicted probability of supporting Trump goes from 0.3 to 0.46.

Insert Figure (12) in about here

To test whether these results hold for all targeted subgroups we proceed with the same approach used to analyze the effect on turnout. In table 7, columns (2) and (3), we re-estimate the model with gender and ideology targeting for the subgroups of female and male voters, respectively. In columns (4) to (6) we look at neutral, republican and democrat voters. Finally, in columns (7) and (8), we re-estimate the model for education and ideology targeting (table 6, column (5)), this time for the subgroups of voters with high

or low educational levels²⁵ .

When it comes to targeting based on the combination of gender and ideology, we find large and significant effects for the group of neutral (uninterested and uninformed) voters (column (4)). The effects on the other ideologically identified groups are not significant. We find large effects also for the subgroup of male voters, however these are weakly significant. Figures 13 and 14 help us interpret the effects of the model's interaction terms for these alternative subgroups of respondents. Figure 13 shows that a common trend is that respondents targeted by highly intensive campaigns are less likely to cast their vote in favour of the Republican candidate. Male respondents who use Facebook on a regular basis are the only subgroup for which this trend is reverted, but the difference with the effect for male respondents who never use Facebook is not statistically significant. The bottom right plot shows that there is some significant difference between the support for Trump among women who use Facebook daily and those who don't. For medium level of campaign intensity, this difference being approximately equal to 2 percentage points.

From figure 14, top graph, we see that, while partisan voters are hardly affected by exposure to the social media, neutral voters who access Facebook on a daily basis are up to twice more likely to support the Republican candidates than neutral voters who do not have a Facebook account. The bottom graph shows that for high levels of campaign intensity, neutral voters who are exposed to Facebook on a daily basis are up to 5 percentage points more likely to cast their vote in favor of Trump. A possible explanation could be that individuals who rarely use Facebook are exposed to alternative forms of campaign, which are possibly effective in convincing voters against choosing Trump. Online micro-targeted politics do very little for partisan voters.

Insert Figures ?? and ?? in about here

²⁵These are defined by the highest educational level attained at the time of the interview. Respondents are divided between those who have at least a college degree (high) and those who have attained only lower educational levels (low)

Finally, we turn to the effect of targeting based on education and ideology. Columns (7) and (8) from table 7 tell us that there is a significant effect on subgroups of voters with low- and high- educational levels. These effects have the same sign, but they are much stronger for the former, than for the latter group. Figure 15 graphically represents these results. The top graph shows that for voters who rarely access Facebook the predicted probability of supporting Trump is orthogonal to the educational level. At the same time, this probability displays a mirroring trend, for respondents with low education who access Facebook on a daily basis. The bottom graph indicates that, for respondents who achieved less than a college degree, accessing Facebook on a regular basis, versus not using it as a tool to gather political news, can explain up to ten percentage points in the support for the Republican candidate.

These results are consistent with the idea that Trump concentrated his targeting effort both on voters who are male, and who are not particularly interested in politics. both categories were naturally key to the candidate’s support base, particularly so for the neutral voters, who were by nature more likely to swing their vote in favor of a persuading politician. There might have also been some targeting based on voters educational level, and it is again not surprising that the low educated voters are the more responsive to the type of political advertising typically displayed on social media. This is also consistent with previous findings showing that the effects of advertising are conditional on having differential impact 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 casted 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 8, this is true for all types of targeting, and holds for subgroups of voters (table 9). Figure 16 graphically confirms these findings. The only positive significant effect, as shown

in figure 17, is that on democrat voters. We find that the average predicted probability of casting a vote for the Democrat candidate increases with campaign intensity for the democrat voters who do not have access to the social media platform, and decreases for those who daily gather political news on it. This effect is not very significant, however it is interesting to see that it mirrors what we had found for Trump.

4.3 Political Behavior Changes: Political Polarization vs Information

As highlighted in Section 2, recent studies have linked the rise of political polarization to the growing use of social media as a vector to convey political messages. It is debated that political micro-targeting may have substantially contributed to this phenomenon. Contrary to traditional media outlets (such as TV, radio, printed press), Facebook is not subject to any regulation on advertising for political campaigns. So targeted voters may be discriminated, or simply be unaware of being the target of political advertising, which comes presented as legitimate news and, as a consequence, may give distorted, biased, or even fake, information. As a possible consequence voters may be "trapped" in filter bubbles, receive less varied content and ultimately lose the ability to critically process information. This might, in turn, make them more unlikely to change their initial voting choices, and possibly make them more sensitive to extreme ideologies (*echo-chambers*).

To investigate this hypothesis, we would need to observe advertising content, and gradual changes in respondents political alignment. These information are unobserved (to the econometrician), and can hardly be approximated. Nevertheless, with our data we can make a first step toward investigating this hypothesis. We can simply analyze and compare Facebook users and non-users' change in voting strategies that may take place during the window between the pre-election interview and the actual vote. If the hypothesis is that

frequent exposure to political news delivered on social media, increases polarization (due to the effect of echo chambers), then we should find that Facebook (regular) users who are highly targeted by political campaigns are less likely to revise their initial voting choices during the electoral campaign, than similar voters who do not use their Facebook account. The time between the interview date and the election identifies a (randomly assigned) window during which candidates are exposed to the “treatment” of online political advertising.

We calculate the probability of changing political preferences by estimating equation(3), where our dependent variable is a zero-one-dummy calculated on the basis of the difference between respondent’s voting intention and actual vote in the 2016 election.

Insert Table (10) in about here

Table 10) shows that, again, targeting down the lines of gender and education, jointly with ideology, is an effective way to anchor voters to their initial voting choice. If we look at the results displayed both in columns two and five of the table we can see that those voters, who are highly targeted on the basis not only of their political preferences but also on their gender and education and consistently rely on Facebook for information, are less likely to change their mind compared to those who are not. Also note that reading the news on Facebook *per se* , as well as being a political target *per se* both marginally affect vote change in a positive way. We proceed our analysis by plotting the effect of gender-based micro-targeting on vote change.

Insert Figure (18) in about here

The inspection of the plots in Figure 18 confirms that there is a strong difference between the behavior of Facebook users and non-users, when they are treated with high level of campaign intensity. When the fight on the political campaign battle field is more intense, an increase

in exposure is associated with a decrease in the probability of changing the initial political view.

Insert Table (11), in about here

Moving to the effect of micro-targeting on specific audiences we re-estimate equation(3) differentiating across genders, political interests and level of educations. Interesting Table (11) shows that there is a divide in the effects of vote change across groups in term of the gender, ideology and education.

Insert Figure 19, in about here

We explore this further with the help of Figure19 and Figure20; motivated by results displayed in Table (11) we define the audiences by gender and ideology. Figure19 looks at gender differences, from the inspection of the top panel of Figure19 it emerges that men relying on Facebook for information, when are treated by low intensity campaign (i.e. for level of intensity around one and below) are about five percent likely to switch their initial voting choices but, when subject to high campaign intensity, they do almost entirely stuck to their initial choice. For women we observe the same trend but this is much steeper and less significant. The figures in the bottom panel show that for low intensity campaign Facebook users change their opinion more than non-users and that this trend is reversed and widened for high intensity campaign. The plots also show that this effect is stronger and more significant for man. It is also to look at the distribution of intensity across gender, overall men appear to have been higher targeted than women as the bar charts at the bottom of the graphs show.

Insert Figure 20, in about here

Finally we investigate patterns in the behaviors of Facebook users and non-user with

respect to their political views and the level of education with the aid of figure 20, reporting the difference in the predicted marginal effects between users and non-users. Inspecting the top panel of the figure, we can infer that the difference in the probability of switching between the two groups due to reading political ads on the social media drastically goes down with campaign intensity for Republican users, the effect is similar but weaker and less significant for moderate voters, no effect at all is registered for Democrat voters instead.

From the comparison of the bar charts reported at the bottom of the graphs we can deduce that politically uninterested/ moderates and Republicans were subjected to slightly higher targeting than democrats. From the bottom panel we can observe again a reduction in the probability of changing mind for Facebook users compared to non-users among voters who did not go to college or university, no effect for those with a degree instead. So education matters too. While for people with a college or university degree reading political news on Facebook does not change the likelihood to revise their views, less educated and largely relying on Facebook for information reduced the probability of changing voting preferences when being highly targeted of at least five percentage points. Comparing the distribution of the intensity across groups along the level of education line we can clearly see that individuals without a university degree were subject overall by higher targeting compared to those with a degree, which may shed light on politicians targeting strategies. This is consistent with the conjecture that political campaigns were largely fought among less educated or informed and this may have had an impact on voting choices in the direction consistent with political polarization.

4.3.1 Media and Political Knowledge

To understand the channels through which reading news on Facebook reduces the likelihood to change individuals' initial voting strategies we need to understand how Facebook operates to affect this behavior. One possibility is that those who are constantly keeping up with the political news on the social media may become better informed compared to those who do not. In this case, it is possible that Facebook users, if more knowledgeable than their counterparts, are 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, aiming to assess respondent political knowledge. During the first interview respondents are asked "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?".

²⁶ During the post-election survey they are asked to "... recall the office of: Joe Biden, Paul Ryan, Angela Merkel, Vladimir Putin?, Justice Roberts?". ²⁷ So we take the ratio of the number of correct post-election answers over the number of correct pre-election as an indicator (albeit a bit crude one) of political knowledge improvement, *PolKnow*, and we estimate the following model

$$Pol - know_{iat} = \omega EXP_i + \phi INT_{at} + \chi EXP_i \times INT_{at} + \pi TV_i + \lambda News_i + \phi Talk_i + \theta Internet_i + \zeta Radio_i + u_{iat} \quad (4)$$

Where *TV*, *News*, *Talk*, *Internet* and *Radio* are dummies taking values of 1 if the respon-

²⁶These are questions V161513, V161514,V161515,V161515 in ANES(2017)

²⁷ Answers: Vice-President Biden, Speaker of the House Ryan, Chancellor of Germany Merkel,President of Russia Putin,US Supreme Ct Chief Justice Roberts .

dent use either of these platforms to learn about the elections. In all our specifications we also control for Political Ideology, Turnout at the 2012 Elections, Vote at the 2012 Election, party registration, vote at the Primary, gender, religion, age, marital status, education, race, children, employment status, home ownership and gun ownership and we include Income Class FE, State FE, Day of Week FE and Weekly FE. Where ω , ϕ , χ , π , λ , ϕ , θ , ζ are the parameters to estimate and u is the error term.

Insert Table (12), in about here

The result for this exercise are displayed in Table12; in the first column we ignore the interaction between exposure and intensity and we only include *EXP* as a regressor to control the effect of Facebook, we then augment this basic specification with various measures for *INT* and its interaction with *EXP*. From the inspection of the table it appears quite clearly that reading about the political campaign on Facebook does not improve respondents' knowledge about politics, the same results for radio, television and talk shows. A positive effect on political knowledge comes from using Internet and reading newspapers instead.

In summary, our results suggest quite clearly that becoming more knowledgeable is not the key to understand why Facebook people are less likely to change their mind and are glued to their initial decisions. Political polarization at this stage may be the more obvious answer, but a deeper understanding of what is going on on this front goes beyond the scope of this paper and is left to future research.

5 Conclusions

This paper has looked at the role played by social media, and in particular Facebook, in shaping political campaign and affecting election outcomes. The logic behind our arguments relies on the fact that social media exploit extensive quantities of user-generated data, which can be used to help politicians personalize their election campaigns and to target with increasing precision those voters who may be decisive in determining the outcomes of elections.

There is already evidence that this new way of campaigning may have large, and possibly unwanted, consequences on election results and on the functioning of our democratic institutions, as the rise of a new wave of populism across the western world suggests, especially in the absence of any regulations in the same way that old media outlets are.

The center piece of our paper is bringing into use a new and unique dataset which allow us to assess the effects and the power of political “micro-targeting” by using daily Facebook advertising prices, collected during the course of the 2016 U.S. presidential campaign to exploit the variation across political ideologies, and to propose a measure for the intensity of online political campaigns. We then employ this novel measure to investigate two fundamental questions: (i) To what extent did political campaigns use the social media platform micro-target voters? and (ii) How large was the effect, if any, of such campaigns on voters who relied on social media to access political news?.

Our results indicate that Facebook matters and advertising was an effective way in persuading Republican and moderate supporters to go to vote, and in swaying their votes towards Trump. Users that are exposed to highly intensive political campaigns, and who do not hold a college degree, are also persuaded to support the Republican candidate. Interestingly, our results show that Clinton also failed to convince her natural core voters to increase their support or participation in the elections. This is consistent with the claim that Trump invested in Facebook as a campaign vehicle considerably more than Clinton

or that his messages were better designed. Our results also indicated that learning about politics on Facebook does not make voters more politically informed but make them less likely to update their voting choices, which is consistent with political polarization. This effect is particularly strong for men, Republican voters and less educated individuals, which are possibly easier to "capture".

This paper constitutes one of the first attempts to shed light on how social media can affect politics. We do believe that our measures of intensity represents a step forward in advancing knowledge. Of course, this measure is not perfect, for example the content of the online messages or the identity of who pays for the ads are not observed by us; so for these reasons we are only able to observe the net outcomes of political campaigning. However we resort to a revealed preference argument, to describe the political strategy that might explain the patterns found in the variation of intensity of political campaign on the social media platform. Given the role that some audiences have in specific states (e.g. moderate voters in swing states), we can infer which political candidates would be willing to compete to secure the votes of these audiences. There still a lot to be done but this can be left to future research.

References

- [1] Allcott, H. and M. Gentzkow (2017). “Social Media and Fake News in the 2016 Election.” *Journal of Economic Perspectives* 31(2), 211-236.
- [2] American National Election Studies, University of Michigan, and Stanford University, (2017), “ANES 2016 Time Series Study. ICPSR36824-v2.” Ann Arbor, MI: Inter-university Consortium for Political and Social Research.
- [3] Barberá, P. (2015). “How Social Media Reduces Mass Political Polarization. Evidence from Germany, Spain, and the U.S.” Working Paper, NYU.
- [4] Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D., Marlow, C., Settle, J. E., and Fowler, J. H. (2012). “A 61-million-person experiment in social influence and political mobilization.” *Nature* 489(7415), 295.
- [5] Bloomberg (2018) <https://www.bloomberg.com/news/articles/2018-09-20/facebook-to-give-less-direct-support-to-trump-in-2020-campaign>.
- [6] Boxell, L., Gentzkow, M. and J.M. Shapiro (2017). “Is the internet causing political polarization? Evidence from demographics.” Working Paper, Stanford University.
- [7] Bright, J., Hale, S. A., Ganesh, B., Bulovsky, A., Margetts, H., and Howard, P. (2017). “Does Campaigning on Social Media Make a Difference? Evidence from candidate use of Twitter during the 2015 and 2017 UK Elections.” arXiv preprint arXiv:1710.07087. <https://www.technologyreview.com/s/609229/how-tweets-translate-into-votes/>
- [8] DellaVigna, S. and E. Kaplan (2007). “The Fox News effect: Media bias and voting.” *The Quarterly Journal of Economics* 122(3), 1187-1234.
- [9] Dubois, A., Zagheni, E., Garimella, K. and I. Weber (2018). “Studying Migrant Assimilation Through Facebook Interests.” Available at arXiv:1801.09430v2
- [10] Eurobarometer, S. 78, (2012). “Media use in the European Union report”. European Commission, pp: 8.
- [11] Falck, O., Gold, R., and S. Heblich (2014). “E-lelections: Voting Behavior and the Internet.” *American Economic Review* 104(7), 2238-65.
- [12] Fletcher, D., and S. Slutsky (2011). “Campaign allocations under probabilistic voting.” *Public Choice* 146(3-4), 469-499.
- [13] Gabler, N.(2016). ”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.

- [14] Gainous, J., and Wagner, K. M. (2014). "Tweeting to power: The social media revolution in American politics." Oxford University Press.
- [15] 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* DOI: 10.1073/pnas.1717781115
- [16] Gentzkow M. (2006). "Television and Voter Turnout." *Quarterly Journal of Economics* 121(3), 931-972.
- [17] Gerber, A. S. , J. G. Gimpel, D. P. Green, and D. R. Shaw (2011), "How Large and Long-lasting Are the Persuasive Effects of Televised Campaign Ads? Results from a Randomized Field Experiment," *American Political Science Review*, 105(1), 135-150.
- [18] Gold, M., and E. Dwoskin (2017) "Trump Campaign's Embrace of Facebook Shows Company's Growing Reach in Elections." *The Washington Post*
- [19] González Cabañas, J., Cuevas, A. and R. Cuevas (2018). "Unveiling and Quantifying Facebook Exploitation of Sensitive Personal Data for Advertising Purposes." *USENIX Security Symposium 2018*, Baltimore.
- [20] Gordon, B.R., Hartmann, W.R. (2016). "Advertising competition in presidential elections". *Quantitative Marketing and Economics* 14(1), 1-40.
- [21] Hillygus, D., Sunshine and T. Shields (2008). "The Persuadable Voter: Wedge Issues in Presidential Campaigns." Princeton, NJ: Princeton University Press.
- [22] Huber, G. A., Arceneaux, K. (2007). Identifying the persuasive effects of presidential advertising. *American Journal of Political Science* 51, 957-977.
- [23] Krasno, J. S. and D. P. Green (2008). "Do Televised Presidential Ads Increase Voter Turnout? Evidence from a Natural Experiment." *Journal of Politics* 70, 245-261.
- [24] Johnston, R., Hagen, M. G., and K.H. Jamieson (2004). "The 2000 presidential election and the foundations of party politics." Cambridge University Press.
- [25] Lindbeck, A., and J.W. Weibull (1987), "Balanced-Budget Redistribution as the Outcome of Political Competition", *Public Choice* 52: 273.
- [26] Liu, Y., Kliman-Silver, C., Bell, R., Krishnamurthy, B. and A. Mislove (2014). "Measurement and Analysis of OSN Ad Auctions." *Proceedings of the Second ACM Conference on Online Social Networks*. Dublin, Ireland.
- [27] Lovett M. and M. Peress (2015), "Targeting Political Advertising on Television", *Quarterly Journal of Political Science*. 10(3), 391-432. Moshary, S. (2017). "Price Discrimination in Political Advertising: Evidence from the 2012 US Presidential Elections". Working Paper, University of Chicago Booth School of Business.

- [28] Nichter, S. (2008). "Vote Buying or Turnout Buying? Machine Politics and the Secret Ballot." *The American Political Science Review* 102(1), 19-31.
- [29] Pariser, E. (2011). "The filter bubble: What the internet is hiding from you". New York, NY: Penguin Press.
- [30] Petrova, M., Sen, A. and P. Yildirim (2017), "Social Media and Political Donations: New Technology and Incumbency Advantage in the United States" CEPR Discussion Paper No. DP11808.
- [31] Pew Research Center, (May 2016), "News Use Across Social Media Platforms 2016" <http://www.journalism.org/2016/05/26/news-use-across-social-media-platforms-2016/>.
- [32] Pew Research Center, (October 2016), "The Political Environment on Social Media", <http://www.pewinternet.org/2016/10/25/the-political-environment-on-social-media/>.
- [33] Pew Research Center, (November 2016)2, "Social Media Update 2016", <http://www.pewinternet.org/2016/11/11/social-media-update-2016/>
- [34] Ridout, T. N. and M. Franz (2011). "The persuasive power of campaign advertising". Philadelphia, PA: Temple University Press.
- [35] Ridout, T. N., Franz, M., Goldstein, K. M., and J.W. Feltus (2012). "Separation by television program: Understanding the targeting of political advertising in presidential elections." *Political Communication*, 29(1), 1-23.
- [36] Saez-Trumper, D., Liu, Y., Baeza-Yates, R., Krishnamurthy, B. and A. Mislove (2014). "Beyond CPM and CPC: Determining the Value of Users on OSNs." *Proceedings of the Second ACM Conference on Online Social Networks*. Dublin, Ireland.
- [37] Shaw, D. R. (2006). "The Race to 270: The Electoral College and the Campaign Strategies of 2000 and 2004". Chicago: University of Chicago Press.
- [38] Stromberg, D. (2008), "How the Electoral College Influences Campaigns and Policy: The Probability of Being Florida," *American Economic Review*, 98(3), 769-807.
- [39] Sunstein, C. R. (2001). *Republic.com*. Princeton, NJ: Princeton University Press.
- [40] Sunstein, C. R. (2009). *Republic.com 2.0*. Princeton, NJ: Princeton University Press.
- [41] Sunstein, C. R. (2017). *Republic: Divided democracy in the age of social media*. Princeton, NJ: Princeton University Press.
- [42] Williams C. B., G. J. Gulati (2017), "Digital Advertising Expenditures in the 2016 Presidential Election", *Social Science Computer Review* , 36, 4, 406 - 421.

- [43] Zagheni E., Weber I. and G. Krishna (2017). “Leveraging Facebook’s Advertising Platform to Monitor Stocks of Migrants.” *Population and Development Review* 43, 721-734
- [44] Zaller, J. (1992). The nature and origins of mass opinion. New York, NY: Cambridge University Press.

6 Figures

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




Figure 2: *Effect of targeting on audience-specific ad prices*




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




Figure 4: *Facebook CPM Prices in Democratic Countries with Fair Elections*




Figure 5: *Facebook CPM Prices across differently aligned states*




Figure 6: *Intensity of Political Campaign across States Type - Partisan vs. Moderate Users*




Figure 7: *Intensity of Political Campaign across States Type - Very Partisan vs. Moderate Users*




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




Figure 9: *Overall Predicted Turnout*


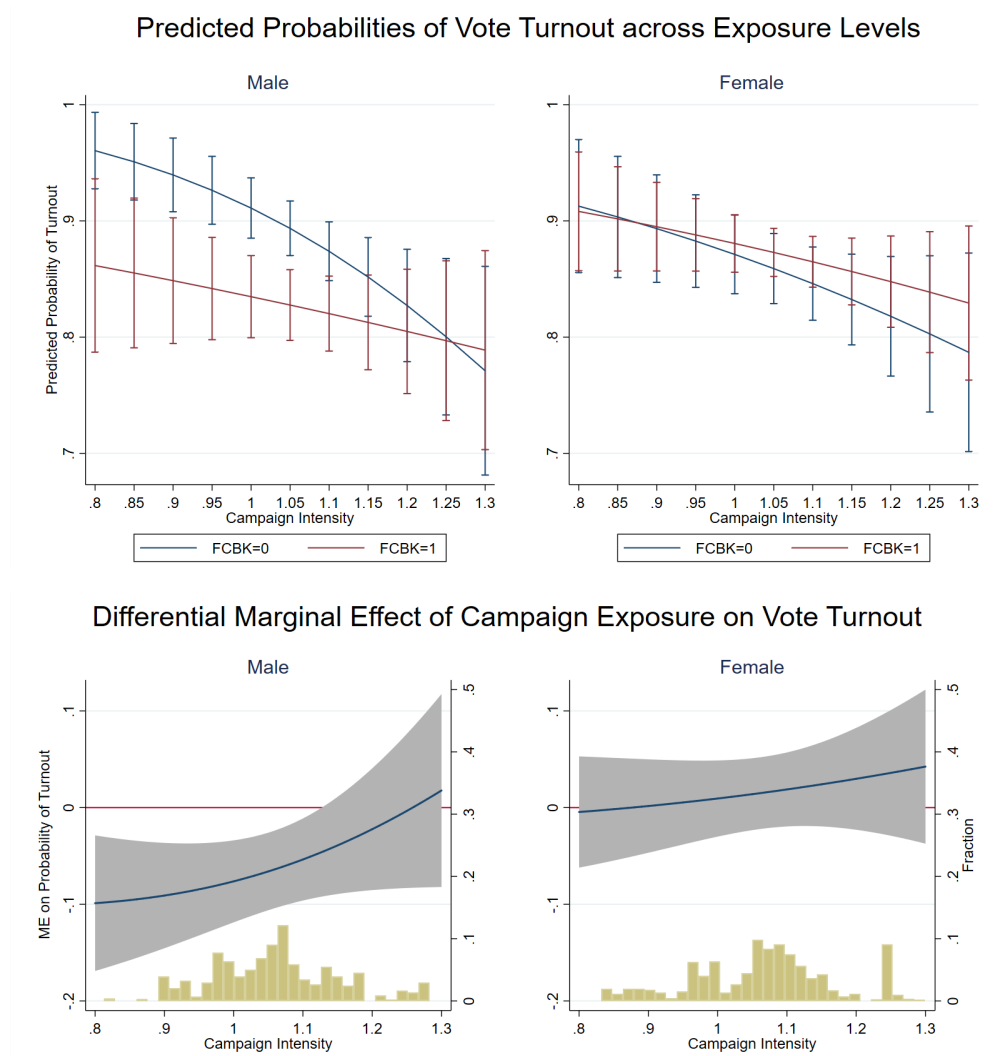


Figure 10: *Effect of Facebook micro-targeted campaigns on Turnout: Gender*



Notes. Marginal Effects w.r.t. Low versus High usage of Facebook, as a tool to access political news

Figure 11: *Effect of Facebook micro-targeted campaigns on Turnout: Ideology*

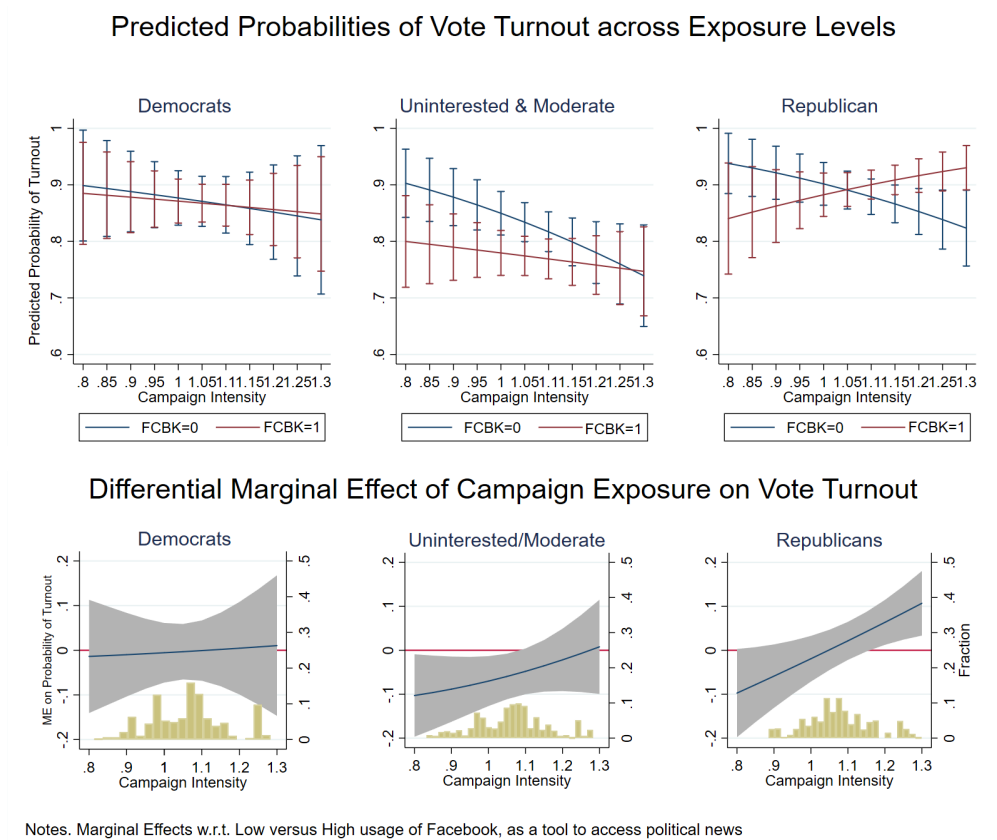


Figure 12: *Overall Predicted Vote for Trump*


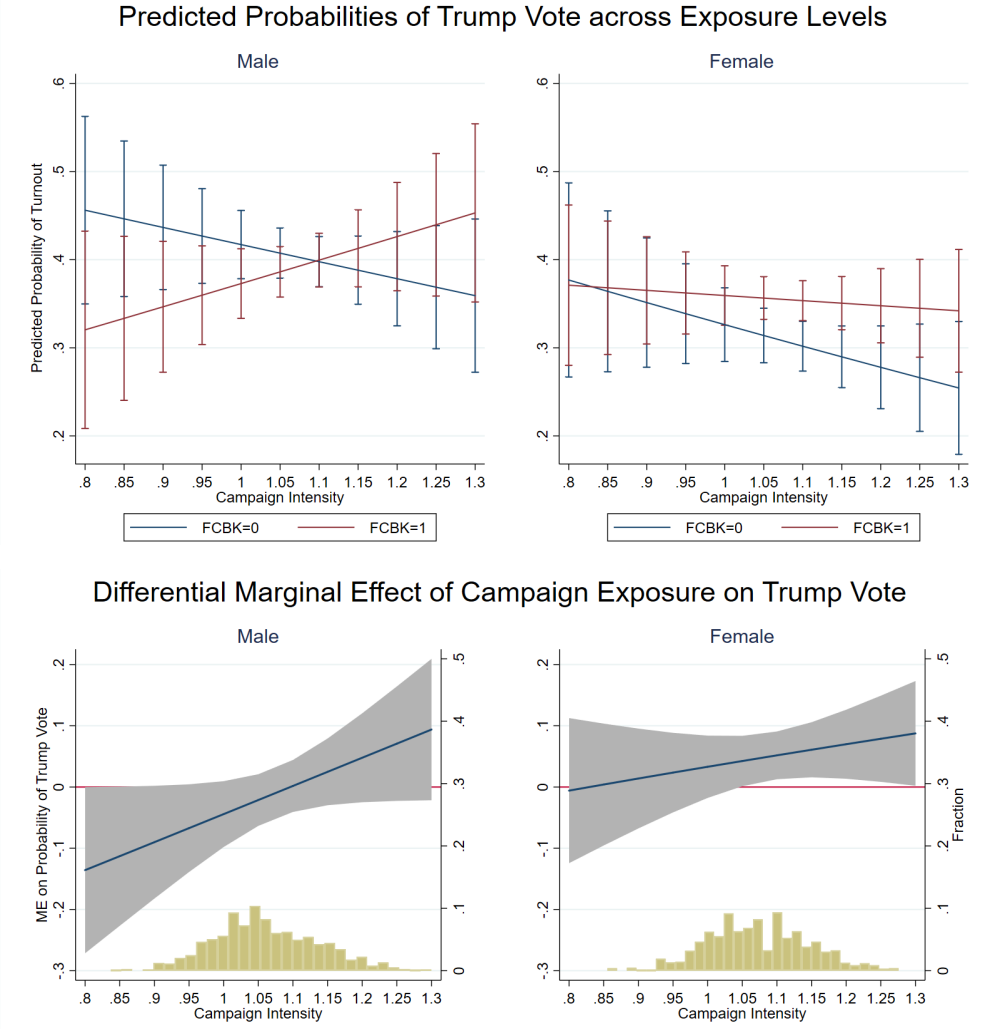
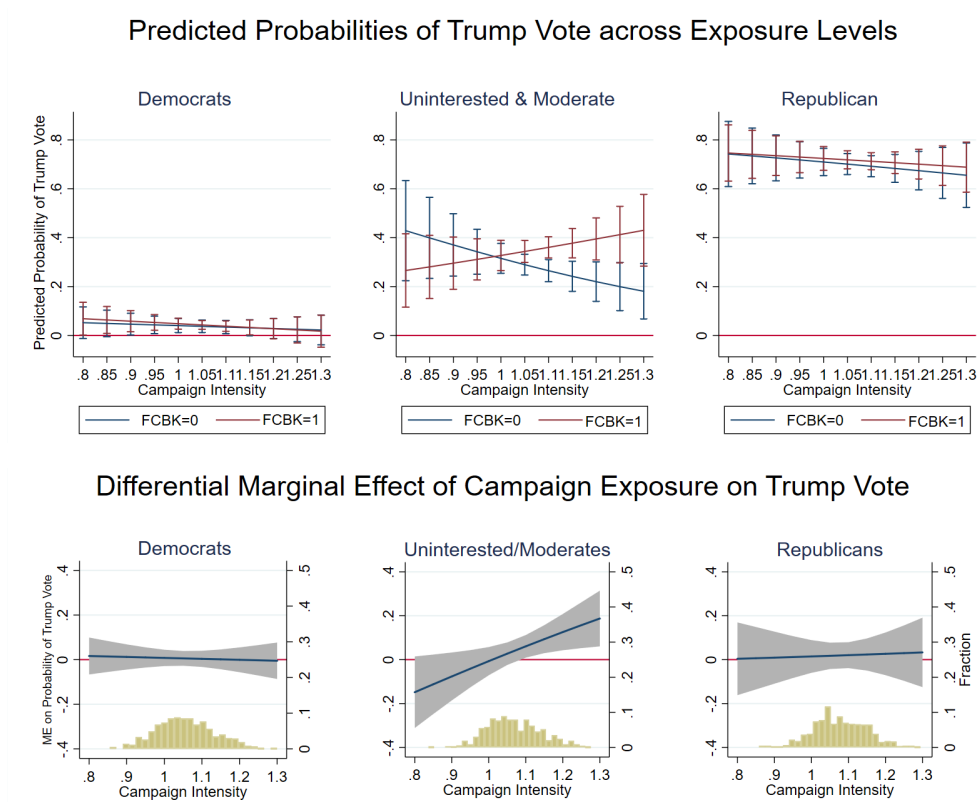


Figure 13: *Effect of Facebook micro-targeted campaigns on Trump Vote: Gender*



Notes. Marginal Effects w.r.t. Low versus High usage of Facebook, as a tool to access political news

Figure 14: *Effect of Facebook micro-targeted campaigns on Trump Vote: Ideology*



Notes. Marginal Effects w.r.t. Low versus High usage of Facebook, as a tool to access political news.

Figure 15: *Effect of Facebook micro-targeted campaigns on Trump Vote: Education*

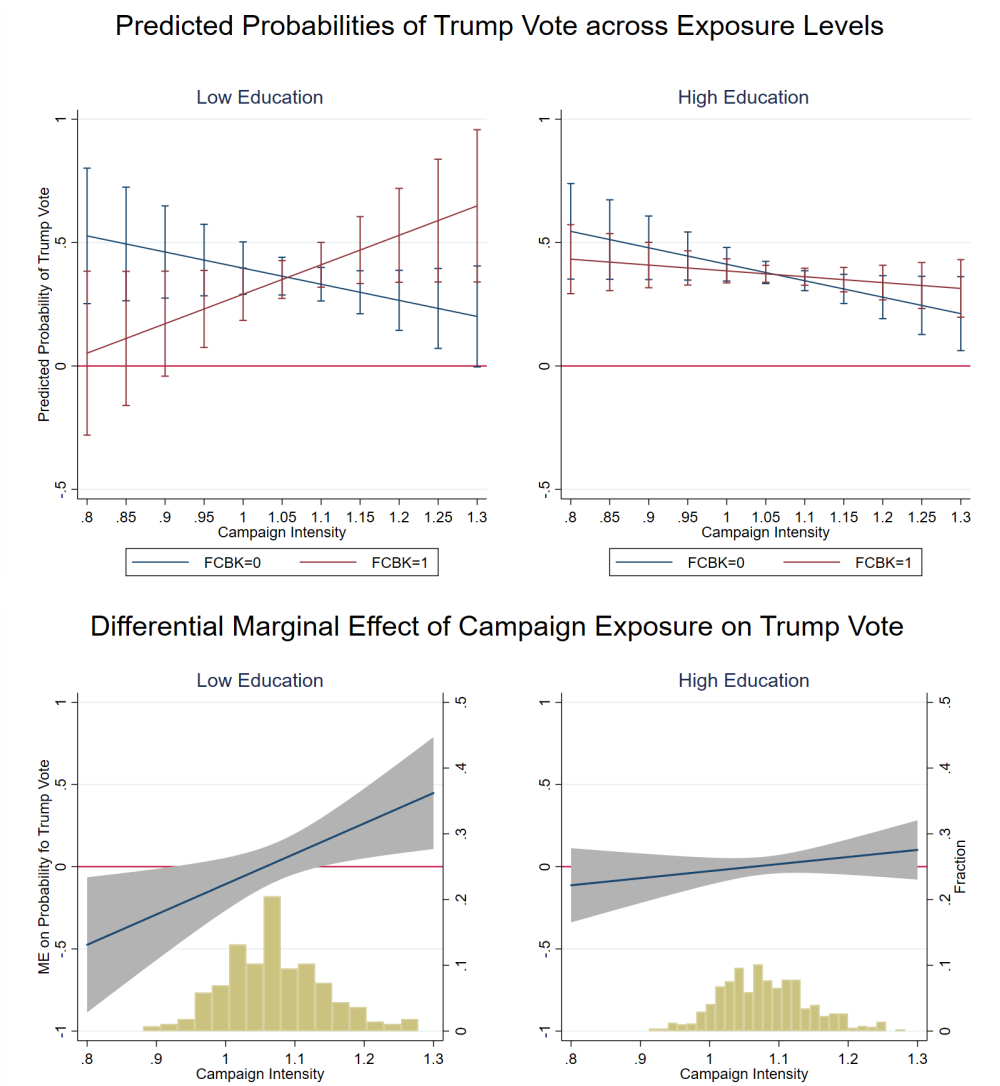


Figure 16: *Overall Predicted Vote for Clinton*


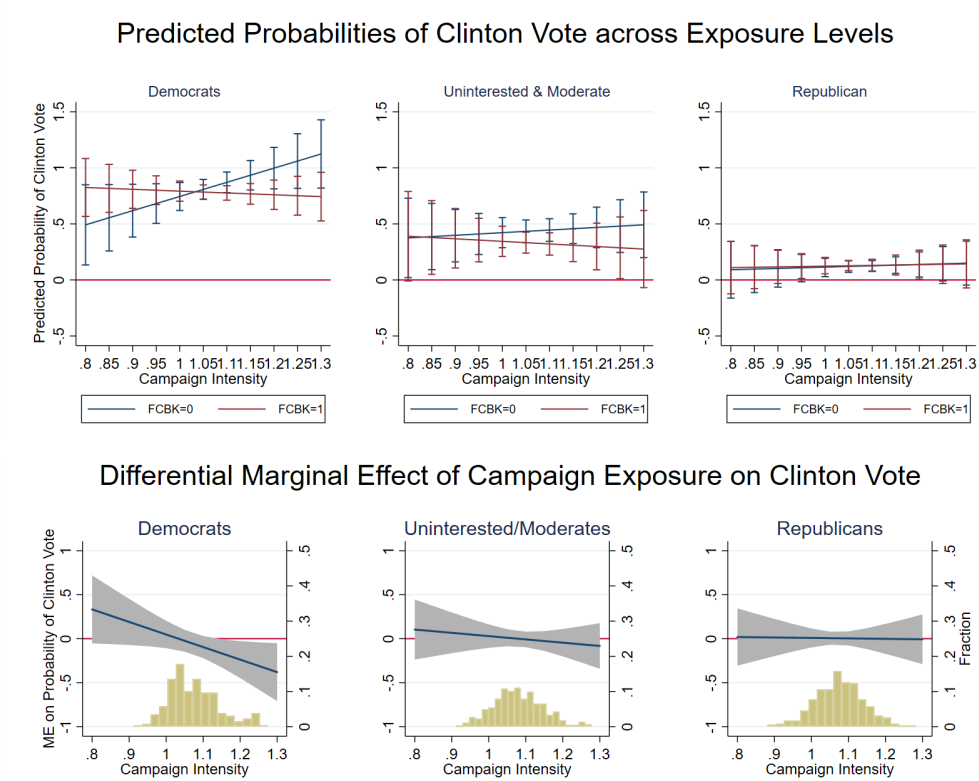


Figure 17: *Effect of Facebook micro-targeted campaigns on Clinton Vote: Ideology*



Notes. Marginal Effects w.r.t. Low versus High usage of Facebook, as a tool to access political news

Figure 18: *Overall Predicted Percentage of Vote Change*




Figure 19: *Effect of Facebook micro-targeted campaigns on Vote Change: by Gender*

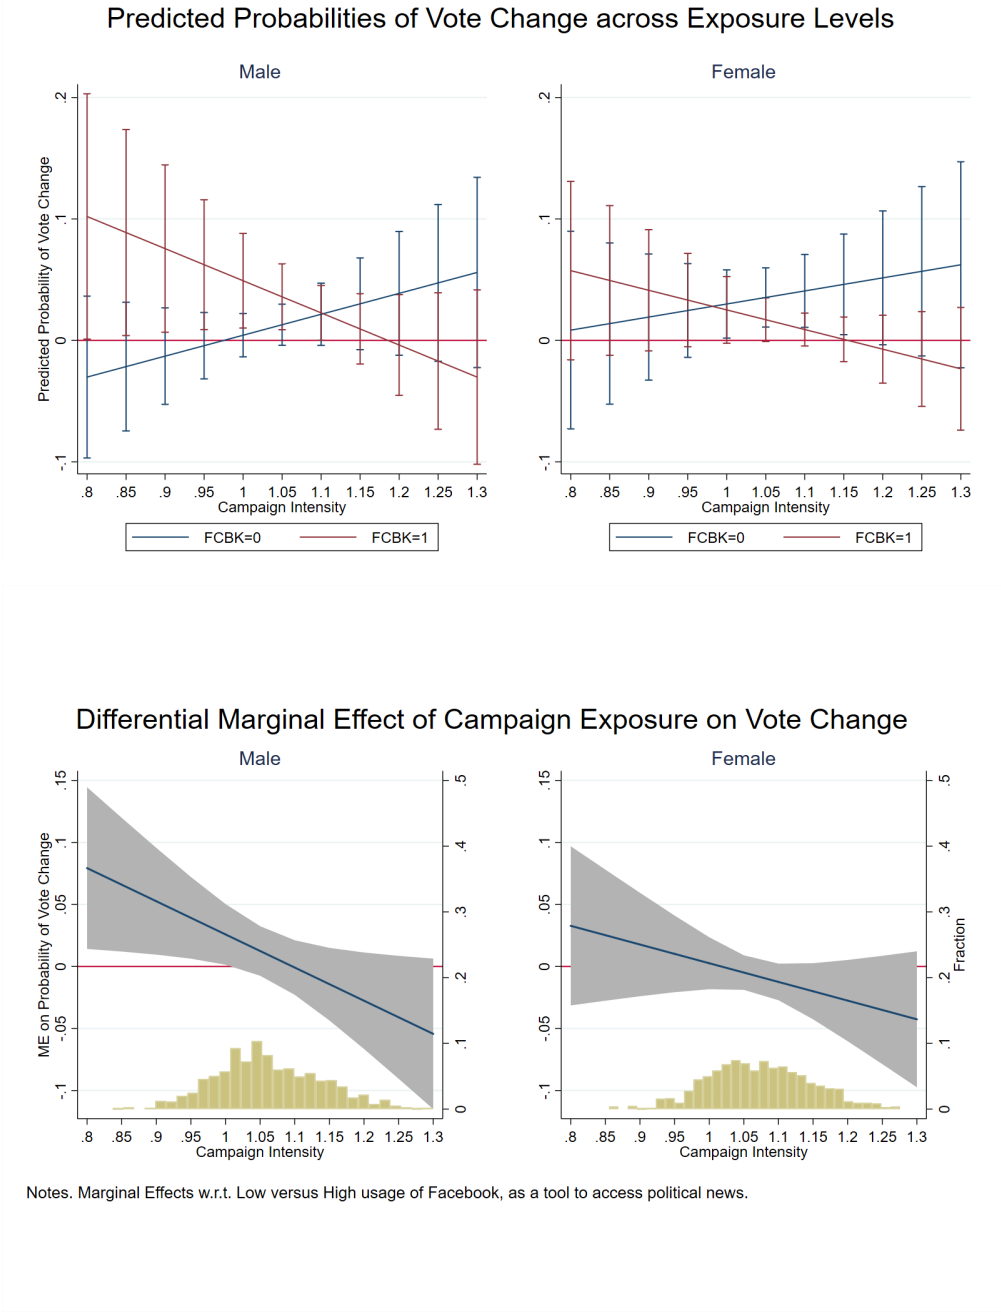





Figure 20: *Effect of Facebook micro-targeted campaigns on Vote Change:by Ideology and Education*



7 Tables

Table 1: Candidates Winning Probabilities and CPM Prices

Dep. Variable: Relative CPM Median	Republican Partisan States States	Democrat Partisan States 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 $cpm_{s,t} = \alpha cpm_{s,t-1} + \beta_1 PDIFF_{s,t-1} + \beta_2 (PDIFF_{s,t-1} \times TADV_{s,t-1}) + e_{s,t}$ with $PDIFF_{s,t} = |(P_{s,t}^T - P_{s,t}^C)|$, 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 of the dependent variable, as well as of the predetermined and endogenous variables for use as instruments. Number of Observations is 1,977 for Republican Partisan States, 1,376 for Democrat and 859 for Swing States. State grouping is based on historical election results.

Table 2: Political Microtargeting by State/Ideology/Gender

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

Table 3: The sample of Facebook users and non-users

Variable	No Facebook		Facebook 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
Democrat	997	0.472	1,430	0.504
Swing	997	0.109	1,430	0.095
Republican	997	0.418	1,430	0.401
Female	997	0.444	1,430	0.591
Age	997	55.580	1,430	46.071
Degree	997	0.538	1,430	0.610
White Resp.	997	0.713	1,430	0.760

Table 4: The effect of Facebook on Voter Turnout

Targeting Type:	Political Ideology only (1)	Political Ideology & Gender (2)	Political Ideology & Age (3)	Political Ideology & Race (4)	Political Ideology & Education (5)
News Access on FCBK (Exposure)	0.0251 (0.108)	-0.237** (0.107)	0.0628 (0.0873)	-0.0454 (0.108)	0.0468 (0.220)
Relative CPM (Intensity 1)	-0.0945 (0.0881)	-0.209** (0.0834)	-0.0623 (0.0563)	-0.131 (0.0858)	0.00858 (0.151)
News Access on FCBK x Relative CPM	-0.0331 (0.0994)	0.209** (0.0986)	-0.0686 (0.0792)	0.0324 (0.0998)	-0.0562 (0.205)
Internet Access	0.0238 (0.0430)	0.0210 (0.0427)	0.0245 (0.0428)	0.0235 (0.0429)	0.0531 (0.0663)
R^2	0.176	0.178	0.177	0.177	0.228
News Access on FCBK (Exposure)	-0.0988 (0.157)	-0.224 (0.174)	0.0188 (0.129)	-0.254 (0.160)	-0.212 (0.310)
CPM Price Ratio (Intensity 2)	0.105 (0.118)	-0.171 (0.127)	-0.0790 (0.0757)	-0.0601 (0.117)	0.0198 (0.211)
News Access on FCBK x CPM Price Ratio	0.0814 (0.143)	0.199 (0.162)	-0.0272 (0.118)	0.226 (0.147)	0.184 (0.286)
Internet Access	0.0213 (0.0429)	0.0246 (0.0428)	0.0246 (0.0428)	0.0244 (0.0429)	0.0527 (0.0661)
R^2	0.177	0.176	0.176	0.176	0.229
Observations	2,422	2,422	2,422	2,422	1,059

Note: Top panel uses Relative CPM, as measured during the week of the Interview. Bottom Panel uses Ratio of Relative CPM during Election Week to Relative CPM during Interview Week. All specifications include Income Class FE, State FE, Day of Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, Vote at the 2012 Election, party registration, vote at the Primary, gender, religion, age, marital status, education, race, children, employment status, home ownership and gun ownership.

Table 5: Heterogeneous effect on Voter Turnout

	All	Female Voters	Male Voters	Democrat Voters	Uninterested & Moderate Voters	Republican Voters	2012 Democrat Voters	2012 Republican Voters
News Access on FCBK (Exposure)	-0.237** (0.107)	-0.0704 (0.137)	-0.416** (0.186)	-0.156 (0.190)	-0.300 (0.227)	-0.281* (0.159)	-0.256 (0.167)	-0.167 (0.158)
Relative CPM (Intensity 1)	-0.209** (0.0834)	-0.227 (0.141)	-0.405*** (0.157)	-0.0511 (0.158)	-0.296* (0.175)	-0.191 (0.126)	-0.295** (0.130)	-0.0371 (0.125)
News Access on FCBK x Relative CPM	0.209** (0.0986)	0.0847 (0.124)	0.349** (0.173)	0.146 (0.183)	0.243 (0.209)	0.279** (0.138)	0.227 (0.160)	0.127 (0.137)
Internet Access	0.0210 (0.0427)	0.00537 (0.0541)	0.0131 (0.0743)	0.138 (0.111)	-0.0519 (0.0596)	0.118 (0.0864)	0.0291 (0.0567)	0.106 (0.0764)
Observations	2,422	1,284	1,138	691	917	807	1,166	784
R-squared	0.178	0.210	0.222	0.251	0.236	0.239	0.175	0.205

Note: Intensity is measured as Ratio of Relative CPM during Election Week to Relative CPM during Interview Week. All specifications include Income Class FE, State FE, Day of Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, Vote at the 2012 Election, party registration, vote at the Primary, gender, religion, age, marital status, education, race, children, employment status, home ownership and gun ownership.

Table 6: The Effect of Facebook on Trump Vote

Targeting Type:	Political Ideology only (1)	Political Ideology & Gender (2)	Political Ideology & Age (3)	Political Ideology & Race (4)	Political Ideology & Education (5)
News Access on FCBK (Exposure)	0.0275 (0.121)	-0.125 (0.122)	0.0990 (0.0991)	-0.0289 (0.123)	-0.207 (0.223)
Relative CPM (Intensity 1)	-0.0323 (0.100)	-0.157 (0.0958)	0.0874 (0.0585)	-0.0724 (0.100)	-0.0700 (0.162)
News Access on FCBK x Relative CPM	0.000553 (0.113)	0.142 (0.114)	-0.0650 (0.0886)	0.0531 (0.116)	0.205 (0.208)
Internet Access	0.0233 (0.0436)	0.0212 (0.0434)	0.0231 (0.0435)	0.0234 (0.0435)	0.0161 (0.0631)
R^2	0.536	0.537	0.537	0.536	0.558
News Access on FCBK (Exposure)	-0.169 (0.171)	-0.378** (0.186)	0.0518 (0.141)	-0.172 (0.182)	-0.775** (0.346)
CPM Price Ratio (Intensity 2)	-0.218* (0.127)	-0.284** (0.128)	-0.0164 (0.0830)	-0.173 (0.131)	-0.590** (0.259)
News Access on FCBK x CPM Price Ratio	0.182 (0.157)	0.379** (0.174)	-0.0219 (0.129)	0.186 (0.169)	0.729** (0.319)
Internet Access	0.0269 (0.0435)	0.0253 (0.0431)	0.0235 (0.0436)	0.0256 (0.0434)	0.0253 (0.0635)
R^2	0.537	0.537	0.536	0.537	0.560

Observations 2,421 2,421 2,421 2,421 1,059

Note: Top panel uses Relative CPM, as measured during the week of the Interview. Bottom Panel uses Ratio of Relative CPM during Election Week to Relative CPM during Interview Week. All specifications include Income Class FE, State FE, Day of Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, Vote at the 2012 Election, party registration, vote at the Primary, gender, religion, age, marital status, education, race, children, employment status, home ownership and gun ownership.

Table 7: Heterogeneous Effect on Trump Vote

	All Sample	Female Voters	Male Voters	Uninterested & Uninformed Voters	Republican Voters	Democrat Voters	Voters with High Education	Voters with Low Education
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
News Access on FCBK (Exposure)	-0.378** (0.186)	-0.287 (0.268)	-0.508* (0.281)	0.0503 (0.162)	-0.715** (0.351)	0.0551 (0.375)	-1.953** (0.780)	-0.457 (0.432)
Relative CPM (Intensity 1)	-0.284** (0.128)	-0.353 (0.224)	-0.182 (0.207)	-0.0598 (0.116)	-0.496** (0.218)	-0.0498 (0.271)	-0.653 (0.465)	-0.667* (0.341)
News Access on FCBK x Relative CPM	0.379** (0.174)	0.320 (0.249)	0.475* (0.264)	-0.0425 (0.153)	0.707** (0.328)	-0.0186 (0.344)	1.847** (0.723)	0.430 (0.399)
Internet Access	0.0253 (0.0431)	0.0157 (0.0523)	0.00153 (0.0748)	0.0209 (0.0311)	-0.0241 (0.0613)	0.136 (0.0933)	0.00751 (0.0961)	0.0221 (0.106)
Observations	2,417	1,282	1,135	691	914	806	275	782
R-squared	0.537	0.576	0.536	0.279	0.428	0.382	0.642	0.588

Note: Intensity is measured as Ratio of Relative CPM during Election Week to Relative CPM during Interview Week. All specifications include Income Class FE, State FE, Day of Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, Vote at the 2012 Election, party registration, vote at the Primary, gender, religion, age, marital status, education, race, children, employment status, home ownership and gun ownership.

Table 8: The Effect of Facebook on Clinton Vote

Targeting Type:	Political Ideology only (1)	Political Ideology & Gender (2)	Political Ideology & Age (3)	Political Ideology & Race (4)	Political Ideology & Education (5)
News Access on FCBK (Exposure)	-0.0472 (0.121)	-0.124 (0.110)	0.0282 (0.101)	-0.0741 (0.119)	0.410* (0.216)
Relative CPM (Intensity 1)	-0.0741 (0.0963)	-0.117 (0.0835)	-0.144** (0.0653)	-0.117 (0.0938)	0.165 (0.159)
News Access on FCBK x Relative CPM	0.0137 (0.111)	0.0851 (0.0995)	-0.0572 (0.0931)	0.0388 (0.109)	-0.403** (0.202)
Internet Access	0.0536 (0.0417)	0.0515 (0.0416)	0.0546 (0.0411)	0.0533 (0.0416)	0.112* (0.0585)
R^2	0.557	0.557	0.559	0.557	0.586
News Access on FCBK (Exposure)	-0.0337 (0.175)	0.183 (0.180)	0.174 (0.149)	-0.162 (0.177)	0.626* (0.347)
CPM Price Ratio (Intensity 2)	0.291** (0.129)	0.0763 (0.131)	-0.0699 (0.0883)	0.152 (0.127)	0.450* (0.251)
News Access on FCBK x CPM Price Ratio	0.00127 (0.161)	-0.201 (0.167)	-0.189 (0.135)	0.120 (0.163)	-0.599* (0.319)
Internet Access	0.0475 (0.0417)	0.0528 (0.0415)	0.0557 (0.0415)	0.0514 (0.0416)	0.105* (0.0594)
R^2	0.558	0.557	0.558	0.557	0.586
Observations	2,417	2,417	2,417	2,417	1,057

Note: Top panel uses Relative CPM, as measured during the week of the Interview. Bottom Panel uses Ratio of Relative CPM during Election Week to Relative CPM during Interview Week. All specifications include Income Class FE, State FE, Day of Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, Vote at the 2012 Election, party registration, vote at the Primary, gender, religion, age, marital status, education, race, children, employment status, home ownership and gun ownership.

Table 9: Heterogenous Effect of Facebook on Clinton Vote

	All	Low Education	High Education	Democrat Voters	Uninterested & Moderate Voters	Republican Voters
News Access on FCBK (Exposure)	0.626* (0.347)	1.009 (0.719)	0.559 (0.440)	1.476** (0.736)	0.403 (0.637)	0.0599 (0.642)
Relative CPM (Intensity 1)	0.450* (0.251)	0.400 (0.419)	0.411 (0.329)	1.265* (0.648)	0.384 (0.438)	0.118 (0.447)
News Access on FCBK x Relative CPM	-0.599* (0.319)	-0.954 (0.670)	-0.537 (0.403)	-1.428** (0.684)	-0.373 (0.584)	-0.0513 (0.600)
Internet Access	0.105* (0.0594)	0.00269 (0.0856)	0.205** (0.0845)	0.499** (0.223)	0.0822 (0.0811)	0.117 (0.0774)
Observations	1,057	275	782	261	444	349
R-squared	0.586	0.688	0.618	0.606	0.586	0.595

Note: Intensity is measured as Relative CPM, as measured during the week of the Interview, based on political and education targeting. All specifications include Income Class FE, State FE, Day of Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, Vote at the 2012 Election, party registration, vote at the Primary, gender, religion, age, marital status, education, race, children, employment status, home ownership and gun ownership.

Table 10: Effect of Facebook on Change in Voter Behaviour

Targeting Type:	Political Ideology only (1)	Political Ideology & Gender (2)	Political Ideology & Age (3)	Political Ideology & Race (4)	Political Ideology & Education (5)
News Access on FCBK (Exposure)	0.0299 (0.0482)	0.0440 (0.0518)	-0.0353 (0.0368)	0.00503 (0.0521)	0.0140 (0.0954)
Relative CPM (Intensity 1)	0.0191 (0.0382)	0.0444 (0.0413)	0.00579 (0.0295)	0.0239 (0.0386)	0.0288 (0.0740)
News Access on FCBK x Relative CPM	-0.0259 (0.0437)	-0.0389 (0.0484)	0.0345 (0.0338)	-0.00282 (0.0485)	-0.0162 (0.0895)
Internet Access	-0.0487** (0.0247)	-0.0480* (0.0245)	-0.0489** (0.0247)	-0.0485** (0.0247)	-0.0733* (0.0385)
R^2	0.068	0.068	0.068	0.068	0.145
News Access on FCBK (Exposure)	0.0430 (0.0652)	0.227** (0.0889)	0.0646 (0.0598)	-0.0250 (0.0686)	0.340*** (0.128)
CPM Price Ratio (Intensity 2)	-0.0246 (0.0465)	0.140** (0.0578)	0.0447 (0.0482)	-0.0284 (0.0511)	0.140 (0.0992)
News Access on FCBK x CPM Price Ratio	-0.0379 (0.0595)	-0.210** (0.0834)	-0.0575 (0.0555)	0.0251 (0.0642)	-0.317*** (0.120)
Internet Access	-0.0480* (0.0246)	-0.0494** (0.0246)	-0.0485** (0.0247)	-0.0481* (0.0246)	-0.0755** (0.0383)
R^2	0.068	0.071	0.068	0.068	0.150
Observations	2,422	2,422	2,422	2,422	1,059

Note: Top panel uses Relative CPM, as measured during the week of the Interview. Bottom Panel uses Ratio of Relative CPM during Election Week to Relative CPM during Interview Week. All specifications include Income Class FE, State FE, Day of Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, Vote at the 2012 Election, party registration, vote at the Primary, gender, religion, age, marital status, education, race, children, employment status, home ownership and gun ownership.

Table 11: Heterogeneous Effect on Change in Voters Behaviour

	All Sample	Female Voters	Male Voters	Democrat Voters	Uninterested& Uninformed Voters	Republican Voters	Voters with No Degree	Voters with University Degree
News Access on FCBK (Exposure)	0.227** (0.0889)	0.153 (0.127)	0.293** (0.129)	0.0602 (0.0808)	0.214 (0.187)	0.348** (0.151)	0.514 (0.418)	0.303* (0.161)
Relative CPM (Intensity 2)	0.140** (0.0578)	0.211* (0.109)	0.0214 (0.0727)	-0.0327 (0.0613)	0.209** (0.106)	0.217* (0.119)	0.0245 (0.221)	0.254* (0.145)
News Access on FCBK x Relative CPM	-0.210** (0.0834)	-0.151 (0.118)	-0.267** (0.122)	-0.0509 (0.0729)	-0.209 (0.175)	-0.317** (0.140)	-0.445 (0.386)	-0.291* (0.154)
Internet Access	-0.0494** (0.0246)	-0.0633* (0.0339)	0.000110 (0.0303)	0.0181 (0.0151)	-0.0655* (0.0359)	-0.00887 (0.0462)	0.00799 (0.0301)	-0.194** (0.0833)
R^2	0.071	0.138	0.121	0.157	0.154	0.221	0.476	0.197
Observations	2,422	1,284	1,138	691	917	807	275	784

Note: Intensity is measured as Ratio of Relative CPM during Election Week to Relative CPM during Interview Week. All specifications include Income Class FE, State FE, Day of Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, Vote at the 2012 Election, party registration, vote at the Primary, gender, religion, age, marital status, education, race, children, employment status, home ownership and gun ownership.

Table 12: Effect of Facebook on Information

	(1) None	(2) Political Ideology Only	(3) Political Ideology & Gender	(4) Political Ideology & Race	(5) Political Ideology & Age	(6) Political Ideology & Education
News Access on FCBK (Exposure)	0.00996* (0.00592)	0.0249 (0.0659)	0.0295 (0.0640)	-0.0805* (0.0484)	-0.0120 (0.0682)	-0.0450 (0.131)
Relative CPM (Intensity 2)		0.372 (0.290)	0.00297 (0.284)	-0.251 (0.182)	0.436 (0.304)	0.330 (0.681)
News Access on FCBK x Relative CPM		-0.0208 (0.0610)	-0.0255 (0.0597)	0.0759* (0.0438)	0.0133 (0.0634)	0.0448 (0.120)
TV	0.0313 (0.0508)	0.0152 (0.0467)	0.0161 (0.0467)	0.0166 (0.0467)	0.0154 (0.0467)	0.0631 (0.0770)
Newspapers	0.00873 (0.0359)	0.0570* (0.0324)	0.0548* (0.0324)	0.0551* (0.0324)	0.0561* (0.0324)	0.0975* (0.0549)
TalkShows	0.00193 (0.0399)	0.0273 (0.0363)	0.0282 (0.0363)	0.0273 (0.0364)	0.0269 (0.0363)	-0.00579 (0.0614)
Internet	0.0967** (0.0419)	0.138*** (0.0381)	0.135*** (0.0382)	0.137*** (0.0381)	0.137*** (0.0381)	0.0418 (0.0611)
Radio	-0.0202 (0.0376)	-0.00338 (0.0335)	-0.00491 (0.0334)	-0.00275 (0.0335)	-0.00338 (0.0335)	0.0493 (0.0569)
Observations	2,191	2,191	2,191	2,191	2,191	938

Note: Top panel uses Relative CPM, as measured during the week of the Interview. Bottom Panel uses Ratio of Relative CPM during Election Week to Relative CPM during Interview Week. All specifications include Income Class FE, State FE, Day of Week FE and Weekly FE. They also control for Political Ideology, Turnout at the 2012 Elections, Vote at the 2012 Election, party registration, vote at the Primary, gender, religion, age, marital status, education, race, children, employment status, home ownership and gun ownership.

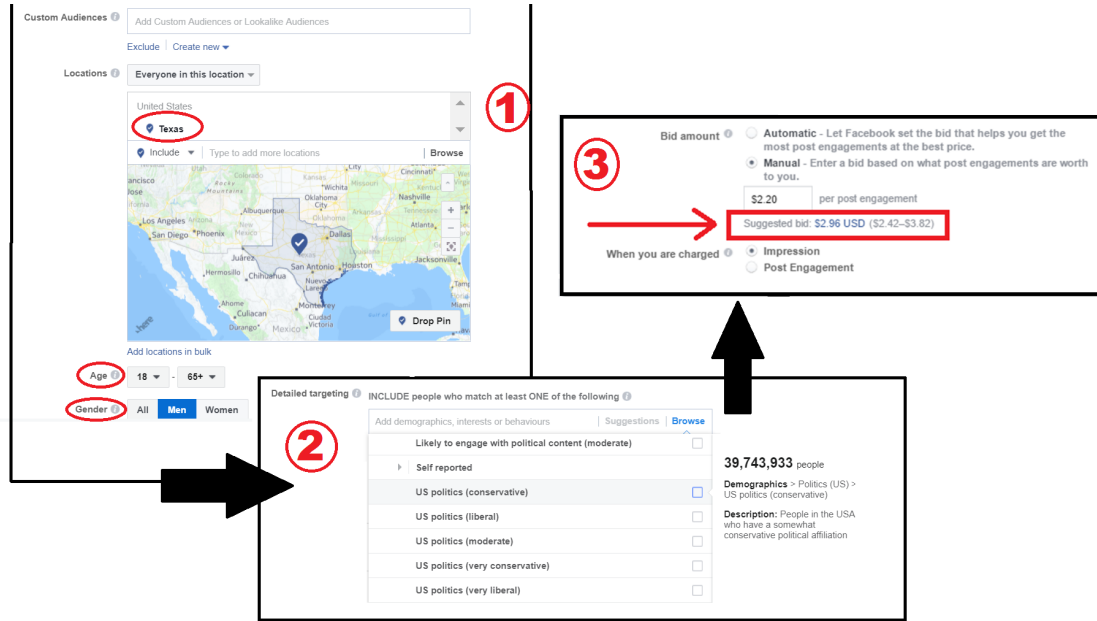
A Appendix

A.1 Selection of Targeted Audience

Advertisers on Facebook compete for ad slots to the *desired* audience. Facebook offers multiple layers of targeting, so that the advertiser can direct their ads to the exact audience that maximizes the return on their investment. Figure 21 shows images of the Audience Selection Dashboard available to Facebook Advertisers. Audiences can be chosen by a number of criteria, among which 1) geographical location and 2) political ideology. Once the audience is selected, Facebook indicates the number of active users and a range of prices (minimum, median and maximum) corresponding to a spectrum of bids that are winning auctions in real time to reach the same audience. Note that political ideology falls under the “Demographics” characteristics, in the Facebook audience selection criteria. This means that political alignment is determined on the basis of users activity on the site, and it is rather constant over time. It depends on pages liked, political preferences stated in the profile, but also on the average political alignment of those who like the same pages of the user. There are two ways to pay for ads, depending on the objective of the campaign: for conversion-oriented objectives advertisers usually pay per click (CPC) or per action (CPA)²⁸, whereas for exposure objectives, like political campaigns, advertisers choose to pay for impressions (CPM). We scrape Daily Prices for State/Ideology specific Audiences. We extract both CPM and CPC. We further do this for Age, Gender, Race and Education State/Ideology specific audiences. Once the bids are placed, Facebook allocates ad spaces according to a mechanism based on the Vickrey-Clarke-Groves auction described below.

²⁸CPC stands for cost-per-click. This is the rate that websites charge advertisers every time someone clicks on an ad. CPA stands for cost-per-action. This is the rate that websites charge advertisers every time a user clicks and does a specific action (such as page likes, video views, mobile app installs, etc.).

Figure 21: *Facebook Audience Selection Dashboard*



A.2 Auctions for Advertising on Facebook

To allocate advertising opportunities on its website, Facebook adopts an auction system based on the Vickrey-Groves-Clark (VCG) mechanism. We briefly present this mechanism (in an informal way), and discuss the implications for our analysis. 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 considers all bids and calculates for each 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, consider the simple case where 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.

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 Facebook.

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, the logic can be described 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 "opportunity cost" 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).

According to Facebook, this system minimizes the probability that ads are shown to uninterested audiences. Furthermore, it increases the price advertisers can expect to pay when attempting to reach an audience that is also targeted by other relevant ads. This induces advertisers to properly target ads at lucrative targets. Facebook favors ads that are directly

relevant to the person browsing at the time. Less relevant ads cause higher "social" harm, so the price the advertiser pays increases as well. This feature implies that users who are interested in politics are more likely to be exposed to politics-related ads. Finally, we expect that the prices paid to reach audiences who are targeted by a greater number of candidates should also increase, because, all else equal, a greater number of bidders interested in an audience makes the marginal loss (and hence the final price) go up.

A.3 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 GDPR, one of the most conservative data protection regulations, 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, we can confirm there is no privacy/data protection risks for

Facebook users derived from our research.