

The Welfare Effects of Facebook: Pre-Analysis Plan

Hunt Allcott, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow*

November 7, 2018

*Allcott: New York University, NBER, and JPAL. hunt.allcott@nyu.edu. Braghieri: Stanford University. lucabrag@stanford.edu. Eichmeyer: Stanford University. saraeich@stanford.edu. Gentzkow: Stanford University and NBER. gentzkow@stanford.edu. We thank Nancy Baym, Moira Burke, Alex Leavitt, and Todd Rogers for helpful comments. We thank Kelly Buckley and Raj Bhargava for exceptional research assistance. We are grateful to the Sloan Foundation for grant funding. The study was approved by IRBs at Stanford (eProtocol #45403) and NYU (IRB-FY2018-2139). Disclosures: Allcott is a paid employee of Microsoft Research. Gentzkow does paid consulting work for Amazon and is a member of the Toulouse Network for Information Technology, a research group funded by Microsoft. Braghieri and Eichmeyer have no relevant or material disclosures.

I Introduction

The rise of social media has profoundly influenced the way many people communicate and get information. Most social media users visit the platforms at least once a day, and the average user spends nearly an hour per day on Facebook alone. It may be that no new technology since television has so profoundly altered the way people spend their time, communicate with others, and get information about the world.

There has been growing concern that social media may have unintended negative impacts. Many correlation studies find that higher social media use is associated with lower subjective well-being and mental health.¹ Negative outcomes such as suicide and depression seem to have risen sharply over the same period that the use of smartphones and social media has expanded (Twenge, Sherman, and Lyubomirsky 2016; Twenge and Park 2017; Twenge, Martin, and Campbell 2018; Twenge et al. 2018). It is alleged that by creating ideological “echo chambers” among like-minded friend groups, Facebook has increased political polarization (Sunstein 2001, 2017). Social media are the primary medium spreading fake news and other types of misinformation (Allcott and Gentzkow 2017), and there is concern that coordinated disinformation campaigns can affect elections in the U.S. and abroad. Alter (2018), along with other academics and prominent Silicon Valley executives in the “time well-spent” movement, argue that digital media devices and social media apps are harmful and addictive. Central to this discussion is Facebook, which despite continuous changes in the social media landscape is by far the largest social media platform in terms of ad revenue, total users, and total time spent.

To test these concerns, we carry out a randomized experiment in a large sample of U.S. Facebook users in the run-up to the 2018 midterm elections. We recruit sample via Facebook ads and measure participants’ willingness-to-accept (WTA) to deactivate Facebook for four weeks using an incentive-compatible Becker-DeGroot-Marschak (1964, “BDM”) survey technique. A Treatment group is randomly assigned to receive an offer price of \$102, which (in our pilots) is sufficient to induce most of them to deactivate, while a Control group is assigned to receive an offer of \$0 and thus remain active. With informed consent from all participants, we measure a suite of outcomes using text messages, online surveys, responses to emails about political topics and apps to social media use, online data measuring participants’ activity on Twitter and Facebook, and administrative records on voting and electoral contributions.

Our design allows us to address two types of questions highlighted above. First, we can carry out a standard impact evaluation, measuring the treatment effects of four-week Facebook deactivation on news knowledge, political engagement, political polarization, subjective well-being, social interaction, opinions about Facebook, and post-experiment Facebook use. Second, using several

¹See, for example, Burke and Kraut 2016, Frison and Eggermont 2015, Kross et al. 2013, Satici and Uysal 2015, Shakya and Christakis 2017, and Tandoc, Ferrucci, and Duffy 2015. See Appel, Gerlach, and Crusiu 2016 and Baker and Algorta 2016 for reviews.

BDM elicitation, we can measure the distribution of consumer valuations of Facebook, as well as how valuations change in expected and unexpected ways after a period of deactivation.

Sections II and III detail the experimental design, outcome variables, and empirical strategy. Section IV presents table and figure shells for the impact evaluation.

II Experimental Design

II.A Experiment Overview

Table 1 presents the experiment timeline and sample sizes. We timed the experiment so that the main period of Facebook deactivation would end shortly after the midterm elections on November 6th. The experiment has seven parts: recruitment, pre-screen, baseline, midline, and endline surveys, daily text messages, and post-endline communications.

We recruited participants using Facebook ads. Our ad said, “Participate in an online research study about internet browsing and earn an easy \$30 in electronic gift cards.” Appendix Figure A1 presents the ad. To minimize sample selection bias, the ad did not further hint at our research questions or suggest that the study was related to social media or Facebook deactivation. We targeted the ads by demographic cells in an attempt to gather an initial sample that was approximately representative of Facebook users on gender, age, college completion, and political ideology. 1,690,076 unique users were shown the ad, of whom 30,064 clicked on it. This 1.8 percent click-through rate is about twice the average click-through rate on Facebook ads across all industries (Irvine 2018).

Clicking on the ad took the participant to a brief pre-screen survey, which included several background demographic questions and the consent form. The pre-screen, baseline, midline, and endline surveys were hosted on the site stanforduniversity.qualtrics.com. 14,324 people passed the pre-screen, by reporting being a U.S. resident born between the years 1900 and 2000 who uses Facebook more than 15 minutes and no more than 600 minutes per day. Of those people, 5,974 people consented to participate in the study.

After the consent form, participants began the baseline survey. The baseline recorded email address, additional demographics, and outcome variables. Finally, we asked for each participant’s name, zip code, Twitter handle, the URL to their Facebook Profile page (which we would use “solely to observe whether your Facebook account is active”), and their phone number (“in order for us to send you text messages during the study”). Finally, at the end of the baseline survey, we informed people that “As part of this study, we will ask you to deactivate your Facebook account twice for a period of 24 hours. You will keep your access to Facebook messenger. If you deactivate, you can choose to come back whenever you want with your content and friends network unchanged,” and asked, “Are you willing to deactivate your Facebook account twice for 24 hours (once after Survey 2 and once after Survey 3)?”

3,234 people finished the baseline survey. Of those, 183 were dropped from the experimental

sample because of invalid data (for example, invalid Facebook Profile URLs) or low-quality baseline responses (for example, discrepancies between average daily Facebook usage reported in the pre-screen vs. baseline survey, completing the survey in less than ten minutes, no text in short-answer boxes, and other patterns suggesting careless responses). The remaining 3,051 participants had valid baseline data and were randomized.

On October 11th, we sent email invitations to the midline survey. First, we asked people to deactivate their Facebook account for 24 hours, and we guide them through the process. Second, we used a Becker-DeGroot-Marschak (BDM) mechanism to elicit willingness-to-accept (WTA) to stay deactivated for four weeks instead of 24 hours.² We then revealed the BDM price offer. Participants whose WTA was strictly less than the price draw were informed that they should deactivate for four weeks. Third, we reminded people that we would again ask them to deactivate for 24 hours after the endline survey, and used a second BDM mechanism to elicit WTA to stay deactivated for the four weeks after endline instead of just 24 hours. We informed people that “We will check continuously whether your account is deactivated for the entire [24 hours/4 weeks] in which it is supposed to be by pinging the URL associated with your profile.”

On November 8th, two days after the midterm election, we will send an email invitation to the endline survey. We measure the same outcome variables as in the baseline. All questions are identical, with several exceptions described in Section II.C below, such as rephrasing questions about the midterm to be in the past tense. We then ask people to again deactivate their Facebook accounts for the next 24 hours, and elicit WTA to stay deactivated for the next four weeks instead of the next 24 hours. Participants were told, “With a 50% chance we will require you to abide by the decision you made 4 weeks ago; with 50% chance we will ignore the decision you made 4 weeks ago and we will require you to abide by the decision you make today.”

We also gather data from three post-endline communications. After the endline, we will ask participants to report how many minutes per day they had used Facebook in the past seven days. We ask participants with iPhones to report the time recorded by their phone in the Settings app, and we ask all other participants to just estimate. We later send an email with links to information on ways to limit smartphone social media use, and finally we send an email with links to donate, volunteer for, or sign petitions related to political causes. Appendix Figures A2-A3 present the emails.

For the approximately six weeks between baseline and endline, we send daily text message surveys to measure aspects of subjective well-being that were best measured daily instead of in a retrospective endline survey. We rotate three types of questions, measuring happiness, the primary

²The survey explained, “The computer has randomly generated an amount of money to offer you to deactivate your Facebook account for the next 4 weeks. Before we tell you what the offer is, we will ask you the smallest offer you would be willing to accept. If the offer the computer generated is above the amount you give, we will ask you to deactivate for 4 weeks and pay you the offered amount if you do. If the offer is below that amount, we will not ask you to deactivate.” We then asked several comprehension questions to make sure that participants understood the mechanism. We did not tell participants the distribution from which the offer prices would be drawn.

emotion felt over the past ten minutes, and loneliness. Appendix Figure A4 presents the three questions.

We verify deactivation by checking each participant’s Facebook Profile page URL regularly at random times. While a user can limit how much content other people see when they browse to the user’s Profile page, the Profile URL returns an error message within the Facebook website stating “Sorry, this content isn’t available right now” if and only if the account has been deactivated.³ This is thus our measure of deactivation. For all participants, we verify deactivation at least once per day for the seven days before midline and the 28 days after endline. Between midline and endline, we verify deactivation at least twice per day for the Treatment group and at least once every four days for the Control group. During the post-midline and post-endline 24-hour deactivation periods, we verify deactivation within a few hours of when they completed the survey.

If participants were not deactivated when they were supposed to be, we immediately send an email informing them that they should again deactivate as soon as possible, along with a survey asking them to explain why they were not deactivated.

All participants receive \$5 per completed survey, paid via gift card immediately upon completion. All participants were told that they would receive a \$15 “completion payment” if they completed all surveys, responded to 75 percent of text messages, and, if in the Treatment group, kept their accounts deactivated for 75 percent of the time between midline and endline.

II.B Randomization

We use the BDM mechanism described above to randomly assign participants to Facebook deactivation. Figure 1 illustrates the randomization. Participants with valid baseline data were randomized into three groups: Treatment (approximately 35 percent of the sample, with offer price for post-midline deactivation $p = \$102$), Control (approximately 65 percent of the sample, with $p = \$0$), and a third group (approximately 0.2 percent of the sample, with p drawn from a uniform distribution on $[\$0, \$170]$). Treatment and Control were balanced within 48 strata defined by age, average daily Facebook use, heavy vs. light news use (those who get news from Facebook fairly often or very often vs. never, hardly ever, or sometimes), active vs. passive Facebook use, and Democrat, Republican, or independent party affiliation.

For the post-endline deactivation, 0.2 percent of participants were randomly selected to a BDM offer price drawn randomly from $p' \in [0, 170]$, while the remaining 99.8 percent receive offer $p' = 0$ for post-endline deactivation. The post-endline BDM offer price was balanced between Treatment

³By default, Facebook profile URLs end in a unique number, which is the numeric ID for that person in the Facebook system. Users can update their default URL to be something customized and they can change their customized URL as often as they want. At the start of the experiment, we checked if a Facebook profile URL (it could be the customized version) was valid by pinging it and looking in the page source for the string containing the numeric ID for the person. If the numeric ID existed, we knew that the URL was valid. We could also use the ID to reconstruct the default numeric URL that always navigates to that person even if they change their customized URL.

and Control, so two Treatment group members and four Control group members were assigned to positive post-endline offers $p' \in [0, 170]$.

This approach allows us to maintain incentive compatibility in the BDM mechanism, assure balance between Treatment and Control groups, and use a simple regression to estimate treatment effects of post-midline deactivation.

II.C Outcome Variables

For the impact evaluation, we consider the outcome variables in the nine families described below. We further collect these families into three groups: substitution (sections II.C.1 and II.C.2), “news and political” outcomes (sections II.C.3-II.C.5), and “well-being” outcomes (sections II.C.6-II.C.9). Appendix B presents details on each outcome variable and moderator, grouped by family. We also construct indices that combine the outcome variables within each family, weighting by the inverse of the covariance between variables, as described in Anderson (2008). In constructing these indices, we orient the outcome variables so that more positive values have the same meaning (e.g. more positive means “more polarized” in all cases). Outcomes to be multiplied by -1 are prefaced by a “-” sign in Appendix B .

II.C.1 Substitute time uses

At baseline and endline, we ask participants how many minutes per day they spent on Facebook on the average day in the past four weeks. At baseline, we also ask participants to report how much of their free time on the average day in the past four weeks they spent on various activities, ranging from using social media apps other than Facebook to spending time with friends and family in person. At endline, we ask parallel questions about free time activities, but phrased to elicit *changes* in how much time they spent on these things in the past four weeks (between endline and midline) relative to what is normal for them. We phrased the questions in this way in order to more precisely detect changes in self-reported time use caused by the deactivation.

II.C.2 Substitute news sources

At baseline, we ask participants how often they get news from different sources, including Facebook, cable TV, print, and radio news, borrowing a standard survey question from the Pew Center (2018). At endline, we ask parallel questions, but again phrased to ask about *changes* in how often they got news from these sources over the past four weeks. For the participants who report having a Twitter handle, we gather data on number of tweets for the four weeks before baseline began and the four weeks between midline and endline. This allows a non-self-reported measure of one kind of potential substitution away from Facebook.

II.C.3 News knowledge

In order to detect broad changes in news exposure, we ask participants how closely they followed politics, how closely they followed news about President Trump, and how many minutes per day they spent watching, reading, or listening to the news (including on social media) over the past four weeks.

In order to measure specific news knowledge, we included a 15-question news knowledge quiz. For each question, we gave a statement from the news in the past four weeks and asked participants to indicate if they thought the statement was true or false. The order of the 15 statements was randomized. Seven of the statements were from news stories covered in the past four weeks in six news websites: New York Times, Wall Street Journal, Fox News, CNN, MSNBC, and US News & World Report, such as “The Trump administration set the maximum number of refugees that can enter the country in 2019 to 30,000.” Three of the headlines were false modifications of articles from those same two sources, such as “President Trump spoke at the funeral of former Arizona Senator John McCain, honoring the late McCain’s wish.” The *news knowledge* variable reflects respondents’ scores on those ten questions. The final five statements were from fake news stories—rated false by third-party fact-checkers snopes.com and factcheck.org—that circulated heavily within a four-week period before the survey. The *fake news knowledge* variable reflects respondents’ scores on those five questions.

II.C.4 Political engagement

We include two measures of political engagement in the political engagement index. First, we measure whether participants voted in the 2018 midterm election, by matching on name, birth year, and zip code to a voting database supplied to Stanford by L2, a voting data provider. Second, we measure whether participants clicked on any of the links in the post-endline politics email.

II.C.5 Political polarization

There are a variety of ways to measure political polarization (Gentzkow 2016), and we use both standard and novel measures. First, we include standard “feeling thermometer” questions capturing how “warm or cold” participants feel toward the Democratic and Republican Parties and President Trump. The *party affective polarization* variable is the respondent’s thermometer warmth toward her own party minus her warmth toward the other party. Second, the *Trump affective polarization* variable is the thermometer warmth toward President Trump (for Republicans) or minus one times the thermometer warmth toward President Trump (for Democrats). For these and all other polarization variables, we include independents who lean toward a party, and we drop independents who do not lean toward either party. Third, we ask respondents to list recent news events that made them angry at the Republican or Democratic Party. *Party anger* is the natural log of one plus the

length (in characters of text) of her response about the other party minus the natural log of one plus the length of her response about her own party. Fourth, *party understanding* is the number of times the respondent reports seeing news that made her better understand the point of view of her own party minus the number of times she saw news that helped her understand the other party.

Fifth, we asked nine opinions about current political issues, such as “To what extent do you think that free trade agreements between the U.S. and other countries have been a good thing or a bad thing for the United States?” The *issue polarization* variable reflects the extent to which the respondent’s issue opinions align with the opinions of other members of her party in our sample. Sixth, *belief polarization* reflects the extent to which the respondent’s beliefs about current news events (from the news knowledge quiz described above) align with the beliefs of the average other co-partisan. Finally, *vote polarization* measures the strength of preferences for the congressional candidate of the same party measured in a generic ballot question described below.

II.C.6 Subjective well-being

There is a vast literature on measuring subjective well-being (see, e.g. Kahneman et al. 2006), and we adopt standard measures from the literature. We modified existing scales in two ways. First, we asked questions in reference to the past four weeks, so as to increase our ability to detect changes as a result of Facebook deactivation. Second, in some cases we chose a subset of questions from standard multi-question scales in order to focus on areas of subjective well-being that might be most affected by Facebook.

The *happiness* variable uses two questions from the Subjective Happiness Scale (Lyubomirsky and Lepper 1999), asking how happy participants were over the past four weeks as a person and “compared to most of my peers.” *Life satisfaction* is three questions from the Satisfaction with Life Scale (Diener et al. 1985), such as the level of agreement with the statement, “During the past 4 weeks, I was satisfied with my life.” *Loneliness* is the Three-Item Loneliness Scale (Hughes et al. 2004). Finally, *depressed*, *anxious*, *absorbed*, and *bored* are four emotions measured by the European Social Survey well-being module (Huppert et al. 2009).

The daily text messages allowed us to measure the aspects of subjective well-being that are most important to record in the moment instead of retrospectively. This approach builds on the Experience Sampling Method of Csikszentmihalyi and Larson (1987) and Stone and Shiffman (1994). The variable *SMS happiness* is the answer to the question, “Overall, how happy do you feel right now on a scale from 1 (not at all happy) to 10 (completely happy)?” The variable *SMS positive emotion* is an indicator variable for whether the participant reports a positive emotion when asked, “What best describes how you felt over the last ten minutes? Please text back the corresponding number.” Finally, *SMS loneliness* is the answer to the question, “How lonely are you feeling right now on a scale from 1 (not at all lonely) to 10 (very lonely)?”

II.C.7 Social interaction

We have three measures of social interaction. The *friends met in person* variable is the natural log of one plus the number of friends seen in person in the last week. Specifically, the survey question asked participants to “list the first names of as many friends you met in person last week that you can think of in 1 minute.” *Offline activities* is the number of offline activities (such as going out to dinner, spending time with your kids, etc.) that the person did at least once last week. *Diverse interactions* is the indicator for whether the respondent interacted with someone who voted the opposite way in the last presidential election plus the indicator for whether the respondent interacted with someone from another country in the last week.

II.C.8 Opinions about Facebook

We ask eight questions eliciting people’s opinions about Facebook, including “To what extent do you think Facebook is good or bad for society?” and “To what extent do you think Facebook makes people more or less politically polarized?” Each of these eight responses is on a ten-point scale. In the endline survey only, we also asked, “As part of this study, you were asked to deactivate your Facebook account for [24 hours/4 weeks]. To what extent do you think that deactivating your account was good or bad for you?” We also have two open answer text boxes in which we ask people to write out the most important positive and negative impacts that Facebook has on their lives. The *positive impacts* and *negative impacts* variables are the natural log of one plus the count of characters in the respective text box.

II.C.9 Post-experiment Facebook use

We use four measures of planned and actual post-experiment Facebook use. First, *planned use*, which is measured in the endline survey, is the extent to which participants plan to use Facebook more or less than they had before they started the study. Second, *clicked time limit email* is an indicator for whether the respondent clicks any of the links in the post-endline social media time limit email. Third, *time to reactivation* is the natural log of one plus the number of days that the participant’s account remains inactive after the post-endline 24-hour deactivation period. Fourth, *post-endline mobile use* is the natural log of one plus the number of minutes that the participant reports using Facebook on their phone in the post-endline text message.

II.C.10 Secondary outcomes

We also plan to include a figure with effects on secondary outcomes.

First, we consider the standard generic ballot question: “If the elections for U.S. Congress were being held today, would you vote for the Republican Party’s candidate or the Democratic Party’s candidate for Congress in your district?” To increase precision, we then ask, “How convinced are

you about whether to vote for the Republican or Democratic candidate?” At endline, we ask these questions in past tense, about whom the respondent did vote for in the 2018 midterm (or whom the respondent would have voted for had she voted, to avoid potentially selective non-response). The *voted Republican* variable is the strength of preferences for the Republican candidate. This variable is secondary because we expect the estimates to be too imprecise to be of interest.

Second, we also ask self-reported voting in the 2018 midterms. This is secondary because it is superseded by the administrative voting data from L2. If for some reason we have problems with the administrative data, we will promote this to a primary outcome.

Third, we gather contributions to political campaigns from the Federal Election Commission database, both in the four years before baseline and during the midline-endline period. This is secondary because very few Americans contribute to political campaigns, and we do not expect to be able to detect effects of four weeks of deactivation.

We may also use this figure to add ex-post exploratory analyses.

III Empirical Strategy

III.A Empirical Strategy

For our impact evaluation, we estimate the treatment effects of the 27 days of additional deactivation in Treatment relative to Control. While the great majority of the treatment group stayed deactivated between midline and endline, our data suggest that many of those who failed to deactivate did so intentionally, and may have been using Facebook. Our primary analyses therefore consider treatment-on-the-treated effects instead of intent-to-treat effects.

Define Y_i as some outcome measured at endline, and \mathbf{Y}_i^b as a vector including the baseline value of the outcome and the baseline value of the index that includes the outcome.⁴ Define D_i as the percent of deactivation checks between 24 hours post-midline and November 7th (the day before endline) that person i is observed to be deactivated. Define $T_i \in \{1, 0\}$ as a Treatment group indicator, and ν_s as the vector of stratum dummies. We can estimate local average treatment effects of deactivation using the following regression:

$$Y_i = \tau D_i + \rho \mathbf{Y}_i^b + \nu_s + \varepsilon_i, \tag{1}$$

instrumenting for D_i with T_i . In Equation (1), τ measures the local average treatment effect of deactivation for people induced to deactivate by the experiment.

Define H_i as person i 's average daily hours of Facebook use reported at baseline, winsorized at 120 minutes. Facebook deactivation should have a larger impact for people who use Facebook

⁴ \mathbf{Y}_i^b excludes the baseline value of the outcome for outcomes such as clicks on post-endline emails that do not have a baseline value. \mathbf{Y}_i^b excludes the baseline index when Y_i is not included in an index. When Y_i is an index, \mathbf{Y}_i^b is simply the baseline value of the index.

more. Thus, we may have more power if our primary right-hand-side variable is D_iH_i instead of D_i , as in the following regression:

$$Y_i = \tau D_iH_i + \beta H_i + \rho \mathbf{Y}_i^b + \nu_s + \varepsilon_i, \quad (2)$$

where we analogously instrument for D_iH_i with T_iH_i . In Equation (2), τ measures the local average treatment effect of deactivation *per hour of daily Facebook use avoided*.

If H_i is measured with significant error, then the latter specification could reduce precision. We will make either Equation (1) or Equation (2) our primary specification, depending on which delivers more power.

For all outcome variables other than willingness-to-accept, we exclude from the sample all participants with WTA of \$102 or higher, as none of these people were assigned to deactivate. When analyzing willingness-to-accept, we do not exclude these people, as sample selection on the outcome variable would bias our results.

We anticipate that we might have attrition rates that differ statistically between Treatment and Control. If that is the case, we will use a version of the procedure introduced by Behagel et al. (2015). We use all observations in the group (either Treatment or Control) with the lower response rate. We then exclude all participants in the other group that started the endline survey after time t^* , where t^* is set to equalize the response rates. Equalizing attrition in this way delivers unbiased estimates under the monotonicity assumption described in Behagel et al. (2015).

III.B Changes to Initial Pre-Analysis Plan

We submitted our initial pre-analysis plan on October 12. This revised version makes the following changes.

- On the basis of data on non-compliance during the study, we now plan to make the intent-to-treat estimates our primary results. We thus changed Equations (1) and (2) to be instrumental variables regressions.
- We now plan to also report effects on indices of substitute time use and substitute news sources.
- We renamed *party thermometer* and *Trump thermometer* to be *party affective polarization* and *Trump affective polarization*, respectively.
- We made a handful of other non-substantive editorial changes.

IV Empirical Results

IV.A Probability of Being Deactivated

Figure 2 presents the probability of being deactivated in Treatment and Control, by day of the experiment.

IV.B Substitutes for Facebook

Figure 3 presents treatment effects on substitutes for Facebook: substitute news sources and substitute time uses. This and the remaining figures use simulated data, as the endline data have not yet been collected.

IV.C Effects on News and Political Outcomes

Figure 4 presents treatment effects on news and political outcomes: news knowledge, political engagement, and political polarization.

IV.D Effects on Well-Being Outcomes

Figure 5 presents estimates on what we broadly call well-being outcomes: subjective well-being, social interactions, opinions about Facebook, and post-experiment Facebook usage.

IV.E Heterogeneous Treatment Effects

We will test the following moderators:

1. We present separate estimates of effects on subjective well-being text message surveys for text messages sent during the time of day when the respondent reported using Facebook the most.
2. For the news and political outcomes in Section IV.C, we present separate estimates after splitting the sample into *heavy news users* vs. *light news users* (those who get news from Facebook fairly often or very often vs. never, hardly ever, or sometimes).
3. For the well-being outcomes in Section IV.D, we present separate estimates after splitting the sample into *active users* vs. *passive users*. We measure this using two questions: share of active vs. passive browsing using a question based on the Passive and Active Facebook Use Measure (Gerson, Plagnol, and Corr 2017), and “what share of your time on Facebook do you spend interacting one-on-one with people you care about,” which was suggested by Facebook’s well-being team as a key measure of the type of usage that they think improves

well-being. Active vs. passive users are defined as having above- vs. below-median sum of their two responses to these questions.

4. For estimates of Equation (1), we present separate estimates after splitting the sample into above- vs. below-median daily Facebook use.
5. In secondary analyses, we also consider two secondary moderators:
 - (a) age, for all outcomes
 - (b) political party, limited to the news and political outcomes

V Measuring the Consumer Surplus from Facebook

Figure 6 presents the average of WTA in treatment and control at midline for post-midline deactivation, at midline for post-endline deactivation, and at endline for post-endline deactivation.

We will extend this analysis significantly as we develop our theory model.

References


- Allcott, Hunt and Matthew Gentzkow. 2017. "Social Media and Fake News in the 2016 Election." *Journal of Economic Perspectives* 31 (2):211–236.
- Alter, Adam. 2018. *Irresistible: The Rise of Addictive Technology and the Business of Keeping Us Hooked*. Penguin Random House.
- Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak. 1964. "Measuring Utility by a Single-Response Sequential Method." *Economic Information, Decision, and Prediction* .
- Behagel, Luc, Bruno Crépon, Marc Gurgand, and Thomas Le Barbanchon. 2015. "Please Call Again: Correcting Nonresponse Bias in Treatment Effect Models." *The Review of Economics and Statistics* 97 (5):1070–1080.
- Csikszentmihalyi, Mihaly and Reed Larson. 1987. "Validity and reliability of the Experience-Sampling Method." *The Journal of Nervous and Mental Disease* .
- Diener, Ed, Robert A. Emmons, Randy J. Larsen, and Sharon Griffin. 1985. "The Satisfaction With Life Scale." *Journal of Personality Assessment* 49 (1).
- Gentzkow, Matthew. 2016. "Polarization in 2016." Toulouse Network for Information Technology Whitepaper.
- Gerson, Jennifer, Anke C. Plagnol, and J. Corr, Philip. 2017. "Passive and Active Facebook Use Measure (PAUM): Validation and relationship to the Reinforcement Sensitivity Theory." *Personality and Individual Differences* 117:81–90.
- Hughes, Mary Elizabeth, Linda J. Waite, Louise C. Hawkey, and John T. Cacioppo. 2004. "A Short Scale for Measuring Loneliness in Large Surveys: Results From Two Population-Based Studies." *Res Aging* 26 (6):655–672.
- Huppert, Felicia A., Nic Marks, Andrew Clark, Johannes Siegrist, Alois Stutzer, Joar Vittersø, and Morten Wahrendorf. 2009. "Measuring Well-being Across Europe: Description of the ESS Well-being Module and Preliminary Findings." *Social Indicators Research* 91:301–315.
- Kahneman, Daniel, Alan B. Krueger, David Schkade, Norbert Schwarz, and Arthur A. Stone. 2006. "Would You Be Happier If You Were Richer? A Focusing Illusion." *Science* 312:1908–1910.
- Lyubomirsky, Sonja and Heidi S. Lepper. 1999. "A Measure of Subjective Happiness: Preliminary Reliability and Construct Validation." *Social Indicators Research* 46:137–155.
- Stone, Arthur A. and Saul Shiffman. 1994. "Ecological momentary assessment (EMA) in behavioral medicine." *Annals of Behavioral Medicine* 16 (3):199–202.

- Sunstein, Cass. 2001. *Republic.com*. Princeton University Press.
- . 2017. *Republic: Divided Democracy in the Age of Social Media*. Princeton University Press.
- Twenge, Jean M., Thomas E. Joiner, Megan L. Rogers, and Gabrielle N. Martin. 2018. “Increases in Depressive Symptoms, Suicide-Related Outcomes, and Suicide Rates Among U.S. Adolescents After 2010 and Links to Increased New Media Screen Time.” *Clinical Psychological Science* 6 (1).
- Twenge, Jean M., Gabrielle N. Martin, and W. Keith Campbell. 2018. “Decreases in psychological well-being among American adolescents after 2012 and links to screen time during the rise of smartphone technology.” *Emotion* 18 (6):765–780.
- Twenge, Jean M. and Heejung Park. 2017. “The Decline in Adult Activities Among U.S. Adolescents, 1976–2016.” *Child Development* .
- Twenge, Jean M., Ryne A. Sherman, and Sonja Lyubomirsky. 2016. “More Happiness for Young People and Less for Mature Adults: Time Period Differences in Subjective Well-Being in the United States, 1972–2014.” *Social Psychological and Personality Science* 7 (2).

Table 1: **Experiment Timeline and Sample Sizes**


Phase	Date	Sample size
Recruitment and baseline	September 24 - October 3	N=1,690,076 shown ads N=30,064 clicked on ads N=14,324 passed pre-screen N=5,974 consented N=3,234 finished baseline N=3,051 had valid baseline, were randomized
Midline	October 11	N=□ began midline N=□ finished midline
Endline	November 8	N=□ began endline N=□ finished endline
Post-endline	November-December	N=□ responded to time use text message
End of second deactivation	December 6	

Figure 1: **Treatment Groups and Randomization**



Notes: This figure presents the share of the sample assigned to each price p offered for post-midline deactivation and to each price p' offered for post-endline deactivation.

Figure 2: **Probability of Being Deactivated**



Notes: The figure shows mean percent of checks at which participants had their Facebook accounts deactivated across treatment/control individuals, by day of the experiment.

Figure 3: **Substitutes for Facebook**




Figure 4: **Effects on News and Political Outcomes**




Figure 5: Effects on Subjective Well-Being, Social Interactions, and Opinions about Facebook





Figure 6: Average Valuation of Facebook in Treatment and Control



Online Appendix: Not for Publication

The Welfare Effects of Facebook

Hunt Allcott, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow

A Experimental Design Appendix

Figure A1: Facebook Advertisement Used for Recruitment



Figure A2: **Post-Endline Social Media Time Limit Email**

Stanford
University

Dear ,

Thank you for participating in our study about Facebook use!

With the end of the study, we wanted to provide resources that can help you manage the technology in your life. We know how difficult it can be to stay away from or limit the time you spend on social media and other technologies.

There are tools that can help you track and limit your social media usage on your smartphone. If you would like to learn about them, [click here](#) if you have an iPhone and [click here](#) if you have an Android device. If you would like further information about other ways to curb smartphone use, here is a TIME magazine article we hope you find useful:

[Learn more about ways to limit smartphone use](#)

Thanks again for your participation in the study, and we wish you all the best.

Sarah, Luca, Kelly, Hunt, Matt, and Raj

The Stanford Online Experience Research Study Research Team

Figure A3: **Post-Endline Politics Email**

The 2018 midterm elections are over, but you can still participate and make your voice heard in important national, state, and local issues. There are many ways you might make a difference in future elections.

If you support Democratic candidates you might:

- Volunteer to help Democratic candidates in your community [here](#).
- Sign a petition denouncing recent voter ID laws [here](#).

Donate to the Democratic Party to help prepare for future elections




If you support Republican candidates you might:

- Volunteer to help Republican candidates in your community [here](#).
- Sign a petition to encourage people to stand for the National Anthem [here](#).

Donate to the Republican Party to help prepare for future elections

We hope you find these resources useful and engage with issues that matter to you.

The Stanford Online Experience Research Study Research Team

Figure A4: **Subjective Well-Being Text Messages**(a) **Happiness**(b) **Primary Emotion**(c) **Loneliness**

B Variable Definitions by Family

Variable name	Question text
<i>Substitute time uses</i>	
Facebook minutes	On an average day in the past 4 weeks, how many minutes would you say you spent on Facebook, including through the Facebook app on your phone? (<i>not included in substitute time uses index</i>) On an average day in the last 4 weeks, how much free time (i.e. excluding work) did you spend... [0 minutes, Between 1 and 30 minutes, Between 31 minutes and 1 hour, Between 1 and 2 hours, Between 2 and 3 hours, More than 3 hours]
Non-FB social media	...using social media apps other than Facebook?
Non-social online	...online (on your computer, tablet, smartphone, etc.) for things other than social media?
TV alone	...watching TV or movies by yourself?
Non-screen alone	...on non-screen activities (e.g. cooking, reading books, exercising – anything without an electronic screen in front of you) by yourself?
Friends and family	...doing anything with friends and family (in person)?
<i>Substitute news sources</i>	
	Over the past four weeks, how often did you... [Never, Hardly Ever, Sometimes, Fairly Often, Very Often]
Facebook	...get news from Facebook (<i>not included in substitute news sources index</i>)
Print	...read any newspapers in print?
Radio	...listen to the news on the radio?
Local TV	...watch local television news?
Network TV	...watch national evening network television news (such as ABC World News, CBS Evening News, or NBC Nightly News)?
Cable TV	...watch cable television news (such as CNN, the Fox News cable channel, or MSNBC)?
Non-FB social media	...get news from social media sites other than Facebook (e.g. Twitter or Snapchat)?
Non-social online	...get news from news websites or apps other than social media?
Number of tweets	$\ln(1+\text{number of tweets in past four weeks})$

<i>News knowledge</i>	
Follow politics	Thinking back over the last 4 weeks, how closely did you follow US politics? [Not at all closely, somewhat closely, rather closely, very closely]
Follow Trump	Thinking back over the last 4 weeks, how closely did you follow news about President Trump? [Not at all closely, somewhat closely, rather closely, very closely]
News minutes	On an average day of the last 4 weeks, how many minutes did you spend watching, reading or listening to the news (including news via social media)? [text box]
News knowledge (at baseline)	Of the following news events, which ones do you think are true, and which ones do you think are false? [True, False, Unsure] Tension in trade negotiations escalated between the United States and China, with President Trump announcing tariffs on \$200 billion worth of goods. An off-duty Dallas police officer entered the apartment of an African-American neighbor and shot and killed the unarmed neighbor. Deputy Attorney General Rod Rosenstein early in his tenure suggested secretly recording President Trump and recruiting cabinet members to remove him from office. The Trump administration set the maximum number of refugees that can enter the country in 2019 to 30,000. Michael Cohen, President Donald Trump's former personal attorney, agreed to cooperate with the Mueller investigation team and discuss Trump's business dealings with Russia. President Trump blasted Attorney General Jeff Sessions for the indictments of two lawmakers who supported Trump during the 2016 election. CBS chief executive Les Moonves resigned after multiple sexual misconduct allegations. President Trump's former campaign chairman Paul Manafort refused deal to cooperate with the Mueller investigation team in exchange for legal charges against him being dropped. President Trump spoke at the funeral of former Arizona Senator John McCain, honoring the late McCain's wish. Hurricane Florence caused more than 300 deaths.
(at endline)	A prominent Saudi Arabian journalist who was critical of the country's government was killed inside the Saudi Arabian consulate in Istanbul.

In the weeks preceding the midterm elections, several high-profile Democrats, including Barack Obama and Hillary Clinton, were sent packages containing explosive devices.

A mass shooting fueled by anti-Semitic sentiment took place in a synagogue in Pittsburgh.

President Trump announced he plans to sign an executive order to prevent second-generation immigrants born in the United States from automatically being granted U.S. citizenship.

The Department of Justice charged a Russian national allegedly involved in a wide-ranging online disinformation campaign aimed at influencing the Midterm elections.

One of the women who made allegations against Supreme Court Justice Brett Kavanaugh has admitted to investigators that the allegations were fabricated. Attorney General Jeff Sessions resigned at President Trump's request.

Harvard University recently stood trial for allegedly discriminating against African-American applicants in its admission process.

Far-right candidate Jair Bolsonaro recently won an election to become the President of Argentina.

Senator Elizabeth Warren's DNA test results show that she has no native American ancestry.

Fake news
knowledge
(at baseline)

After researcher Dr. Christine Blasey Ford accused Supreme Court nominee Brett Kavanaugh of sexual assault, it is revealed that Kavanaugh's mother once ruled against Dr. Blasey Ford's parents in a foreclosure case.

CNN's Anderson Cooper reported deceptively on Hurricane Florence, standing in a ditch to create the misleading impression that he was filming amidst waist-deep floodwaters.

Mayor Carmen Yulín Cruz of San Juan was arrested for misappropriating \$3 million in disaster relief funds intended for the victims of Hurricane Maria in Puerto Rico.

Clerk refused to sell gas to a man fleeing hurricane Florence over a Trump bumper sticker.

WikiLeaks released an email showing that Hillary Clinton's presidential campaign bribed prominent Republicans to oppose Donald Trump during the 2016 election.

(at endline)

Billionaire George Soros was revealed to be one of the funders of a caravan of Central American emigrants traveling through Mexico to the U.S. border.

A Russian feminist activist poured bleach on men who were “manspreading” on the train (“manspreading” refers to men sitting in public transport with legs wide apart, thereby covering more than one seat).

In a recent vote, all Democrats in Congress voted against a 2.8% cost of living allowance in Social Security benefits.

Cesar Sayoc, suspect in an act of domestic terrorism directed at vocal critics of President Trump, was a registered Democrat.

None of the 154 mass shootings in 2018 was committed by a black man, illegal alien, or woman.

Political engagement

Voted	Takes value 1 if recorded as having voted in 2018 midterm, and 0 otherwise
Clicked politics email	Takes value 1 if clicked on any link in the post-endline politics email, and 0 otherwise

Political polarization

Party affective polarization	Thinking back over the last 4 weeks, how warm or cold did you feel towards the parties and the president on the feeling thermometer?
Trump affective polarization	Thinking back over the last 4 weeks, how warm or cold did you feel towards the parties and the president on the feeling thermometer?
Party anger	List as many recent (last 4 weeks) news events you can think of that made you angry at the Republican Party. (If more than 5, just list those 5 that left you most angry. If less than 5, list less. If none, enter "none" in the first textbox.) List as many recent (last 4 weeks) news events you can think of that made you angry at the Democratic Party. (If more than 5, just list those 5 that left you most angry. If less than 5, list less. If none, enter "none" in the first textbox.)
- Party understanding	Thinking back over the last 4 weeks, how often did you see news that made you better understand the point of view of the Republican Party? [Never, Once, Two or three times, Four times or more] Thinking back over the last 4 weeks, how often did you see news that made you better understand the point of view of the Democratic Party? [Never, Once, Two or three times, Four times or more]
Issue polarization	To what extent do you think that free trade agreements between the U.S. and other countries have been a good thing or a bad thing for the United States? Overall, would you say that blacks or whites are treated more fairly in dealing with the police?

	Do you think that employers firing men who have been accused of sexual harassment or assault before finding out all the facts is a major or a minor problem?
	As you may know, Brett Kavanaugh is a federal judge who has been nominated to serve on the Supreme Court. Would you like to see the Senate vote in favor of Kavanaugh serving on the Supreme Court, or not?
	On the whole, do you think immigration is a good thing or a bad thing for this country today?
	How confident, if at all, are you that the Justice Department special counsel Robert Mueller will conduct a fair investigation into Russian involvement in the 2016 election?
	In general, do you feel that the laws covering the sale of firearms should be made less strict, more strict, or kept as they are now?
	In presenting the news dealing with political and social issues, do you think that news organizations deal fairly with all sides, or do they tend to favor one side?
	To what extent do you think President Trump is honest and trustworthy?
Belief polarization	Level of agreement with co-partisans on beliefs questions
Vote polarization	Strength of generic ballot preference for co-partisan candidate (see Voted Republican question)

Subjective well-being

Happiness	Over the last 4 weeks, I think I was [1 (not a very happy person) ... 7 (a very happy person)]
	Over the last 4 weeks, compared to most of my peers, I think I was [1 (less happy) ... 7 (more happy)]
Life satisfaction	Below are three statements that you may agree or disagree with. Indicate your agreement with each item and please be open and honest in your responding. [Strongly disagree, Disagree, Slightly disagree, Neither agree nor disagree, Slightly agree, Agree, Strongly agree]
	In most ways my life during the past 4 weeks was close to ideal.
	The conditions of my life during the past 4 weeks were excellent.
	During the past 4 weeks, I was satisfied with my life.
- Loneliness	How often did you feel that you lacked companionship over the past four weeks [Hardly ever, Some of the time, Often]
	How often did you feel left out over the past four weeks [Hardly ever, Some of the time, Often]

	How often did you feel isolated from others over the past four weeks [Hardly ever, Some of the time, Often]
	Below are some ways you might have felt or behaved in the past 4 weeks. Please tell us how much of the time during the past 4 weeks: [1 None or almost none of the time, 2, 3, 4 All or almost all of the time]
- Depressed	... you felt depressed.
- Anxious	... you felt anxious.
Absorbed	... you were absorbed in doing something worthwhile.
- Bored	... you felt bored.
SMS happiness	Overall, how happy do you feel right now on a scale from 1 (not at all happy) to 10 (completely happy)?
SMS positive emotion	What best describes how you felt over the last 10 minutes? Please text back the corresponding number. [1: Lonely/left out 2: Shameful/guilty 3: Absorbed in doing something worthwhile 4: Sad 5: Loving/tender 6: Bored 7: Happy 8: Angry 9: Worried 10: Other positive feeling 11: Other negative feeling 12: Other neutral feeling]
- SMS loneliness	How lonely are you feeling right now on a scale from 1 (not at all lonely) to 10 (very lonely)?

Social interaction

Friends met in person	List the first names of as many of the friends you met in person last week that you can think of in 1 minute (if none, enter "none"). Separate the names using commas (",").
Offline activities	Which of the following activities did you do at least once last week? Check all that apply Go out for dinner Go to the cinema Talk to friends on the phone Go to a party Get together with friends Go to a shopping mall Spend time with your parents Spend time with your kids
Diverse interactions	Interact with someone who voted the opposite way as you in the last presidential election Interact with someone from another country

<i>Opinions about Facebook</i>	
Improves social life	To what extent do you think Facebook improves or worsens people's social lives?
Good for you	To what extent do you think Facebook is good or bad for you?
Good for society	To what extent do you think Facebook is good or bad for society?
Makes people happy	To what extent do you think using Facebook makes people more or less happy?
- Wouldn't miss it	To what extent do you agree or disagree with the following statement: "If people spent less time on Facebook, they would soon realize that they don't miss it."?
Follow news better	To what extent do you think Facebook helps people follow the news better?
- Clickbait, fake news	To what extent do you think Facebook exposes people to clickbait or false news stories?
- More polarized	To what extent do you think Facebook makes people more or less politically polarized?
- Deactivation good	As part of this study, you were asked to deactivate your Facebook account for [24 hours/4 weeks]. To what extent do you think that deactivating your account was good or bad for you?
Positive impacts	What are the most important positive impact(s) that Facebook has on your life? [text box]
- Negative impacts	What are the most important negative impact(s) that Facebook has on your life? [text box]

<i>Post-experiment use</i>	
Planned use	After going through this study, how much more or less time do you plan to spend on Facebook compared to before you started the study?
- Clicked time limit email	Takes value 1 if clicked on any link in the post-endline social media time limit email, and 0 otherwise
- Time to reactivation	$\ln(1 + \text{number of days deactivated after 24-hour post-endline deactivation period})$
Post-endline mobile use	[if have an iPhone] Please write down the amount of screen time you spent on the Facebook app according to your battery report. [if do not have an iPhone] How many hours would you say you spent on the Facebook app on your phone in the past seven days, in total?

<i>Secondary outcomes</i>	
---------------------------	--

Voted Republican	If the elections for U.S. Congress were being held today, would you vote for the Republican Party's candidate or the Democratic Party's candidate for Congress in your district? [Republican candidate, Democratic candidate, Other/don't know] [If would vote for Republican or Democratic candidate] How convinced are you about whether to vote for the Republican candidate or the Democratic candidate? [slider from 0 to 100]
Voted (self-report)	Did you [midline: Do you plan to] vote in the midterm elections on November 6th, 2018?
Contributions	ln(1+FEC contributions between October 12 and November 10)

Moderators

Time of day	At what times of day do you usually use Facebook the most? [Morning (6AM-12 noon), Afternoon (12 noon-5PM), Evening (5-9PM), Night (9PM-midnight), Late night/early morning (midnight-6AM)]
Active browsing	People talk about two different ways to use Facebook: "Active" users often post status updates, comment on other people's walls and pictures, post photos, etc. "Passive" users mostly check out their news feeds and/or other people's photos and profiles but don't comment or interact much with others on the site. Which would you say describes your Facebook use best? What share of your time on Facebook do you spend interacting one-on-one with people you care about (for example, commenting on their posts or sending them private messages)?
Get news from Facebook	Over the past four weeks, how often did you ... get news from Facebook [Never, Hardly Ever, Sometimes, Fairly Often, Very Often]
Facebook minutes	On an average day in the past 4 weeks, how many minutes would you say you spent on Facebook, including through the Facebook app on your phone?
