## UCLA UCLA Electronic Theses and Dissertations

### Title

Digital Campaign Advertising and Turnout: Evidence about Framing and Dosage from a Randomized Field Experiment

### Permalink

https://escholarship.org/uc/item/88b4w48k

**Author** Bakr, Adam

Publication Date

### Supplemental Material

https://escholarship.org/uc/item/88b4w48k#supplemental

Peer reviewed|Thesis/dissertation

#### UNIVERSITY OF CALIFORNIA

Los Angeles

Digital Campaign Advertising and Turnout:

# Evidence about Framing and Dosage from a Randomized Field Experiment

A dissertation submitted in partial satisfaction of the requirements for the degree Doctor of Philosophy in Political Science

by

Adam Bakr

2025

© Copyright by Adam Bakr

2025

#### ABSTRACT OF THE DISSERTATION

# Digital Campaign Advertising and Turnout: Evidence about Framing and Dosage from a Randomized Field Experiment

by

Adam Bakr

Doctor of Philosophy in Political Science University of California, Los Angeles, 2025 Professor Lynn Vavreck Lewis, Chair

A recent wave of studies on political advertising has demonstrated the small (and sometimes absent) persuasive effects of television and digital advertising in American elections. However, less is known about how the *types* of messages and *dosage* of political advertisements affect people and even less about how dosage interacts with messaging. This dissertation reports the results of an individual-level randomized field experiment on 778,882 registered voters conducted over Facebook during the 2018 Republican gubernatorial primary in Texas. The ads, produced by the incumbent governor's campaign, were designed to increase turnout among likely Republican voters in the state. I find small but discernible increases in voter turnout, with some message strategies resonating more than others. In particular, a social pressure message emphasizing civic duty outperformed a partisan identity appeal. I further demonstrate that among the more effective messages, delivering a higher volume of advertisements yields larger returns in voter turnout (up to about 1.2 percentage points increase for 20 exposures of a social pressure ad). Finally, when considering race as a moderating variable, the effects were not significantly different across racial groups when the advertisements emphasized partian identity instead of racialized policies, suggesting that non-racial appeals had uniformly modest effects across demographics. These findings contribute to our understanding of digital campaign effects by highlighting the roles of message framing and exposure frequency in voter mobilization. The dissertation of Adam Bakr is approved.

Chad J. Hazlett

Daniel Thompson

Efrén Osvaldo Pérez

Lynn Vavreck Lewis, Committee Chair

University of California, Los Angeles

2025

### TABLE OF CONTENTS

List of Figures		
$\mathbf{Li}$	st of	Tables
1	Intr	$\operatorname{oduction}$
<b>2</b>	Lite	rature Review
	2.1	The Evolution of Political Campaign Advertising
	2.2	Advantages and Limitations of Digital Advertising
	2.3	Voter Mobilization and Message Framing
	2.4	Demographic Moderators of Responsiveness
	2.5	Challenges in Measuring Digital Ad Effects
3	Met	hodology $\ldots \ldots 10$
	3.1	Research Design Overview
	3.2	Data and Variables
	3.3	Analytical Strategy
4	Res	earch Design
	4.1	Context: 2018 Texas Primary and Sample 15
	4.2	Treatment Conditions – Message Content
	4.3	Treatment Conditions – Dosage Levels
	4.4	Implementation and Integrity
5	Emj	pirical Findings

	5.1	Turnout Overview	21
	5.2	Main Treatment Effects vs. Control	22
	5.3	Effects by Message Content	26
	5.4	Effects of Ad Dosage	29
	5.5	Interaction Effects Between Message Type and Dosage	32
	5.6	Election Day vs. Early Voting Effects	36
	5.7	Subgroup Analysis: Does Treatment Effect Vary by Demographics?	40
		5.7.1 Race/Ethnicity	41
		5.7.2 Gender	53
		5.7.3 Age	58
6	Dis	cussion	64
	6.1	Revisiting Expectations and Literature	64
	6.2	Implications for Campaign Strategy	66
	6.3	Limitations and Considerations of External Validity	68
	6.4	Directions for Future Research	70
7	Cor	nclusion	72
	7.1	Summary of Findings	72
	7.2	Theoretical Contributions	73
	7.3	Practical Implications	74
	7.4	Limitations and Scope Conditions	75
	7.5	Future Research Directions	76
	7.6	Concluding Reflections	77
$\mathbf{A}$	Арі	pendix A: Supplementary Tables and Robustness Checks	78

vi

В	Appendix	B: Advertisement Transcripts	83
	A.0.3	Regression Diagnostic Plots	80
	A.0.2	Heteroscedasticity Tests	79
	A.0.1	Multicollinearity Diagnostics	78

### LIST OF FIGURES

5.1	Turnout Increase by Dosage	32
5.2	Turnout Effect by Message Type and Dosage	35
5.3	Estimated marginal treatment effects on voter turnout by age group. Effects	
	represent percentage-point changes from baseline turnout	63
A.1	Residuals vs. Fitted Values for OLS Turnout Model. Each point represents an	
	individual voter's residual (actual minus predicted turnout) plotted against	
	the model's fitted turnout probability for that voter. The red dashed line at	
	0 indicates perfect prediction. The spread of residuals is relatively constant	
	around zero, though slightly wider in the mid-range of fitted values, consistent	
	with heteroscedasticity from the binary outcome	80
A.2	Normal Q–Q Plot of OLS Residuals. The quantiles of the residuals from	
	the turnout regression are plotted against theoretical quantiles from a normal	
	distribution. Deviations from the $45^{\circ}$ line indicate departures from normality.	
	The plot shows moderately heavier tails than a normal distribution (points at	
	the extremes lie off the line), reflecting the excess of large residuals inherent	
	in a linear model for a binary outcome. Overall, residuals are approximately	
	normal in the center of the distribution but with some heavy-tail behavior at	
	the extremes.	81

### LIST OF TABLES

5.1	Turnout among Experiment Subjects by Method of Voting, 2018 Texas Primary	21
5.2	Treatment Effects on Voter Turnout for All Participants	25
5.3	T-Test Results – Turnout Differences by Message Content Treatment	27
5.4	T-Test Results – Turnout Differences by Ad Dosage	30
5.5	Treatment Effects on Voter Turnout – Election Day	38
5.6	Treatment Effects on Voter Turnout for White Voters	42
5.7	Treatment Effects on Voter Turnout for Black Voters	45
5.8	Treatment Effects on Voter Turnout for Latino Voters	49
5.9	Treatment Effects on Voter Turnout for AAPI Voters	51
5.10	Treatment Effects on Voter Turnout – Female Voters	54
5.11	Treatment Effects on Voter Turnout – Male Voters (OLS Models) $\ldots$ .	56
5.12	OLS Regression: Seniors (65+) × Message × Dosage $\ldots \ldots \ldots \ldots$	59
5.13	OLS Regression: Older Adults (50–64) × Message × Dosage	60
5.14	OLS Regression: Middle-Aged Adults (30–49) × Message × Dosage	61
5.15	OLS Regression Model: Young Adults (20–29) $\times$ Message $\times$ Dosage Level $% {\rm (Intermation of the set of the $	62
A.1	Variance Inflation Factors for Key Independent Variables	79

#### ACKNOWLEDGMENTS

Completing this dissertation marks not only the culmination of years of academic research and personal growth but also reflects the invaluable support and encouragement of many remarkable individuals. It is an honor and a pleasure to acknowledge those whose guidance, friendship, and unwavering belief have made this milestone achievable.

Foremost, I extend my deepest gratitude to my dissertation chair, Professor Lynn Vavreck. Her tireless support, insightful mentorship, and intellectual rigor have significantly shaped my scholarly trajectory. Beyond guiding my research, Professor Vavreck has profoundly influenced my approach to teaching and pedagogy, enriching my academic identity. Her enthusiasm and encouragement empowered me to explore my research interests with confidence and determination.

I am equally indebted to Professor Efrén Osvaldo Pérez, whose steadfast support and thoughtful advice have continuously propelled my academic development. Efren provided indispensable guidance on implementing moderators in this study and earlier projects, deeply informing my research methodology. Most notably, Efren was the first faculty member to offer me a research opportunity, sparking a passion for inquiry that ultimately culminated in this dissertation.

I am also grateful for the feedback and support from colleagues within UCLA's American Politics Lab, Political Psychology Lab, and Race, Ethnicity, and Politics Reading Group. The insights gained from presenting and discussing my work in these settings have been instrumental in refining my ideas and methodology.

To my dear friend Jenny Hamilton, I extend my heartfelt appreciation. Jenny's consistent encouragement, unwavering belief in my potential, and compassionate friendship have provided me with immense strength, both within academia and beyond. Her support has helped sustain me through moments of uncertainty and reminded me always of my value and contributions.

I also thank my close friend Derek Holliday, whose companionship was especially invalu-

able during the isolating and challenging periods of the COVID-19 pandemic. Derek continually reminded me of the importance of balance, the restorative power of music, sports, and friendship, grounding me in optimism when it was most needed.

To my dear friends and chosen family, Cordell Crismon and Diana Rivera, I owe immense gratitude. Their warmth, encouragement, and steadfast friendship provided crucial emotional sustenance and laughter throughout my graduate studies. Whether ensuring I was well-nourished, keeping me grounded in the realities and joys outside academia, or enthusiastically championing my efforts, their presence was truly indispensable. Additionally, I extend a special acknowledgment to their cats, Albus and Luna, whose affectionate antics, from turning the pages of my books because I was reading too slowly for them or gently insisting on breaks by napping on my keyboard, made them unforgettable study companions.

I also owe special acknowledgment to my dog, Snoop Doggie, whose boundless joy, affection, and companionship significantly enriched my life during the often stressful years of graduate school. Snoop intuitively provided comfort when I felt down and playful distraction when I most needed relief, filling these years with more joy and laughter than I ever imagined possible.

I owe a special thanks also to my in-laws, Shamim and Seema Pirzada, who graciously welcomed me into their home during the past year as my wife, Samah, and I navigated the early stages of marriage. Their unwavering support, hospitality, and heartfelt words of encouragement served as a guiding light, profoundly enriching my life. I feel extraordinarily fortunate to have not just one, but two sets of parents whose kindness and generosity I will always cherish.

I would further like to extend my heartfelt gratitude to my sisters, Kenzie and Alia, who have consistently provided me with a safe space filled with unconditional love, friendship, and unwavering support throughout this journey. Their presence in my life is a daily reminder that no matter the challenges I face, I am never alone. Beyond their support, Kenzie and Alia inspire me profoundly through their own notable successes in both their careers and personal lives, and through their continuous pursuit of new knowledge and skills. Their resilience, curiosity, and determination serve as guiding examples that continually motivate me to strive toward greater achievement and personal growth.

My deepest appreciation and admiration belong to my parents, Medhat and Nagwa Bakr, whose unconditional love, unwavering support, and profound sacrifices have shaped who I am today. To my father, whose steadfast motivation and firm belief in my academic potential continually inspired me toward excellence, both as a student and an educator, I remain sincerely grateful. To my mother, whose boundless encouragement and emotional strength form the very foundation of my character, I dedicate a special acknowledgment. My mother's Alzheimer's diagnosis during my graduate studies has been deeply challenging, knowing that she will neither read this dissertation nor attend my graduation. Yet, her invaluable lessons about belonging, perseverance, dedication, and unconditional love remain vivid and deeply inspiring. Her frequent encouragement—"Go show them what Bakrs are made of!"—still resonates powerfully with me, motivating me always to offer my very best, trusting that the world will reciprocate. While I dearly wish I could tell her, "Mom, I showed them!," I am comforted by the knowledge that she remains proud, wholeheartedly rooting for my progress, success, and happiness from afar.

Finally, words fall short in expressing the immense gratitude and admiration I hold for my wife, Samah Pirzada. Her presence and partnership have fundamentally transformed my life's trajectory in immeasurably positive ways. From the beginning, Samah constantly reminded me of my capabilities, especially during moments of doubt, always believing in my potential even when I struggled to see it myself. Her consistent celebration of milestones, no matter how small—such as insisting we mark the completion of my Master's degree, despite my own reluctance—provided me with the motivation and confidence to persevere.

Samah's dedication and hard work, both professionally through her impactful career in Diversity, Equity, and Inclusion, and personally in the care and support she generously provides to me and our broader family, continue to inspire me daily. She has helped me understand the importance of balancing professional ambition with meaningful family experiences, often reminding me that while work remains constant, our opportunities for adventures and cherished moments are fleeting. Her encouragement to be fully present, whether traveling abroad or enjoying local experiences, has significantly enriched our life together.

Samah, my love and appreciation for you are unwavering. Without your continuous compassion, strength, and belief in me, I would not be where I am today. This dissertation, and every milestone I achieve, is infinitely more meaningful because of your presence and support. For all of this and so much more, thank you from the bottom of my heart.

To all those mentioned here, and the many others whose kindness, friendship, and guidance have supported me throughout this journey, I offer my heartfelt appreciation. Your contributions are forever embedded in this dissertation and in my continued growth as a scholar and a person.

### VITA

2011 B.A. (	Political Science	and History)	, UCLA
-------------	-------------------	--------------	--------

- 2020 M.A. (Political Science), UCLA
- 2021 Ph.D. Candidate (Political Science), UCLA

### CHAPTER 1

### Introduction

The advent and proliferation of digital media technologies have profoundly transformed the landscape of political campaigning, reshaping the strategies, tools, and channels through which candidates and political organizations communicate with voters. Historically, political campaigns relied heavily on broadcast media such as television and radio to disseminate their messages to the electorate. These traditional communication channels, while effective in their era, offered limited scope for personalized messaging and direct voter interaction. In contrast, contemporary campaigns operate within a highly digitized environment characterized by social media platforms such as Facebook, X (formerly Twitter), YouTube, and Instagram, enabling political messages to reach voters with unprecedented precision, personalization, and frequency (Stromer-Galley 2014; Kreiss 2016; West 2013).

The strategic shift toward digital advertising is underscored by its rapid growth in prominence and financial investment over recent election cycles. In the 2020 U.S. presidential election alone, campaigns spent approximately \$1.65 billion on digital advertisements, with platforms like Facebook and Google commanding the majority of these expenditures (Roberts 2020; Project 2021). This expenditure accounted for roughly 24 percent of total campaign advertising budgets, second only to broadcast television, underscoring the centrality of digital media in contemporary electoral strategies (Baum 2021; Williams and Gulati 2013).

Despite the growing prominence of digital advertising, scholarly understanding of its effectiveness in mobilizing voter turnout remains limited and contested. Early academic research in campaign communications primarily concentrated on persuasion—whether messages delivered via traditional media could change voters' candidate preferences or attitudes

(Broockman and Kalla 2018; Coppock, Hill, and Vavreck 2020). These studies generally found modest or negligible effects, prompting scholars to shift attention toward voter mobilization—the capacity of campaigns to encourage voter participation itself (Gerber and Green 2012). Studies exploring mobilization via traditional methods such as canvassing, direct mail, and telephone outreach have demonstrated clear, albeit often modest, effects on turnout (Nickerson 2006; Gerber, Green, and Larimer 2008). However, research examining the mobilization potential of digital campaign advertising remains relatively sparse and yields mixed results (Haenschen and Jennings 2019; Hager 2019; Coppock, Green, and Porter 2022; Collins 2018).

This dissertation seeks to address this gap in the literature through an individual-level randomized field experiment conducted during the 2018 Texas Republican gubernatorial primary election. Collaborating with the incumbent governor's campaign, this study systematically tests whether variations in message framing and exposure frequency—referred to as "dosage"—affect voter turnout among targeted Republican voters identified via voter-file matching on Facebook. The experiment examined four distinct message frames: a social pressure message emphasizing civic duty and accountability, a partisan identity appeal invoking group solidarity and partisan loyalty, a traditional reminder message emphasizing the election date, and a hybrid approach combining identity and social pressure cues. Voters in treatment conditions were randomly assigned to receive varying dosages of these advertisements (10, 15, or 20 exposures). The study thus provides empirical evidence about the relative effectiveness of different message framings and the optimal frequency of digital advertisement exposure necessary to stimulate voter participation.

Theoretically, this dissertation draws from foundational models of voter turnout, notably Downs' (1957) rational-choice calculus of voting, expanded upon by Riker and Ordeshook (1968), who introduced intrinsic psychological and normative motivations such as civic duty into the calculus. The existing literature highlights two particularly effective strategies: social pressure appeals, which leverage voters' adherence to social norms and accountability, and partisan identity appeals, which activate voters' psychological affiliations with political parties and group loyalties (Gerber, Green, and Larimer 2008; Michelson 2005; Brader 2005). Digital platforms amplify the potential impact of these strategies by providing highly granular targeting capabilities, real-time feedback mechanisms, and scalable, cost-effective communication channels (Coleman and Freelon 2015; Kim et al. 2018; Tucker et al. 2018).

While digital advertising offers significant theoretical and practical promise, its effectiveness is not uniform across all voter segments. Prior research suggests that demographic factors—including party identification, race and ethnicity, gender, and age—serve as moderators influencing how individuals perceive and respond to campaign appeals (Ramakrishnan and Baldassare 2004; Burns, Schlozman, and Verba 2001; Eagly 1987; Malhotra, Michelson, and Valenzuela 2012). This dissertation explicitly evaluates these potential moderating effects, providing a nuanced understanding of whether and how message efficacy and dosage sensitivity vary across different demographic groups.

Conducting randomized field experiments in digital environments introduces unique methodological complexities, including challenges of dosage compliance, platform opacity, and interference effects (Bond et al. 2012a; Gerber and Green 2012). For example, online users' engagement with digital content can differ dramatically, complicating efforts to precisely measure exposure and treatment fidelity. Moreover, digital platforms continuously adjust their content delivery algorithms, further complicating attribution and generalization of findings over time (Tucker et al. 2018). Recognizing these challenges, the present study integrates methodological rigor with practical relevance, using verified voter turnout data rather than intermediate digital engagement metrics to robustly evaluate the efficacy of digital advertising in real-world electoral contexts.

In summary, this dissertation advances the literature by providing robust causal evidence about digital advertising's mobilization potential and the conditions under which digital messages effectively influence electoral behavior. By systematically testing message framing and dosage in a large-scale randomized field experiment, the research offers actionable insights for campaigns seeking to optimize digital strategies and contributes substantively to theoretical debates surrounding voter mobilization and the behavioral impact of digital media. The following chapters thoroughly elaborate on the literature, methodology, empirical findings, and implications of these results for both scholarship and electoral practice.

### CHAPTER 2

### Literature Review

The emergence of digital technologies has transformed political campaign communication, reshaping the strategies and channels that candidates use to engage voters. The proliferation of platforms such as Facebook, Twitter (now X), YouTube, and Instagram provides campaigns unprecedented opportunities for continuous, personalized voter engagement, far surpassing traditional broadcast methods (Stromer-Galley 2014; Kreiss 2016). These platforms enable targeted political advertising at scale, real-time adjustments based on voter interactions, and detailed audience analytics, marking digital advertising as a distinct component of modern electoral strategy. While early digital studies focused on persuasion, recent scholarship has increasingly turned toward mobilization—examining whether and under what conditions digital ads can effectively increase voter turnout (Bond et al. 2012b; Haenschen and Jennings 2019). This dissertation contributes to this literature by investigating how message framing and ad frequency (dosage) affect turnout in a real-world election experiment conducted via Facebook during the 2018 Texas Republican gubernatorial primary.

This chapter reviews the theoretical and empirical foundations guiding this inquiry. Section 2.1 traces the evolution from broadcast media to the current digital paradigm, emphasizing key shifts in strategy and message delivery. Section 2.2 discusses digital advertising's distinct advantages and inherent limitations. Section 2.3 connects digital advertising to foundational theories of voter mobilization and examines previous findings regarding effective GOTV message framing. Section 2.4 reviews demographic moderators influencing responsiveness to campaign messages. Finally, Section 2.5 addresses methodological challenges in evaluating digital advertising's causal effects, particularly concerning dosage and platform dynamics.

#### 2.1 The Evolution of Political Campaign Advertising

Political advertising has always reflected prevailing communication technologies. Early campaigns leveraged mass media (newspapers, radio, television) primarily for one-way communication, prioritizing broad reach but lacking interactivity and personalization. Televised advertising dominated from the 1960s through the 1990s, characterized by highly choreographed messages aimed at mass audiences (West 2013). Digital media's rise marked a qualitative shift; platforms like Facebook allowed campaigns to engage voters continuously, personally, and interactively (Stromer-Galley 2014; Kreiss 2016). Obama's 2008 presidential campaign pioneered data-driven online strategies for voter outreach, setting a new precedent (Kreiss 2016). By 2016, digital methods were integral to campaign infrastructure, exemplified by the Trump campaign's extensive use of Facebook's micro-targeting capabilities (Boghani 2018). By 2020, digital campaign expenditures exceeded \$1.5 billion, solidifying digital media as an essential electoral battleground (Roberts 2020; Project 2021).

#### 2.2 Advantages and Limitations of Digital Advertising

Digital political advertising offers distinct advantages including cost-effectiveness, scalability, precise targeting, personalization, and interactive engagement. Platforms employ auctionbased pricing models enabling tight budget control and rapid scalability, allowing campaigns flexibility in testing and adjusting strategies (Coleman and Freelon 2015).

Precision targeting allows campaigns to segment voters based on detailed demographics and behaviors, delivering personalized content that research suggests enhances message reception and engagement (Kim et al. 2018). However, precise targeting raises concerns about reinforcing polarization and echo chambers, and risks potential voter backlash if messages are misdirected or overly aggressive (Tucker et al. 2018; Cadwalladr 2018). Digital ads also facilitate real-time engagement, offering immediate voter feedback through likes, shares, and comments. Campaigns utilize this feedback loop, continuously refining messages—a practice termed "analytic activism" (Karpf 2012). Yet, high online engagement does not always correlate with actual voter turnout, necessitating careful interpretation of these metrics (Anspach 2017).

#### 2.3 Voter Mobilization and Message Framing

Digital ads must be understood within broader voter mobilization frameworks. Classic rational-choice theories emphasize psychological incentives like civic duty ("D-term") as critical motivators given the negligible individual instrumental benefit of voting (Downs 1957b; Riker and Ordeshook 1968). Field experiments consistently find personal GOTV contacts, such as door-to-door canvassing, most effective, while impersonal methods (mail, digital reminders) yield modest results (Gerber and Green 2012; Green and Gerber 2015). Nonetheless, digital platforms offer affordable scalability compared to personal contacts.

Among GOTV message strategies, *social pressure* (emphasizing civic norms and accountability) has proven highly effective, leveraging voters' intrinsic motivation to comply with social expectations (Gerber, Green, and Larimer 2008; Green and Gerber 2015). In contrast, *social identity* appeals, emphasizing group loyalty or partisan stakes, yield mixed mobilization outcomes, potentially due to voter fatigue or lower relevance in primary contexts (Panagopoulos 2010). This dissertation experimentally compares these frames within a digital advertising context, examining effects across varying exposure frequencies (dosages)—an understudied aspect of digital mobilization.

#### 2.4 Demographic Moderators of Responsiveness

Demographic characteristics—party identification, race, gender, and age—significantly moderate responses to campaign communication. Party identification strongly influences receptivity, with partisans inclined toward messages reinforcing in-group loyalties, though risk of backlash from misdirected partisan cues remains (Arceneaux and Kolodny 2009; Campbell et al. 1980).

Race and ethnicity condition responsiveness through shared experiences, community consciousness, and historical marginalization. Ethnically tailored GOTV strategies can enhance mobilization among minority voters, though generic messages still moderately mobilize across racial groups (Michelson 2005; Enos, Fowler, and Vavreck 2014; Ramakrishnan and Baldassare 2004). Gender also moderates responses, potentially influencing message effectiveness based on differing socialization experiences (Burns, Schlozman, and Verba 2001; Eagly 1987). Younger voters, though traditionally less engaged, can respond positively to digital strategies due to their digital-native habits, but often require stronger, repeated stimuli due to high content saturation (Haenschen and Jennings 2019; Malhotra, Michelson, and Valenzuela 2012).

#### 2.5 Challenges in Measuring Digital Ad Effects

Evaluating digital advertising effects involves significant methodological challenges. Content saturation and voter fatigue complicate determining optimal exposure frequency. Digital engagement metrics—clicks, views—often poorly correlate with actual turnout, requiring rigorous validation through official voter records to ensure accurate measurement of true mobilization effects (Anspach 2017; Hersh 2020).

Digital field experiments face practical hurdles in treatment assignment and compliance; intended exposures often differ from actual exposures due to platform algorithms, reducing precision in causal attribution (Gerber and Green 2012). Additionally, rapidly evolving platform policies and algorithms limit study replicability and generalizability over time, highlighting the importance of theoretical grounding to ensure lasting insights despite changing technological contexts (Tucker et al. 2018).

This dissertation addresses these challenges using a rigorous randomized field experi-

ment design, employing validated turnout measures, clearly defined treatment dosages, and careful moderation analyses to offer robust conclusions about digital political advertising's mobilization potential.

### CHAPTER 3

### Methodology

#### 3.1 Research Design Overview

To test the hypotheses and questions derived above, I implemented a randomized controlled trial (RCT) in the field. Field experiments are considered a gold standard for causal inference in voter mobilization research. By randomly assigning voters to different advertising conditions (or a control condition with no ads), we can ensure that any differences in turnout observed between groups can be attributed to the exposure to the ads, rather than preexisting differences between voters. This methodology follows in the tradition of large-scale voter experiments conducted by Gerber and Green (2012) and others, but applies it to the domain of digital advertising.

The study was carried out during the lead-up to the 2018 Republican gubernatorial primary in Texas. This context was selected for several reasons. First, primaries in Texas are open, meaning voters choose which party's primary to participate in on election day, and turnout tends to be moderate (not as high as a general presidential election, but high enough to measure effects on turnout). Second, the campaign of the incumbent governor was actively interested in boosting Republican turnout and was willing to collaborate by providing ads and allowing random assignment of a subset of their digital outreach. This partnership ensured ecological validity—these were real campaign ads delivered in a real election.

All experimental procedures were approved and complied with relevant ethical guidelines for research. Voter data (such as demographics and past vote history) were provided by the campaign and supplemented with publicly available voter file information. Randomization and outcome measurement were conducted by the researcher independently to maintain integrity; the campaign did not selectively choose who saw which ads beyond the random assignment provided.

#### 3.2 Data and Variables

The primary outcome of interest is whether an individual voted in the 2018 primary election (any method of voting). This turnout data was obtained from the official state voter file after the election, which records whether each registered voter cast a ballot and by what method (early in-person, mail, election day in-person, etc.). From this, I constructed a binary variable *Voted* (1 if the person voted, 0 if not). Additionally, I have an indicator for voting in the Republican primary specifically (versus the Democratic primary), since in Texas an individual can choose either primary. In our analysis, any vote (Republican or Democratic primary) counts as turnout, but we also examine Republican primary voting as the focal behavior the campaign cared about. The vast majority of voters who cast ballots from our sample did so in the Republican primary (about 93% of voters in the sample, with 7% opting for the Democratic primary).

Key independent variables include:

- Treatment Group: A categorical variable with multiple levels corresponding to each experimental condition (including Control). This captures both message content and dosage, e.g., Social Pressure 10, Social Pressure 15, Social Pressure 20, Social Identity 10, etc., as well as a Control group. In total, there were 1 (control) + 5 (message conditions) × 3 (dosages) = 16 groups.
- Message Type: A factor indicating the content of the ad(s) shown: Social Pressure, Traditional GOTV, Social Pressure & Traditional (combined/mixed content), Social Identity, Social Identity & Social Pressure (combined identity and pressure content).

Control has no message.

- **Dosage**: A numeric variable for the intended number of ad impressions: 0 (control), 10, 15, or 20. These represent approximate exposure counts per person in the treatment period.
- **Demographics**: Variables for each individual's race/ethnicity (categorized generally as White, Black, Hispanic/Latino, or AAPI), gender (Male, Female, with a very small number listed as unknown or other, which were excluded or categorized as other), and age (as of the election, in years).
- Voting History: Although randomization handles selection bias, we record each individual's prior general election turnout history as context. The sample was drawn from voters with a recent history of Republican primary voting or likely Republican partisans, so by design these individuals have above-average baseline turnout propensities.

One important note is that the randomization was stratified or blocked on certain variables to ensure balance. In particular, assignment to treatments was roughly balanced across major race groups and regions of the state. This was to guard against any chance imbalances given the large sample size.

#### **3.3** Analytical Strategy

The experimental analysis is straightforward: compare turnout rates between groups. The primary estimands are the Average Treatment Effects (ATEs) of each treatment condition relative to control. Because the sample size is very large (over 778,000), differences of even a few tenths of a percentage point can be estimated with statistical significance. We rely on ordinary least squares (OLS) regression for most estimates, modeling the turnout outcome as a function of treatment indicators. Although turnout is binary, OLS provides a straightforward interpretation in terms of percentage-point differences, and with large samples, OLS and logistic regression tend to yield very similar inferences.

We estimate several regression models:

- 1. A model with a single indicator for *Any Treatment* (pooling all those who saw ads vs. the control group). This provides an overall sense of whether being exposed to *any* ad affected turnout.
- 2. Models with indicators for each *Message Type*, ignoring dosage (pooling across dosages within a message category). This tests which content was more effective on average.
- 3. Models with indicators for each *Dosage* level (10, 15, 20 exposures, pooling across content). This tests dosage effects irrespective of content.
- 4. Fully interacted models (factorial design) with indicators for each specific combination of message type and dosage (e.g., Social Pressure 10, Social Pressure 15, ... Social Identity & Social Pressure 20, etc.), using the Control group as the omitted baseline. This yields an estimate for each treatment cell's effect.
- 5. Interaction models to test moderation, for example including treatment indicators interacted with a demographic indicator, to see if treatment effects differ by subgroup (e.g., adding treatment \* Black to see if Black voters responded differently than White voters).
- 6. Separate regressions by subgroup (e.g., running the main models on only female voters, only male voters, only white voters, etc.) to directly observe treatment effects within those groups.
- 7. Simple difference-in-means tests (t-tests) between selected pairs of groups (particularly to compare message types and dosages pairwise).

Standard errors in all regressions are robust to heteroskedasticity. Clustering was not necessary since randomization was individual-level and there was no hierarchical grouping in treatment delivery (ads were delivered individually via Facebook's system). However, as a precaution, we also checked that results were similar when clustering by zip code (to account for any localized turnout shocks like a concurrent election or weather event) and found virtually identical standard errors given the random geographic distribution of treatments.

All hypothesis tests use conventional significance levels (\*p < 0.1, \*\*p < 0.05, \*\*\*p < 0.01). Given the large sample, the emphasis is on substantive significance as well—i.e., is an effect of 0.5 percentage points meaningful in context? We interpret effects in terms of percentage point changes in turnout.

Having described the methodology, the next chapter details the specific context of the study—the experimental setting, treatments, and implementation on Facebook.

### CHAPTER 4

### **Research Design**

#### 4.1 Context: 2018 Texas Primary and Sample

The field experiment took place during the run-up to the March 6, 2018 Texas primary, focusing on the Republican gubernatorial primary. While the incumbent Governor was a heavy favorite in his primary, his campaign was interested in driving up Republican turnout to demonstrate enthusiasm and also to potentially benefit down-ballot primary races. Texas primaries typically see turnout in the range of 10-20% of registered voters in midterm years; however, our sample was a targeted subset more likely to vote (because it was drawn from prior GOP voters and modeled Republicans). Indeed, in our control group we observe a turnout rate of about 40% (as we'll see in the results), which is higher than the statewide average, reflecting the campaign's targeting strategy.

The sample of 778,882 individuals was drawn from the campaign's voter contact list on Facebook. These were matched voter file records that Facebook could target via custom audiences (using identifiers like names, emails, phone numbers, etc. to match to Facebook accounts). All individuals in the sample were intended to be eligible voters residing in Texas. Most had voted in a recent Republican primary or were identified by the campaign's data analytics as likely Republicans (in a state with no party registration, campaigns use statistical models and past primary participation as proxies for party affiliation).

Random assignment was done at the individual level: each person was randomly assigned to one of the 16 groups (15 treatment combinations or the control) with equal probability. The control group (approximately 48,674 people, about 6.25% of the sample) did not receive any ads from this campaign during the experiment. The treated groups (the remaining  $\sim 730,208$  people) each received a specific set of ads and frequency.

The experiment was fielded over roughly a three-week period leading up to Election Day (March 6). Early voting in Texas ran from Feb 20 to March 2, meaning a significant portion of votes would be cast before Election Day. The ads ran throughout this period. Importantly, some voters in our sample likely cast their ballots during early voting; they would still have been shown ads until they voted (since the campaign did not have real-time data on who had already voted early). This means a voter could potentially have voted on, say, Feb 27 but still seen ads on March 1. Such cases should be randomly spread across groups and thus not bias comparisons, though it could dilute measurable effects (because an ad cannot mobilize someone who has already voted).

#### 4.2 Treatment Conditions – Message Content

The campaign produced a set of video advertisements (each about 15 seconds long) for use in this experiment. We categorized them into a few messaging strategies:

- 1. Social Pressure (Ad 1): Emphasized that a study was observing who does or does not vote, implying public knowledge of one's voting behavior. It used a stern "civic duty" appeal and highlighted that voting records are public. This ad was designed to induce a feeling of accountability and social pressure to vote.
- 2. Traditional GOTV (Ad 2): A straightforward, positive reminder to vote, featuring Texas First Lady Cecilia Abbott encouraging people to vote on Election Day. The tone was informative and upbeat (e.g., "Don't miss your chance to vote. Find your voting location and cast your ballot.") without any element of social pressure or negativity.
- 3. Social Pressure & Traditional Combined (Ad 5): A mix of the above two approaches effectively showing both the social pressure message and a traditional reminder in the course of the treatment. In practice, voters assigned to this condition saw

a rotation of Advertisement 1 and Advertisement 2 (the pressure ad and the reminder ad) across their impressions. The messaging combined the accountability aspect with a helpful reminder.

- 4. Social Identity (Ad 4): Framed voting in terms of group competition. This ad depicted a conversation where one person expresses frustration about the state of the country, and another responds with a partisan call to action: "They're energized and they're voting... We need to go vote so they don't take over.". Without explicitly naming parties, it clearly implied Republicans ("we") versus Democrats ("they"). The appeal was to in-group loyalty and a fear of the out-group winning due to higher turnout.
- 5. Social Identity & Social Pressure (Hybrid) (Ad 3): A combined approach where elements of the identity framing were mixed with a social pressure message. For instance, a line in the ad might chastise someone for not voting ("You haven't voted in years... our neighbors could know via public records.") thus injecting a mild social pressure, while also invoking the idea of "our side vs. their side". This hybrid was intended to see if blending partisan identity with pressure yields any synergy or if it muddles the message.

These five messaging strategies (pure pressure, pure reminder, pressure+reminder, pure identity, identity+pressure) correspond to the message factor in our design. Each treated subject was assigned to see only one of these message types (or two ads alternating, in the hybrid cases) throughout the experiment.

### 4.3 Treatment Conditions – Dosage Levels

Within each message condition, subjects were further assigned to a dosage: 10, 15, or 20 impressions (ad exposures) on average. In practical terms, the campaign set up the Facebook advertising such that each user in, say, the *Social Pressure 15* group would be targeted

enough times to achieve about 15 impressions of the ad by the end of early voting. Not every user sees exactly the same number of ads due to how Facebook's pacing and delivery work, but the averages were close to those targets. The cost per impression was low enough that reaching 20 impressions per user over a few weeks was feasible within the campaign's budget.

The rationale for these dosage levels was to test a range from a modest number of exposures (10) to a relatively high number (20). Prior campaign experience suggested that fewer than 10 exposures might not register with voters (the message could be missed or forgotten), while more than 20 might be overkill for the timeframe. We also wanted to see if there's an inflection point in the 10-20 range.

The control group received 0 impressions of any of these campaign ads. It's worth noting they were not prevented from seeing other political content on Facebook (for instance, other campaigns' ads or organic political posts), only that this particular campaign did not show them the governor's ads. Thus, control reflects a status quo baseline of typical political exposure minus the governor's ads.

#### 4.4 Implementation and Integrity

The campaign's digital team uploaded the custom audience of our sample to Facebook's Ads Manager and created separate ad campaigns corresponding to each experimental group. Each campaign used the same overall audience but with different creative (ad content) and frequency settings. Random assignment was done outside of Facebook; effectively, individuals were tagged with codes that assigned them to experimental groups, and those tags determined which campaign bucket they fell into. The campaign then delivered the ads accordingly. As a result, randomization was not done by Facebook's system but by our own procedure, ensuring true random assignment.

Throughout the early voting period and up to Election Day, the campaign monitored delivery to make sure the intended number of impressions were being served. There was

some natural variance (some people might log in less often and see slightly fewer ads, some more), but follow-up data from Facebook indicated that the median impressions for the 10group was about 9, for the 15-group about fourteen, and for the 20-group about nineteen, confirming a good separation between conditions.

Because this was a live campaign environment, there was always a small risk of contamination: for example, if a control group member followed the Governor's Facebook page, they might see a posted video of the ad in their feed organically. However, such instances would be rare and, importantly, not systematically different across groups. Similarly, if people shared the ads, theoretically a control person could see it via a friend's share. These spillovers are possible but were assumed to be minimal given the scale and the fact that the ads were not especially viral content.

After the election, official turnout data were matched back to the experimental subjects. The primary turnout (and which primary they chose) is public record in Texas. This matching was done with high accuracy using voter IDs.

In summary, the research design achieved a clean implementation of a multi-factor field experiment in a real election, with the campaign environment lending authenticity to the treatments. With this setup, we can now analyze how the different advertisements and dosages influenced voter turnout.
# CHAPTER 5

## **Empirical Findings**

This chapter presents the empirical results of the randomized controlled field experiment conducted during the 2018 Texas Republican gubernatorial primary. The analyses reported herein address the central question of whether, and to what extent, targeted digital advertisements delivered via Facebook influenced individual-level voter turnout. The chapter proceeds in a structured fashion. It begins by establishing baseline patterns of political participation in the experimental sample, offering a descriptive overview of turnout rates and methods of voting across all 778,882 participants.

Subsequent sections investigate the effects of the experimental treatments on turnout, beginning with a comparison between treated and control groups. These analyses examine whether exposure to any form of advertisement increased the likelihood of voting, and further explore whether the nature of the message (e.g., social pressure, traditional GOTV, partisan identity) or the number of ad exposures (10, 15, or 20 impressions) conditioned the magnitude of the treatment effect. We also evaluate the interactive effects of message content and dosage to determine whether certain combinations were particularly effective or ineffective.

In addition to average treatment effects, we analyze how the timing of vote casting whether an individual voted early or on Election Day—may have moderated responsiveness to the treatment. Because the advertisement campaign spanned the full early voting window and culminated just before Election Day, treatment effects may have differed based on when voters cast their ballots and how much exposure they had accumulated by that point.

The chapter concludes with a set of subgroup analyses that assess treatment heterogeneity by key demographic characteristics, including race/ethnicity, gender, and age. These analyses enable us to determine whether the intervention produced differential effects across subpopulations and whether such differences were statistically and substantively meaningful.

### 5.1 Turnout Overview

Before analyzing treatment effects, it is important to establish the context in which the experimental interventions were delivered. Table 5.1 provides a summary of voter turnout among all individuals included in the study, disaggregated by method of voting. Of the 778,882 subjects in the sample, 43.61% participated in the primary election, while 56.39% abstained. This overall turnout rate is relatively high for a midterm primary election in Texas, which is consistent with the composition of the sample. Subjects were drawn from a list of individuals either known to have previously voted in Republican primaries or modeled by the campaign's analytics as likely Republican identifiers. As such, the sample was enriched with individuals possessing relatively high baseline propensities to vote.

Among those who voted, approximately 18.28% of the full sample voted early in person, 19.79% voted on Election Day, and 5.51% submitted absentee (mail-in) ballots. A negligible fraction of individuals cast provisional ballots, of which only 0.02% were accepted and 0.01% rejected.

Vote Method	Vote Count	Percent of Sample
Absentee (Mail)	42,943	5.51%
Early (In-Person)	$142,\!390$	18.28%
Election Day (In-Person)	$154,\!143$	19.79%
No Vote (Did Not Vote)	439,240	56.39%
Provisional – Accepted	122	0.02%
Provisional – Rejected	44	0.01%
Total	778,882	100.0%

Table 5.1: Turnout among Experiment Subjects by Method of Voting, 2018 Texas Primary

*Note:* "No Vote" includes all who did not cast a ballot in this primary. Turnout (voted vs. not) is 43.61% overall among subjects.

These descriptive statistics provide the empirical baseline against which the effects of the treatments must be evaluated. The fact that nearly 44% of the sample voted suggests that the individuals targeted by the campaign were, on average, moderately engaged. At the same time, the 56% who abstained represent a substantial portion of the sample with room for potential mobilization. Importantly, this means that treatment effects would have to be sufficiently strong to move individuals who, although selected for their relatively high vote propensity, remained uncommitted or undecided.

It is also worth noting that because the campaign ads were delivered throughout the early voting period and continued until Election Day, variation in when individuals voted may have influenced their exposure to the treatment. For instance, individuals who voted early may have done so after limited exposure to the ads, or even prior to receiving their full dosage. By contrast, individuals who waited until Election Day may have had more time to be exposed to the assigned treatments, particularly those assigned to higher-frequency conditions. Therefore, in addition to estimating the intent-to-treat (ITT) effect on overall turnout, the analysis that follows will also examine turnout effects disaggregated by voting method.

In summary, this initial turnout overview highlights a context in which a large portion of the sample was reachable and potentially mobilizable. The remaining analyses evaluate whether and how the experimental treatments influenced turnout beyond what would have occurred in the absence of the intervention.

#### 5.2 Main Treatment Effects vs. Control

The most general and foundational research question guiding this study is whether exposure to digital advertisements—specifically, those delivered via Facebook—had a statistically and substantively meaningful effect on voter turnout during the 2018 Texas Republican gubernatorial primary. This question assesses the overall efficacy of the campaign's digital outreach, independent of message type or exposure dosage, by comparing the average turnout rate among all treated individuals to that of the untreated control group.

Model 1 in Table 5.2 presents this estimate. The specification is a simple ordinary least squares (OLS) regression in which the dependent variable is a binary indicator for whether an individual voted, and the independent variable is a dichotomous treatment indicator coded 1 for individuals assigned to any of the advertising conditions (regardless of message content or dosage) and 0 for those in the control group. The estimated coefficient for the treatment indicator is +0.644 percentage points (p < .01), suggesting that, on average, voters who received any treatment were 0.64 percentage points more likely to vote than those who received no ads.

To contextualize this effect, the control group's baseline turnout rate was approximately 40.0%. Thus, the estimated treatment effect reflects a relative increase of approximately 1.6%. While the effect size is modest in absolute terms, it is statistically significant and consistent with prior literature on get-out-the-vote (GOTV) interventions, which frequently find that well-designed treatments—especially those delivered at scale—tend to produce small but measurable gains in turnout (see Gerber, Green, and Larimer (2008); Green and Gerber (2015)). The 95% confidence interval for the estimate ranges from approximately 0.19 to 1.10 percentage points, and the standard error is approximately 0.23, owing to the large sample size of 778,882 observations.

These findings suggest that, even in an information-rich electoral environment characterized by a saturated media landscape, a well-targeted digital ad campaign can exert a measurable mobilizing effect. It is noteworthy that this result emerges despite the relatively low intensity of the intervention: each treated individual received between 10 and 20 ad impressions over a period of several weeks. Given the brevity of the video ads (approximately 15 seconds each) and the impersonal nature of the delivery medium, the estimated effect of 0.64 percentage points provides meaningful evidence that digital mobilization efforts, even when modest in scope, can contribute to measurable increases in political participation.

Nonetheless, it is important to emphasize that the average effect masks potential variation across types of messages and dosages. The digital ad campaign tested five distinct message framings (e.g., social pressure, traditional GOTV, partisan identity) and three dosage levels (10, 15, and 20 ad impressions). Accordingly, the next stage of the analysis unpacks this heterogeneity by estimating treatment effects separately by message type (Model 2), by dosage level (Model 3), and by the full set of message-by-dosage combinations (Model 4). These models help identify which specific combinations of message content and exposure frequency were most effective at increasing turnout and whether there were interaction effects between these two dimensions of the intervention.

	Model 1 (Any Treat)	Model 2 (Message)	Model 3 (Dosage)	$\begin{array}{c} \textbf{Model 4} \\ (\text{Message} \times \text{Dosage}) \end{array}$
Any Treatment	$0.644^{***}$ (0.230)	_	_	
Social Pressure		$0.867^{***}$ (0.257)	_	_
Traditional GOTV	-	0.709***	-	-
	-	(0.257)	-	—
Social Pressure & Traditional	_	$(0.915^{***})$	_	_
Social Identity	_	(0.201) $0.426^{*}$	_	-
	_	(0.257)	-	—
Social Identity & Social Pressure	_	(0.302)	_	_
10 Impressions	_	(0.201)	0.657***	_
10 1.0000000	_	_	(0.244)	_
15 Impressions	_	-	0.516**	_
-	_	-	(0.244)	_
20 Impressions	_	-	$0.758^{***}$	-
	_	-	(0.244)	-
Social Pressure 10	_	-	-	$0.572^{*}$
	_	-	-	(0.315)
Social Pressure 15	—	—	—	0.867***
	-	-	-	(0.315)
Social Pressure 20	_	-	-	1.163***
	—	-	-	(0.315)
Traditional 10	-	—	_	0.826***
	—	—	—	(0.315)
Traditional 15	—	-	-	0.618**
	—	—	-	(0.315)
Traditional 20	—	—	—	0.684**
	_	-	-	(0.315)
Social Pressure & Iraditional 10	_	-	-	$(0.053^{**})$
Social Processo & Traditional 15	_	_	_	(0.315)
Social Pressure & Traditional 15	_	_	_	(0.315)
Social Pressure & Traditional 20	_	_	_	1 161***
Social Pressare C Thaillonal 20	_	_	_	(0.315)
Social Identity 10	_	_	_	0 714**
Social Including 10	_	_	_	(0.315)
Social Identity 15	_	_	_	0.074
	_	_	_	(0.315)
Social Identity 20	_	_	_	0.489
	—	-	-	(0.315)
Social Identity & Social Pressure 10	_	-	-	$0.522^{*}$
	_	—	_	(0.315)
Social Identity & Social Pressure 15	_	-	-	0.089
	_	—	—	(0.315)
Social Identity & Social Pressure 20	_	—	—	0.295
	_	_	-	(0.315)
Intercept (Control Mean)	40.038***	40.038***	40.038***	40.038***
• • • /	(0.223)	(0.223)	(0.223)	(0.223)
N (subjects)	778,882	778,882	778,882	778,882

## Table 5.2: Treatment Effects on Voter Turnout for All Participants

p < .1; p < .05; p < .01

*Note:* OLS coefficients represent percentage-point changes in turnout relative to the control group mean of 40.04%. Models 1–4 estimate effects by treatment status, message frame, dosage, and message-by-dosage combinations. Positive and statistically significant effects appear for social pressure, traditional GOTV, and hybrid messages, particularly at higher dosage levels. Ads were uniformly distributed and not demographically targeted. Robust standard errors appear in parentheses.

As Table 5.2 shows, the analysis proceeds by disaggregating the average treatment effect along two key dimensions—message content and dosage—before examining their interaction in a fully saturated factorial model. The following sections unpack each of these dimensions in greater detail, beginning with message content (Model 2), followed by dosage level (Model 3), and then the combined interaction model (Model 4), to identify which elements of the digital ad campaign were most effective in increasing turnout among Republican primary voters.

### 5.3 Effects by Message Content

Model 2 in Table 5.2 provides a disaggregated analysis of the treatment effects by message type, averaging across dosage levels. In this specification, each coefficient represents the average treatment effect of a particular message frame—regardless of dosage—relative to the control group. The results indicate that all five tested message conditions produced positive turnout effects, although the magnitude and statistical significance of these effects varied across message types:

- Social Pressure Message: +0.867 percentage points (p < .01)
- Traditional GOTV Message: +0.709 percentage points (p < .01)
- Combined Social Pressure & Traditional Message: +0.915 percentage points (p < .01)
- Social Identity Message: +0.426 percentage points (p < .10)
- Combined Social Identity & Social Pressure Message: +0.302 percentage points
   (p = .20)

These results strongly support the hypothesis that normative appeals emphasizing civic obligation and accountability (i.e., social pressure messages) are more effective than identitybased appeals in mobilizing voter turnout. The standalone social pressure message yielded a nearly 0.87-point increase in turnout relative to the control group and was the most effective individual treatment condition when controlling for dosage. Notably, the combination of social pressure with a traditional reminder also produced a robust effect (+0.915), suggesting that integrating complementary frames—civic duty and informational nudges—may enhance turnout effects.

In contrast, the social identity message yielded only a marginally significant increase in turnout (+0.426), and the hybrid identity-pressure message was not statistically distinguishable from zero. These results imply that the partisan in-group vs. out-group framing used in the identity message did not effectively motivate voter participation in this context. This underperformance could stem from the relative weakness of the identity framing in the absence of a salient political threat or from the possibility that the tone of the message—implicitly invoking partisan conflict—failed to resonate with the targeted audience.

To test whether these differences between message types were statistically significant, we conducted a series of pairwise post-hoc t-tests comparing average turnout across message conditions (pooled across dosages). The most relevant comparison—between the Social Pressure and Social Identity messages—revealed a statistically significant difference of 0.44 percentage points (p < .05). That is, the average voter exposed to the social pressure ad was approximately 0.44 points more likely to vote than a voter exposed to the identity message. Differences between the Social Pressure and Traditional messages, as well as between Social Identity and Traditional, were not statistically significant at conventional levels, indicating that the traditional GOTV message performed at a level intermediate between the two.

Tab	le $5.3$ :	T-Test	Results –	Turnout	Differences	by	Message	Content	Treatment
-----	------------	--------	-----------	---------	-------------	----	---------	---------	-----------

Treatment Group 1	Treatment Group 2	Diff. in Turnout	t-stat	95% CI Low	95% CI High
Social Identity & Social Pressure	0.44**	2.43	0.09	0.80	
Social Identity & Traditional	0.28	1.56	-0.07	0.64	
Social Pressure & Traditional	-0.16	-0.87	-0.51	0.20	

Note: \*\*p < 0.05. Positive differences represent higher turnout in Group 1 relative to Group 2. For instance, the first row indicates that individuals exposed to the Social Pressure message had turnout rates that were 0.44 percentage points higher, on average, than those exposed to the Social Identity message.

These findings yield several interpretive insights. First, they reinforce the consistent

advantage of social pressure appeals in voter mobilization campaigns, as documented in prior studies such as Gerber, Green, and Larimer (2008). Messages that emphasize civic norms and social accountability appear to trigger a psychological response that enhances voting behavior, even in low-salience electoral contexts. In contrast, messages grounded in partisan identity—particularly those that evoke out-group threat or competition—may require a more specific or emotionally salient political environment to be effective. The partisan identity message in this experiment, which warned that "they're energized and voting" and framed the act of voting as a means of preventing the other side from "taking over," may have lacked the immediate relevance or perceived urgency needed to compel action.

Second, the relative ineffectiveness of the combined Social Identity & Social Pressure message merits closer examination. Rather than reinforcing one another, the dual framing may have undermined the motivational force of each component. The social pressure element relies on internalized civic norms and the desire to conform to expectations of responsible behavior. The identity element, by contrast, appeals to group-based loyalty and out-group antagonism. When combined, these cues may introduce cognitive dissonance or dilute the clarity of the message, reducing its impact. In this case, the hybrid condition underperformed both the pure identity and pure pressure messages, though not significantly.

In summary, these results underscore that **message content plays a critical role in shaping the effectiveness of digital political advertisements**. In this experimental setting, normative appeals to civic duty and accountability—especially when conveyed through the lens of social pressure—proved more effective at increasing turnout than appeals rooted in partisan identity. While traditional reminder messages also generated significant, albeit smaller, effects, identity-based messaging failed to produce consistently meaningful gains and may, in some configurations, diminish the clarity or credibility of the call to action. These findings have direct implications for the design of digital GOTV strategies and suggest that campaigns should prioritize psychologically grounded and positively framed appeals over confrontational or overly politicized messaging.

### 5.4 Effects of Ad Dosage

We now turn to an analysis of treatment dosage—specifically, whether the number of times a voter was exposed to a political advertisement influenced the likelihood of turnout. The theoretical premise for exploring dosage effects is rooted in the concept of message salience and repetition: increased exposure may reinforce key messages, enhance memory retention, or elevate perceived importance, thereby increasing the probability of behavioral change. However, excessive repetition may also produce diminishing returns or trigger disengagement due to ad fatigue or perceived intrusiveness.

Model 3 in Table 5.2 estimates the impact of three distinct dosage levels—10, 15, and 20 impressions—on voter turnout, relative to the untreated control group. Each dosage indicator represents the assigned average number of ad exposures during the campaign period. The results are as follows:

- 10 Impressions: +0.657 percentage points (p < .01)
- 15 Impressions: +0.516 percentage points (p < .05)
- 20 Impressions: +0.758 percentage points (p < .01)

All three dosage conditions yielded statistically significant increases in turnout compared to the control group, suggesting that even modest exposure to campaign advertisements on Facebook can positively influence voter participation. However, the magnitude and shape of the relationship between dosage and turnout are more nuanced. Contrary to expectations of a strictly linear or monotonic trend, the 15-impression group exhibited a slightly lower estimated effect (+0.516) than the 10-impression group (+0.657), suggesting a potential nonlinear relationship. The highest dosage level—20 impressions—produced the largest effect, at approximately 0.76 percentage points above control, reinforcing the idea that more exposure can lead to greater mobilization, at least up to this point. To formally assess whether these differences across dosage levels were statistically meaningful, we conducted a series of pairwise t-tests (Table 4). These tests compared average turnout across the three dosage groups:

- 10 vs. 15 impressions
- 10 vs. 20 impressions
- 15 vs. 20 impressions

Table 5.4: T-Test Results – Turnout Differences by Ad Dosage

Dosage Group 1	Dosage Group 2	Diff. in Turnout	t-stat	95% CI Low	95% CI High
10 Impressions	15 Impressions	-0.14	-1.01	-0.41	0.13
10 Impressions	20 Impressions	+0.10	0.72	-0.17	0.38
15 Impressions	20 Impressions	$+0.24^{*}$	1.72	-0.03	0.52

Note: \* indicates p < 0.1. No pairwise differences are statistically significant at the 0.05 level. The 15 vs. 20 comparison is marginally significant at the 0.10 level ( $p \approx 0.086$ ). Positive differences indicate higher turnout in Group 1.

Although none of the differences are statistically significant at conventional thresholds (p < .05), the comparison between 15 and 20 impressions approaches marginal significance at the 10% level, with the higher dosage group exhibiting a 0.24 percentage-point greater turnout rate  $(p \approx 0.086)$ . The difference between 10 and 20 impressions is smaller (+0.10) and not statistically meaningful. Interestingly, the 10 vs. 15 comparison favors the lower dosage, albeit not significantly, reinforcing the suggestion of a mild dip in the mid-range.

Taken together, these findings imply that dosage matters, though the effects are not uniformly increasing across levels. The most consistent inference is that all three dosage levels produced significant and positive effects relative to the control, and that the 20-impression condition yielded the strongest effect. However, the marginal gain from increasing exposure beyond 10 impressions was relatively modest, and not statistically significant in most pairwise comparisons.

One plausible explanation for the slight underperformance of the 15-impression group is the onset of ad fatigue or cognitive saturation. Repeated exposure to the same or similar advertisements may initially increase attention and perceived urgency, but beyond a certain point may lead to habituation or irritation, reducing effectiveness. Alternatively, the observed pattern could simply reflect random variation or sampling noise, particularly since all estimates lie within a fairly narrow band of 0.5 to 0.75 percentage points.

From a campaign strategy perspective, these results suggest that even relatively modest exposure—on the order of 10 impressions—can generate meaningful increases in turnout. Incremental gains from additional exposure appear to be limited, though not negligible. The highest dosage group (20 impressions) yielded the largest effect, indicating that for campaigns with sufficient resources, higher repetition may still be justifiable, particularly in competitive electoral environments where small turnout differences can be decisive.

To visually represent the relationship between dosage and turnout, Figure 5.1 presents the estimated treatment effects by dosage level, with 95% confidence intervals. The figure illustrates a generally positive trajectory from 10 to 20 impressions, with a slight flattening or dip at 15, followed by a rebound at 20. As subsequent sections will demonstrate, this pattern may vary depending on the specific message content, highlighting the importance of analyzing content-dose interactions. Figure 5.1: Turnout Increase (pp) by Ad Dosage (All Messages Pooled). *Note:* Dots indicate point estimates of turnout difference from control for each dosage group, averaged across all message types. Vertical lines represent approximate 95% confidence intervals ( $\pm 1.96$  SE). All dosage conditions yielded statistically significant increases relative to control, though differences between them are not significant at the 95% level.



In sum, the analysis of dosage effects reveals a weakly positive association between the number of exposures and the magnitude of the turnout effect. While the 20-impression condition yielded the highest observed impact, the marginal gains beyond 10 impressions were relatively small and inconsistent across comparisons. These findings provide evidence in support of digital campaign strategies that prioritize at least moderate exposure levels while tempering expectations regarding the benefits of ad saturation. The full implications of these patterns become clearer when analyzed in conjunction with message content, as discussed in the next section.

#### 5.5 Interaction Effects Between Message Type and Dosage

Up to this point, the analysis has considered the effects of message content and advertisement dosage independently. However, theoretical considerations suggest that message content and exposure frequency may interact in important ways. For instance, cognitively demanding or affectively complex messages may require repeated exposures to exert an influence on behavior, whereas straightforward reminders may be effective after only minimal exposure. The factorial structure of this experiment enables an examination of how different message types perform across varying levels of dosage.

Model 4 in Table 5.2 incorporates a full set of indicators for each unique treatment condition, defined by the interaction of message type and dosage. In total, the model estimates the turnout effect of each individual treatment cell relative to the control group. While interpreting every coefficient is impractical in this context, a number of noteworthy patterns emerge that merit discussion.

The most pronounced treatment effects were concentrated in the conditions that combined social pressure content with higher dosage levels:

- Social Pressure (20 exposures): +1.163 percentage points (p < .01)
- Social Pressure & Traditional (20 exposures): +1.161 percentage points (p < .01)</li>
- Social Pressure & Traditional (15 exposures): +0.931 percentage points (p < .01)</li>
- Social Pressure (15 exposures): +0.867 percentage points (p < .01)
- Traditional GOTV (10 exposures): +0.826 percentage points (p < .01)

Conversely, the weakest effects were associated with identity-based treatments, particularly at higher dosage levels:

- Social Identity (15 exposures): +0.074 (not significant)
- Social Identity & Social Pressure (15 exposures): +0.089 (not significant)
- Social Identity (20 exposures): +0.489 (not significant;  $p \approx .20$ )

- Social Identity & Social Pressure (20 exposures): +0.295 (not significant)
- Social Identity (10 exposures): +0.714 (p < .05), the only identity-based condition that achieved statistical significance

These results yield two principal insights:

- 1. The effectiveness of social pressure messages increases with frequency. The turnout effect associated with the social pressure message grew steadily across dosage levels: from +0.572 at 10 exposures (significant at the 10% level), to +0.867 at 15 exposures (p < .01), and ultimately to +1.163 at 20 exposures (p < .01). This suggests that the persuasive mechanism activated by social pressure—namely, a heightened sense of civic obligation or reputational accountability—is reinforced with repeated exposure.
- 2. The traditional GOTV message exhibits a diminishing returns pattern. The 10-exposure version of the traditional reminder ad was effective (+0.826\*\*\*), but additional exposures yielded smaller marginal gains (+0.618 at 15 exposures and +0.684 at 20 exposures). These findings imply that while traditional reminders are useful, their utility plateaus after a certain threshold. In contrast to social pressure messages, the efficacy of standard reminders does not appear to scale with repetition and may even diminish slightly, possibly due to message fatigue or habituation.

The hybrid message combining social pressure and traditional GOTV content also performed well, especially at higher dosages. At 20 exposures, this treatment condition produced a turnout increase nearly identical to that of the pure social pressure condition (+1.161 vs. +1.163). The hybrid was also effective at lower dosages, suggesting that combining normative appeals with positive reinforcement may create a robust mobilization strategy capable of sustaining effectiveness across dosage levels.

In contrast, the performance of the social identity message was consistently underwhelming. While a modest positive effect was observed at 10 exposures (+0.714, p < .05), this impact dissipated at 15 and 20 exposures. The lack of improvement—or even decline—in performance at higher frequencies suggests that identity-based appeals of the sort deployed here may not benefit from repeated exposure. It is possible that the adversarial framing of the identity message ("us vs. them") did not resonate strongly with recipients or that its emotional impact eroded with repetition. The hybrid identity and social pressure treatment was similarly ineffective, indicating that simply layering a sense of accountability onto a weak identity frame does not substantially enhance its mobilizing power.

These findings reinforce the conclusion that **message content and frequency interact meaningfully to shape voter response**. Social pressure messaging, in particular, appears to operate cumulatively, with higher frequencies amplifying its psychological salience. Conversely, traditional and identity-based appeals appear more sensitive to overexposure, with either diminishing or inconsistent returns at higher dosage levels.

Figure 5.2 visualizes these patterns, plotting estimated turnout effects by message type across the three dosage levels.



Figure 5.2: Turnout Effect by Message Type and Dosage

As shown in Figure 5.2:

- The Social Pressure and Social Pressure & Traditional lines display a clearly increasing trend, with the largest effects observed at 20 exposures.
- The **Traditional** line begins high at 10 exposures and declines slightly at 15 and 20, suggesting a saturation point.
- The Social Identity and Social Identity & Social Pressure lines are generally flat, with only the 10-exposure Social Identity condition showing a significant effect.

Taken together, these findings underscore the central conclusion that the effectiveness of digital GOTV ads depends critically on both message content and delivery frequency. Social pressure appeals are uniquely well-suited to high-frequency deployment, consistently producing the largest turnout gains in the study. In contrast, identity-based appeals appear less effective regardless of dosage, and traditional reminders may provide limited marginal returns beyond initial exposure. These results provide important guidance for the design and targeting of digital mobilization campaigns, particularly in the final weeks of an election cycle when message saturation and voter fatigue become salient considerations.

#### 5.6 Election Day vs. Early Voting Effects

Because the Facebook advertisements were delivered throughout the early voting period and up to Election Day, one would expect the treatment effects to be concentrated among voters who had not yet voted early. Early voters, once having cast their ballots, were effectively immune to further mobilization efforts, limiting potential treatment impact among that subset. To investigate this, we restricted the analysis to individuals who had not voted during early voting and assessed whether the treatment increased their likelihood of voting specifically on Election Day. Table 5.5 summarizes these findings.

Among voters who had abstained throughout early voting, Election Day turnout in the control group was roughly 66%. Within this subset, the treatment effects were considerably

larger than those observed in the full sample. Assignment to any advertisement increased the probability of voting on Election Day by approximately 1.13 percentage points (p < 0.01), nearly double the average effect of about 0.64 percentage points observed for overall turnout in the entire sample. This indicates the digital advertisements exerted their strongest mobilizing impact on voters who delayed participation until Election Day.

Table 5.5 further disaggregates these Election Day effects through OLS regression models that separately estimate the influence of message framing, dosage, and their interactions. Across these models, clear patterns emerge:

	Model 1	Model 2	Model 3	Model 4
	(1)	(2)	(3)	(4)
Any Treatment	$1.131^{***}$	_	_	_
a	(0.000)	1.12.0111		
Social Pressure	_	1.426***	_	-
Traditional	_	(0.437) 1.202***	_	_
Traditional	_	$1.392^{(1)}$	_	_
Social Prossura & Traditional	_	(0.437) 1.600***	_	_
Social i ressure & rraditional	_	(0.437)	_	_
Social Identity	_	0.616	_	_
2	_	(0.436)	_	_
Social Identity & Social Pressure	_	0.531	_	_
U U	_	(0.437)	_	_
10 Dose Targets	_	_	1 111***	_
TO Dose Targets	_	_	(0.414)	_
15 Dose Targets	_	_	0.957**	_
10 2000 1005000	_	_	(0.414)	_
20 Dose Targets	_	_	1.325***	_
C C	_	_	(0.414)	_
Social Prossure 10	_	_	_	0.035*
Social i lessure 10	_	_	_	(0.935)
Social Pressure 15	_	_	_	(0.000) 1 440***
Social Pressure 19	_	_	_	(0.535)
Social Pressure 20	_	_	_	1.903***
	_	_	_	(0.535)
Traditional 10	_	_	_	1.586***
	_	_	_	(0.535)
Traditional 15	_	_	_	1.229**
	-	_	-	(0.535)
Traditional 20	_	_	_	$1.362^{**}$
	_	_	-	(0.535)
Social Pressure & Traditional 10	_	—	—	1.404***
	—	_	_	(0.535)
Social Pressure & Iraditional 15	_	_	_	$1.430^{-1.4}$
Social Programs & Traditional 20	_	—	_	(0.000) 0.000***
Social i lessure & fraditional 20	_	_	_	(0.535)
Social Identity 10	_	_	_	$(0.912^{*})$
Social Identity 10	_	_	_	(0.535)
Social Identity 15	_	_	_	0.046
U U	_	_	_	(0.535)
Social Identity 20	_	_	_	$0.890^{*}$
	_	_	_	(0.535)
Social Identity & Social Pressure 10	-	_	_	0.718
	_	_	—	(0.535)
Social Identity & Social Pressure 15	-	_	-	0.631
	_	—	—	(0.535)
Social Identity & Social Pressure 20	-	-	_	(0.525)
	_	—	—	(0.535)
Control Group Mean	$\begin{array}{c} 66.335^{***} \\ (0.378) \end{array}$			
Observations $(N)$	502,063	502,063	502,063	502,063

Table 5.5: Treatment Effects on Voter Turnout – Election Day

p < 0.1; p < 0.05; p < 0.05; p < 0.01

Notes: Table presents OLS regression estimates for treatment effects on voter turnout, limited to voters who cast ballots on Election Day. Estimates and standard errors are multiplied by 100 to facilitate interpretation as percentage-point effects. Model 1 compares any treatment versus control. Model 2 differentiates treatment by message content. Model 3 evaluates treatment effects by dosage targets. Model 4 jointly considers message content and dosage targets. Robust standard errors are provided in parentheses. Cells without values indicate parameters not included in the respective model.

Model 1 shows a significant effect of exposure to any advertisement (b = 1.131, SE = 0.390, p < .01), confirming the general mobilization capability of digital GOTV messages among Election Day voters.

Model 2 reveals substantial variation by message content. Specifically, the social pressure (b = 1.426, SE = 0.437, p < .01), traditional reminder (b = 1.392, SE = 0.437, p < .01), and combined social pressure and traditional messages (b = 1.690, SE = 0.437, p < .01) all had robust positive effects. Conversely, purely identity-based messages (*Social Identity*, b = 0.616, and *Social Identity & Social Pressure*, b = 0.531) yielded smaller, statistically insignificant impacts. These results underscore that normative (social pressure) and informative (traditional reminder) appeals were particularly effective at mobilizing voters who had not voted early.

Model 3 highlights a clear dose-response relationship, indicating increased effectiveness with higher frequency of exposure: 10 (b = 1.111, SE = 0.414, p < .01), 15 (b = 0.957, SE = 0.414, p < .05), and notably, 20 ad impressions (b = 1.325, SE = 0.414, p < .01). Thus, additional exposures continued to drive incremental increases in turnout, with no signs of voter fatigue within the tested range.

Model 4, examining interactions between message content and dosage, provides the most detailed insights. The largest gains in turnout were observed in conditions featuring high-frequency exposure combined with normative or traditional appeals. The Social Pressure  $\mathscr{C}$  Traditional 20 condition produced the largest overall effect (b = 2.228, SE = 0.535, p < .01), followed closely by Social Pressure 20 (b = 1.903, SE = 0.535, p < .01). Purely identity-based conditions performed inconsistently, generating modest or negligible turnout increases even at high dosages, further reinforcing the superiority of normative messaging and election reminders.

Overall, these results demonstrate that the observed digital GOTV mobilization effect was driven predominantly by influencing voters who had not yet cast ballots before Election Day. Early voters, in contrast, showed limited room for additional mobilization, as their decisions had already been finalized by the time they received subsequent ads. This finding aligns logically with campaign targeting strategies, suggesting digital interventions are most productive when directed toward late-deciding or procrastinating voters. The results further reinforce the conclusion that effective mobilization hinges on both carefully tailored message framing (particularly emphasizing civic duty and social accountability) and adequate dosage (frequency of ad exposure), especially in mobilizing voters who remain available through Election Day. These modest but substantively meaningful mobilization effects underscore the strategic value of targeted digital advertising in contemporary electoral campaigns.

# 5.7 Subgroup Analysis: Does Treatment Effect Vary by Demographics?

In the final stage of the analysis, we examine whether the effects of the treatment conditions varied systematically across key demographic subgroups, with particular attention to race, gender, and age. Given the randomized nature of the experimental design and the exceptionally large overall sample size, the study maintains sufficient statistical power to conduct disaggregated analyses within each subgroup. For example, the dataset includes over 650,000 White voters, approximately 32,000 Black voters, roughly 67,000 Hispanic voters, and approximately 27,000 individuals classified as AAPI. These subgroup sizes enable us to estimate treatment effects with reasonable precision for each demographic category.

To assess potential treatment heterogeneity, we employ two complementary approaches. First, we estimate separate ordinary least squares (OLS) regression models within each subgroup to identify within-group responsiveness to the various treatment conditions. Second, we estimate pooled models that include interaction terms between the treatment indicators and subgroup identifiers (e.g., Treatment  $\times$  Black, Treatment  $\times$  Female, etc.) in order to formally test whether differences in treatment effects across groups are statistically significant. This dual strategy allows us to determine both the magnitude of treatment effects within each demographic category and whether these effects differ in meaningful ways across categories.

#### 5.7.1 Race/Ethnicity

Although this study did not design or deliver race-specific messaging, race remains a critical axis for understanding variation in political behavior and responsiveness to campaign interventions. Decades of scholarship in political science have documented the ways in which racial identity, experiences with political institutions, and patterns of group-based socialization shape voting behavior and perceptions of political messaging. Accordingly, this project treats race not as a treatment condition but as a potential moderator of treatment effects. The ads deployed in this experiment were race-neutral in content and uniformly distributed across the entire treatment population, yet differential responsiveness by racial subgroup may still emerge due to baseline differences in turnout propensity, levels of trust in political messaging, or the salience of digital advertisements among different communities. In line with this theoretical framing, the following section disaggregates the treatment effects by race—beginning with White voters—and estimates a set of parallel OLS models to examine whether the impact of message content and ad dosage varied across demographic lines. While the aggregate analysis revealed modest treatment effects overall, these subgroup models provide additional granularity and explore whether digital campaign ads mobilized, suppressed, or left unaffected voters from distinct racial backgrounds.

Table 5.6 presents the results of subgroup-specific OLS regression models estimating the effects of digital advertisements on turnout among White voters (N = 652,471). The analysis follows the same modeling strategy as in earlier sections: Model 1 estimates the average treatment effect of receiving any advertisement; Model 2 includes indicator variables for each message condition; Model 3 includes dosage conditions averaged across messages; and Model 4 introduces full interaction terms between each message condition and dosage level. All outcome measures are scaled to reflect percentage-point changes in turnout relative to the control group, with standard errors shown in parentheses.

Model 1 reveals a negative point estimate for the average treatment effect (b = -1.967, SE = 1.607), suggesting that, on average, exposure to digital advertisements may have

	Model 1	Model 2	Model 3	Model 4
	(Any Ireat)	(Message)	(Dosage)	(Message × Dosage)
Any Treatment	-1.967 (1.607)	_	_	_
Social Pressure		-2.104	_	_
	_	(1.802)	_	_
Traditional GOTV	_	-2433	_	_
	_	(1.804)	_	_
Social Pressure & Traditional	_	-1 409	_	_
	_	(1.795)	_	_
Social Identity	_	-2.932	_	_
	_	(1.800)	_	_
Social Identity & Social Pressure	_	-0.959	_	_
Social Including C Social I locoure	_	(1.803)	_	_
10 Dose Taraets	_	(11000)	$-2.872^{*}$	_
10 D000 10/9000	_	_	(1.708)	_
15 Dose Targets	_	_	-1 244	_
10 D000 10/9000	_	_	(1.709)	_
20 Dose Targets	_	_	-1 781	_
20 Dose Targets	_	_	(1.703)	_
Social Procesure 10	_	_	(11100)	-3.276
Social Pressure 10		_	_	(2, 235)
Social Procesure 15	_	_	_	-0.324
50Ciui 1 105Sure 15	_	_		(2, 228)
Social Processing 20				2.220)
Social Tressure 20				(2.187)
Traditional 10				(2.107)
11aanonai 10	_	_	_	(2.034)
Traditional 15				2.221)
Traditional 15	_	_	_	-3.048 (2.218)
Traditional 00				(2.210)
Traditional 20	_	_	_	(2.215)
Control Decouver 64 The dition of 10				(2.210)
Social Fressure & Truattional 10	_	_	_	-3.338 (2.201)
Social Processor & Traditional 15				(2.201)
Social Pressure & Traditional 15				(2.207)
Social Processing & Traditional 20				(2.221)
Social Pressure & Traditional 20				(2.171)
Social Identity 10				2.171)
Social Identity 10	_	_	_	(2.211)
Social Identity 15				2.211)
Social Identity 15	_	_	_	(2.218)
Social Identity 20				2 127
Social Identity 20	_	_	_	(2.205)
Social Identity & Social Processing 10				(2.200)
Social Identity & Social Tressure 10				(2.222)
Social Identity & Social Processing 15				0.266
Social Identity & Social Tressure 15				(2.222)
Social Identity & Social Processor DO	_	_	_	(2.222)
Social Taentity & Social Fressure 20	_	_	_	(2.213)
	-		_	(2.210)
Control Mean	32.720***	32.720***	32.720***	32.720***
	(1.777)	(1.777)	(1.777)	(1.777)
N (subjects)	652,471	652,471	652,471	652,471

Table 5.6: Treatment Effects on Voter Turnout for White Voters

p < .1; p < .05; p < .01

Note: OLS coefficients represent percentage-point changes in turnout relative to a control group mean of 32.72%. Models 1–4 estimate treatment effects by overall exposure, message frame, dosage level, and message-by-dosage interactions, respectively. All ads were race-neutral. Most estimates are negative and not statistically significant, suggesting limited or potentially demobilizing effects among White voters. Robust standard errors appear in parentheses.

reduced turnout among White voters. However, the estimate is not statistically significant and should be interpreted with caution. While not definitive, this result runs counter to the assumption that digital GOTV messages necessarily produce neutral or positive effects in all subpopulations.

Model 2 disaggregates the treatment effects by message frame. None of the five message conditions produce statistically significant results. Moreover, all point estimates are negative. The *Social Identity* message yields the largest negative effect (b = -2.932, SE = 1.800), followed by *Traditional* (b = -2.433, SE = 1.804), and *Social Pressure* (b = -2.104, SE = 1.802). These estimates suggest that, rather than mobilizing White voters, digital ads—particularly those invoking identity, civic obligation, or even neutral reminders—may have had a demobilizing effect on this group, or at minimum failed to spur additional participation.

Model 3 introduces dosage conditions across all message types and reveals similar trends. All dosage-level estimates are negative, and the effect for 10 Dose Targets is statistically significant at the p < .1 level (b = -2.872, SE = 1.708). Neither 15 Dose Targets (b = -1.244, SE = 1.709) nor 20 Dose Targets (b = -1.781, SE = 1.703) are significant, but their consistent directionality—downward—lends weight to the hypothesis that repeated exposure to these messages may have generated fatigue, annoyance, or disengagement among White voters.

Model 4 presents interaction terms between each message condition and dosage level. Although none of the coefficients reach conventional levels of statistical significance, the consistent direction of the estimates is striking. With only two exceptions (*Social Pressure*  $\mathscr{C}$  Traditional 15, b = 0.305; and Social Identity  $\mathscr{C}$  Social Pressure 15, b = 0.366), every message-dosage combination is associated with a decrease in turnout. Some of the largest negative effects occur at lower dosage levels, such as Social Pressure 10 (b = -3.276, SE = 2.235) and Traditional 15 (b = -3.648, SE = 2.218). The uniformity and magnitude of these negative point estimates, even though individually imprecise, suggest that exposure to these digital ads may have backfired among White voters. This pattern contrasts sharply with findings among other demographic groups (e.g., male voters, Election Day voters), for whom the same messages—especially at higher dosages produced small but consistent turnout gains. The divergence suggests the possibility of conditional or context-sensitive backfire effects. For example, ads emphasizing civic norms or identity may be interpreted differently depending on baseline turnout propensity, partisan alignment, trust in digital media, or perceived relevance of the message source. Alternatively, White voters may have experienced higher levels of message saturation, ad fatigue, or skepticism toward online political content, any of which could mitigate or reverse the intended mobilizing effect.

The analysis of White voters in Table 5.6 reveals little evidence that digital GOTV messaging increased turnout for this group. On the contrary, across all four models, the estimated treatment effects are consistently negative. While these findings are not uniformly statistically significant, their consistency in direction and magnitude raises important questions about heterogeneous treatment effects. These results suggest that campaign strategists should be cautious in assuming uniform efficacy across racial subgroups, and that additional research is needed to understand why digital outreach may suppress rather than activate participation among certain voters.

Table 5.7 presents the results of subgroup-specific OLS regression models estimating the effects of digital GOTV advertisements on turnout among Black voters (N = 31,875). Consistent with earlier sections, Model 1 reports the average treatment effect across all ad exposures, Model 2 introduces message-specific estimates, Model 3 assesses average effects across dosage levels, and Model 4 includes fully interacted message-by-dosage terms. All estimates are expressed in percentage-point changes relative to the control group mean, which is 32.72%, and standard errors are shown in parentheses.

Model 1 reveals a negative and non-significant average treatment effect among Black voters (b = -2.134, SE = 1.943), suggesting that exposure to any digital ad—regardless of message content or frequency—did not increase turnout in this group and may have slightly decreased it. This estimate, while imprecise, raises the possibility that treatment

	Model 1 (Any Treat)	Model 2 (Message)	Model 3 (Dosage)	Model 4 (Message × Dosage)
Any Treatment	-2.134	_	_	-
Any Traimeni	(1.943)	_	_	_
Social Pressure	(	-2.085	_	_
	_	(2.178)	_	_
Traditional GOTV	_	-1.458	_	_
	_	(2.182)	_	_
Social Pressure & Traditional	_	-1.348	_	_
	_	(2.173)	-	-
Social Identity	_	$-4.850^{**}$	-	-
	-	(2.178)	-	-
Social Identity & Social Pressure	_	-0.936	—	_
	-	(2.180)	—	-
10 Dose Targets	-	-	-3.153	-
	—	_	(2.066)	_
15 Dose Targets	—	-	-1.218	—
	-	-	(2.066)	-
20 Dose Targets	-	-	-2.029	-
	-	—	(2.061)	-
Social Pressure 10	-	-	-	-2.879
	-	-	-	(2.697)
Social Pressure 15	-	-	—	-1.249
	—	—	-	(2.686)
Social Pressure 20	-	_	-	-2.086
	-	_	_	(2.648)
Traditional 10	_	_	—	-2.537
	_	_	_	(2.688)
Traditional 15	—	—	-	-2.541
	—	—	_	(2.083)
Traditional 20	-	_	_	0.688
Control Decomo & Traditional 10	_	_	_	(2.005)
Social Pressure & Traditional 10	—	—	_	-4.037 (2.671)
Social Procesure & Traditional 15				(2.071)
Social Pressure O Traditional 15	_	_	_	(2.685)
Social Procesure & Traditional 20	_	_	_	(2.000)
bociai i ressure e i radiconai so	_	_	_	(2.641)
Social Identity 10	_	_	_	-3 223
Decide Tachibity 10	_	_	_	(2.681)
Social Identity 15	_	_	_	-4.063
	_	_	_	(2.676)
Social Identity 20	_	_	_	-7.269***
0	_	_	_	(2.675)
Social Identity & Social Pressure 10	_	_	_	-2.496
U U	_	_	-	(2.679)
Social Identity & Social Pressure 15	_	_	_	0.306
U U	_	_	-	(2.682)
Social Identity & Social Pressure 20	_	-	-	-0.589
	_	-	-	(2.679)
Control Mean	32 720***	32 720***	32 720***	32 720***
Control Hilder	(1.777)	(1.777)	(1.777)	(1.777)
N (subjects)	31,875	31,875	31,875	31,875
· · · /	,	,	, - · -	,

Table 5.7: Treatment Effects on Voter Turnout for Black Voters

 $^*p<.1;\,^{**}p<.05;\,^{***}p<.01$  Note: OLS regression estimates show treatment effects on voter turnout among Black voters.  $\label{eq:coefficients} Coefficients \ represent \ percentage-point \ changes \ relative \ to \ a \ control \ group \ mean \ of \ 32.72\%.$ Models 1-4 estimate treatment effects by exposure, message frame, dosage levels, and combined message-by-dosage interactions. All ads were race-neutral. Social Identity treatments significantly reduced turnout at higher doses, indicating possible backlash effects. Robust standard errors in parentheses.

effects among Black voters differ meaningfully from those observed in other demographic subgroups.

Model 2 disaggregates treatment effects by message type. Four of the five message frames produce negative, non-significant coefficients: Social Pressure (b = -2.085, SE = 2.178), Traditional (b = -1.458, SE = 2.182), Social Pressure & Traditional (b = -1.348, SE = 2.173), and Social Identity & Social Pressure (b = -0.936, SE = 2.180). However, one estimate does reach conventional significance: the Social Identity message yields a large and statistically significant negative effect (b = -4.850, SE = 2.178, p < .05). This finding indicates that appeals to group identity may have backfired among Black voters in this context, producing turnout suppression rather than mobilization. The direction and magnitude of this effect warrant careful consideration, especially given the race-neutral design of the ads and their deployment in a Republican primary context.

Model 3 shows uniformly negative effects across dosage levels. None of the estimates are statistically significant, but the consistent pattern is striking: 10 Dose Targets (b = -3.153, SE = 2.066), 15 Dose Targets (b = -1.218, SE = 2.066), and 20 Dose Targets (b = -2.029, SE = 2.061) all suggest that more exposure did not enhance, and may have reduced, the likelihood of turnout. These results contrast with the dosage-response curves observed among Election Day voters and male voters, where increased exposures generated modest turnout boosts. Among Black voters, repetition appears to have had no consistent mobilizing effect and may even have reduced engagement.

Model 4 explores interaction effects between message type and dosage. As in Model 3, the overwhelming majority of coefficients are negative. The only statistically significant effect is for the *Social Identity 20* condition, which shows a large and significant negative effect on turnout (b = -7.269, SE = 2.675, p < .01). This represents the largest treatment effect in magnitude observed in any subgroup and suggests a strong demobilizing response to repeated exposure to the identity-based message among Black voters. While many other message-by-dosage combinations also yield large negative coefficients—e.g., *Social Identity 15* (b = -4.063), *Social Pressure 10* (b = -2.879), and *Traditional 15* (b = -2.541)—none

are statistically significant, though their consistency in direction is notable. One notable exception is the *Social Pressure & Traditional 10* condition, which also shows a marginally significant negative effect (b = -4.637, SE = 2.671, p < .1), suggesting a possible early exposure backfire in that condition as well.

Unlike the White voter results, where the effects were mostly negative but imprecise, the Black voter findings include several large, statistically significant suppression effects, particularly among identity-based and high-dosage exposures. These patterns are difficult to reconcile with the intention of the ads, which were designed to mobilize rather than demobilize. One potential explanation is that certain frames—especially those invoking group identity—may not have resonated in a meaningful or culturally congruent way for Black voters in the context of a Republican primary. Alternatively, repeated digital exposure may have been perceived as intrusive or manipulative, particularly when the content lacked cultural specificity or contextual relevance. Given the race-neutral design of the messages, these effects suggest that race-specific political context and message reception may play a more critical role than anticipated in digital mobilization campaigns.

Table 5.7 shows that the digital GOTV campaign tested in this study failed to mobilize Black voters and, in several conditions—most notably *Social Identity 20*—appears to have actively reduced turnout. These results contrast sharply with findings from other subgroups and reinforce the importance of testing for heterogeneous treatment effects. While the underlying mechanisms are not directly observable in this dataset, the magnitude and statistical significance of several negative coefficients among Black voters suggest caution when applying standardized messaging across diverse communities. Future campaigns should consider the contextual, cultural, and electoral dimensions of message targeting to avoid unintended backlash.

Table 5.8 presents the results of subgroup-specific OLS regression models assessing the effects of digital advertisements on voter turnout among Latino voters (N = 66,829). The table includes four models that incrementally estimate the impact of treatment assignment (Model 1), message frame (Model 2), dosage intensity (Model 3), and the interaction between

message and dosage (Model 4). All coefficients are scaled as percentage-point changes in turnout relative to the control group mean of 32.72%, with standard errors in parentheses.

Model 1 reports the average effect of receiving any digital ad, regardless of message content or dosage. The point estimate is negative (b = -1.506, SE = 1.770), though not statistically significant. This suggests that, on average, the digital advertisements tested in this study did not successfully mobilize Latino voters—and may even have slightly reduced turnout. While this estimate should be interpreted cautiously, its direction mirrors the negative point estimates found among White and Black voters, raising early questions about digital ad effectiveness in this subgroup.

Model 2 disaggregates the treatment effect by message content. Every message frame yields a negative coefficient, though none are statistically significant. The largest negative effect comes from the *Traditional* message (b = -2.437, SE = 1.985), followed by the *Social Identity* message (b = -2.296, SE = 1.983), and the *Social Pressure* message (b = -1.862, SE = 1.984). The combined *Social Pressure & Traditional* and *Social Identity & Social Pressure* messages yield smaller negative effects, but again with no statistical significance. The consistent negativity across frames suggests that none of the message strategies—whether duty-based, identity-based, or blended—resonated effectively with Latino voters in this campaign.

Model 3 examines dosage effects independent of message content. As with the messagebased results, all three dosage tiers yield negative point estimates. The 10 Dose Targets condition is the most negative (b = -2.743, SE = 1.881), followed by 20 Dose Targets (b = -1.437) and 15 Dose Targets (b = -0.340), neither of which are statistically significant. This pattern stands in contrast to earlier results for Election Day and male voters, where dosage appeared to enhance message effectiveness. Among Latino voters, repetition does not appear to increase mobilization, and may even generate disengagement.

Model 4 evaluates message-by-dosage interaction terms. Again, the results show a broad pattern of negative point estimates. The largest negative effects appear in the *Traditional* 10 condition (b = -3.901, SE = 2.447) and the *Social Pressure* 10 condition (b = -3.047,

	Model 1	Model 2	Model 3	Model 4
	(Any Ireat)	(Message)	(Dosage)	(Message × Dosage)
Any Treatment	-1.506 (1.770)	_	_	_
Social Pressure	_	-1.862	_	_
	_	(1.984)	_	-
Traditional GOTV	_	-2.437	_	_
	_	(1.985)	-	-
Social Pressure & Traditional	_	-0.214	_	_
	_	(1.978)	-	-
Social Identity	_	-2.296	_	_
U U	_	(1.983)	-	-
Social Identity & Social Pressure	_	-0.714	_	_
0	_	(1.985)	_	-
10 Dose Taraets	_	_	-2.743	_
	_	_	(1.881)	_
15 Dose Taraets	_	_	-0.340	_
	_	_	(1.882)	_
20 Dose Taraets	_	_	-1437	_
20 2000 1 argono	_	_	(1.876)	_
Social Pressure 10	_	_	_	-3.047
50000 170500rC 10	_	_	_	(2.455)
Social Pressure 15	_	_	_	0.292
50000 170500rC 10	_	_	_	(2.448)
Social Pressure 20	_	_	_	-2 786
Social Pressare 20	_	_	_	(2.100)
Traditional 10	_	_	_	-3 001
Traditional 10	_	_	_	(2.447)
Traditional 15				2.206
Traditional 15	_	_	_	(2.440)
Traditional 20				0.122
Traditional 20	_	_	_	(2.436)
Social Processor & Traditional 10				2.130)
Social Tressure O Traditional To	_	_	_	(2.420)
Social Procesure & Traditional 15	_	_	_	2.421)
Social Tressure O Traditional 15	_	_	_	(2.203)
Social Processor & Traditional 20				0.552
Social Tressure O Traditional 20	_	_	_	(2.397)
Social Identity 10	_	_	_	(2.001) -2.024
Social Identity 10	_	_	_	(2.437)
Social Identity 15	_	_	_	-1.834
Social Identity 15	_	_	_	(2.440)
Social Identity 20	_	_	_	-3.014
Social Identity 20	_	_	_	(2.427)
Social Identity & Social Pressure 10	_	_	_	(2.121)
Social Identity C Social I lessure 10	_	_	_	(2.445)
Conial Identity & Conial Processon 15				(2.440)
Social Identity & Social Tressure 15	_	_	_	(2.442)
Social Identity & Social Processing 20				1 742
Social Identity & Social Fressure 20	_	_	_	-1.140 (2.422)
				(2.100)
Control Mean	32.720***	32.720***	32.720***	32.720***
	(1.777)	(1.777)	(1.777)	(1.777)
N (subjects)	66,829	66,829	66,829	66,829

Table 5.8: Treatment Effects on Voter Turnout for Latino Voters

 $^*{\rm p}<.1;\,^{**}{\rm p}<.05;\,^{***}{\rm p}<.01$  Note: OLS regression estimates represent percentage-point changes in voter turnout among Latino voters relative to a control mean of 32.72%. Models estimate effects by overall treatment (Model 1), message framing (Model 2), dosage (Model 3), and message-by-dosage interactions (Model 4). All coefficients are statistically insignificant or negative, suggesting minimal or slightly demobilizing effects. Ads were race-neutral. Robust standard errors in parentheses.

SE = 2.455). These early-exposure conditions suggest that initial contact with ads—rather than later saturation—may have driven down engagement. Other low-performing combinations include *Traditional 15* and *Social Identity 20*, both exceeding three-point reductions in turnout. Although some higher-dosage conditions such as *Social Pressure & Traditional* 15 (b = 2.283) and *Social Identity & Social Pressure 15* (b = 0.919) produce small positive estimates, none are statistically significant, and the directionality is inconsistent.

Taken together, these results reveal a striking uniformity of null or negative effects across message types, dosage levels, and their interactions. Latino voters in this experiment did not appear to be mobilized by the tested advertisements, and some exposure conditions suggest the possibility of backlash or disengagement. These findings highlight the limits of race-neutral digital messaging when targeting racially diverse constituencies. They also underscore the need for cultural specificity and audience resonance in designing GOTV strategies for Latino communities, particularly when engaging them in the context of partisan primaries where base alignment may be more complex.

In sum, Table 5.8 indicates that the digital advertisements tested in this study failed to generate significant increases in turnout among Latino voters. While the underlying mechanisms cannot be directly observed, the pattern of negative point estimates suggests that message content and dosage—at least in the formats deployed here—did not meaningfully engage this demographic. These findings reinforce the broader conclusion that digital outreach campaigns require audience-specific tailoring to avoid null or counterproductive effects.

Table 5.9 reports the results of subgroup-specific OLS regression models estimating the effects of digital GOTV advertisements on turnout among Asian American and Pacific Islander (AAPI) voters. The table includes four models: Model 1 estimates the average treatment effect across all treated individuals, Model 2 disaggregates effects by message content, Model 3 examines dosage levels regardless of message, and Model 4 includes fully interacted message-by-dosage conditions. Coefficients are expressed as percentage-point changes in turnout relative to a control group mean of 37.49%, with robust standard errors in parentheses.

Model 1 estimates the overall effect of being exposed to any digital advertisement. The

	Model 1 (Any Treat)	Model 2 (Message)	Model 3 (Dosage)	Model 4 (Message × Dosage)
Any Treatment	1.622	(	(8-)	(88-)
Any Treatment	(2.428)	_	_	_
Social Pressure	()	-1.161	_	_
	_	(2.725)	_	_
Traditional GOTV	_	-2.302	_	_
	_	(2.722)	_	_
Social Pressure & Traditional	_	-0.749	_	_
	-	(2.723)	-	-
Social Identity	_	-3.598	_	_
	-	(2.723)	-	-
Social Identity & Social Pressure	_	-0.345	-	—
	-	(2.719)	-	-
10 Dose Targets	_	_	-2.106	-
	-	—	(2.580)	-
15 Dose Targets	-	-	-1.654	-
	—	-	(2.581)	_
20 Dose Targets	-	-	-1.140	-
	-	-	(2.577)	-
Social Pressure 10	-	—	—	-3.842
	-	-	—	(3.348)
Social Pressure 15	_	-	-	0.178
	-	—	—	(3.360)
Social Pressure 20	-	-	—	0.387
	—	-	-	(3.352)
Traditional 10	_	-	-	-4.387
	-	-	—	(3.352)
Traditional 15	_	-	-	-4.573
The line last	_	-	-	(3.362)
Traditional 20	-	-	—	1.901
	—	_	_	(3.320)
Social Pressure & Traditional 10	_	-	-	-2.194
Control Decomo & Traditional 15	—	—	_	(0.000)
Social Fressure & Traditional 15	_	_	_	2.450 (2.378)
Conial Processing & Traditional 00				0.010)
Social Tressure & Traditional 20	_	_	_	(3.314)
Social Identity 10	_	_	_	-0.507
Social Identity 10	_	_	_	(3.351)
Social Identity 15	_	_	_	-6.024*
Social Identity 10	_	_	_	(3.348)
Social Identity 20	_	_	_	-4.115
	_	_	_	(3.346)
Social Identity & Social Pressure 10	_	_	_	0.543
	_	_	_	(3.341)
Social Identity & Social Pressure 15	_	_	_	-0.076
	_	_	_	(3.326)
Social Identity & Social Pressure 20	_	_	_	-1.474
	_	-	-	(3.350)
Control Maan	27 100***	97 100***	97 100***	97 100***
Control mean	(1.536)	(1.536)	(1.536)	(1.536)
N (subjects)	778 882	778 889	778 889	778 882
· (caplece)	110,004	110,002	110,002	110,002

Table 5.9: Treatment Effects on Voter Turnout for AAPI Voters

 $<sup>^*{\</sup>rm p}<.1;\,^{**}{\rm p}<.05;\,^{***}{\rm p}<.01$  Note: OLS regression estimates represent percentage-point changes in turnout among AAPI voters relative to a control mean of 37.49%. Models estimate effects by overall treatment (Model 1), message framing (Model 2), dosage (Model 3), and message-by-dosage interactions (Model 4). Most coefficients are statistically insignificant or negative, indicating minimal mobilization effects. Social Identity 15 shows a marginally significant demobilizing effect. Ads were race-neutral. Robust standard errors in parentheses.

result is negative and statistically insignificant (b = -1.633, SE = 2.428), suggesting that, on average, exposure to the tested advertisements had no meaningful mobilizing effect among AAPI voters and may have slightly reduced turnout.

Model 2 disaggregates this average effect by message type. None of the message frames produced statistically significant results. Four of the five tested messages yielded negative coefficients, including *Social Identity* (b = -3.598, SE = 2.723), *Traditional* (b = -2.302, SE = 2.722), and *Social Pressure* (b = -1.161, SE = 2.725). The *Social Pressure & Traditional* and *Social Identity & Social Pressure* messages yielded smaller negative estimates. While none of these coefficients are statistically significant, the consistent directionality suggests that these message frames were either ineffective or mildly demobilizing within this subgroup.

Model 3 examines the average effect of dosage across all message types. Again, all coefficients are negative and statistically insignificant: 10 Dose Targets (b = -2.106, SE = 2.580), 15 Dose Targets (b = -1.654, SE = 2.581), and 20 Dose Targets (b = -1.140, SE = 2.577). Unlike other subgroups (e.g., Election Day or male voters), the AAPI subgroup shows no evidence of a dosage-response pattern. Instead, repetition appears to have either no effect or a slightly suppressive effect on turnout.

Model 4 includes interaction terms between message frame and dosage level. As in previous models, nearly all coefficients are negative, with the exception of a few conditions that yield small positive but statistically insignificant results. Among the most negative estimates are *Traditional 15* (b = -4.573, SE = 3.362), *Traditional 10* (b = -4.387, SE =3.352), and *Social Pressure 10* (b = -3.842, SE = 3.348). The *Social Identity 15* condition yields the only statistically significant effect across the entire model: a substantial and negative impact on turnout (b = -6.024, SE = 3.348, p < .1). This suggests that repeated exposure to group-based identity messaging may have triggered disengagement or message fatigue among AAPI voters. Other Social Identity conditions—such as *Social Identity 20* and *Social Identity & Social Pressure 20*—also show negative but statistically insignificant estimates. A few positive estimates emerge, such as Social Pressure 15 (b = 0.178), Social Pressure 20 (b = 0.387), and Social Pressure & Traditional 15 (b = 2.430), but none reach significance. The pattern overall is one of statistical imprecision and general ineffectiveness.

These findings echo those for Black and Latino voters: race-neutral digital GOTV messaging did not improve turnout among AAPI voters, and some message-dosage combinations appear to have had unintended suppressive effects. Although the sample size for AAPI voters is smaller than for other groups, the consistent directionality and the statistical significance of one of the most negative estimates (*Social Identity 15*) highlight the importance of tailoring content to specific audiences. Messages that do not meaningfully resonate or reflect the political and cultural priorities of the targeted group may not simply fail—they may backfire.

Table 5.9 shows that digital GOTV advertisements tested in this study were not effective at mobilizing AAPI voters. While many point estimates are imprecise, the majority are negative, and at least one (Social Identity 15) reaches statistical significance in the demobilizing direction. These findings reinforce broader conclusions about the limitations of race-neutral digital outreach and the potential importance of culturally relevant messaging in political mobilization efforts.

#### 5.7.2 Gender

Tables 5.10 and 5.11 present subgroup-specific ordinary least squares (OLS) models estimating the effect of digital campaign advertisements on voter turnout among female and male registered voters, respectively. Each model follows a consistent structure: Model 1 estimates the average treatment effect of receiving any advertisement, Model 2 examines message frame effects by pooling dosage levels, Model 3 evaluates dosage frequency effects across all treated individuals, and Model 4 includes interaction terms between each message frame and dosage, capturing the conditional effects of repeated exposure.

	Model 1 (Any Treat)	Model 2 (Message)	Model 3 (Dosage)	$\begin{array}{c} \textbf{Model 4} \\ (\text{Message} \times \text{Dosage}) \end{array}$
Any Treatment	1.837 (1.120)	_	_	-
Social Pressure		$2.403^{*}$	-	-
Traditional COTV	_	1.600	_	_
	_	(1.260)	_	_
Social Pressure & Traditional	_	1.258	_	_
	_	(1.255)	-	-
Social Identity	_	$2.386^{*}$	_	_
	_	(1.255)	_	_
Social Identity & Social Pressure	_	1.536 (1.257)	_	_
10 Impressions	_		1 579	_
10 1107000000	_	_	(1.190)	_
15 Impressions	_	_	1.909	_
10 110010000	_	-	(1.190)	_
20 Impressions	_	_	2.022*	_
	_	_	(1.190)	_
Social Pressure 10	_	_	_	3.048**
	_	_	-	(1.547)
Social Pressure 15	_	_	_	2.377
	_	-	-	(1.536)
Social Pressure 20	_	_	_	1.780
	_	—	-	(1.544)
Traditional 10	_	_	_	$2.863^{*}$
	—	—	-	(1.547)
Traditional 15	_	_	_	-0.384
	_	-	-	(1.553)
Traditional 20	—	—	-	2.325
	-	-	-	(1.560)
Social Pressure & Traditional 10	_	_	—	0.118
	-	—	—	(1.537)
Social Pressure & Traditional 15	-	-	-	1.607
	-	—	—	(1.544)
Social Pressure & Traditional 20	_	_	-	2.070
	-	—	-	(1.547)
Social Identity 10	_	—	—	1.998
	-	—	_	(1.549)
Social Identity 15	-	-	-	3.346**
	—	-	-	(1.547)
Social Identity 20	-	—	_	1.834
	—	-	-	(1.534)
Social Identity & Social Pressure 10	-	—	_	-0.114
	—	-	-	(1.555)
Social Identity & Social Pressure 15	_	—	_	$2.541^{*}$
	_	—	-	(1.540)
Social Identity & Social Pressure 20	—	-	-	2.129
	_	—	_	(1.550)
Control Mean	41.000***	41.000***	41.000***	41.000***
	(1.033)	(1.033)	(1.033)	(1.033)
N (Female subjects)	363,582	$363,\!582$	363,582	363,582

Table 5.10: Treatment Effects on Voter Turnout – Female Voters

 $<sup>^*</sup>p<.1;\,^{**}p<.05;\,^{***}p<.01$  Note: OLS coefficients represent percentage-point increases in turnout relative to a control mean of 41.0% for female voters. Model 1 pools all treatments. Model 2 distinguishes effects by message content. Model 3 distinguishes effects by dosage level. Model 4 fully interacts message content and dosage. Positive, significant effects observed for Social Pressure (10 doses), Traditional (10 doses), Social Identity (15 doses), and Social Identity & Social Pressure (15 doses). Robust standard errors in parentheses.

Among female voters (N = 363,582), the overall treatment effect in Model 1 is a positive but statistically insignificant increase of 1.84 percentage points (b = 1.837, SE = 1.120), suggesting weak evidence that exposure to any ad increased turnout. However, when the analysis is disaggregated by message content (Model 2), two message frames stand out: the *Social Pressure* message (b = 2.403, SE = 1.255, p < .1) and the *Social Identity* message (b = 2.386, SE = 1.255, p < .1), both of which yield statistically significant or marginally significant positive effects. These findings suggest that female voters may be responsive to both normative appeals invoking civic obligation and group-based appeals that emphasize shared identity or values.

In Model 3, a monotonic pattern emerges across dosage levels. Turnout increases slightly at 10 exposures (b = 1.579, SE = 1.190), rises further at 15 exposures (b = 1.909, SE =1.190), and peaks at 20 exposures with a statistically significant effect (b = 2.022, SE =1.190, p < .1). These results indicate that increased message frequency yields modest but potentially meaningful gains in turnout among female voters.

Model 4 reveals that this dosage responsiveness is not evenly distributed across message types. For example, Social Pressure 10 generates a statistically significant increase in turnout (b = 3.048, SE = 1.547, p < .05), whereas Social Pressure 15 and 20 produce smaller, nonsignificant effects. The Traditional 10 and 20 conditions also yield turnout boosts, with the latter approaching statistical significance (b = 2.325, SE = 1.560). The Social Identity 15 condition yields the highest significant estimate among identity-based treatments (b = 3.346, SE = 1.547, p < .05), suggesting that repeated exposure to group-referent appeals may build cumulative mobilizing effects. Importantly, some conditions—such as Traditional 15—produce negative coefficients, underscoring the importance of dosage-message alignment.

Overall, these results suggest that female voters respond most favorably to civic duty and identity-based messages when delivered at low to moderate frequencies. However, gains are modest, and few effects are statistically significant beyond the .10 threshold, implying that digital ads alone are insufficient to drive large-scale behavioral change among women.
	Model 1 (Any Treat)	Model 2 (Message)	Model 3 (Dosage)	$\frac{Model \ 4}{(Message \times Dosage)}$
Any Treatment	2 073*	_	_	_
ing ireachent	(1.118)	_	_	_
Social Pressure		$2.344^{*}$	_	_
	_	(1.253)	_	_
Traditional GOTV	_	1 990	_	_
	_	(1.257)	_	_
Social Pressure & Traditional	_	1.483	_	_
	_	(1.253)	_	_
Social Identity	_	2.768**	_	_
	_	(1.253)	_	_
Social Identity & Social Pressure	_	1.775	_	_
	_	(1.255)	_	_
10 Impressions	_	_	1.810	_
1	_	_	(1.188)	_
15 Impressions	_	_	$2.015^{*}$	-
<u>r</u>	_	_	(1.187)	_
20 Impressions	_	_	2.391**	_
<u>r</u>	_	_	(1.188)	_
Social Pressure 10	_	_	_	2.443
	_	-	-	(1.544)
Social Pressure 15	_	_	_	3.002*
	_	-	-	(1.532)
Social Pressure 20	_	_	_	1.584
	_	-	-	(1.541)
Traditional 10	_	_	_	2.932*
	_	-	_	(1.544)
Traditional 15	_	-	-	-0.103
	-	-	-	(1.550)
Traditional 20	_	-	-	3.143**
	-	-	-	(1.557)
Social Pressure & Traditional 10	_	-	_	0.297
	-	-	-	(1.534)
Social Pressure & Traditional 15	-	-	-	1.839
	_	-	-	(1.540)
Social Pressure & Traditional 20	-	-	-	2.334
	-	-	-	(1.544)
Social Identity 10	-	-	-	3.314**
	-	-	-	(1.546)
Social Identity 15	_	-	—	2.892*
	-	-	-	(1.544)
Social Identity 20	-	-	-	2.122
	_	—	—	(1.531)
Social Identity & Social Pressure 10	-	-	_	0.081
	_	-	-	(1.551)
Social Identity & Social Pressure 15	_	—	—	2.382
	_	—	—	(1.537)
Social Identity & Social Pressure 20	—	-	-	$2.804^{*}$
	_	—	—	(1.546)
Control Mean	41.000***	41.000***	41.000***	41.000***
	(1.033)	(1.033)	(1.033)	(1.033)
N (Male subjects)	380,443	380,443	380,443	380,443

# Table 5.11: Treatment Effects on Voter Turnout – Male Voters (OLS Models)

<sup>\*</sup>p < .1; \*\*p < .05; \*\*\*p < .01 Note: OLS coefficients represent percentage-point increases in turnout relative to a control mean of 41.0% for male voters. Model 1 pools all treatments. Model 2 distinguishes effects by message content. Model 3 distinguishes effects by dosage level. Model 4 fully interacts message content and dosage. Positive, significant effects observed for Social Identity and certain dosage combinations. Robust standard errors in parentheses.

Among male voters (N = 380,443), the average treatment effect in Model 1 is statistically significant (b = 2.073, SE = 1.118, p < .1), suggesting that digital advertisements collectively raised turnout by just over two percentage points relative to control. This contrasts with the ambiguous effect found among women and implies a stronger baseline responsiveness to digital outreach among men.

In Model 2, message-level disaggregation again highlights the Social Pressure frame (b = 2.344, SE = 1.253, p < .1) and the Social Identity frame (b = 2.768, SE = 1.253, p < .05) as relatively more effective than others. Interestingly, the Traditional message—while positive (b = 1.990, SE = 1.257)—does not attain statistical significance. These patterns suggest that male voters, like female voters, are most influenced by psychologically charged messages that invoke civic duty or collective identity.

Dosage effects in Model 3 exhibit a clear upward trend, similar to that observed among women: turnout increases progressively from 10 to 15 to 20 impressions, with statistically significant effects at both 15 exposures (b = 2.015, SE = 1.187, p < .1) and 20 exposures (b = 2.391, SE = 1.188, p < .05). These findings provide further evidence for a dosageresponse curve in digital campaign effectiveness, at least for men.

Model 4 introduces interaction terms that offer deeper insight. Among the most notable findings is the high turnout associated with *Social Pressure 15* (b = 3.002, SE = 1.532, p < .1), *Traditional 20* (b = 3.143, SE = 1.557, p < .05), and *Social Identity 10* (b =3.314, SE = 1.546, p < .05). These results reinforce the idea that certain message types become more effective at higher frequencies—but not universally so. For example, *Traditional* 15 yields a negative coefficient (b = -0.103, SE = 1.550), while *Social Identity & Social Pressure 10* is essentially null. This variability suggests that some messages wear well with repetition while others decay or lose salience over time.

In summary, the male-specific findings suggest that men may be more consistently responsive to digital mobilization efforts than women, particularly when exposed to targeted messages multiple times. Though most individual effects remain small in absolute terms, the consistency and statistical significance of dosage-linked turnout increases among men are notable. These results bolster the case for repetition-based targeting in digital campaigns, at least within certain messaging frameworks.

Together, Tables 6.5 and 6.6 suggest that message content and dosage interact meaningfully with gender in shaping voter responsiveness to digital political advertising. Whereas male voters exhibit a modest but statistically significant mobilization response—particularly to identity and social pressure appeals—female voters show more varied and weaker effects. Importantly, the effectiveness of these messages hinges not only on their psychological content but also on the frequency with which they are delivered. These patterns underscore the importance of audience segmentation and message calibration in the design of scalable, gender-sensitive digital GOTV strategies.

### 5.7.3 Age

To evaluate whether age moderates the effect of digital GOTV advertisements, Tables 5.12 through 5.15 present four stratified OLS regression models by age cohort. Each model estimates the interaction between message frame and dosage level within a specific age group. The results provide insight into how digital messaging may differentially influence political behavior across the life cycle. Figure 5.3 visually summarizes the marginal treatment effects derived from these models.

Seniors (65+). Table 5.12 reveals that for seniors, none of the main message frames or dosage levels produce statistically significant effects on turnout. However, several positive coefficients in the interaction terms—particularly at higher dosages—suggest a potential for modest mobilization. For example, the interaction *Social Pressure & Traditional × Dosage* = 20 yields a small positive estimate (b = 0.010, SE = 0.006), though it falls just short of statistical significance. Overall, the results indicate that seniors were largely unresponsive to digital advertisements, with only weak signs of persuasion when exposed to high-frequency norm-based messages.

Variable	Coefficient	Std. Error
Message: Social Identity & Social Pressure	-0.005	(0.004)
Message: Social Pressure	-0.001	(0.004)
Message: Social Pressure & Traditional	0.001	(0.004)
Message: Traditional	0.003	(0.004)
Dosage = 15	-0.005	(0.004)
Dosage = 20	-0.004	(0.004)
Message: Social Identity & Social Pressure $\times$ Dosage = 15	0.007	(0.006)
Message: Social Pressure $\times$ Dosage = 15	0.007	(0.006)
Message: Social Pressure & Traditional $\times$ Dosage = 15	0.004	(0.006)
Message: Traditional $\times$ Dosage = 15	0.003	(0.006)
Message: Social Identity & Social Pressure $\times$ Dosage = 20	0.003	(0.006)
Message: Social Pressure $\times$ Dosage = 20	0.007	(0.006)
Message: Social Pressure & Traditional $\times$ Dosage = 20	0.010	(0.006)
Message: Traditional $\times$ Dosage = 20	-0.001	(0.006)
Intercept (Seniors, Control)	0.449***	(0.003)
Observations	391	,517

Table 5.12: OLS Regression: Seniors  $(65+) \times \text{Message} \times \text{Dosage}$ 

p < .1; p < .05; p < .01

*Note:* OLS coefficients for Seniors (65+), interacting message type and dosage level. The reference category is untreated Seniors at dosage = 10. Coefficients represent percentage-point changes in turnout. Standard errors in parentheses.

**Older Adults (50–64)**. Table 5.13 shows that older adults are slightly more responsive than seniors, though effect sizes remain small and mostly insignificant. The main effect for being in this age group is significantly negative (b = -0.030, SE = 0.005, p < .01), suggesting that, on average, older adults were less likely to vote than seniors in the absence of treatment. Among the interactions, the most promising is *Older Adults × Social Pressure* & *Traditional × Dosage = 15* (b = 0.008, SE = 0.010), which—while not significant—points in a direction consistent with the hypothesis that norm-based appeals paired with reminders may increase participation in this group. The small magnitude of effects overall implies that digital outreach may have limited influence among older adults, despite their relatively higher turnout rates.

Variable	Coefficient	Std. Error
Main Effect: Older Adults	$-0.030^{***}$	(0.005)
Older Adults × Message: Social Identity & Social Pressure	0.004	(0.007)
$Older \ Adults \  imes \ Message: \ Social \ Pressure$	0.001	(0.007)
Older Adults × Message: Social Pressure & Traditional	-0.0003	(0.007)
$Older \ Adults \ \times \ Message: \ Traditional$	-0.005	(0.007)
$Older \ Adults \times Dosage = 15$	-0.0004	(0.007)
$Older \ Adults \times Dosage = 20$	0.004	(0.007)
Older Adults $\times$ Social Identity & Social Pressure $\times$ Dosage = 15	-0.004	(0.010)
$Older \ Adults \times Social \ Pressure \times Dosage = 15$	-0.0004	(0.010)
$Older \ Adults \ \times \ Social \ Pressure \ & Traditional \ \times \ Dosage = 15$	0.008	(0.010)
$Older \ Adults \ \times \ Traditional \ \times \ Dosage = 15$	0.005	(0.010)
$Older \ Adults \times Social \ Identity \ & Social \ Pressure \ \times \ Dosage = 20$	-0.004	(0.010)
$Older \ Adults \times Social \ Pressure \times Dosage = 20$	-0.002	(0.010)
$Older \ Adults \ \times \ Social \ Pressure \ \& \ Traditional \ \times \ Dosage = 20$	-0.004	(0.010)
$Older \ Adults \  imes \ Traditional \  imes \ Dosage = 20$	0.008	(0.010)
Intercept (Seniors, Control)	$0.449^{***}$	(0.003)
Observations	258	,728

Table 5.13: OLS Regression: Older Adults  $(50-64) \times \text{Message} \times \text{Dosage}$ 

p < .1; p < .05; p < .01

*Note:* OLS model for Older Adults (50–64), interacted with message type and dosage. The reference group is untreated Seniors at dosage = 10. Standard errors in parentheses.

Middle-Aged Adults (30–49). The results in Table 5.14 provide some of the strongest evidence of differential responsiveness by message and dosage. Notably, the two-way interaction *Middle-Aged* × Social Identity & Social Pressure is statistically significant and positive (b = 0.018, SE = 0.011, p < .1), indicating that middle-aged voters may respond favorably to blended messages that combine identity and norm cues. However, this mobilization appears to decline at higher levels of exposure: *Middle-Aged* × Social Identity & Social Pressure × Dosage = 20 yields a negative and marginally significant estimate (b = -0.025, SE = 0.015, p < .1), suggesting that overexposure may diminish the persuasive effect. Other interaction terms in this model, such as those involving traditional and norm-based frames, remain small and statistically indistinct. These results imply that middle-aged voters are selectively responsive to certain combinations of content and dosage, but that the relationship is non-linear.

Table 5.14: OLS Regression: Middle-Aged Adults  $(30-49) \times \text{Message} \times \text{Dosage}$ 

Variable	Coefficient	Std. Error
Main Effect: Middle-Aged Adults	-0.011	(0.008)
Middle-Aged × Social Identity & Social Pressure	$0.018^{*}$	(0.011)
$Middle-Aged \times Social \ Pressure$	-0.007	(0.011)
Middle-Aged  imes Social Pressure & Traditional	0.002	(0.011)
$Middle-Aged \times Traditional$	-0.003	(0.011)
$Middle-Aged \times Dosage = 15$	0.0003	(0.011)
$Middle-Aged \times Dosage = 20$	0.014	(0.011)
$Middle-Aged \times Social \ Identity \ & Social \ Pressure \times Dosage = 15$	-0.014	(0.015)
$Middle-Aged \times Social \ Pressure \times Dosage = 15$	0.013	(0.015)
$Middle-Aged \times Social \ Pressure \ & Traditional \times Dosage = 15$	0.001	(0.015)
$Middle-Aged \times Traditional \times Dosage = 15$	-0.005	(0.015)
$Middle-Aged \times Social \ Identity \ & Social \ Pressure \times Dosage = 20$	$-0.025^{*}$	(0.015)
$Middle-Aged \times Social \ Pressure \times Dosage = 20$	0.002	(0.015)
$Middle-Aged \times Social \ Pressure \ & Traditional \times Dosage = 20$	-0.018	(0.015)
$Middle-Aged \times Traditional \times Dosage = 20$	-0.014	(0.015)
Intercept (Seniors, Control)	$0.449^{***}$	(0.003)
Observations	76,	052

p < .1; p < .05; p < .01

*Note:* OLS regression model for Middle-Aged Adults (30–49), including message and dosage interactions. Reference group is untreated Seniors at dosage

= 10. Standard errors in parentheses.

Young Adults (20–29). Table 5.15 reveals the most negative and statistically meaningful treatment effects across any subgroup. The main effect for being a young adult is strongly negative (b = -0.125, SE = 0.030, p < .01), confirming that this group has the lowest baseline likelihood of voting. Furthermore, interaction terms involving normbased messages consistently suggest a demobilizing effect. For example, Young Adults × Social Pressure yields a marginally significant negative effect (b = -0.085, SE = 0.044, p < .1), and other interaction terms—such as Young Adults × Dosage = 20 (b = -0.038, SE = 0.043)—continue the trend. While several message-by-dosage combinations do yield small positive coefficients, none reach statistical significance. These findings suggest that digital ads may backfire among young adults, possibly due to overexposure, skepticism of political messaging, or poor resonance with the frames employed.

Table 5.15: OLS Regression Model:	Young Adults (	(20 - 29)	) × Message ×	Dosage Level
-----------------------------------	----------------	-----------	---------------	--------------

Variable	Coefficient	Standard Error
Main Effect: Young Adults	$-0.125^{***}$	(0.030)
Young Adults × Message: Social Identity & Social Pressure	-0.050	(0.042)
Young Adults $\times$ Message: Social Pressure	$-0.085^{*}$	(0.044)
Young Adults × Message: Social Pressure & Traditional	-0.016	(0.044)
Young Adults $\times$ Message: Traditional	0.006	(0.044)
Young Adults $\times$ Dosage = 15	-0.029	(0.043)
Young Adults $\times$ Dosage = 20	-0.038	(0.043)
Young Adults $\times$ Message: Social Identity & Social Pressure $\times$ Dosage = 15	0.053	(0.061)
Young Adults $\times$ Message: Social Pressure $\times$ Dosage = 15	0.087	(0.062)
Young Adults $\times$ Message: Social Pressure & Traditional $\times$ Dosage = 15	0.030	(0.062)
Young Adults $\times$ Message: Traditional $\times$ Dosage = 15	-0.046	(0.062)
Young Adults $\times$ Message: Social Identity & Social Pressure $\times$ Dosage = 20	0.045	(0.061)
Young Adults $\times$ Message: Social Pressure $\times$ Dosage = 20	0.089	(0.062)
Young Adults $\times$ Message: Social Pressure & Traditional $\times$ Dosage = 20	-0.004	(0.063)
Young Adults $\times$ Message: Traditional $\times$ Dosage = 20	-0.039	(0.062)
Observations		3,911

p < .1; p < .05; p < .01

*Note:* OLS regression estimates for Young Adults (20–29), interacted with message type and dosage level. The reference group is untreated Seniors (65+) with dosage = 10. Coefficients represent percentage-point changes in turnout. Standard errors in parentheses.

Marginal Effects Summary. Figure 5.3 presents estimated marginal treatment effects across age groups. The bar plot shows that while seniors and older adults experienced small positive average treatment effects (approximately +0.8 and +0.9 percentage points, respectively), middle-aged adults show near-zero impact (-0.4 points), and young adults exhibit a substantially negative marginal treatment effect (-5.6 points). This visual reinforces the age-stratified models' conclusions: responsiveness to digital GOTV efforts is minimal or mildly positive among older voters, but may be neutral or even counterproductive among the youngest.



Figure 5.3: Estimated marginal treatment effects on voter turnout by age group. Effects represent percentage-point changes from baseline turnout.

Taken together, these findings highlight clear heterogeneity in digital ad effectiveness across age groups. Seniors and older adults appear modestly responsive to civic and reminderbased appeals, particularly at higher dosages. Middle-aged adults show some responsiveness to hybrid messages, though these effects diminish at high levels of exposure. By contrast, young adults exhibit patterns of resistance or demobilization, raising concerns about digital campaign strategies that apply uniform messaging across age cohorts. These results underscore the importance of age-specific targeting and message tailoring in the design of digital GOTV campaigns.

# CHAPTER 6

# Discussion

This experimental study offers a rare, large-scale, individual-level examination of the efficacy of digital campaign advertisements in influencing voter turnout. Unlike many prior studies that rely on aggregate or observational data, the randomized controlled design employed here allows for a robust causal inference regarding the impact of specific ad content and dosage on electoral participation. In doing so, the study contributes important empirical evidence to debates within political science regarding the mobilizing capacity of digital communication strategies, particularly in the context of a low-salience, intra-party primary election.

This chapter interprets the findings within the framework of existing theories of campaign communication and voter mobilization. It compares the results to prior research, identifies practical lessons for political strategists, outlines the study's methodological and contextual limitations, and proposes avenues for future research.

### 6.1 Revisiting Expectations and Literature

The results of this experiment resonate with the established literature on campaign communications, particularly the "minimal effects" tradition. Beginning with early empirical studies and extending into contemporary scholarship, researchers have frequently found that most campaign outreach has limited persuasive or mobilizing power—especially in high-salience electoral contexts where voters are already engaged, polarized, or inundated with competing messages (Broockman and Kalla 2018). In this study, the average treatment effect of digital ads on voter turnout ranged from 0.5 to just over 1 percentage point, depending on dosage and message type. While modest in absolute terms, these effects are statistically significant and substantively meaningful in tightly contested elections or low-turnout primaries, where even marginal gains can shift outcomes.

Of particular importance is the finding that message framing significantly moderated the effectiveness of the advertisements. Ads that emphasized civic norms and individual accountability—social pressure messages—consistently outperformed those that invoked partisan identity. This supports the theoretical framework developed in prior GOTV literature, notably Gerber, Green, and Larimer (2008), which demonstrates that social pressure interventions—by invoking social norms and the observability of voting behavior—can effectively increase participation. Our study extends this finding into the digital realm, suggesting that even impersonal, algorithmically distributed messages on Facebook can activate these psychological mechanisms.

Conversely, the relatively weak performance of the identity-based message underscores a limitation in relying solely on partian appeals. In the context of a primary election—where inter-party competition is absent and the stakes may be less salient—messages that frame voting as an in-group response to out-group threats may lack sufficient urgency or relevance. This aligns with theoretical expectations from political psychology: for identity-based mobilization to be effective, the group identity invoked must be salient, threatened, and mean-ingful within the specific political moment. In the absence of these conditions, such appeals may be ignored or discounted.

Equally important is the study's contribution to the understudied question of advertising dosage. While many prior field experiments in political communication assess only one or two treatment doses, this study systematically varied exposure at three levels: 10, 15, and 20 impressions. The findings suggest diminishing marginal returns to increased frequency but indicate that higher doses may yield greater effects for certain message types—especially those invoking social pressure. The non-linear trend, with a slight dip in effectiveness at 15 exposures, may reflect noise or threshold effects, where moderate repetition fails to reinforce the message further, while higher repetition breaks through voter apathy or digital clutter.

For practitioners, this suggests that while 10 exposures may suffice in many cases, a 20impression strategy could produce greater returns for high-performing messages.

Finally, the consistency of treatment effects across key demographic subgroups—race, gender, and age—warrants attention. The absence of significant heterogeneity in responsiveness supports the proposition that well-designed GOTV messaging, particularly normative appeals, may have broad applicability across diverse constituencies. This stands in contrast to the assumptions underpinning microtargeting strategies, which often emphasize hyper-segmentation and tailored appeals. While such strategies may be more effective in persuasion contexts, our findings suggest that in turnout-focused campaigns, generic but normatively resonant messages can perform reliably across demographic lines. Moreover, the fact that minority voters did not exhibit backlash or disengagement—despite the lack of culturally specific targeting—reinforces the value of message neutrality and inclusivity, particularly in polarized environments.

# 6.2 Implications for Campaign Strategy

The results of this study carry several implications for the design and implementation of digital GOTV campaigns:

- 1. Message Content Matters Substantially: Campaigns seeking to mobilize voters via digital advertisements should consider emphasizing civic duty, accountability, and participation norms over partian conflict. Social pressure messages that invoke a sense of personal responsibility and social expectation were among the most effective treatments in this study. Conversely, identity-based appeals—framed around partian rivalry—were largely ineffective in a primary context, suggesting such messages may be better suited for general elections where partian stakes are more salient.
- 2. Dosage Should Be Optimized, Not Maximized: The findings provide empirical support for a strategic approach to message repetition. While 10 exposures proved

sufficient to generate a statistically significant boost in turnout, the highest dosage (20 exposures) yielded the largest effect—especially when paired with social pressure content. However, the gains from increased exposure are marginal and may not justify the additional expense in all contexts. Campaigns should weigh the cost of higher frequency against the relatively small incremental returns, and consider focusing higher doses on key voter segments where marginal gains are most valuable (e.g., low-propensity voters in battleground districts).

- 3. Hybrid Messaging Strategies Can Be Effective: The combined pressure-plustraditional message—featuring both a normative cue and a friendly reminder—performed consistently well. This suggests that campaigns need not choose between emotional and informational appeals. Alternating or layering message types can enhance effectiveness by engaging multiple cognitive and emotional pathways. Such hybrid strategies also reduce the risk of message fatigue by providing variation without sacrificing consistency.
- 4. Demographic Segmentation May Be Unnecessary for Turnout-Focused Ads: While demographic microtargeting remains a powerful tool in campaign strategy—particularly for persuasion or fundraising—our results suggest that when the objective is to increase turnout, broadly effective messages can be deployed with equal efficacy across subgroups. This reduces complexity and cost for campaign planners and offers reassurance that generalized civic appeals are unlikely to alienate minority or marginalized communities, provided the messaging avoids racialized or culturally exclusive language.
- 5. Timing Is Critical for Maximizing Impact: The strongest treatment effects were observed among Election Day voters, suggesting that digital advertisements may be most effective when deployed as last-minute reminders. Voters who had not yet cast a ballot during the early voting window were more susceptible to digital nudges, while early voters were likely unreachable or already committed. Accordingly, campaigns should consider intensifying ad delivery in the days leading up to Election Day and

using voter file integrations to retarget those who have not yet participated.

Together, these insights support a pragmatic, evidence-informed approach to digital GOTV strategy—one that emphasizes message quality, dosage efficiency, and timing optimization over maximal frequency or demographic tailoring.

## 6.3 Limitations and Considerations of External Validity

While the study offers meaningful insights, several limitations should temper the scope of its conclusions:

- Context Specificity: The experiment took place in a midterm Republican primary in Texas—a relatively low-salience election with a dominant incumbent. Turnout in the control group was approximately 40%, suggesting a sample of relatively high-propensity voters. In more competitive or high-salience elections, where baseline turnout is higher and voter engagement is more intense, the marginal effect of digital advertisements may be smaller. Conversely, in lower-turnout municipal or special elections, similar treatments may produce larger effects. Generalization to other electoral contexts should therefore be approached with caution.
- Platform Limitation: The study evaluated only Facebook advertisements. Although Facebook remains one of the most widely used platforms for political communication, voters interact with content differently across platforms. Instagram, YouTube, TikTok, and X (formerly Twitter) each have distinct user demographics, interface designs, and content norms. The efficacy of similar messages on other platforms—especially those preferred by younger voters—may differ, and future research should explore cross-platform comparisons.
- Asymmetry of Campaign Effort: This experiment was conducted in partnership with a single campaign. Control group members were not exposed to the campaign's

ads, but they were not isolated from other political communication. In more contested races where multiple campaigns are simultaneously advertising, the relative impact of a single campaign's ads may be attenuated or counteracted. As Broockman and Kalla (2018) demonstrate, campaign effects in such competitive environments can often cancel each other out, resulting in null net effects. However, GOTV messaging—unlike persuasion—is less likely to be zero-sum, as multiple campaigns may simultaneously increase participation.

- Lack of Persuasion Metrics: The study focuses exclusively on turnout and does not examine vote choice, candidate favorability, or message persuasiveness. While these outcomes were not the campaign's primary objective, it remains possible that repeated exposure to certain messages could affect attitudes, issue salience, or trust in political institutions. Future research should include attitudinal and perceptual measures to assess broader consequences of repeated digital exposure.
- Subgroup Coverage and Missing Data: While the study assessed variation by race, gender, and age, it did not examine other potentially important moderators, such as educational attainment, political ideology, or geographic location (e.g., rural vs. urban). Moreover, while Facebook reported impressions data, we could not verify whether individuals actively viewed or engaged with the ads. The analysis relies on an intent-to-treat framework, meaning effects are averaged over all treated individuals regardless of actual exposure. As a result, our estimates are conservative and may underestimate the impact of message content among fully exposed voters.

Despite these limitations, the study's alignment with prior GOTV findings—particularly regarding the effectiveness of social pressure cues—reinforces confidence in the robustness of its conclusions.

### 6.4 Directions for Future Research

The results of this experiment suggest a number of promising avenues for additional inquiry:

- Explore Additional Message Types: Future experiments should test a wider array of message framings, including policy-based appeals, identity-based appeals tied to specific group memberships (e.g., veterans, students), and emotionally evocative content. Prior work by Brader (2005) demonstrates that emotional content—particularly fear and enthusiasm—can shape voter attitudes and behavior, and its effectiveness in the digital realm remains underexplored.
- Examine Explicitly Negative or Contrast Ads: Our identity message was relatively neutral and avoided inflammatory language. Future studies should evaluate whether more aggressive or explicitly negative messaging (e.g., anti-opponent or fearbased appeals) can enhance mobilization or trigger backlash. It is especially important to examine how such messages affect different subgroups and whether they polarize or depress participation.
- Measure Longer-Term Engagement: While this study measured turnout in a single election, it would be valuable to examine whether digital GOTV ads have spillover effects on subsequent participation or civic engagement. Do repeated exposures encourage habit formation, or are effects confined to the immediate electoral cycle? Longitudinal tracking using voter file data could provide answers.
- Compare Platforms and Content Formats: The behavioral effects of digital advertisements may depend not only on content but also on the medium and format. Comparing identical messages delivered via Facebook, YouTube, TikTok, and email, for example, would provide insight into platform effects. Similarly, comparing static images to video content or text-based appeals could reveal differences in cognitive processing and motivational impact.

- Develop Cost-Effectiveness Models: While we demonstrate turnout gains from digital advertising, the question of cost-per-vote remains. Future research should calculate the return on investment for different message types, frequencies, and audience segments, enabling campaigns to allocate limited budgets more strategically.
- Unpack Behavioral Mechanisms Through Qualitative Inquiry: Finally, combining experimental methods with surveys or interviews would help uncover the psychological pathways activated by different messages. Do recipients feel guilt, civic pride, partisan loyalty, or annoyance? Understanding these mechanisms would allow researchers to design more effective, targeted, and ethically grounded communication strategies.

In sum, while this dissertation answers key questions about digital GOTV effectiveness, it also opens the door to deeper theoretical and applied exploration. As campaign technologies evolve and voter behavior adapts, continual empirical investigation will be essential for refining our understanding of how political communication shapes democratic participation.

# CHAPTER 7

# Conclusion

This dissertation has presented the results of a large-scale, individual-level randomized field experiment designed to assess the causal effects of digital campaign advertising—specifically, Facebook advertisements—on voter turnout in the context of a real-world electoral setting. The findings of this study add theoretical nuance and empirical depth to the literature on political communication, campaign effects, and digital mobilization. Importantly, the study engages with enduring questions in political science concerning the extent to which electoral behavior can be shaped by campaign contact, the mechanisms through which voters are mobilized, and the broader implications of new media technologies for democratic participation.

# 7.1 Summary of Findings

The results demonstrate that digital campaign advertising, when executed with appropriate message content and exposure frequency, can yield small but statistically significant increases in voter turnout. The effect sizes—ranging from 0.5 to 1.2 percentage points depending on message type and dosage—are consistent with prior research emphasizing the limited but nontrivial role that campaign outreach plays in mobilizing voters. These findings affirm the central tenets of the "minimal effects" paradigm (Broockman and Kalla 2018), while simultaneously offering new evidence that digital GOTV efforts can overcome some of the constraints that have historically limited the impact of campaign communication.

Crucially, the results reveal that the effectiveness of digital ads is highly contingent on message framing. Appeals to civic duty, accountability, and social norms—commonly referred to as social pressure messages—proved to be significantly more effective in increasing turnout than identity-based messages emphasizing partian in-group loyalty. This outcome aligns with the theoretical insights of behavioral political science, which highlight the power of norm activation and reputational concerns in motivating political behavior (Gerber, Green, and Larimer 2008). By contrast, the relative ineffectiveness of partian identity appeals—especially in the context of a primary election—underscores the conditional nature of identity salience. Without a clearly defined out-group threat or partian stakes, voters may fail to internalize such messages as personally motivating.

Additionally, the study sheds light on the role of exposure frequency, or "dosage," in shaping campaign effectiveness. While all dosage levels produced positive effects relative to the control group, higher exposure levels (particularly 20 impressions) generated the largest turnout gains. These findings suggest a weakly increasing relationship between repetition and behavioral impact, at least within the frequency bounds tested. Importantly, there was no evidence of backlash or diminishing returns within the tested range, a result that runs counter to concerns about ad fatigue or voter annoyance in saturated media environments. This reinforces the notion that digital advertisements, when carefully designed, can be repeated without adverse effects and may benefit from cumulative exposure.

# 7.2 Theoretical Contributions

Theoretically, this study contributes to multiple strands of literature. First, it reinforces the principle that campaign effects are context-dependent and that message content is critical. The comparative advantage of social pressure appeals supports theories of conditional motivation, which posit that behavior is often shaped more by social cues and perceived expectations than by abstract ideology or partian identification. Moreover, this study extends the validity of those theories into the domain of digital media, demonstrating that messages traditionally delivered through mailers or canvassing can be adapted to digital formats without losing their psychological efficacy. Second, the study speaks to ongoing debates about microtargeting and identity politics in digital campaigning. While the technological capacity for hyper-personalization has increased dramatically, the empirical case for its necessity in turnout campaigns is far from conclusive. The lack of significant treatment heterogeneity by race, gender, or age suggests that broad-based, normatively framed messages can be effective across diverse demographic segments. This finding challenges the assumption that only demographically tailored messages can achieve behavioral impact and supports a more parsimonious approach to digital mobilization—one that focuses on cognitive universals such as duty, obligation, and social accountability.

Third, the findings have implications for models of political participation. The turnout effects observed here were concentrated among Election Day voters, indicating that digital ads may be most effective in activating individuals who delay decision-making or require lastminute prompts. This insight is consistent with dual-process theories of political behavior, which distinguish between habitual and deliberative participation. For habitual voters, ads may serve as benign reminders; for non-habitual or uncertain voters, they may provide the motivational nudge necessary to convert latent intention into action.

### 7.3 Practical Implications

From a practitioner's standpoint, this study offers a set of concrete recommendations for political campaigns seeking to increase voter turnout via digital advertising:

- **Prioritize Normative Message Framing:** Messages emphasizing civic duty, public accountability, and the normativity of voting behavior outperformed partian messages in all tested conditions. Campaigns should consider foregrounding these themes in their digital GOTV efforts.
- Invest in Adequate Dosage: The highest levels of ad exposure (20 impressions) generated the largest turnout effects, suggesting that repeated contact can amplify

message salience. When budget permits, increasing frequency for promising messages is a viable strategy, particularly in the final days of a campaign.

- Use Hybrid Strategies Thoughtfully: The combination of a social pressure cue with a friendly reminder performed well in this experiment, demonstrating that message synergy is possible. Campaigns should test blends of emotional and informational appeals to maximize effectiveness.
- Simplify Targeting in GOTV Campaigns: Given the lack of significant treatment heterogeneity, campaigns may not need to develop highly segmented GOTV messages based on race, gender, or age. Generalized messages that resonate broadly can be deployed more efficiently and cost-effectively.
- Time Ads to Peak Voter Decision Windows: The data suggest that voters who had not yet voted by Election Day were more susceptible to digital nudges. Campaigns should consider intensifying their digital outreach in the final week before Election Day and utilizing voter file integrations to retarget likely voters who have not yet cast ballots.

### 7.4 Limitations and Scope Conditions

While the findings of this study are robust and statistically well-powered, several limitations must be acknowledged. First, the study was conducted in a Republican primary in Texas—a relatively low-salience election with an uncompetitive incumbent. The generalizability of the results to general elections, high-stakes contests, or other partian environments remains uncertain. Future research should replicate this experiment across different electoral contexts to assess external validity.

Second, the platform of interest—Facebook—represents only one node in the broader digital ecosystem. While Facebook remains a dominant player in political advertising, voter behavior and ad engagement can vary across platforms (e.g., Instagram, YouTube, TikTok), especially as user demographics evolve. Additional research is needed to test the portability of these findings across platforms and media formats.

Third, the experimental design measured turnout but not persuasion, candidate favorability, or long-term engagement. Thus, while the ads were effective in mobilizing voters in the short term, their broader political consequences remain unknown. Incorporating attitudinal measures or longitudinal follow-up could illuminate whether digital GOTV campaigns foster deeper forms of political engagement or merely serve as one-time nudges.

Finally, while the use of an intent-to-treat framework enhances causal inference, it may also obscure variation in exposure and attention. Not all individuals assigned to treatment would have seen or cognitively processed the ads, and the actual "dose received" likely varied. Without detailed engagement data from the platform (e.g., views, clicks, watch time), perprotocol effects cannot be estimated with confidence.

## 7.5 Future Research Directions

Building on the results of this study, future research should address several open questions:

- Comparative Testing Across Elections and Platforms: Replicating this design across different electoral environments (e.g., general elections, municipal contests) and digital platforms would improve our understanding of context- and platform-specific effects.
- Behavioral Mechanism Analysis: Qualitative surveys or embedded experiments that measure intermediate outcomes—such as recall, motivation, perceived pressure, or emotion—would clarify why particular messages work and how voters interpret them.
- Content Variation and Message Framing: Future studies should test a wider array of message types, including appeals based on policy, representation, negative

partisanship, and emotional content. Prior work by Brader (2005) suggests that emotion, especially enthusiasm and fear, may play a critical role in mobilization.

- **Cost-Effectiveness Modeling:** Estimating the cost per additional vote generated through digital campaigns remains a high priority for practitioners. Future work should integrate cost data and explore diminishing marginal returns across budget scenarios.
- Longitudinal Outcomes: Measuring whether digital mobilization has persistent effects on voter behavior over multiple election cycles would provide insight into whether these interventions foster habit formation or civic learning.

# 7.6 Concluding Reflections

This dissertation offers one of the most comprehensive individual-level evaluations to date of how digital advertising—specifically on Facebook—affects voter turnout. By combining experimental rigor with practical relevance, it contributes to both academic and applied conversations about the role of campaign communications in modern democracy. In an era of political polarization, misinformation, and declining trust in institutions, the ability to increase participation—even marginally—through scalable, cost-effective, and normatively framed interventions is a finding of genuine importance.

While the road to more engaged and equitable participation will not be paved by digital advertising alone, this study shows that such tools can be part of the solution. In the increasingly crowded digital marketplace of ideas, it appears that simple appeals to civic responsibility—and perhaps a subtle reminder that voting is both expected and observable—can still cut through the noise.

# APPENDIX A

# Appendix A: Supplementary Tables and Robustness Checks

This appendix presents robustness checks including tests for multicollinearity (using VIF), heteroscedasticity (Breusch-Pagan), and visual diagnostics like residual and leverage plots. I'll also explore subsample analyses (e.g., by voting propensity or timing) to assess the consistency of treatment effects.

#### A.0.1 Multicollinearity Diagnostics

One potential concern in the regression analyses is multicollinearity among the treatment indicators, message framing dummies, dosage levels, and demographic covariates. To assess this, we computed variance inflation factors (VIF) for the OLS model specifications including the full set of predictors (treatment assignment or message-type indicators, dosage, age, gender, and race). Table A.1 reports the VIF values for the key independent variables. All VIF values are well below the common heuristic threshold of 10 for serious multicollinearity problems. The highest VIF is around 24 for the indicator of white voters, which reflects the fact that the race dummies are highly correlated (e.g., being White is the complement of the other racial categories in the data). Similarly, the VIF for the linear dosage variable is about 14 and for age about 22, indicating some collinearity between those covariates and the treatment/message indicators (since dosage is zero for the control group and age is somewhat correlated with race in this sample). Importantly, the core experimental variables—such as the message-frame dummies and the treatment indicator—exhibit only moderate multicollinearity (VIF  $\approx 6-7$ ), suggesting that any inflation of standard errors due to multicollinearity is minimal. We conclude that multicollinearity is not a serious concern and does not distort the OLS estimates of treatment effects.

Predictor	VIF
Social Identity message (dummy)	6.27
Social Identity & Social Pressure message	6.24
Social Pressure message	6.31
Social Pressure & Traditional message	6.33
Traditional message	6.32
Female (gender dummy)	9.46
Male (gender dummy)	9.87
Asian (race dummy)	1.44
Black (race dummy)	2.14
Hispanic (race dummy)	3.31
White (race dummy)	24.57
Dosage (number of ad exposures)	14.54
Age (in years)	22.08

 Table A.1: Variance Inflation Factors for Key Independent Variables

*Note:* VIF values were calculated for an OLS model including all message conditions, dosage, and demographic covariates. No VIF exceeds the rule-of-thumb threshold of 10 for high multicollinearity. The relatively higher VIF for the White race dummy is a byproduct of the categorical coding of race (with White being the largest group and thus highly negatively correlated with the other race dummies). Overall, the results indicate that multicollinearity is not severe.

#### A.0.2 Heteroscedasticity Tests

We next examined whether the linear probability models violated the assumption of homoscedastic residuals. Because the dependent variable (voter turnout) is binary, some degree of heteroscedasticity is expected in an OLS framework (the variance of the error term can depend on the predicted turnout probability). We conducted Breusch–Pagan tests for heteroskedasticity on the core regression models. In each case, the Breusch–Pagan test strongly rejects the null hypothesis of constant error variance (p-value < 0.001). This indicates the presence of heteroscedasticity in the OLS models, likely due to the binary nature of the outcome and minor group-level variance differences. As a result, all OLS regression results reported in the dissertation employ robust (Huber–White heteroskedasticity-consistent) standard errors to ensure valid inference despite the non-constant variance.

#### A.0.3 Regression Diagnostic Plots

Figures A.1–A.3 present diagnostic plots for the main OLS regression model (linear probability model for turnout with full covariates). These plots help evaluate model fit and the OLS assumptions.

Figure A.1 plots the residuals versus fitted values. The residuals appear symmetrically distributed around zero across the range of predicted turnout values, with no obvious nonlinear patterns. We do observe a slight "funnel" shape – the spread of residuals is somewhat larger for moderate fitted values than for extreme fitted values – which is consistent with heteroscedasticity due to the binary outcome (predicted turnout around 50% yields higher variance in residuals, whereas predictions near 0 or 1 constrain the possible residuals).



Figure A.1: Residuals vs. Fitted Values for OLS Turnout Model. Each point represents an individual voter's residual (actual minus predicted turnout) plotted against the model's fitted turnout probability for that voter. The red dashed line at 0 indicates perfect prediction. The spread of residuals is relatively constant around zero, though slightly wider in the mid-range of fitted values, consistent with heteroscedasticity from the binary outcome.

Figure A.2 shows a normal Q–Q plot of the residuals. The residual distribution exhibits mild deviations from normality at the tails: there are slightly more extreme residuals than a normal distribution would predict (the points stray from the 45-degree line at the extremes). This heavy-tailed pattern is again attributable to the binary outcome – many observations have residuals of large magnitude (near  $\pm 0.5$  or beyond) when the predicted probability differs from the actual 0/1 outcome. Nonetheless, in the central portion of the distribution the residuals roughly follow a straight line, suggesting no severe departures from normality aside from the expected binary-outcome effects.



Figure A.2: Normal Q–Q Plot of OLS Residuals. The quantiles of the residuals from the turnout regression are plotted against theoretical quantiles from a normal distribution. Deviations from the 45° line indicate departures from normality. The plot shows moderately heavier tails than a normal distribution (points at the extremes lie off the line), reflecting the excess of large residuals inherent in a linear model for a binary outcome. Overall, residuals are approximately normal in the center of the distribution but with some heavy-tail behavior at the extremes.

Finally, Figure A.3 plots residuals against leverage (Hat values) for each observation. We do not find any data points with exceedingly high leverage or influence. Given the very large sample size  $(N \approx 7.78 \times 10^5)$ , each individual observation has negligible influence on the overall fit. The leverage values are uniformly very small (all below 0.001), and no observation has an outsized Cook's distance. In sum, the diagnostic plots do not reveal any worrisome model mis-specification: the linear probability model fits without systematic bias, and no



single observations or outliers unduly affect the results.

Figure A.3: Residuals vs. Leverage for OLS Turnout Model. Each point shows an observation's residual plotted against its leverage (Hat value). No observation exhibits extremely high leverage or an unusually large residual. The absence of points with both high leverage and large residual suggests there are no influential outliers unduly driving the results. Cook's distance values for all observations were near zero (not shown), further indicating no single data point has disproportionate influence on the OLS estimates.

# APPENDIX B

# **Appendix B: Advertisement Transcripts**

The following are full transcripts of the four digital advertisements shown to participants in the treatment and control groups. Each advertisement was approximately 15 seconds in length and concluded with the same closing bumper card: "VOTE! Tuesday, March 6th". The content was designed to reflect a variety of message frames, including civic duty, identity, and social pressure.

### Advertisement 1: Social Pressure – Analyst Warning Frame

**Voice-Over (Male)** [0:00]: You are being studied. A statewide study is being conducted to see if people do or do not vote.

**Bumper Card** [0:00]: ALERT! AN IMPORTANT STATEWIDE STUDY IS BEING CON-DUCTED.

**Voice-Over (Male)** [0:06]: They will be analyzing public records of who turns out in the March 6th primary election. Do your civic duty. Vote.

**Bumper Card** [0:06]: ALERT! ANALYSTS WILL BE REVIEWING PUBLIC RECORDS OF THE MARCH 6TH VOTER TURNOUT.

Voice-Over (Male) [0:10]: Do your civic duty.

Bumper Card [0:10]: ALERT! DO YOUR CIVIC DUTY...

Voice-Over (Male) [0:12]: Vote.

Bumper Card [0:12]: VOTE! TUESDAY, MARCH 6TH

#### Advertisement 2: Traditional GOTV – Texas First Lady Message

Voice-Over (Female) [0:00]: This is Texas First Lady Cecilia Abbott...

Bumper Card [0:03]: PRIMARY VOTING IS MARCH 6TH!

Voice-Over (Female) [0:03]: ...reminding you to vote in the March Primary Election.

Voice-Over (Female) [0:07]: Don't miss your chance to vote. Find your voting location and cast your ballot.

Bumper Card [0:12]: VOTE! TUESDAY, MARCH 6TH

#### Advertisement 3: Social Pressure – Voting Records Public Frame

#### Dialogue

Man [0:00]: It's unbelievable what goes on in this country.

Woman [0:03]: (scoffs) You can't complain. You haven't voted in years.

**Man** [0:06]: Nobody knows that.

Woman [0:07]: Oh yeah? Our neighbors could. You know voting records are public.

**Man** [0:11]: (Looks at woman with concern across his face.)

Bumper Card [0:12]: VOTE! TUESDAY, MARCH 6TH

#### Advertisement 4: Identity-Based – In-Group Mobilization Frame

### Dialogue

**Man** [0:00]: It's unbelievable what goes on in this country.

Woman [0:03]: And what's worse? They're energized and they're voting.

**Man** [0:06]: So, this means...

Woman [0:07]: This means we need to go vote so they don't take over.

Bumper Card [0:12]: VOTE! TUESDAY, MARCH 6TH

84

### BIBLIOGRAPHY

- Anspach, Nicolas M. 2017. "The new personal influence: How our Facebook friends influence the news we read." *Political Communication* 34 (4): 590–606.
- Arceneaux, Kevin, and Robin Kolodny. 2009. "Educating the Least Informed: Group Endorsements in a Grassroots Campaign." American Journal of Political Science 53 (4): 755–770.
- Baum, Matthew A. 2021. Campaign Advertising and Democratic Accountability. Cambridge, UK: Cambridge University Press.
- Boghani, Priyanka. 2018. Brad Parscale, Trump's 2020 campaign manager, calls Facebook ad policy a 'gift'. https://www.pbs.org/wgbh/frontline/article/brad-parscale-trumps-2020-campaign-manager-calls-facebook-ad-policy-a-gift/. Accessed: 2024-03-26.
- Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D. I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. 2012a. "A 61-million-person experiment in social influence and political mobilization." *Nature* 489 (7415): 295–298.
- ———. 2012b. "A 61-million-person experiment in social influence and political mobilization." *Nature* 489 (7415): 295–298.
- Brader, Ted. 2005. "Striking a Responsive Chord: How Political Ads Motivate and Persuade Voters by Appealing to Emotions." American Journal of Political Science 49 (2): 388– 405.
- Broockman, David E., and Joshua L. Kalla. 2018. "The Minimal Persuasive Effects of Campaign Contact in General Elections: Evidence from 49 Field Experiments." American Political Science Review 112 (1): 148–166.

- Burns, Nancy, Kay Lehman Schlozman, and Sidney Verba. 2001. The Private Roots of Public Action: Gender, Equality, and Political Participation. Cambridge, MA: Harvard University Press.
- Cadwalladr, Carole. 2018. Cambridge Analytica: 'We know we changed votes and swung elections'. https://www.theguardian.com/membership/2018/sep/29/cambridgeanalytica-cadwalladr-observer-facebook-zuckerberg-wylie. Accessed: 2024-03-26.
- Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. 1980. *The American Voter*. University of Chicago Press.
- Coleman, Stephen, and Deen Freelon. 2015. *Digital Politics: Mobilization, Engagement, and Participation*. Polity Press.
- Collins, Jason. 2018. "Campaigning for Votes Online: Facebook Ads, Clicks, and Effects on Political Behavior." *Journal of Political Marketing* 17 (2): 163–181.
- Coppock, Alexander, Donald P. Green, and Ethan Porter. 2022. "Is your treatment actually too strong? How average treatment effects can be biased when treatment effects are heterogeneous." *Journal of Experimental Political Science* 9 (1): 14–26.
- Coppock, Alexander, Seth J. Hill, and Lynn Vavreck. 2020. "The Small Effects of Political Advertising Are Small Regardless of Context, Message, Sender, or Receiver: Evidence from 59 Real-Time Randomized Experiments." Science Advances 6 (36): eabc4046.
- Downs, Anthony. 1957a. An Economic Theory of Democracy. New York: Harper / Row.
- ———. 1957b. An Economic Theory of Democracy. New York: Harper / Row.
- Eagly, Alice H. 1987. Sex Differences in Social Behavior: A Social-Role Interpretation. Hillsdale, NJ: Lawrence Erlbaum Associates.

- Enos, Ryan D., Anthony Fowler, and Lynn Vavreck. 2014. "Increasing Inequality: The Effect of GOTV Mobilization on the Composition of the Electorate." *The Journal of Politics* 76 (1): 273–288.
- Gerber, Alan S., and Donald P. Green. 2012. Field Experiments: Design, Analysis, and Interpretation. New York: W. W. Norton & Company.
- Gerber, Alan S., Donald P. Green, and Christopher W. Larimer. 2008. "Social Pressure and Voter Turnout: Evidence from a Large-Scale Field Experiment." American Political Science Review 102 (1): 33–48.
- Green, Donald P., and Alan S. Gerber. 2015. *Get Out the Vote: How to Increase Voter Turnout.* 3rd. Brookings Institution Press.
- Haenschen, Katherine, and Jay Jennings. 2019. "Mobilizing Millennial Voters with Facebook Ads: A Field Experiment." *Political Communication* 36 (3): 311–329.
- Hager, Anja. 2019. "Can Facebook really persuade? A study of political advertising online." Journal of Information Technology & Politics 16 (3): 278–292.
- Hersh, Eitan. 2020. Politics Is for Power: How to Move Beyond Political Hobbyism, Take Action, and Make Real Change. New York: Scribner.
- Karpf, David. 2012. The MoveOn Effect: The Unexpected Transformation of American Political Advocacy. Oxford University Press.
- Kim, Young Mie, Jordan Hsu, David Neiman, Colin Kou, Levi Bankston, Soo Yun Kim, Richard Heinrich, Robyn Baragwanath, and Garvesh Raskutti. 2018. "The Stealth Media? Groups and Targets behind Divisive Issue Campaigns on Facebook." *Political Communication* 35 (4): 515–541.
- Kreiss, Daniel. 2016. Prototype Politics: Technology-Intensive Campaigning and the Data of Democracy. Oxford University Press.

- Malhotra, Neil, Melissa R. Michelson, and Ali Adam Valenzuela. 2012. "Emails from Official Sources Can Increase Turnout." Quarterly Journal of Political Science 7 (3): 321–332.
- Michelson, Melissa R. 2005. "Meeting the Challenge of Latino Voter Mobilization." The ANNALS of the American Academy of Political and Social Science 601 (1): 85–101.
- Nickerson, David W. 2006. "Volunteer Phone Calls Can Increase Turnout: Evidence From Eight Field Experiments." *American Politics Research* 34 (3): 271–292.
- Panagopoulos, Costas. 2010. "Affect, Social Pressure, and Prosocial Motivation: Field Experimental Evidence of the Mobilizing Effects of Pride, Shame, and Publicizing Voting Behavior." *Political Behavior* 32 (3): 369–386.
- Project, Wesleyan Media. 2021. Political Advertising in the 2020 Election. https://mediapr oject.wesleyan.edu/releases/2020-election-final-report/. Accessed: 2024-03-26.
- Ramakrishnan, Karthick, and Mark Baldassare. 2004. "The Ties That Bind: Changing Perspectives on Ethnicity and Political Participation." *Political Behavior* 26 (4): 349–378.
- Riker, William H., and Peter C. Ordeshook. 1968. "A Theory of the Calculus of Voting." American Political Science Review 62 (1): 25–42.
- Roberts, Kyle. 2020. OnAir: Breaking Down Online Advertising Trends. https://www. cookpolitical.com/analysis/national/political-advertising/onair-breaking-down-onlineadvertising-trends. Accessed: 2024-03-26.
- Stromer-Galley, Jennifer. 2014. Presidential Campaigning in the Internet Age. Oxford University Press.
- Tucker, Joshua A., Andrew Guess, Pablo Barberá, Cristian Vaccari, Alexandra Siegel, Sergey Sanovich, Denis Stukal, and Brendan Nyhan. 2018. "Social Media, Political Polarization, and Political Disinformation: A Review of the Scientific Literature." *Political Science Research and Methods* 6 (4): 511–530.

- West, Darrell M. 2013. Air Wars: Television Advertising and Social Media in Election Campaigns, 1952-2012. 6th. Washington, D.C.: CQ Press.
- Williams, Christine B., and Girish J. Gulati. 2013. "Social Networks in Political Campaigns: Facebook and the Congressional Elections of 2006 and 2008." New Media & Society 15 (1): 52–71.