

Scaling the Field Program in Modern Political Campaigns

Investigating Determinants of Capacity in Mobilizing and Organizing

By Ben McGuire

Harvard Kennedy School of Government

Master in Public Policy Candidate

Expected Graduation: May 30, 2019

Policy Analysis Exercise

Democratic National Committee Analytics Team

Advisor: Professor Marshall Ganz

Seminar Leader: Professor Thomas Patterson

Submission Date: April 2, 2019

Table of Contents

<i>Executive Summary</i>	3
<i>Literature Review and Background</i>	5
What Is A Field Program?	5
The Impact of Field Program Outcomes on Voter Mobilization.....	7
The Research Gap in the Organizing Logic Model	9
Mobilizing Versus Organizing: Incentives and Strategies	11
Program Quality and Relationships: The Problem Is Scale	14
<i>Data, Methodology, and Methods</i>	16
Data Sources	16
Interviews and Acknowledgements	17
<i>Volunteerism and Individual Participation</i>	18
Measuring Volunteerism.....	18
Types of Shift Completed	19
Timing of First Engagement	19
<i>Scaling the Field Program for Mobilization</i>	27
Measuring Field Program Capacity.....	27
Aggregate Shift Volume and Type	27
How Timing Affects the Impact and Value of pre-GOTV Shifts	30
<i>Scaling the Field Program for Organizing Capacity</i>	35
Measuring Organizing Capacity.....	35
How Timing Affects Shift Impact on Organizing Capacity.....	38
<i>Synthesizing Findings</i>	41
<i>Recommendations Based on Initial Findings</i>	43
Field Experiments in Future Cycles.....	43

Executive Summary

This Policy Analysis Exercise explores factors that influence mobilizing and organizing capacity within field programs of political campaigns. Previous research has identified clear links between voter contact programs in the final days of a campaign and turnout among those who the campaign contacts. However, comparatively little research to-date has explored how activities within the cycle influence the ability of the campaign organization to drive capacity.

Part of the research challenge is availability; most campaigns do some kind of get out the vote (GOTV) activity, but relatively few build organizing programs large, coherent, and early enough in the cycle to conduct rigorous experiments. This report uses volunteer engagement data from the last two presidential cycles to identify potential relationships between volunteer engagement, and measures of capacity in a GOTV operation.

These findings are observational, not experimental or quasi-experimental. It is difficult and, in many areas, impossible to escape endogeneity problems within the volunteer population as well as within the distribution and activities of field offices based purely on findings of what happened in the past. For this reason, these initial findings should not be taken as rigid prescriptions, but rather as promising recommendations to inform future experimentation and analysis. The final section of this report proposes some guidelines for field experiments which could help to isolate the impact of these factors and inform future field management.

Engaging volunteers early in the cycle is associated with greater levels of aggregate activity, but late-engaged volunteers are in general more likely to participate in get-out-the-vote (GOTV) activities. One in six volunteers who had their first shift completion in July participated in GOTV; over one in four who began in October participated in GOTV. **On average, completing at least one door-to-door canvass shift during the cycle is associated with much higher levels of individual volunteerism (i.e., completing shifts as well as showing up for GOTV) than completing either a one-on-one shift or phone bank.** Sixty percent of those who did at least one canvass shift during the cycle participated in GOTV, compared to 11 percent of those who never canvassed.

Completing more one-on-one shifts in the aggregate is associated with scheduling more GOTV shifts in total; completing more canvass shifts is associated with scheduling more canvass shifts during GOTV. A one percent increase in total unique cycle volunteers is associated with roughly a one percent increase in GOTV canvass shifts scheduled. Having one additional volunteer in the office has an unclear impact on capacity development when controlling for time effects across the cycle; the estimated impact of both one-on-one shifts and canvass shifts on GOTV capacity fall at the end of the cycle.

Using both total GOTV canvass shifts scheduled as well as the share of cycle volunteers engaged in GOTV as output metrics, the incremental estimated value of shifts across the cycle follows an interesting and consistent pattern.

- **One-on-one shifts are on average most valuable in developing mobilizing and organizing capacity roughly three months before GOTV begins.** An additional one-on-one shift completed in the July-August timeframe is associated with close to a one percentage point increase in the share of cycle volunteers who participate in GOTV, and an additional 7.5 canvass shifts scheduled.
- **Canvass shifts have a small estimated impact on capacity very early and very late in the cycle, and peak in estimated impact just over two months before GOTV begins; the statistical significance of this relationship increases throughout the cycle.** An additional canvass shift completed in the August-September timeframe is roughly associated with 0.8 percentage point increase in the share of cycle volunteers who participate in GOTV.
- **Phone shifts have a consistently weaker relationship with GOTV capacity at the turf level than either one-on-one or canvass shifts.**

These data suggest that in order to develop a field program that has maximum ability to build capacity both in terms of mobilization volume and organizing strength, programs should front-load one-on-one meetings with volunteers as a way to drive mid-cycle capacity in canvass shifts that provide offices with the ability to practice logistics and teamwork that are central to GOTV. Phone shifts should remain a large part of daily operations but should be evaluated more in their ability to drive other kinds of recruitment and action than in their ability to drive capacity independently.

However, these figures are historical observation, not tested in conditions that can allow us to avoid endogeneity problems. To more rigorously test these theories, I recommend that the Democratic National Committee use small-scale pilots of intentionally-designed programs and use randomized control trials of different models against outcomes that are low-risk (e.g., summer organizing conferences) but which can provide early-stage evidence.

Literature Review and Background

What Is A Field Program?

What are the purposes and values of a political campaign? Who does it serve, and what should we expect from it? These are not idle questions. Thanks to huge infusions of money and technology in the last three decades, a modern Presidential campaign has become either an impressive juggernaut of optimized technology delivering relevant messages to citizens or a cold, clinical machine that has lost touch with the beating heart of American democracy. And in the wake of a Presidential election that saw a well-funded, well-organized, and well-planned campaign lose to an ad-hoc, upstart, chaotic outsider, every conventional wisdom about the correct way to organize and manage a campaign is rightly being questioned.

At the center of the debate lies the field program—phone and door-to-door contact with voters to convince them to support a candidate, volunteer, and show up to vote on or before Election Day. In a time when campaigns can leverage technology tools and the Internet to influence voters at massive scale, the field organization of a political campaign—once at the heart of strategy and operations—can feel very antiquated. And because field programs are such a visible piece of the campaign, they are the subject of substantial armchair punditry: a field program was lauded for handing two elections to President Obama¹; critics pilloried Secretary Clinton’s 2016 campaign for not doing enough in the field.² Both claims are likely out-sized relative to the real impact that voter contact can have in swinging public opinion or turnout.

Field programs have moved in and out of central importance in American elections; face-to-face connection was a critical component of elections controlled party bosses and precinct captains from the early Republic up through the 1960’s.³ But Democratic Party electoral reforms in 1976 and the rise of mass media technologies hobbled once-preeminent local party leaders in the Presidential process, while making name recognition and message broadcasting for all candidates dramatically easier.⁴ Field came back into vogue around the turn of the 21st century, when the 2004 Democratic primary effort for Vermont Governor Howard Dean and re-election campaign of President George W. Bush invested heavily in individually-targeted field outreach programs. The difference in these campaigns was not that targeting and individual optimization was new; going back to the mid-nineteenth century, President Lincoln’s campaign plan to generate lists of voter targets at the local level based on past behavior “could be mistaken for a modern field strategy.”⁵ Rather, it was the combination of old organizing techniques for local turnout, married to new technologies for data management, analysis, and

¹ See for example: <https://fivethirtyeight.com/features/obama-leads-better-than-31-in-field/>.

² See for example: <https://www.progressivepolicy.org/blog/clinton-lost-ground-game-view-trenches/>

³ Forthal, Sonya. *Cogwheels of Democracy: A Study of the Precinct Captain*. Praeger. 1946.

⁴ Han, Hahrie and McKenna, Liz. *Groundbreakers: How Obama’s 2.2 Million Volunteers Transformed Campaigning in America*. Oxford University Press. 2014.

⁵ Han, Hahrie and McKenna, Liz. *Groundbreakers*. Page 31.

voter contact (technologies which were also critical for Dean's low-dollar fundraising). The Dean campaign in New Hampshire, outmatched by well-funded opponents, embraced an approach focused on shared commitments and organizing: house meetings, one-on-one conversations,⁶ and close relationships between local organizers that drove expanding networks of volunteers.⁷ The George W. Bush reelection campaign focused on targeted contacts and mobilization: they matched volunteers to voters through shared characteristics (e.g., veteran status, age of children) to enhance the relevance of voter contact messaging.⁸ Both campaigns demonstrated the power of local volunteers in bringing voter contact programs not only to a large scale, but also connecting with people in a way that felt relevant and personal. As subsequent research summarizes: "voters are less likely to be influenced by campaigners from other communities."⁹

One critical innovation came before the 2008 Democratic presidential primary, when a large-scale study by Marshall Ganz and Ruth Wageman on leadership development and organizational capacity in the Sierra Club demonstrated a potential framework by which campaigns could structure their organizing approach. Rather than think of volunteer organizations and local chapters as distributed units doing low-level production labor, the authors found a huge value came from "investment in creating the individual and collective organizational capacity required to accomplish specific strategic goals."¹⁰ This approach directly informed the work of President Obama's 2008 field campaign, which developed a 'snowflake' model for organizers and volunteer teams to develop capacity not only for completing more activities, but also for building leadership and team connections which would help a local organization be more than the sum of its individual parts.

With this model as a rough template and a data infrastructure in place to manage operations, field programs for Democratic general elections since then have followed a somewhat predictable pattern. Early on, the priority is shifts to develop volunteers through one-on-one meetings, recruit volunteers over the phone, and register voters. Over time, door-to-door canvass shifts are added as the office develops capacity and organizational strength, until by the end of the campaign voter all types of shifts but phone bank and canvass are abandoned to drive GOTV capacity. This report includes analysis of one-on-one meetings and two of the three main voter-contact shifts: phone bank and canvass shifts. Voter registration shifts are excluded because they do not have GOTV capacity implications, but this should not be construed as devaluing their importance. Indeed, many of the insights about developing capacity around GOTV are relevant for increasing scale in registration.

⁶ I.e., 45-60-minute conversations between staff and volunteers to discuss goals and roles for the volunteer in the field program.

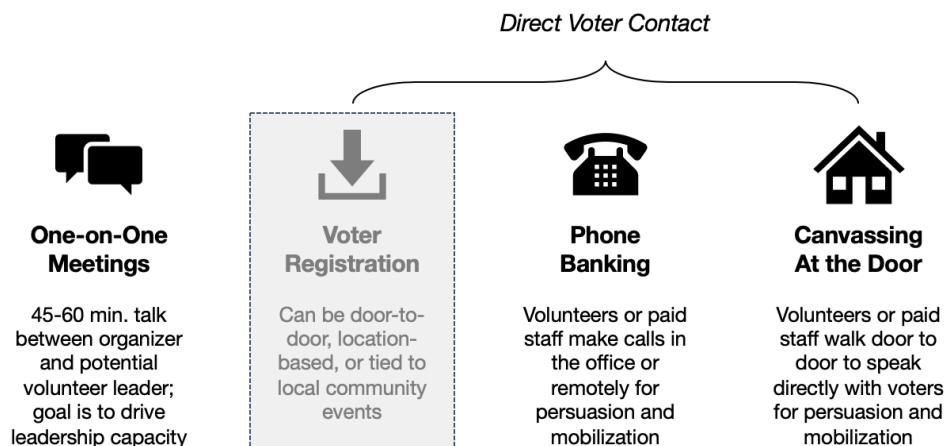
⁷ Author interview with Karen Hicks.

⁸ Han and McKenna, *Groundbreakers*, 41.

⁹ Nickerson David W., Feller Avi. *Can Voter Turnout Contaminate a Neighborhood?* Paper presented at Harvard Networks in Political Science Conference, June 13, 2008. Cambridge, MA, USA.

¹⁰ Marshall Ganz and Ruth Wageman. *Sierra Club Leadership Development Project: Pilot Project Report and Recommendations*. Leadership Development Project. May 2008.
<http://marshallganz.usmblogs.com/files/2012/11/LDP-Final-Report.pdf>.


Shift Definitions Included In This Report



The graph below shows the average distribution of shifts in each week in the 2012 and 2016 general-election cycles; with some minor differences, the shape of this distribution is likely similar across a range of campaigns at the Congressional, state-wide, and national level.

Average Share of Volunteer Shifts Completed by Week

Averaging 2012 and 2016 General Election Campaigns



The Impact of Field Program Outcomes on Voter Mobilization

A number of randomized field experiments over the last several decades have confirmed the impact of door-to-door and phone contacts with voters in boosting turnout. The literature is mixed on if these tasks are best done by volunteers or if it is possible to achieve the same

kinds of impact with paid professionals. Conventional wisdom prioritizes local volunteers as always and substantially better than paid canvass or phone calls. Many studies have identified geographic proximity of canvasser as a key factor in determining the effectiveness of contact in increasing turnout.¹¹ However, David Nickerson has demonstrated that the relationship may be more complex. Volunteer operations can be ineffective; robust investments in paid outreach can generate high-quality programs.¹² But if generating high-quality paid programs are a question of levels of resource investment, the best way to improve the quality of volunteer operations remains an open question. The consequences of these programs and the urgency of finding good answers are quite high; Enos and Fowler estimate that the “2012 presidential campaigns increased turnout in highly targeted states by 7–8 percentage points, on average.”¹³

In *Get Out The Vote: How To Increase Voter Turnout*, Donald P. Green and Alan S. Gerber maintain a summary of different campaign programs and the field experiments that demonstrate the effects they have (or do not have). Gerber and Green’s meta-analysis summarizes both qualitative advice about how to launch a field campaign as well as the quantitative evidence that is backed by repeated scientific experiments; findings follow:¹⁴

- **Door-to-door Canvass:** “One vote per 15 contacts plus effects of spillover on housemates;” the effect is statistically significant; “At \$16 per hour and 6 contacts per hour, one vote costs \$31”
- **Volunteer Phonebank:** “One vote per 35 contacts;” the effect is statistically significant; “At \$16 dollars per hour and 16 contacts per hour, one vote costs \$35”

Because door-to-door canvassing has the strongest claim to both impact and value, canvass shifts will be one of the main focuses of this report. That is not to say that phone bank contacts via volunteers are not a critical way to conduct outreach to volunteers as well as voters; many voters (e.g., those who live in far-flung rural areas, apartment buildings, retirement communities) might only be reachable by phone. Rather, it is a recognition that the nature of the technology itself in telephone organizing allows for the development of scale relatively cheaply; scaling the kinds of face-to-face contacts that are most effective requires a different-in-kind approach. There are increasing challenges to phone contact¹⁵ and anecdotal evidence from interviews suggests similar trends in canvass; making volunteer operations more efficient and effective should be a priority for campaigns that are focused on voter contact.

¹¹ Sinclair, Betsy, McConnell, Margaret, and Michelson, Melissa. Local Canvassing: The Efficacy of Grassroots Voter Mobilization. *Political Communication*, 30:42–57, 2013.

¹² Nickerson, David W. Quality Is Job One: *Professional and Volunteer Voter Mobilization Calls*. *American Journal of Political Science*. 2007.

¹³ Enos, Ryan and Fowler, Anthony. *Aggregate Effects of Large-Scale Campaigns on Voter Turnout*. *Political Science Research and Methods*. Vol 6, No. 4, 733–751. The European Political Science Association. 2016.

¹⁴ Gerber, Alan S. and Green, Donald P. *Get Out The Vote: How To Increase Voter Turnout*. Brookings Institution Press, Washington, DC. 2015.

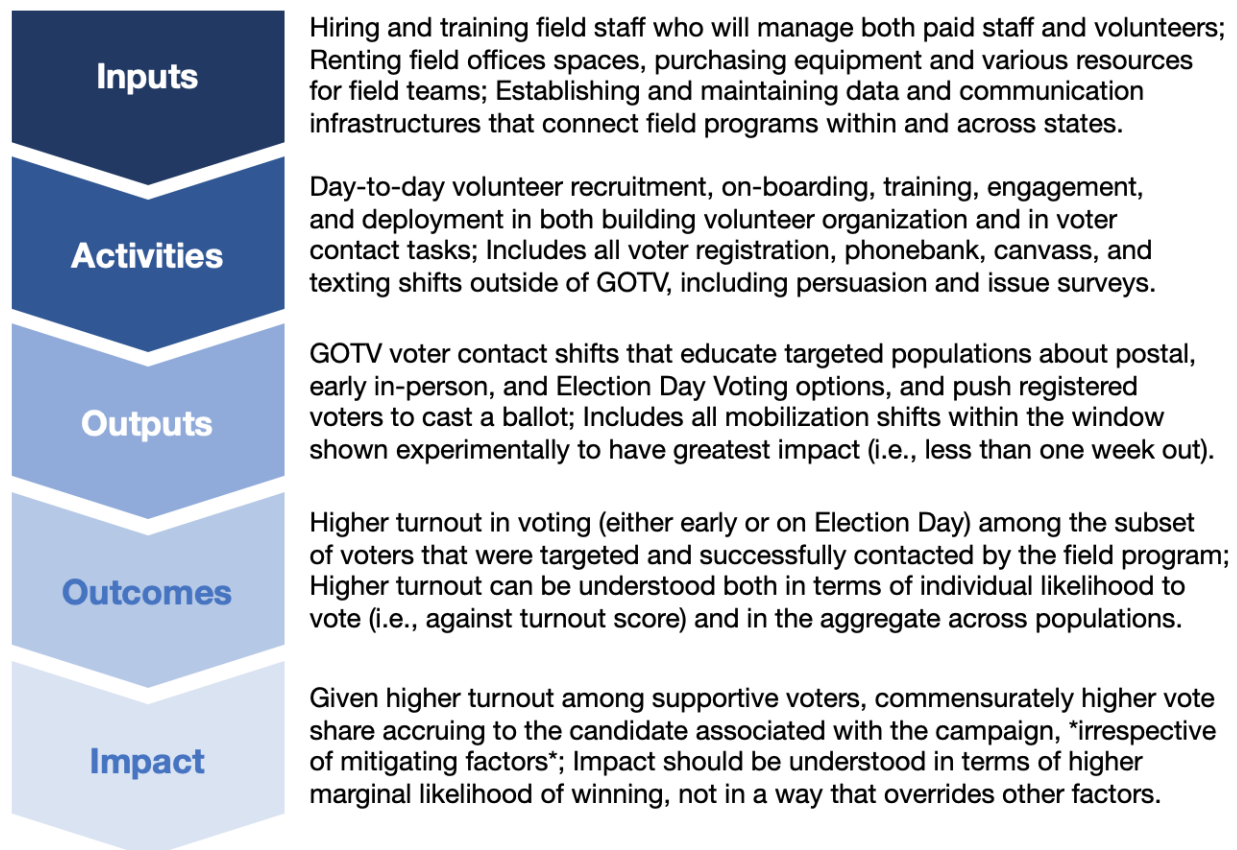
¹⁵ Scott Keeter, Nick Hatley, Courtney Kennedy and Arnold Lau. “What Low Response Rates Mean for Telephone Surveys.” Pew Research Center. May 15, 2017. <https://www.pewresearch.org/methods/2017/05/15/what-low-response-rates-mean-for-telephone-surveys/>.

This report will also focus exclusively on ways to build capacity via volunteers, rather than through paid canvass or commercial calling services. This is because the policy adjustment to increase scale of paid services is essentially ‘Raise and spend more money.’ But even if Democrats raise and spend as much money as possible in future election cycles, there could be substantial value in raising both the efficiency and effectiveness of volunteer operations.

The Research Gap in the Organizing Logic Model

The goal of this work is to inform how campaigns establish and manage volunteer operations. A useful framing device for program evaluation used in many sectors is a logic model, which breaks down the work of an organization into inputs, activities, outputs, outcomes, and impacts. The value of interpreting the complex work of the organization through this framework is that by distinguishing between investments made (inputs), the work done with those resources (activities), the direct effects of that work (outputs), the measurable impact created by outputs (outcomes), and the overall result (impact), organizations can think more coherently about the investments they are making and the value of how they spend time and energy.

Visualizing A Simplified Field Program Logic Model

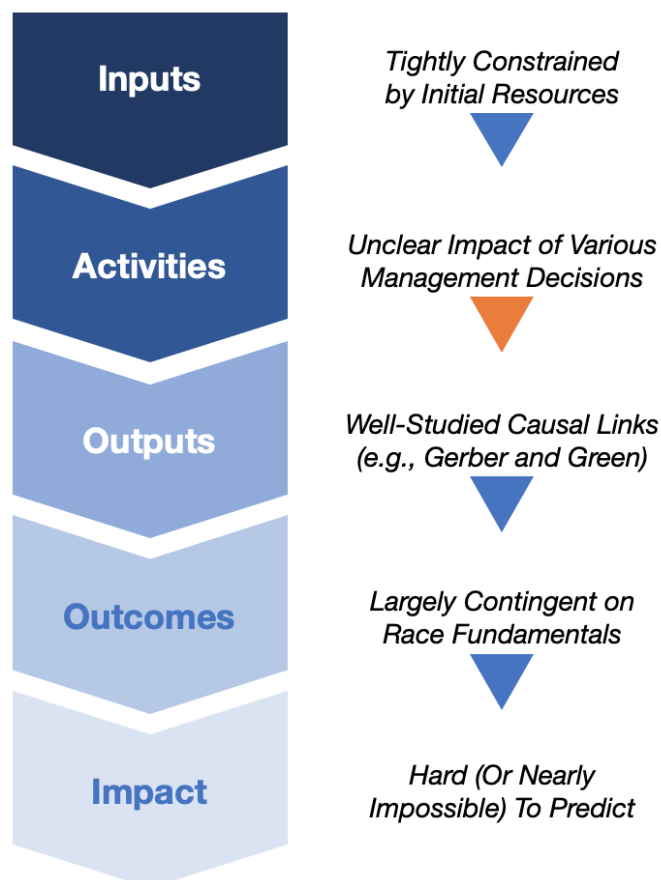


The individual pieces of this model should be familiar to those who have previously participated on political campaigns within or adjacent to the field program. The missing link is

not articulating this logic model to explain activity; it is instead that to-date, research has focused on, defined, explored, and quantified parts of this model while neglecting others.

Field program leaders can draw on libraries of research to know that more activity close to Election Day (e.g., a door contact immediately prior to voting) can move the needle. Managers can quantify a win number based on turnout and support projections, map back the number of additional votes needed, and quantify to a highly granular degree how many doors must be knocked to achieve that number, and therefore how many doors must be attempted, how many shifts scheduled, and how many volunteers engaged in GOTV. But if the link between outputs and outcomes is extremely well-defined, we know comparatively little about how the work of field organizers and field programs can help maximize the capacity of programs to produce those activities. We can know that, in general, GOTV needs to be as big as possible, but there has been much less exploration of how choices in program structure and the nature of in-cycle activities drives the ability of campaigns to create scale effectively.

The Gap in the Field Program Logic Model



In particular, existing research does not adequately address the degree to which breadth (i.e., size of program) versus depth (i.e., strength of individual relationships) in field programs impacts the ability of campaigns to not only mobilize but also organize strong local community

networks. Field organizers know that ‘more is better,’ but do not have adequate tools to understand the differential value of organizing activities across the cycle.

It is useful to make the problem concrete. Campaign strategists develop field program goals through backwards-mapping. They estimate the votes needed for victory (or win number) based on voter registration, estimates of support, and previous election turnout. Then, based on how many voters need to be engaged through this process (using academically-tested assumptions of how many additional votes are gained for every knock and phone call) and field programs in previous cycles, they go backwards through the calendar to allow for growth and organization-building across the cycle. Thus, in June, an organizer might have a relatively high goal set for calling potential volunteers, but a relatively low goal for knocking on voters’ doors. These goals will shift over time to reflect where best estimates about what worked in the past, as well as what is realistic given current resources. The dilemma for organizers and managers of organizers is as follows. For a given shift or voter contact goal, it is not always clear if it is more beneficial to have a few highly-engaged or a large number of low-engagement volunteers complete that work. For a given overall volunteer goal, it is not always clear what those volunteers should do. And it is nearly impossible to understand in what ways those summer and early-fall tasks are meant to connect with the campaign’s overall goals. The first goal of this report is to elucidate some of the quantitative relationships between different activities within the cycle and later capacity generation; the second is to think through how these relationships can be studied experimentally to provide clear management guidance for future campaign management.

Mobilizing Versus Organizing: Incentives and Strategies

At a higher level, there is a more complex question to explore. There is a fundamental distinction between a campaign aiming to turn out the votes of existing supporters (mobilizing) and a campaign which seeks to build internal leadership capacity within a community to enable change (organizing). Another way to analogize this is in terms of investment in support. Mobilizing is ‘spending’ supporter enthusiasm without investing in ways to grow the asset; organizing can be understood as investing to grow the assets while (and through) deploying them in the field.

The difference is frequently framed in terms that have normative implications. After Secretary Hillary Clinton’s loss in 2016, many field staff and volunteers vented to reporters that the explanation was simple. The campaign, in their view, was too focused of mobilization at the expense of persuasion and organizing; they didn’t try to really engage, and so “they didn’t provide people anything to do until GOTV.” One volunteer confided that the operation was “surgical and corporate... maybe they don’t care about us.”¹⁶ The criticism was reminiscent of

¹⁶ Dove, Edward-Isaac. “How Clinton lost Michigan — and blew the election: Across battlegrounds, Democrats blame HQ’s stubborn commitment to a one-size-fits-all strategy. Politico. 14 December 2016. <https://www.politico.com/story/2016/12/michigan-hillary-clinton-trump-232547>.

criticism of many past campaigns; while Howard Dean's campaign in Iowa was exciting, researchers argued that its infusion of huge numbers of outsiders just before the caucus "repelled" voters.¹⁷ The truth in both of these cases, and in the general case of retrospective looks at campaigns, is likely much more nuanced. Undoubtedly, every campaign includes some decisions which favor mobilization, and some which favor organization. It is easy in hindsight to argue that very specific local challenges and frustrations are a general indictment of campaign strategy, just as it is easy to argue the unilateral benefits of any particular piece of a winning campaign. We should be skeptical of the tendency for justification to elide useful analysis. To quote AJP Taylor, the "defect of this explanation is that, since it explains everything, it explains nothing."¹⁸ Electoral politics, rich in tradition and information but lacking clear connections between individual decisions and campaign outcomes, are susceptible to 'superstitious' learning¹⁹ and the misattribution of success or failure to those factors which are proximate in time, physical distance, or personal relevance. Organizing, like every other apparatus of a campaign that persuades and mobilizes voters, makes an impact on the margin.

The tradeoff is better understood as a prioritization choice by managers with limited resources and little time to make allocation decisions, and who feel immense personal and professional pressure to perform. As data sources have not proliferated within campaigns and have become more accessible through analysis and reporting tools, these managers have increasingly turned to large datasets to narrow their decision. Hahrie Han, who directs the P3 Research Lab at the University of California-Santa Barbara and has written extensively on organizing programs in recent political cycles, argues that the danger which results is that factors which are easy to measure are the ones which are managed; "attention is paid to the things that are the most countable, but not the most important. Some of the things we measure are useful proxies to the things we care about—like volunteer engagement and excitement—but we start to allocate and reward and punish based on things that are way upstream."²⁰

As a result, it can be hard to justify spending significant staff time on the kinds of activities which are helpful in building long-term relationships and developing trust within networks (e.g., one-on-one meetings between volunteers and organizers) when those activities do not produce the kinds of metrics which are easy to understand, frame in terms of output, and describe to campaign staff, donors, and members of the media. Instead, managers are frequently incentivized to push organizing teams and volunteers to generate what might be characterized as 'empty activity.' Cold-calling enormous numbers of people to show up for one volunteer shift in which they also cold-call hundreds of voters generates a lot of data that looks good superficially; it has high topline for volunteer engagement and voter contact. The numbers that emerge from this work can easily be visualized and packaged. They can easily be

¹⁷ Keeter, Scott, Funk, Cary, and Kennedy, Courtney. "Deaniacs and Democrats: Howard Dean's Campaign Activists." In *The State of the Parties: The Changing Role of Contemporary American Parties*

¹⁸ Taylor, AJP. *The Origins of the Second World War*. Touchstone. New York. 1961.

¹⁹ Rita Gunther McGrath. On The Pitfalls of Superstitious Learning. Harvard Business Review. July 19, 2011. <https://hbr.org/2011/07/superstitious-learning.html>.

²⁰ Author interview with Dr. Hahrie Hahn

used to congratulate highly-productive offices and chastise those that are doing less activity. Han suggests that “campaigns are trying to super-efficient and are trying to optimize resource allocations based on efficiency, without a way to track how those decisions impact downstream outcomes. So, they look at how effectively persuasion knocks change minds, and unsurprisingly they don’t win you the race in July, so they put money elsewhere. But the act of building a team and training organizers and developing relationships might have a much longer yield curve that you’re just not capturing.”²¹

The important point is not that people are myopic or making bad decisions: far from it. Rather, it is that no one—in the headquarters office or in the states—has a consistent way to predict accurately how these activities map to the outcomes and outputs about which organizers actually care. What could be lost if the focus of in-cycle campaign management falls squarely on topline productivity metrics (i.e., producing activity now) and neglects attention to organizing work that may not produce many tangible goods now but instead drives long-term outcomes? In her study of recent social movements, sociologist Zeynep Tufekci compares movement organizing in the Civil Rights Movement to modern digital movements like the Arab Spring and Occupy Wall Street. While she documents a host of critical gaps and challenges that modern movements face, one set of observations is particularly critical for political campaigns to understand. In the Civil Rights Movement, Tufekci found that organizing to complete “the tasks that are required for organization, logistics, and coordination together over time” created huge value. The “benefits and collective capabilities attained during the process of forming durable networks” however, only became available when working together “poses challenges that must be overcome collectively and require decision making, building of trust, and delegation among a semidurable network of people who interact over time.”²²

While the development of what Tufekci calls ‘network internalities’ in social movements has differences with the strategy and tactics of electoral campaigns, relationships and trust are critical for organizing structures within campaigns. An enormous challenge faced by managers of field programs is that it is very difficult to control the myriad factors that shape volunteer experiences. One bad experience on the phone with an angry voter, an extremely hot day spent canvassing, or even a night when no one picks up the phone might scare the volunteer away from helping in the future. This challenge can be mitigated by relationships, and not only those formed vertically between organizers and volunteers. Lateral relationships—formed between volunteers, either informally or through formal teams—can build separate sources of motivation and strength that supplement the candidate and campaign apparatus.

Organizing capacity and the relationships, connections, and shared purpose that it demands can help to standardize a program’s approach to volunteer training, management, and evaluation. In doing so, they might also help to ‘inoculate’ a field program and its volunteers

²¹ *ibid*

²² Tufekci, Zeynep. *Twitter and Tear Gas: Power and Fragility of Networked Social Movements*. Yale University Press. 2017.

from the vagaries and exigencies of random events within a cycle, creating a basis of collective action that drives sustainable capacity.

We should not expect that the incentives which guide electoral decision makers to focus on mobilization will shift radically in the near term, nor should we want or expect campaigners to unilaterally adopt tactics which will decrease their short-term chances of winning in order to develop local capacity. Rather, we should examine the degree to which incremental changes in program structure in the months well before traditional GOTV windows can help campaigns build an organization which can organize both at greater scale and more effectively.

Program Quality and Relationships: The Problem Is Scale

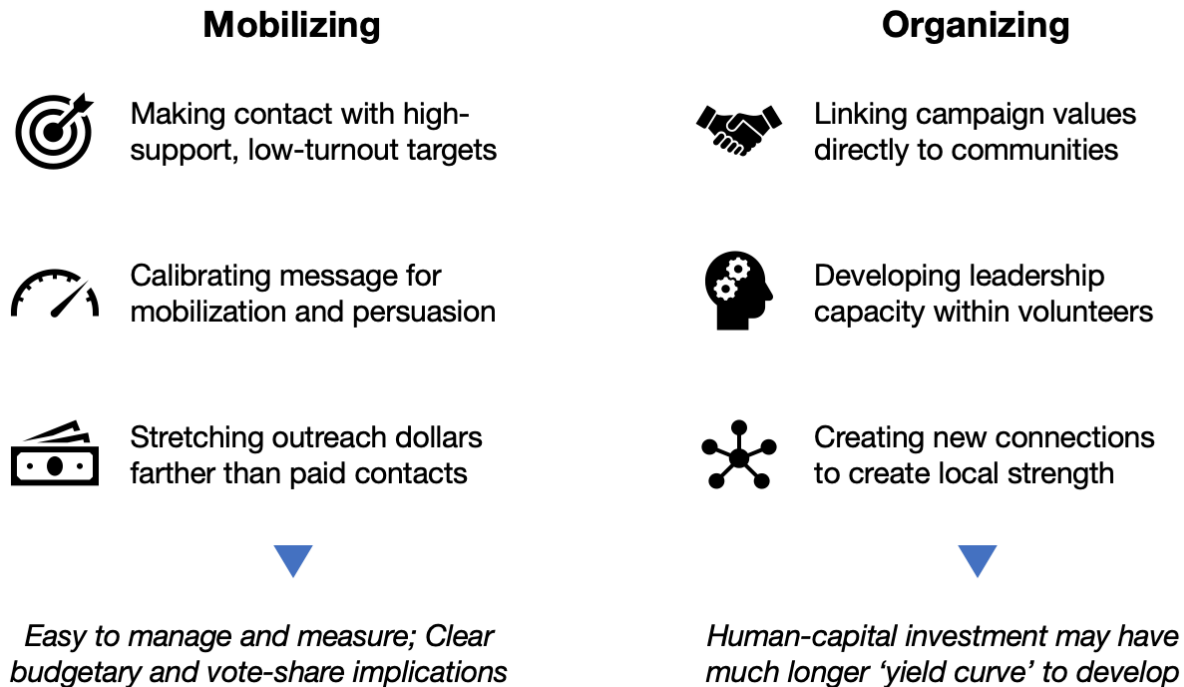
Having a field program increases turnout on the margins; these effects are significant and measurable and accrue to programs that invest in phone calls as well as door-to-door canvassing. One of the key distinctions that we can draw from previous research is that in virtually all cases, voters are more responsive to contact from engaged volunteers with whom they share some kind of connection than either unknown, disconnected volunteers or paid canvassers. Gerber and Green summarize the research in phone bank turnout experiments: “Callers who genuinely believe in the cause that they are calling for tend to be much more effective than those who view this work as just another job.”²³ On the doorstep, knocks are more effective when they come from someone that the voter knows or has some connection with. On the one hand, this is not a surprising finding; if human beings are social to some degree, we should expect a degree of social connection to be of value in motivating them to act. But it also implies a discrete challenge that campaigns must confront.

The “biggest challenge is bringing a door-to-door campaign “to scale.” It is one thing to canvass 3,600 voters; quite another to canvass 36,000 or 360,000.”²⁴ The challenge here can be understood in terms of raw effort (e.g., staff engaged, offices opened), but is perhaps better expressed in terms of time. Field programs that build trust cannot be stood up overnight. It takes time to build relationships that motivate sustained, effective effort and engage a broader community. It takes time to establish the kind of presence within a community that inspires trust and makes people feel that the sacrifice of their own time and energy is worth the effort. It is precisely this deficit, this sense of urgency and time running out, that can drive a management culture responsive to volumetrics and current activity rather than the work that builds a longer-term foundation for organizing capacity. And until there is a deeper exploration of how intra-cycle activities impact mobilizing and organizing scale, it will be difficult for campaigns to develop a more coherent approach to the mobilizing-organizing question.

²³ Gerber and Green, 68.

²⁴ Gerber and Green, 38.

Summarizing Two Meta-Strategies for Field Programs



This dichotomy is useful only to an extent; every general election campaign completes many activities that can be categorized into both groups. The relevant question for managers is to what degree and at what time are these activities most useful for building capacity.

This report explores the challenges involved in bringing the volunteer organization in a modern political campaign to scale. It begins with an exploration of the individual-level factors which tend to impact proxies for 'volunteerism'—that is, the likelihood of individuals to volunteer more or in more critical moments, scaled against their participation in various activities with relevant campaigns. The following section looks at how management decisions in field programs are associated with volumetric measurements of campaign output during GOTV and links mid-cycle activities with mobilization effectiveness. The next section explores how some of those same activities can be tied to better proxies for true organization capacity. Finally, it lays out a framework through which these concepts might be tested in future campaign cycles.

Data, Methodology, and Methods

Data Sources

The data contained in this report come from the Voter Activation Network (VAN) MyCampaign and MyVoter database records collected by political campaigns in the last several election cycles for candidates running for the Presidential General election. Records associated with Presidential elections are chosen for several related reasons:

- **Presidential races engage a larger and more diverse group of volunteers and voters than more localized races (especially in midterm or off-cycle races).** Races that rely on the civic participation of only the most engaged partisans might have few broad lessons to teach advocacy organizations and races in less fraught contexts. But understanding what is effective in motivating and maintaining support of the less-politically engaged has the potential to be valuable for campaigns and activists in a variety of situations (while still holding positive value for those campaigns which engage a smaller group of very committed activists).
- **Data from multiple states decreases the potential for local idiosyncrasies to overwhelm other factors and decreases the risk of overfitting in model specification.** Volunteer engagement strategies and organizational design for a race in New Hampshire differs immensely from best practices in states like Florida or Ohio; organizing metrics that worked well in a very specific context might fail to achieve results in an environment with substantially different voter characteristics. We should be interested in those region-level variations, opportunities, and risks, but also focus attention on identifying those organizing strategies and factors which can inform planning and development for a wide array of campaign types.
- **Data collected by Presidential campaigns benefit from a larger dedicated technology and analytics team; they are more consistently and accurately collected and analyzed, which should minimize the risks presented by using self-reported data.** The degree of oversight and quality control over organizing data diminishes significantly as the office moves from Presidential candidates down-ballot. Top-of-the-ticket campaigns contain their own data and management idiosyncrasies which make data quality a challenge (especially a high degree of turnover and high-powered incentive structures which can reward ‘padding’ numbers associated with volumetric goals), but these challenges are somewhat more predictable and manageable than the myriad of issues which beset smaller campaign teams.

Data are pulled from relevant backend systems maintained by the Democratic National Committee using a system of interconnected identification numbers which are tagged to campaigns and are referred to as ‘committees’ - legal entities that hire staff, collect donations, and fall under the jurisdiction of state and federal election laws. These committees are then

mapped to staff (i.e., battleground state regional directors and organizers), contacts with volunteers, and contacts with voters. A representative listing of the kinds of data points available through this system follows:

- Individual Volunteers
 - Limited demographic details (gender, age, and some race information)
 - Total shifts scheduled and completed, categorized by shift type and shift date
 - Listing of all volunteer activity scheduled for GOTV (completion history unreliable)
- Field Organizations
 - Superstructure of field organizers and regions within battleground states
 - Field organizations where volunteers complete the greatest number of individual shifts (defined here as their 'home' turf)
 - Daily volunteer engagements and shifts completed tagged to individual field organizations

These tables were then turned into separate data tables for analysis using custom SQL. First, hundreds of thousands of individual volunteers were mapped manually to organizations using their geographic location or volunteer activity history. Following the construction of these base mappings, volunteers were categorized for each week throughout the cycle based on their engagement with the campaign (i.e., types of shifts scheduled and completed) and whether they had been an active volunteer during the previous one-to-six weeks. These tables were then merged with measures of voter contact output during GOTV, and exported without identifiable information for statistical analysis using R. The relationships between input and output variables were measured using a combination of aggregated and time-sensitive linear regression models. Regression outputs and statistical significance are reported for all models; control specifications are included as relevant.

Interviews and Acknowledgements

In addition to quantitative analysis, this report relies on the advice and perspective of academic experts and field practitioners who have previously participated in and led field campaigns. Many of these interviews were not for public attribution; when relevant, sources are cited directly. The author thanks Catherine Tarsney, Marshall Ganz, Tom Patterson, Nicco Mele, Robby Mook, Elan Kriegel, Karen Hicks, Hahrie Han, Liz McKenna, Bobby Courtney, Jeremy Bird, Emily Norman, and many others for the gift of their time and perspective.

Volunteerism and Individual Participation

Before studying organizational metrics, it is helpful to build a baseline of how individual volunteers interact with campaigns. The relationship between individual volunteerism and overall capacity-building activities is not simple; if volunteer hours during the summer months are spent recruiting additional volunteers, for example, then the value of that recruitment should not be ignored when considering whether those volunteers who make the calls are extremely likely to volunteer again.

However, as individual volunteers are the most basic unit of analysis and of policy decisions, understanding their engagement provides valuable context for organizational analysis. The following pages unpack initial findings into how the timing and nature of volunteer engagement in a field organization might drive individual volunteerism.

Measuring Volunteerism

Data collected by campaigns offers several ways in which we can measure some proxies for individual volunteerism, narrowly conceived here as the likelihood of an individual to perform a volunteer on a campaign, holding available controls constant. The most obvious way to measure volunteerism under these constraints would be the raw number of volunteer shifts completed: a volunteer that does more is more engaged. We could also look at the regularity of contact: a volunteer who shows up at regular and tight intervals is more engaged than a volunteer who shows up rarely and sporadically.

The methodological trap in these measures of engagement is that it is very difficult to understand whose engagement is organic (e.g., who has made volunteering on campaigns a habit and would have shown up anyway) and which engagement is being driven at least partially as a result of organizing activity. In order to avoid these issues and begin to isolate useful information, it is helpful to try to control for as many factors as possible to identify discrete policy decisions that could be shaped to improve engagement.

For example, if we control for overall shifts completed and basic demographic details, and then look at how timing of first shift in isolation affects the likelihood of an individual volunteer to participate in GOTV activities, we would have a better idea of how field programs should think about scaling the volume of new volunteers coming in during each stage of the campaign. If we control for timing of first engagement and overall activity but select out the impact of discrete kinds of volunteer activities, then we might have a better sense of how different volunteer recruitment and voter contact priorities roll up towards GOTV capacity.

Types of Shift Completed

To begin the analysis of volunteerism, we can start with a very high-level question: If a volunteer completes at least one shift during the cycle, does it increase their likelihood of participating in GOTV? The question is worth answering concretely because if doing more shifts does not have a positive association with GOTV participation, then the overall value of field programs outside of paid-canvass operations would come into question. The table below estimates the impact of an additional shift in general and at least one shift of one-on-one, canvass, and phone bank on the likelihood of a person participating in GOTV.

Impact of Completing at Least One Shift per Shift Type on GOTV Participation

Dependent variable:				
	(1)	At Least One GOTV (2)	(3)	(4)
Additional Shift of Any Kind	0.027*** (0.0001)			
At Least One One-on-One		0.074*** (0.001)		
At Least One Canvass			0.098*** (0.0004)	
At Least One Phonebank				0.039*** (0.0002)
Constant	0.165*** (0.001)	0.221*** (0.001)	0.195*** (0.001)	0.207*** (0.001)
Observations	481,868	481,868	481,868	481,868
Log Likelihood	-246,377.600	-267,920.400	-247,282.500	-256,617.000
Akaike Inf. Crit.	492,759.200	535,844.700	494,569.100	513,238.000

Note:

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

The takeaway from this table is that at a high level, doing *something* during the cycle as a volunteer has on average a positive relationship with showing up to help during GOTV when volunteer capacity is most critical. The organizer's work, then, is to understand to what degree, in what kind of activities, and when the volunteer's engagement is most impactful.


Timing of First Engagement

When it makes sense to bring people into the field program? Too early, and organizers worry about burnout; will the volunteer persist in engagement until Election Day? Too late, and

volunteers might not be well-trained to effectively engage voters, recruit their friends and neighbors, or leverage the network effects at the heart of relational organizing trends.

Targeting and universe construction for likely volunteers is a campaign-specific process, and the individual targeting models and scores that go into recruitment were not available for this analysis. However, we can retrospectively examine the degree to which timing influences two measures of capacity: (1) how many shifts a volunteer completes overall; (2) the volunteer's likelihood to participate in GOTV in some way. The table below groups individual volunteers who participated in the 2016 cycle and completed at least one shift of any kind by the first day in which they completed a volunteer shift and were counted by the campaign as showing up. It then takes the average number of shifts completed across the cycle and the average share scheduling at least one GOTV shift for each day.


When A Volunteer First Engages the Campaign Shapes How They Will Contribute



The high-level answer to the timing question is that volunteers who don't participate until relatively late in the cycle tend to do few shifts—an intuitive result, as they have less time to do things. Those who show up very late in the cycle are much more likely to participate in GOTV than those who help out during the summer months. This initial analysis indicates that there may be some trade-off for field programs to consider when they recruit volunteers; people that join early help to complete tasks but are individually unlikely to persist through to GOTV. However, waiting to recruit late in the cycle when GOTV participation might be higher risks missing opportunities to build capacity to engage those volunteers.

However, it is also worth unpacking how these results hold up for volunteers that engage in qualitatively distinct kinds of activities. After all, volunteers are not brought into offices in general, but to complete specific tasks. In particular, we might be interested in how timing matters for those volunteers that have a higher level of engagement or are taking on some aspect of volunteer leadership.

One-on-One Completers Volunteer More and Do More in GOTV



The obvious challenge in interpreting these findings is the selection bias that we can expect with one-on-one completion. These volunteers are willing to take the time to complete a forty-five to sixty-minute conversation with an organizer, which already indicates that they are relatively engaged. The organizer has identified them as engaged enough to take time to complete one of these conversations and has also identified potential value in such a conversation. This means that the difference between the two lines above should not be interpreted as caused by the participation in one-on-one shifts; rather, we should think of the difference as providing an empirical distinction between volunteers who engage more personally and directly with the organizer and those who do not.

However, there are also some useful empirical results here. Regardless of whether the completion of a one-on-one shift is the result of endogenous volunteer enthusiasm or the consequence of effective organizing, volunteers who complete a one-on-one shift complete 2-3 more shifts across the cycle than those who do not. And regardless of timing, these volunteers are 10 to 20 percentage points more likely to engage in GOTV activity. This provides field managers with some useful benchmarks and potential early-warning flags for the quality of one-on-one engagements. If volunteers that have completed a one-on-one meeting aren't

showing a higher level of engagement during the cycle than those who have not, this could be a useful signal that the organizer is not effectively translating those time-intensive conversations into the commitments that the field program needs to build capacity.

The table below separates volunteers into the first month in which they completed a shift. The estimate of the linear relationship between completing at least one one-on-one shift and participating in GOTV, controlling for the total number of shifts completed, confirms the graphical relationship shown above.


Impact of At Least One-on-One Shift on GOTV Likelihood

Dependent variable:				
	July (1)	At Least One GOTV August (2)	September (3)	October (4)
At Least One 1-on-1	0.033*** (0.003)	0.034*** (0.004)	0.045*** (0.004)	0.085*** (0.009)
Total Shifts Completed	0.026*** (0.0002)	0.031*** (0.0003)	0.042*** (0.0003)	0.055*** (0.001)
Constant	0.090*** (0.002)	0.127*** (0.002)	0.131*** (0.001)	0.186*** (0.001)
Observations	60,778	78,987	115,454	161,624
Log Likelihood	-22,589.940	-37,433.660	-56,631.630	-93,174.220
Akaike Inf. Crit.	45,185.870	74,873.320	113,269.300	186,354.400
Note: *p<0.1; **p<0.05; ***p<0.01				

This table includes a control for total shifts completed; the directionality and magnitude of these relationships were robust to additional controls for age, race, and gender. This relationship is overall positive but not dramatically so; the impact of an additional shift completed in general is not that different from the additional value of shifting between the categories of completing zero one-on-one shifts and completing at least one. That is to say, one-on-one completers are more engaged than those who do not complete a one-on-one, but it is difficult to tell whether this indicates urgency to complete more one-on-one shifts, or if it merely indicates that people who were engaged in general enough to do a one-on-one shift also demonstrate that underlying spirit in other ways.

When we perform a similar analysis but separate volunteers by whether they completed a phone bank shift at any point in the cycle, we see a similar pattern emerge, but with a wider gap between volunteer groups. The graphs below visualize the average total shifts completed and share scheduling at least one GOTV shift by date of first shift completion, separating volunteers by whether they did at least one phone bank shift during the cycle.

Early Phone Volunteers Do Many More Shifts; Steady GOTV Participation Across Time



Volunteers who started early and completed at least one phone bank shift (which at that point in the cycle is more likely a volunteer-recruitment call than a voter contact conversation) do more shifts on average than those who never do a phone bank, although this gap deteriorates rapidly. In contrast, phone bank volunteers on average participate in GOTV between 40 and 50 percent of the time, regardless of when they start.

This more stable relationship in GOTV capacity with respect to the timing of engagement is reflected in the linear estimates included in the table below, which as above tests the likelihood of GOTV participation contingent on completing at least one phone bank and the overall number of shifts completed during the cycle.

Impact of At Least Phone Bank Shift on GOTV Likelihood


Dependent variable:				
	July (1)	At Least One GOTV August (2)	September (3)	October (4)
At Least One Phone Bank	0.175*** (0.004)	0.134*** (0.003)	0.136*** (0.003)	0.154*** (0.003)
Total Shifts Completed	0.021*** (0.0003)	0.027*** (0.0003)	0.038*** (0.0003)	0.047*** (0.001)
Constant	0.071*** (0.002)	0.099*** (0.002)	0.104*** (0.001)	0.167*** (0.001)
Observations	60,778	78,987	115,454	161,624
Log Likelihood	-21,503.700	-36,532.400	-55,359.190	-91,659.420
Akaike Inf. Crit.	43,013.390	73,070.790	110,724.400	183,324.800

Note: *p<0.1; **p<0.05; ***p<0.01

Here, there is comparatively little difference in GOTV propensity between volunteers that start early and start late based on if they did a phone bank shift at any point in the cycle. This has some potential implications for volunteerism. First, the magnitude of the effect size of doing at least one phone bank shift is much higher than doing at least one one-on-one shift is substantially larger. This is unlikely to be the result of some innate characteristics of phone bank volunteers (as this is the default for most field offices). Second, being involved in some kind of voter contact appears to mitigate the timing effects of when people join the campaign.

The picture that emerges from taking a parallel approach to volunteers completing at least one canvass shift during the cycle is even more stark. The graphs below duplicate the methodology above by sorting individual volunteers by the day of their first completed shift, then averages the total number of shifts completed during the cycle and the average share of volunteers who started on that date who scheduled at least one shift during GOTV. The only difference is that now, we have separated volunteers into two different categories; those who did not complete a canvass shift during the cycle (in black) and those completed at least one canvass shift during the cycle (in blue).

Those Who Start Early and Canvass Do More Shifts, More GOTV



We can now see a different story emerging about individual volunteerism with respect to canvass activity. On average, volunteers who started in the August-September timeframe and did at least one canvass shift at some point during the cycle completed between 3 and 4 times as many total shifts as those who never did a canvass shift. Across the cycle, volunteers who completed at least one canvass shift were between 30 and 50 percentage points more likely to participate in GOTV in some way. The former gap narrows to literally no distinction by GOTV, while the latter gap remains wide throughout the cycle but actually declines over time.

How should we interpret these data? It seems that unlike in the case of one-on-one completers, doing a canvass shift during the cycle has its own separable effects on individual volunteerism. What do we think is driving this discrepancy? There are two important factors to consider. The first is that completing a canvass shift is not as simple as blocking time and sitting down; it implies a level of logistical capacity and commitment that demands (and suggests) a more comprehensive organizational structure. Thus, part of this might be simply definitional; those individuals who are completing canvass shifts are simply joining better organizations (these findings are robust to controls for state and region, undermining this as the sole reason for the difference). The second factor is the training element of completing a canvass shift, the formation of “network internalities” that Tufekci describes in an earlier section of this paper. Here, completing canvass shifts as an individual volunteer might be not just selecting for those who are enthused, but also generating a deeper connection with the campaign and with the community via face-to-face contacts with voters.

The table below estimates linear probability models of an individual volunteer participating in GOTV based on whether they completed at least one canvass shift during the cycle and the total number of shifts they completed.

Later Canvass Completers Less Likely to Participate in GOTV

Dependent variable:				
	July (1)	At Least One GOTV August (2)	September (3)	October (4)
At Least One Canvass	0.331*** (0.004)	0.262*** (0.004)	0.261*** (0.003)	0.234*** (0.002)
Total Shifts Completed	0.019*** (0.0002)	0.025*** (0.0003)	0.034*** (0.0003)	0.045*** (0.001)
Constant	0.073*** (0.002)	0.103*** (0.002)	0.106*** (0.001)	0.141*** (0.001)
Observations	60,778	78,987	115,454	161,624
Log Likelihood	-19,675.770	-34,896.060	-53,146.260	-88,663.200
Akaike Inf. Crit.	39,357.540	69,798.110	106,298.500	177,332.400

Note: *p<0.1; **p<0.05; ***p<0.01

The interesting difference between the estimates of the relationships here as compared to one-on-one and phone bank shifts is that the individual volunteer's likelihood to participate in GOTV if they complete a canvass shift gets lower over time. This suggests that the capacity-building value of canvass shifts at the individual level is stronger earlier in the cycle than later. In other words, it is the combination of engaging someone relatively early in the cycle and moving them toward canvassing that drives the greatest bump in GOTV capacity.

Scaling the Field Program for Mobilization

Measuring Field Program Capacity

There are many methods by which we could imagine measuring the organizing capacity of a given region and turf. A useful metric would be created by comparing doors knocked (i.e., contacts made) during GOTV and thereby count the number of conversions into votes. But during the final weeks, most organizing programs explicitly ask organizers and volunteers to stop spending time compiling data; usually, the only complete records that are accessible in this form come from dedicated external staff supporting a randomized control trial in the field.

For this reason, we can use shifts scheduled during GOTV as a good proxy for organizational capacity. We also want to distinguish between general shifts scheduled (as an aggregate measure of capacity) from the specific value of voter contact shifts (especially canvass shifts). These are not the only ways to measure mobilization capacity, but because these figures can easily be compared across field offices and regions and are measured somewhat consistently, they provide a helpful starting place for analysis.

Analyzing shift data to answer these kinds of question introduces two potential causal channels: (1) scheduling and completing shifts during the cycle is the instrument by which capacity is developed; (2) scheduling and completing shifts during the cycle are evidence that capacity has been developed, and are features of effective organizing, not the other way around. Unfortunately, the data collected to complete this report do not include the kinds of qualitative indicators (e.g., quality of conversations between organizers and volunteers, training rigor, clarity of purpose and execution) that would allow differentiation between activity that drives capacity and effective organizing which drives activity as a byproduct of capacity. However, training and management practices can be considered as separate kinds of policy problems for the Democratic National Committee to address; the quantitative signals of effective capacity, whether leading or lagging, remain a valuable input for planning and strategy.

Aggregate Shift Volume and Type

Before exploring the drivers of capacity, it is useful to confirm that GOTV capacity in terms of shifts scheduled is driven by capacity as measured by total volunteers. This point is intuitive; more people are likely to do more activities, all other things being equal. But we also want to differentiate between pure mechanical capacity (i.e., the sheer number of people who are around during GOTV performing various activities) and the overall size of the volunteer organization in a given area (i.e., the total number of volunteers that complete a shift during the cycle). The following regression table estimates a linear relationship between GOTV shifts and volunteer totals for both of these groups:

Bigger is (Generally) Better

	Dependent variable:			
	Total GOTV Shifts Scheduled (1)		Total GOTV Canvass Shifts (3)	
		(2)		(4)
GOTV Volunteers	0.179 (0.203)		0.734*** (0.095)	
Total Cycle Volunteers		-0.150* (0.084)		0.174*** (0.040)
Constant	247.089*** (11.864)	265.824*** (12.182)	96.444*** (5.560)	105.592*** (5.737)
Observations	4,233	4,233	4,233	4,233
R2	0.0002	0.001	0.014	0.005
Adjusted R2	-0.0001	0.001	0.014	0.004
Residual Std. Error (df = 4231)	639.522	639.340	299.697	301.119
F Statistic (df = 1; 4231)	0.783	3.189*	59.846***	19.402***

Note:

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

The simplest interpretation of this table is that having an additional volunteer as part of the GOTV operation (i.e., appearing as a scheduled volunteer in the week before Election Day) is associated with an additional 0.18 total shifts, and an additional 0.734 canvass shifts. We might push this finding to say that, unsurprisingly, having more volunteers translates into more shifts. However, the relationship between the total size of a volunteer operation during the cycle (i.e., in the months leading up to Election Day but not including the final week) and GOTV shifts is much murkier, both in terms of its estimated impact and the degree to which it explains differences between field programs.

We can advance by asking a further question: What is the estimated impact of successfully completing one additional shift at any point in the cycle on a field office's ability to schedule more shifts in GOTV, and how does that additional impact vary by the type of shift completed? Here, we want to separate the overall value of size from individual tasks and discover if the aggregate activities completed have different kinds of value in building toward capacity.

The following table estimates the impact of additional shifts completed in-cycle, after controlling for the total unique volunteers that show up at the field office at any point in the cycle:

Canvass and One-on-One Shifts Associated with GOTV Volume, Even with Controls

	Dependent variable:			
	(1)	Total Shifts Scheduled in GOTV (2)	(3)	(4)
Total Canvass Shifts Completed	3.109*** (0.096)			2.630*** (0.127)
Total One-on-One Shifts Completed		2.903*** (0.357)		1.024*** (0.328)
Total Phone Shifts Completed			2.117*** (0.086)	0.561*** (0.111)
Unique Cycle Vols	-0.778*** (0.078)	-0.282*** (0.085)	-1.000*** (0.086)	-0.953*** (0.082)
Constant	112.951*** (11.872)	227.164*** (12.993)	97.834*** (13.301)	78.365*** (13.011)
Observations	4,233	4,233	4,233	4,233
R2	0.200	0.016	0.125	0.208
Adjusted R2	0.200	0.016	0.125	0.207
Residual Std. Error	572.028 (df = 4230)	634.490 (df = 4230)	598.324 (df = 4230)	569.440 (df = 4228)
F Statistic	529.662*** (df = 2; 4230)	34.586*** (df = 2; 4230)	302.307*** (df = 2; 4230)	277.379*** (df = 4; 4228)

Note:

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

To begin, we can see that completing an individual shift of any kind across the entire cycle is associated positively with GOTV capacity after controlling for the total size of the GOTV operation. In other words, having a field program before GOTV is associated with a larger GOTV program; this is intuitive, but also useful to demonstrate systematically. The other important takeaway from this table is that while phone shifts decline substantially in estimated impact when controlling for other shift types, total canvass and one-on-one shifts are more robust to these controls; any additional one-on-one shift in the months leading up to Election Day is associated with an additional 1.024 GOTV shifts of any kind, and an additional canvass shift was associated with an additional 2.630 GOTV shifts.

However, not all GOTV shifts are created equal. If we accept the proposition from previous academic research that face-to-face canvass shifts are the most effective instrument to turn out voters, then we can recreate the table above, but instead of looking at total shifts, focus specifically on the relationship of in-cycle activity to additional canvass shifts scheduled during GOTV. This is a more nuanced kind of question, because canvass shifts require more investment in training and logistical support; rather than measuring GOTV capacity as a volumetric in terms of raw size of activity, we are getting closer to a measure of operational strength in terms of territory that has been separated into walk packets for individual volunteers, printing and driving logistical support, as well as the training and deployment of volunteers to do a more labor-intensive but also higher-value kind of work.

More Canvass Shifts More Strongly Associated with More GOTV Canvass

Dependent variable:				
	(1)	Total Canvass Shifts Scheduled in GOTV		(4)
		(2)	(3)	
Total Canvass Shifts Completed	1.783*** (0.042)			1.799*** (0.056)
Total One-on-One Shifts Completed		1.658*** (0.168)		0.739*** (0.145)
Total Phone Shifts Completed			1.000*** (0.041)	-0.066 (0.049)
Unique Cycle Vols	-0.186*** (0.034)	0.099** (0.040)	-0.228*** (0.040)	-0.196*** (0.036)
Constant	17.925*** (5.244)	83.505*** (6.097)	26.245*** (6.262)	12.576** (5.757)
Observations	4,233	4,233	4,233	4,233
R2	0.299	0.027	0.129	0.304
Adjusted R2	0.299	0.027	0.129	0.303
Residual Std. Error	252.680 (df = 4230)	297.735 (df = 4230)	281.687 (df = 4230)	251.956 (df = 4228)
F Statistic	903.113*** (df = 2; 4230)	58.789*** (df = 2; 4230)	313.539*** (df = 2; 4230)	460.747*** (df = 4; 4228)

Note:

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

We can now see a more interesting picture emerge with respect to GOTV capacity. There appears to be some value in ‘specialization’ at a high level for field offices; completing more canvass shifts during the cycle is more strongly associated with GOTV canvass shifts than more one-on-one shifts, and the relationship is more statistically significant. An additional canvass shift completed at any point during the General Election cycle is associated with an additional 1.799 canvass shifts scheduled during GOTV, while the estimated impact of an additional phone shift completed is statistically indistinguishable from zero.


How Timing Affects the Impact and Value of pre-GOTV Shifts

The figures above estimate the impact of aggregate volunteers and shifts on an individual field organization’s ability to increase mobilization capacity during GOTV. They provide some useful high-level lessons; namely, that increasing the overall volume of in-cycle volunteer operations has a positive relationship with GOTV capacity, that one-on-one shifts and canvass shifts appear to have a more robust relationship with greater capacity than phone shifts, and that having a higher level of canvass activity during the cycle is associated with specifically higher canvass capacity during GOTV.

However, the critical management and policy question to answer is not only if bigger is better. With limited resources to devote to each part of a campaign, we need to understand where dollars and effort are driving the greatest value. Instead of asking simply if more is good, we want to answer the questions ‘More of what?’, ‘To what extent?’ and ‘When?’ To begin unpacking these questions, we can disaggregate the linear models presented in the above pages into models constructed for every field office for each week of the campaign between the beginning of the General Election (i.e., post-Convention) and GOTV.

The first question to explore is the supposition that larger volunteer organizations are unilaterally better for GOTV capacity. The graphs below report the estimated coefficient for one additional unique volunteer in each week of the campaign cycle, alongside the percentage of variance explained as measured by R-squared. There is substantial variance in the estimate of how valuable a larger organization is across the months leading up to Election Day, ranging from 0.5 to 1.5 additional GOTV canvass shifts per unique volunteer.


Value and Explanatory Power of Simply More Individual Volunteers Is Highly Variable



A critical takeaway from these data is that simply managing by asking for greater volume of volunteers may not be maximally effective by itself; it might matter substantially what these volunteers are doing and how they are being engaged. Put another way, we can imagine that field programs of similar size might be doing substantially different things with the capacity they have developed either innately as a result of organizing work. The value of having additional people come in the door to volunteer, then, hinges on the program's ability to translate enthusiasm into action. The critical question to ask is when that additional organizing activity drives capacity, and which kinds of work tend to be associated with the greatest value.

First, we can characterize the impact of scheduling an additional one-on-one shift in a given week across the cycle. The graphs below estimate a series of linear models calculated for each field program, looking at the associated increase in total GOTV canvass shifts associated with an additional one-on-one shift in each week between late July and late October.

Additional One-on-Ones Most (But Still Very Weakly) Predictive Early in Cycle




The models' outputs show that the highest incremental value as measured by GOTV canvass shifts for additional one-on-one shifts actually comes quite late in the cycle. An assumption that one-on-one shifts are about long-term relationships makes this result non-intuitive, but one-on-one shifts very late in the cycle are usually about identifying, cultivating, and training leaders directly in advance of GOTV. In other words, these shifts are likely designed to fulfill a specific logistical purpose. The other mitigating factor is the number of shifts that occur across the cycle. In the last month of the campaign before GOTV begins, one-on-one shifts account for a tiny share of daily activity (i.e., less than 1 percent of total shifts completed), which is also why the explanatory power of these shifts as measured by R-squared is much lower at the end of the cycle than at the start.

For these reasons, we should not necessarily interpret one-on-one shifts as being more effective late in the cycle. Rather, we should differentiate between the management and capacity value of shifts which happen early and might be used to build team leadership and guide higher levels of engagement, from purely logistical task management happening immediately before GOTV. Both are critical elements for a field program to develop, but the models shown above have a much higher level of confidence for the former and relatively little confidence in the latter.

We can then turn to a parallel analysis, focused on the estimated impact of phone bank shifts and GOTV canvass capacity. The graphs below similarly show the results of linear models estimated for each week of the cycle, examining how an additional phone shift completed in that week is associated with additional GOTV canvass shifts.

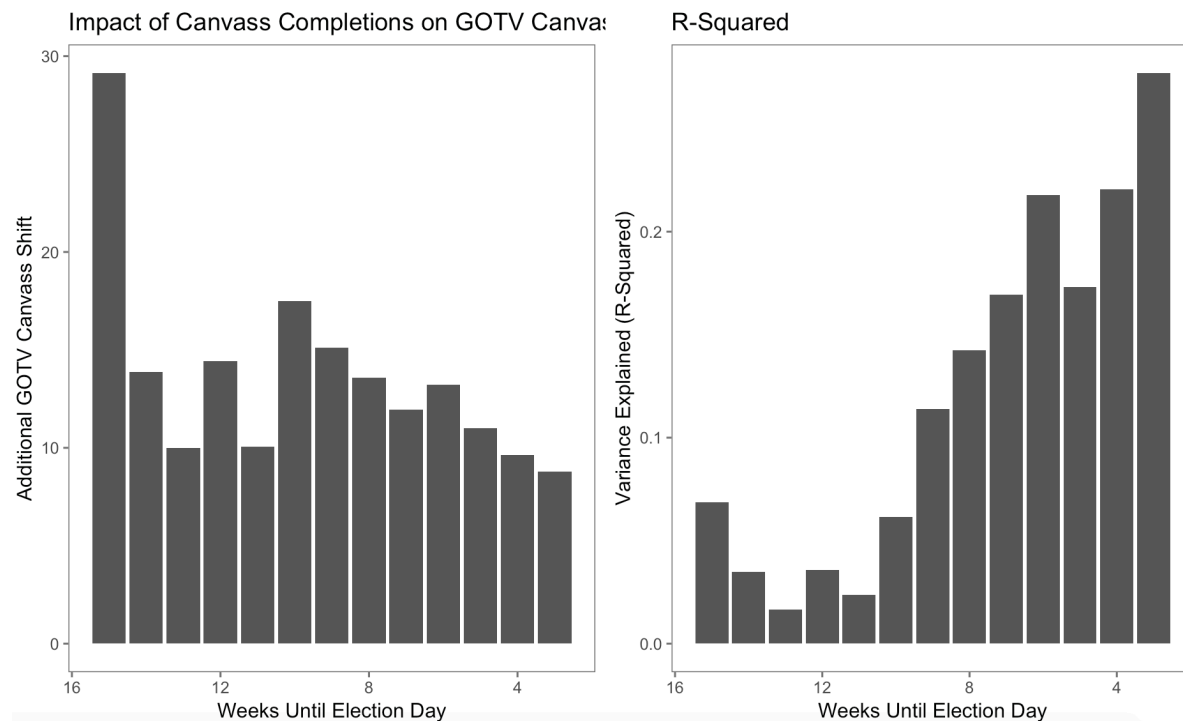
Phone Banks Relatively Less Impactful on Capacity Over Time



The relationship of phone shifts to GOTV canvass capacity appears to be quite different than what we saw previously with one-on-one shifts. Here, the estimated impact of more phone shifts on capacity declines during the cycle, which makes a certain kind of sense. Lots of volunteer recruitment calls early in the cycle might ‘pay off’ to a greater degree than calls relatively late in the cycle, which allow for less time to build relationships and train volunteers in voter contact.

We can take a similar approach to estimate the impact of canvass shifts on additional capacity. Here, the interesting question is the degree to which developing organizational competency and logistical skill in all of the tasks that a canvass shift requires, fulfills not only the mechanical necessities for current shifts but also builds the kinds of durable networks and relationships which drive late-cycle capacity. The graph below mirrors the approach above. Each bar on the left represents a linear model estimate of the impact of an additional canvass shift completed for a given field office (on the left); the bars on right provide the R-squared of this model in explaining the difference between field offices in total GOTV canvass capacity.

Early Impact, But Highest Confidence in Canvass Shift Capacity Impact Late in-Cycle



Perhaps mostly as a result of their rarity early in the cycle, the estimated impact of canvass shifts on canvass GOTV capacity is both quite variable and statistically insignificant in the first few weeks above. However, an extremely simple regression of additional canvass shifts late in the cycle demonstrates serious impact (approximately ten additional GOTV canvass shift per canvass shift completed) and a high degree of confidence (one-third of all variance from this relationship between field offices is explained only by canvass shifts completed in that week).

The tendency for statistical significance of canvass shifts as a capacity driver to increase with the sheer volume of shifts completed across the cycle is intuitive. However, the estimated impact of this work in terms of capacity varies substantially earlier in the cycle, making it difficult to interpret the value of more canvass activity. An aggressive approach to organizing might assume a strong effect of training and team-building in developing capacity; to examine this possibility more closely, we can study the impact of these in-cycle activities in the context of organizational depth, rather than shift volume.

Scaling the Field Program for Organizing Capacity

Activists and scholars like Marshall Ganz have long argued that forming deep relationships with and between volunteers is at the heart of a successful organizing program.²⁵ To better understand the degree that campaigns are accomplishing this work, we should not only look at impact and value in terms of pure mobilizing (i.e., the actual number of activities scheduled or completed), but also the coherence and effectiveness of the organization as a whole.

In addition to a broader philosophical basis for this argument, there is a practical, capacity-oriented argument at play: stronger teams that are better connected with each other and their communities theoretically also are better leaders of new volunteers, stronger advocates with voters, and more effective builders of shared capacity. There is substantial evidence from other industries²⁶ that it takes a long time to develop human capital and train effective teams; we might take a similar economic approach to better understand the degree to which campaigns are developing human capital capacity.

Measuring Organizing Capacity


However, how should we measure ‘depth’ versus ‘breadth’ in on-the-ground organizing programs? The kinds of evidence which we would want to deeply understand the relationships formed during an electoral campaign are unfortunately quite hard to construct from VAN records alone. The quality of conversations between teams, the degree of trust and responsibility shared within an organization, and the maturity of understanding of organizational goals and tasks within an office are not usefully captured in the records of when people showed up and what they did at a high level.

However, there are some useful proxy measures which we can look at more closely. One such proxy is the share of volunteers from across the cycle that schedule at least one shift of any type during GOTV. This serves as a rough measure of how engaged the overall volunteer base of a program is relative to GOTV, as well as some potential measure of how effectively the organizing leadership within an office is stitching together cycle activities and GOTV capacity. The following graph counts field offices in 2016 by the share of all unique volunteers within the cycle who showed up for GOTV.

²⁵ See for example: Ganz, Marshall. “How to Organize to Win.” *The Nation*. March 16, 2018. <https://www.thenation.com/article/how-to-organize-to-win/>.

²⁶ See for example: Gersick, Connie. “Time and Transition in Work Teams: Toward a New Model of Group Development.” *Academy of Management Journal*; Mar 1, 1988; 31, 1

Distribution of Cycle Volunteers Participating in GOTV by Field Office



If we accept the proposition that the share of cycle volunteers participating in GOTV is a useful proxy for organizational strength, then we can frame the work of field management at the region and state level as helping individual offices move from the left side of the distribution above (where very few cycle volunteers show up when their work is most impactful) to the right side.

Before moving forward, it is important to demonstrate that organizational depth as such is not detrimental to overall capacity at a high level. That is, if having a higher share of cycle volunteers showing up for GOTV is negatively correlated with total GOTV capacity, then we would want to be extremely careful about using this metric as a goal to approach. For explanatory ease, the table below includes linear estimates for impact on the log of canvass shifts scheduled in GOTV provided by three input variables: log of GOTV volunteers, the log of total volunteers who show up during the cycle at some point, and the share of cycle volunteers who participate in GOTV.

Impact of Volunteer Base Metrics on Log of Canvass Shifts Scheduled

Dependent variable:			
	Log of Canvass Shifts Scheduled in GOTV		
	(1)	(2)	(3)
Log of GOTV Volunteers	1.034*** (0.011)		
Log of Total Cycle Volunteers		0.997*** (0.026)	
Share of Cycle Volunteers in GOTV			5.708*** (0.166)
Constant	0.300*** (0.043)	-0.651*** (0.129)	2.974*** (0.040)
Observations	2,235	2,235	2,235
R2	0.799	0.392	0.347
Adjusted R2	0.799	0.392	0.347
Residual Std. Error (df = 2233)	0.486	0.845	0.876
F Statistic (df = 1; 2233)	8,888.387***	1,442.637***	1,189.206***
Note: *p<0.1; **p<0.05; ***p<0.01			

The interpretation of these simple regression models is relatively straightforward.

- **GOTV Volunteers:** A one percent increase in total volunteers who are successfully recruited for GOTV (irrespective of their earlier activity) is associated with a 1.034 percent increase in GOTV canvass shifts scheduled.
- **Total Cycle Volunteers:** A one percent increase in total unique volunteers who show up for at least one shift during the cycle is associated with a 0.997 percent increase in GOTV canvass shifts scheduled.
- **Share of Cycle Volunteers in GOTV:** A one percentage point increase in the share of cycle volunteers who also participate in GOTV is associated with a 5.708 percent increase in GOTV canvass shifts scheduled.

At this admittedly high level, there is no initial reason to suspect that increases in this proxy measure for engagement would have a countervailing or negative effect on the ability of a field program to also generate a high volume of GOTV activity. The following pages replicate some of the analyses from the previous section but replace the outcome variable of interest (volume of GOTV and GOTV canvass shifts) with the share of volunteers that participate in GOTV. First, we can look at the aggregate impact of an incremental additional shift completed at any point in the cycle on that office's ability to recruit more in-cycle volunteers to participate in GOTV:

Very Small Bump in Capacity for Shifts Completed; Larger Turfs Struggle in Engagement

	Dependent variable:			
	(1)	Share of Cycle Vols in GOTV (2)	(3)	(4)
Total Canvass Shifts Completed	0.001*** (0.00005)			0.001*** (0.00005)
Total One-on-One Shifts Completed		0.001*** (0.0001)		0.001*** (0.0001)
Total Phone Shifts Completed			0.001*** (0.00003)	0.0003*** (0.00003)
Unique Cycle Vols	-0.0003*** (0.00003)	0.00004* (0.00002)	-0.0002*** (0.00003)	-0.0005*** (0.00003)
Constant	0.201*** (0.004)	0.189*** (0.004)	0.188*** (0.004)	0.189*** (0.004)
Observations	2,261	2,261	2,261	2,261
R2	0.177	0.047	0.131	0.235
Adjusted R2	0.177	0.046	0.131	0.234
Residual Std. Error	0.103 (df = 2258)	0.111 (df = 2258)	0.106 (df = 2258)	0.099 (df = 2256)
F Statistic	243.340*** (df = 2; 2258)	55.336*** (df = 2; 2258)	170.820*** (df = 2; 2258)	173.501*** (df = 4; 2256)


Note:

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

How Timing Affects Shift Impact on Organizing Capacity

With this baseline, we can examine the role of timing within the cycle in capacity-building. To begin, we can again look at the impact and statistical significance of additional one-on-one shifts completed, but with respect to our new outcome variable of interest.


One on One Shifts Peak In 'Organizing' Value Over Three Months Before GOTV



Here, we can detect a much clearer kind of relationship between shifts being completed and capacity generated. An additional one-on-one shift completed between 13 and 14 weeks before the election is associated with a 1 percentage point increase in the share of cycle volunteers that participate in GOTV; this impact is also the point at which we have the most confidence in the relationship between these shifts and GOTV capacity.

We can next run the same sort of analysis, but use phone shifts completed in place of one-on-one shifts as our input variable of interest:

Organizing Value of Phone Calls Steadily Diminishing Across the Cycle




Interestingly, the shape of the distribution of impact and statistical significance is nearly identical for this ‘organizing’ outcome measure as it was for the ‘mobilizing’ outcome measure. Phone shifts tend to have the highest estimated impact on GOTV canvass capacity when they are completed early in the cycle and this estimated impact declines over time. The statistical significance of this relationship moves in the opposite direction; initial-week phone calls explain very little variation in this model compared to later weeks. It’s critical to note here that phones are useful for many things, especially voter contact for hard-to-reach populations and recruitment at scale. These figures should not be interpreted as an aspersion on phone operations at any point in the cycle, but rather a recognition that phone shifts in and of themselves do not have a strong relationship with increased GOTV organizing capacity.

The effect sizes here are not insignificant; an additional phone shift completed 14 weeks before Election Day is associated with a 0.45 percentage point increase in the share of total volunteers across the cycle who also participate in GOTV. The overall takeaway is that in general, earlier phone shifts tend to be more effective than later phone shifts in helping a program build capacity, but the difference is not so stark that it demands urgent reform.

Finally, we can apply the same methodology to canvass shifts. The graphs below duplicate the analysis above, substituting in the estimated impact and statistical significance of an additional canvass shift scheduled in a week against the share of cycle volunteers participating in GOTV.

Canvass Shifts Maximally Impact Organizing Capacity Roughly 10 Weeks from E-Day



Here, a much clearer picture emerges of how canvass shifts are estimated to impact capacity. Very early and very late in the cycle, canvass shifts do not have a strong relationship with additional canvass capacity created. On the early side, we could imagine that this is the result of a combination of factors: very little training; new organizers unfamiliar with the area; few strong relationships between volunteers or organizers; even the summer heat and frustration that could lead to abandonment. On the later side, this might be the result of simply having less time to train and develop new volunteers in time. But in the center of the timeframe, roughly 10-12 weeks from Election Day (i.e., the August-September range), there is a clear peak in the estimated impact of these shifts on organizing capacity, a statistical relationship that also is at the higher end of variance explained.

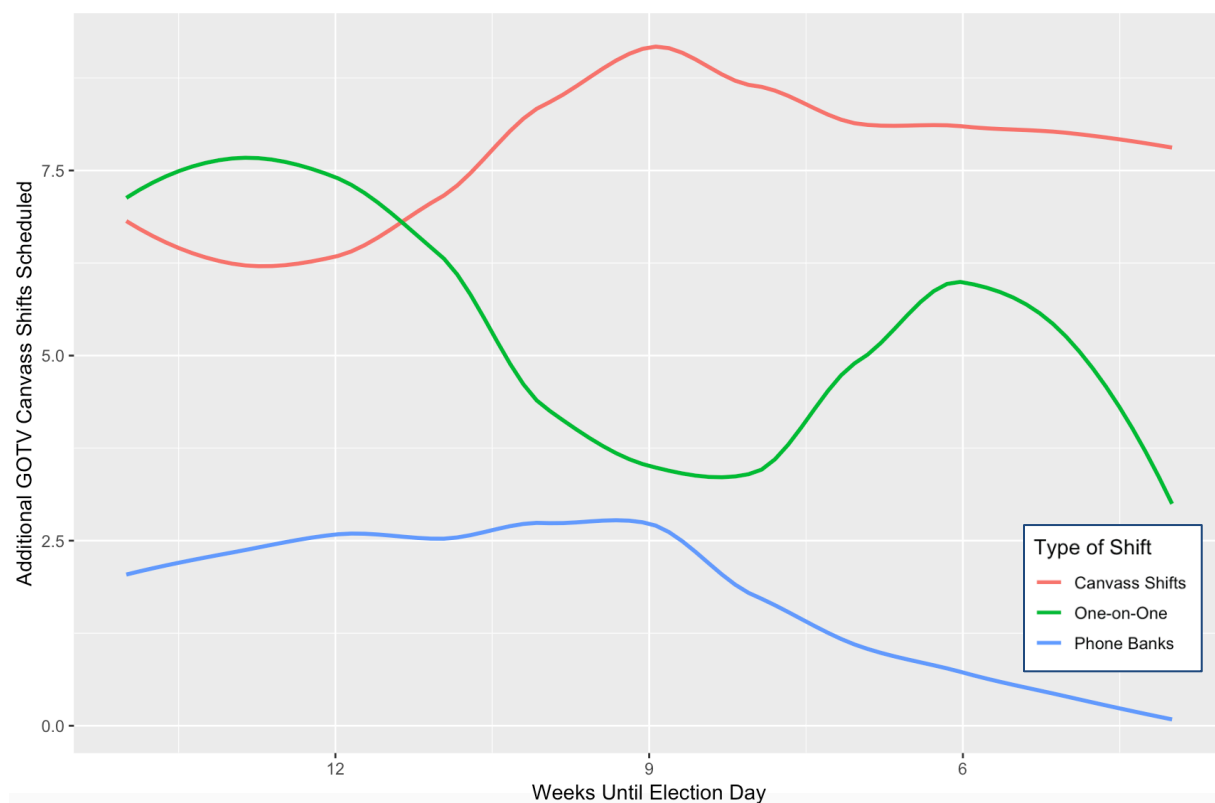
Synthesizing Findings

While the preceding regression output findings might be interesting, we ultimately want to understand not only the estimated impact of these factors in isolation, but rather their interaction within a program. To this end, we can duplicate the methodology used above, but create linear models within each week of the cycle, including as independent variables not only the kinds of shifts completed (i.e., canvass, one-on-one, and phone bank) but also the size and relationship strength of the organization in each week.

We can measure the former by counting the number of total unique volunteers in the door. For the latter, we can categorize volunteers by how active they are as volunteers (i.e., how often they show up to complete a shift); in this case, I measure the share of volunteers showing up that week who had been active volunteers at least one time in the last two weeks. Linear models can then be developed for each week, estimating the relationship between shifts completed and GOTV canvass shifts scheduled, taking these factors into account.

Estimated Impact of Shifts Completed on GOTV Canvass Shifts Scheduled

Controlling for Unique Volunteers and Share Active, Separated by Shift Type




This set of models demonstrates a more useful lesson from the data. One-on-one shifts have the highest estimated impact quite early (i.e., at least three months before Election Day) and

then degrade rapidly; canvass shifts peak in impact between 11 and 8 weeks before Election Day; phone shifts are less predictive of GOTV canvass than either.

We can duplicate the visualization of many linear models above but instead of using canvass capacity as our outcome variable, look more closely at the share of volunteers who participate in GOTV. This gives us a better sense of how ‘organizing’ capacity is created across the cycle.

Estimated Impact of Shifts Completed on GOTV Canvass Shifts Scheduled

Controlling for Unique Volunteers and Share Active, Separated by Shift Type



Very interestingly, the generation of organization capacity follows that of mobilizing capacity quite closely; one-on-one shifts are most useful early in the cycle and degrade, quickly followed by canvass shifts, while phone banks do not have a significant relationship to either.

These visualizations highlight the estimated impact of shifts completed in a retrospective sense, and also highlight a potential opportunity to manage and evaluate field programs in future political cycles. However, they also present substantial questions. What is the right way to manage transitions between shift types? Should all volunteers do the same kinds of shifts, or should there be different-in-kind engagement strategies for some (e.g., volunteer ‘leaders’)? To what degree should these contacts be geared at different strategic priorities of volunteer recruitment, voter contact, and registration? Retrospective data can offer only an increment of insight into these kinds of questions; the following section proposes a structure for a lightweight experiment to test some critical points that emerge from these findings.

Recommendations Based on Initial Findings

With the strong caveat that these observational findings may include significant unobserved endogeneity in the input and output variables (e.g., very experienced volunteers who show up repeatedly and in GOTV regardless of what shifts they are asked to do), they suggest the sketch of an organizing program which fulfills the purposes of not only mobilizing volume, but also organizing a more cohesive and collectively engaged group.

Based on the observable relationships between shifts completed and capacity produced in volume and in engagement, the most effective structure based on these data involves very early July investment of organizer time in one-on-one shifts as a primary goal alongside production of phone shifts for volunteer recruitment. The next stage of production involves a transition from one-on-one shifts directly into canvass shifts in the August-September time frame as phone bank shifts maintain their share of activity. The final month upswing in phone and canvass shifts should then be built on an earlier and larger base of experience in canvass shifts, with volunteer leaders having previously been cultivated and trained in one-on-ones.

In practice, this implies a management model that would be able to train, monitor, measure, evaluate, and manage hundreds of organizers and potentially thousands of volunteer leaders and even more 'line-level' volunteers in complex organizational shifts between modes. The difficulties of this idealized model should not be underestimated; effectively deploying any large organizing program requires immense cooperation and flexibility across diverse regions and states as well as between field offices and headquarters teams. It also demands rigorous data collection, analysis, and reporting, along with the infrastructure and tools that allow field teams to identify and ameliorate risks as well as find and replicate leading practices.

Field Experiments in Future Cycles

The suggestions above are potentially overly simplistic and even naïve in their expectations of smooth, easy recruitment and management of volunteers across fast-moving organizations in battleground states. More importantly, it implies a one-size-fits-all model, when we have every reason to expect that different approaches might work better in different contexts. This is why the recommendation of this report is not to unilaterally apply these as expected outcomes for future campaigns, but rather to identify the quantitative relationships between inputs and capacity that hold potential promise and propose discrete experiments in which these relationships can be tested.

The chief difficulty in making effective use of these findings is that testing volunteer engagement models carries several kinds of inherent risk. Deploying an ineffective engagement model even in a small area risks offending and disengaging volunteers, in the process souring future volunteers in other regions. It also risks wasting potentially critical moments in capacity building in pursuit of what might be very tiny improvements in organizational strength.

For this reason, the experimental approach that I recommend utilizes a model borrowed from the organizing team associated with the 2004 primary campaign of Howard Dean. Attempting to build out a new structure for organizers to engage with volunteers, the Dean New Hampshire team rolled out a volunteer leadership model focused on house meetings alongside several statewide organizing ‘conferences’ that took place weeks and months out. House meetings fed volunteers directly into phone and canvass operations. In this way, the team could motivate action against a discrete and measurable goal, while providing the overall program with adequate time for course correction in time for critical voting periods.

I propose the following simple structure for experimental study of capacity generation in both mobilization and organizing strength, to be conducted in the off-cycle (i.e., 2019 primary and general elections) or in the primary stage of 2020 elections once the initial wave of organizing staff has been hired.

Establish Intermediate Calendar of Capacity Tests

- At six-to-eight-week intervals, establish multiple statewide organizing events to occur on the same day and involve everyone in shared effort. These might be an organizing ‘conference’ set up in a central location for training and best practices sharing. They could also be statewide dry runs of GOTV operations (though the cost both in time and resources of these kinds of practice runs can be substantial).
- These sprints will serve a dual purpose. The primary stated purpose and technical value is that they will provide organizers and volunteers with a chance to ‘practice’, either formally or in the context of a staged event, the logistical and planning efforts that are required for real GOTV. The invite-to-attendance records for all volunteers across the state also provide organizing managers with a clear intra-cycle measure of engagement and effectiveness with plenty of time to course-correct prior to Election Day.

Assign Field Turfs and Organizing Superstructure

- Based on available budget available for field organizers, the state should be cut into regions and assigned to individual organizers, overseen by regional directors. The data and reporting available to all of these staff should be consistent across all management levels to allow for transparency in goals as well as transparency in performance and evaluation of line staff.

Determine Measurement and Management Priorities

- Establish the metrics by which the campaign will focus on both inputs and outputs. These should include volume metrics of pure activity generated as well as specific proxy measures for engagement and organizing quality. These goals should not be determined or communicated directly to organizers until they can be tailored to the particular volunteer demographics and underlying capacities of each turf.
- Initial Recommendations for Program-wide Evaluation:
 - Share and number of one-on-one shifts that lead to a positive designation of the volunteer as a volunteer leader or team captain.

- Share and number of one-on-one shifts that are followed up with an additional shift of any time within two weeks.
- Share and number of one-on-one shifts that are followed up with a canvass shift completion within four weeks.
- Share and number of canvass shifts that are followed up with an additional canvass shift within four weeks.

Allocate Organizing Models to Turfs with Similar Characteristics

- Within each state, similar turfs should be identified for splitting small organizing tests. These turfs should be similar along a set of critical underlying variables:
 - **Voter registration** (Democrats as a share of overall registered voters)
 - **Age distribution** (Over-65 as a share of overall Democrats)
 - **Race and ethnicity** (Groupings of Asian, Black, Hispanic, and White populations within overall Democratic population)
 - **Voting history** (Share of Presidential and Midterm voters within Democratic population)
 - **Previous volunteer activity** (Participation in electoral campaign volunteer operations in either 2018 or 2016)
- There are many aspects of the quantitative findings in this report which are good candidates for further analysis via controlled experiments. Based on the findings in this research, I recommend the following candidate for initial testing across these districts:
 - First canvass shift preparation: Identify whether the inclusion of pre-work by organizers to prepare volunteers for canvass shifts improves throughput capacity, as is indicated by observational findings in this report. Pilot-test one-on-one shift-to-canvass shift transitions in small, controlled environment well before GOTV so that larger-scale organization can roll out program-wide adoption within a window that maximizes effective capacity building.
 - *Groups*:
 - **Control**: Verbal instructions only (Organizer provides volunteer with turf map, list of names, and instructions for what to do)
 - **Control + Written/Video Instructions**: Verbal instructions and detailed handout explaining purpose of canvass shift, best practices, and goals. Written instructions could also be swapped for online/video resource such as Resistance School guide.
 - **Control + Written Instructions + One-on-one Shift**: Volunteers only scheduled for initial canvass shift once they have completed a 30-45-minute one-on-one conversation with the organizer. Volunteers will also be provided with the same written or online resources as the second group.
 - *Intermediate Outcome Measures*:
 - **Ask-to-schedule ratio across offices** (i.e. Of those volunteers who are asked to canvass in each case, how many say 'Yes?'): **Measured continuously**

- **Schedule-to-completion ratio across offices** (i.e., Flake rate for each kind of preparation): **Measured continuously**
- **Re-shift rate across offices** (i.e., Within two and four weeks, how many volunteers volunteer again?): **Measured each week**
- **Attendance rate at statewide organizing events** (i.e., What share of volunteers in each group participate in mock-GOTV activities?): **Measured at six-to-eight-week intervals**

For legibility, these steps are summarized in the graphic below:

Proposal: First Canvass Shift Preparation

