Article

AMERICAN MARKETING ASSOCIATION

Journal of Interactive Marketing
1-15

DOI: 10.1177/10949968241240453
journals.sagepub.com/home/jnm
Sage

Email Campaigns That Suit the Candidate: Leveraging Automated Text Analysis to Increase Political Donations

Simha Mummalaneni, Rebecca Jen-Hui Wang, and Mathew S. Isaac

Abstract

This research employs automated text analysis to explore how textual characteristics in campaign emails affect monetary donations received by political candidates. The authors outline a new methodological framework that combines a machine learning approach for natural language processing with fixed effect regressions, thereby enabling researchers to study and interpret the impact of textual characteristics on donations while also accounting for individual differences across candidates and their email recipients. Using this framework, the authors analyze 764 emails from 19 candidates in the 2020 U.S. Democratic presidential primary election and evaluate how certain textual characteristics (e.g., empathy, vulnerability) in campaign emails affect donation outcomes. Identifying these effects would enable candidates to improve their email text and increase their donations by 9% on average. This research provides a practical and flexible roadmap for automated text analysis in situations where political campaigns do not have clear a priori hypotheses about which textual characteristics will be effective for them.

Keywords

political marketing, email marketing, marketing communications, natural language processing

Online supplement: https://doi.org/10.1177/10949968241240453

Submitted April 5, 2023

Political candidates regularly communicate with citizens directly through email during election cycles to engage with their supporters and to solicit campaign contributions from them. These kinds of individual campaign donations represent a large and growing share of campaign funding; in the 2016 presidential election, political donations from individuals accounted for 71% of Hillary Clinton's fundraising total and 40% of Donald Trump's fundraising total, respectively (Hughes 2017). To tap into the varied motivations for giving, campaigns have historically relied on a range of donation appeals (Hassell 2011; Sabato 1981; Shea and Burton 2006; Verba, Schlozman, and Brady 1995). Given that the content of candidates' appeals exerts influence on donation behavior (Han 2009), it is critical that campaign managers identify the most effective communication strategies for their candidate. In this research, we address this challenge by examining historical campaign emails and the subsequent campaign donations that the campaign received.

We aim to identify specific textual characteristics that different candidates should focus on as they craft their campaign messages. Our intention is to produce high-level insights that candidates can adopt as they proceed through an entire election cycle, rather than generating specific word-by-word recommendations that may be overly granular or idiosyncratic. Our work shows that by leveraging insights gained from a trove of email messages sent earlier in a campaign, candidates can produce superior email messages later in the campaign that will positively influence donation behavior.

From a methodological standpoint, we demonstrate the value of a simple "bottom-up" framework that lets campaigns examine the effectiveness of their email communication over time. Our goal is not to measure all potential reasons why individuals may donate to presidential campaigns, because these are

Simha Mummalaneni is Assistant Professor of Marketing, Foster School of Business, University of Washington, USA (email: simha@uw.edu). Rebecca Jen-Hui Wang is Associate Professor of Marketing, College of Business, Lehigh University, USA (email: rwang@lehigh.edu). Mathew S. Isaac is the Thomas F. Gleed Chair of Business Administration and Professor of Marketing, Albers School of Business and Economics, Seattle University, USA (email: isaacm@seattleu.edu).

varied, idiosyncratic, and often unmeasurable. Instead, we examine email communications and develop a framework that enables us to analyze the effects of email while avoiding the most notable confounds caused by other common explanations. Specifically, as we discuss in detail subsequently, we believe that the timing and context of our research minimizes the likely impact of many factors that might affect donations, including television advertising, online advertising, social media, retail politics, fundraisers, and holidays/major events.

Our framework has three components: a dictionary-based natural language processing approach, a dimension reduction method, and regression analysis. Rather than identifying and labeling potential sources of variation in advance, we look across predefined categories of psychological concepts taken from an established dictionary known as LIWC (Linguistic Inquiry and Word Count) and rely on principal component analysis (PCA) to condense them. Our framework uses machine learning to analyze email text and to generate features that summarize that content (i.e., textual characteristics), and then uses a panel regression to understand how those different textual characteristics affect campaign donations. This combination of methods leads to a set of interpretable textual characteristics that can help political candidates better gauge the effectiveness of various styles and themes used in their emails and make improvements accordingly.

We focus on campaign emails because they are the most effective way for political campaigns to generate campaign donations. In the 2012 presidential campaign, the Obama campaign garnered a record $690 million in online donations, the majority of which could be attributed to fundraising emails (Green 2012; Madrigal 2012). Since then, the value of email lists in political campaigns has continued to grow. Political candidates now routinely pay between $2 and $5 for a single name on an email list, and federal candidates spent at least $6.7 million to rent or acquire lists that included email addresses in the first quarter of 2019 (Evers-Hillstrom and Erickson 2019). However, despite the importance of emails to political candidates, very little academic research has examined the efficacy of political email campaigns. Even this limited work has been largely descriptive rather than normative or predictive (see, e.g., Epstein and Broxmeyer 2020; Williams and Trammell 2005).

Political campaigns ideally would like to send emails with customized content that will resonate with each individual email recipient. However, political campaigns face a number of operational challenges that limit their ability to customize in this way. With few exceptions (e.g., Madrigal 2012), the only customized element in a campaign email is the suggested donation amount. As summarized by a report from the investigative journalism site ProPublica, "the most significant difference between emails in a single message blast is perhaps unsurprising: campaigns change the amount of money they ask for based on how much you've donated in the past" (Larson 2012b). One reason why more sophisticated customization in email messaging is difficult is because political campaign email lists are often compiled from various sources (e.g., lists purchased from or exchanged with other campaigns) and may therefore contain inconsistent or incomplete information about voters in the database. As a result, campaigns send nearly identical emails to all recipients because their lack of reliable information means that they would "mistarget all the time" if they tried to customize each email individually (Hersh quoted in Larson 2012a; see also Hersh 2018; Hersh and Schaffner 2013). This reasoning is also echoed in nonpolitical settings by companies that choose to send identical emails to their customer mailing lists (see, e.g., Neslin et al. 2013).

To verify that presidential campaigns do not send customized emails to individual recipients, we interviewed three experts with high-level experience in digital outreach for presidential campaigns.[1] According to Tobin Van Ostern, a campaign advisor to multiple presidential and gubernatorial campaigns over the last 15 years,

> Emails from presidential campaigns are primarily personalized at the individual level in two ways: they use the names of the individuals in the subject line or salutation to make it seem more personal, and they try to suggest a donation amount that makes sense for the recipient. So for instance, some recipients might be asked to donate $20 while others might be asked to donate $100. Apart from these types of customizations, everything else in the email's text tends to be the same for everyone on the mailing list. In general, the goal is to write a good email to send to everyone on their mailing list, rather than implementing significant differences in language or tone within a given email campaign.

These views were corroborated by other experts we spoke to. Brandon English, who led email and digital outreach for the Democratic Congressional Campaign Committee from 2006 to 2015 as well as the Joe Biden campaign in 2019, said:

> The Obama 2012 campaign briefly experimented with sending emails on different topics to different people, but that led to a drop in donations and they realized they didn't have enough data or staff to do it well. Since then, campaigns have stayed away from that level of email customization and instead tried to send well-written emails to the widest possible audience.

Lauren Miller, former senior advisor and digital director for Elizabeth Warren from 2012 to 2021, also spoke about the difficulties involved with sending different emails to different recipients. When asked why campaigns do not do this, she responded:

> I wish they had that level of sophistication. Most of these digital teams have just a handful of people on staff. It's not feasible to write that many different emails, even apart from the difficulties of figuring out which people should get which emails. From a day-to-day standpoint, the goal is just to write the email that will resonate with people the most.

[1] All three of these interviews took place in September 2023.

Because campaigns typically send the same email to every recipient in their email list, their goal should be to make that email as effective as possible. One way to gauge the effectiveness of various email techniques is to use A/B email tests. These pilot tests have focused primarily on small changes to subject lines, requested donation amounts, or link placements (e.g., "donate now" buttons) in an email message. In the 2020 presidential election cycle, Bernie Sanders's campaign deployed A/B pilot tests to evaluate the effectiveness of providing information about donor history or local events in email correspondence (Viveiros 2019). The goal of A/B pilot tests is to utilize a small portion of the candidate's database to identify the most effective version of an email, which will subsequently be sent to the entire email list (Moth 2013). However, despite these benefits, A/B tests are inherently limited in that they can only reveal a handful of insights at a time (Gallo 2017; Heng et al. 2018). A promising alternate path that campaigns can take to arrive at an overarching email strategy involves the analysis of historical data sets. Recently, researchers adopted this approach and mined 700 emails sent to supporters of Texas governor Greg Abbott in the 2018 election cycle to identify certain themes and email subject lines that were correlated with higher email open rates, click-through rates, and contributions (Gaynor and Gimpel 2021). Although useful as a proof of concept, this prior work has focused on emails from a single candidate in an election, leaving it unclear whether the results would hold for other candidates. To our knowledge, our research is the first to conduct a systematic investigation to assess how textual characteristics of political emails exert a differential impact on donation behavior *depending on the candidate*.

Methodological Framework

We propose a methodological framework to analyze how the textual characteristics of campaign emails affect campaign donations. We apply this framework to emails sent by candidates who participated in the 2020 Democratic presidential primary race, which enables us to derive substantive insights related to those candidates' email strategies. However, our framework is generalizable to future campaigns as well. Since it relies solely on publicly available data, it can be replicated by political campaigns to measure and improve the efficacy of their campaign emails without having to expend significant financial resources.

Data Sources

Our data consist of two parts: campaign email data and campaign donation data. We collected each of these separately and then merged them to create a data set in which each observation is a unique candidate-state-day.[2]

Email data. The campaign email data were collected by the data journalism website FiveThirtyEight, as part of an article focusing on how Democratic presidential candidates were discussing Donald Trump in their emails (Mehta 2019).[3] FiveThirtyEight's data collection method was simple and can easily be replicated by researchers analyzing future campaigns: they simply signed up to be on the email lists for every major Democratic presidential candidate and collected the emails they received.

One potential threat to our data analysis would be if our email data were nonrepresentative, that is, if the emails in our data set were substantially different from the emails that most people received. Although our email data set does not include all the emails that everyone sees, this concern is mitigated for two reasons. First, prior research has shown that most campaign emails are sent to a wide audience and any customization in the email tends to not affect the main body of the email text (Larson 2012a). Second, FiveThirtyEight used a generic name and email address that was not linked with previous campaign donations or any other traceable activity. This means that even if campaigns wanted to customize the emails they sent to FiveThirtyEight, they had no information that would enable them to do so.

Our email data consist of html files of each of the 764 emails that 19 Democratic presidential campaigns sent out between May 25 and June 26, 2019.[4] We focus on the main body of the text in each email and ignore the subject line, the footer/signature, and any images or logos that may appear. This approach enables us to focus on the textual characteristics of the emails rather than visual aspects, and it also ensures that we are focusing on the parts of the campaign emails that are unlikely to be customized for specific email recipients (Larson 2012b; Madrigal 2012).[5]

Donation data. The campaign donation data come from the Federal Election Commission (FEC). Campaigns are legally required to disclose itemized donations from individuals who donate at least $200 to their campaign, and this information is logged by the FEC. The fact that campaigns are not required to disclose itemized donations below $200 usually yields censored data, in which "small donors" are not observed. As a result, most of the prior literature focusing on campaign donor behavior has not been able to analyze the donation behavior of these small donors. Focusing solely on people who donate at least $200 to a particular campaign (e.g., Barber 2016; Heerwig 2016) may result in only about 66% of all donations being represented in the data set (Alvarez, Katz, and Kim 2020).

We avoid this censored data problem by following recent research from Alvarez, Katz, and Kim (2020), who recommend

[2] The data set is on Open Science Framework (https://osf.io/gzu86/).

[3] This particular data set is publicly available on FiveThirtyEight's GitHub page (https://github.com/fivethirtyeight/candidate-emails/).

[4] The original data from FiveThirtyEight include emails from 22 candidates, but three of the candidates (Tulsi Gabbard, Mike Gravel, and Eric Swalwell) were dropped from our analysis because they did not have enough variation in their emails during this time period.

[5] In these emails, the visual aspects are minimal. Typically, the only visual aspects are a small campaign logo and a blue "donate" button, and these two are both held constant across all emails sent by each particular candidate.

including campaign donations that are reported by ActBlue, an online payment processor used by the Democratic presidential candidates' websites. ActBlue reports all these itemized online donations to the FEC, not just ones exceeding $200. Including the ActBlue data adds another layer of data processing but dramatically improves our data coverage: estimates by Alvarez, Katz, and Kim suggest that incorporating small ActBlue donations along with the larger campaign-disclosed donations enables nearly 100% of all campaign donations to be captured.

Dependent Variables

The original donation data from the FEC are itemized: each observation is a unique donation from a specific individual to a specific candidate. The FEC data contain relatively limited information about the donor, all of which is self-reported: their name, their occupation and employer, and their city, state, and zip code. This paucity of individual-level characteristics means that we focus on overall donation behavior patterns rather than trying to understand why specific donors behave differently than others.[6] As a result, we aggregate the FEC donation data to the candidate-state-day level. This level of aggregation enables us to flexibly account for candidates having different levels of popularity and financial support in different parts of the country, without having to model this heterogeneity at the level of the individual donor.

We examine two different dependent variables: (1) the total donation amount and (2) the donation amount coming from small donations (defined here as donations below $200). For a given candidate c, state s, and day t, the former dependent variable represents the total donation amount (in dollars) that candidate c received from donors living in state s on day t. Next, we create the latter dependent variable by repeating this process but only looking at donations below $200. Alvarez, Katz, and Kim (2020) demonstrate that small donors have slightly different demographic characteristics than large donors; in particular, they are more likely to be women, and they are less likely to be White. Keeping track of small donations as a separate dependent variable enables us to examine whether these small donors also respond differently to textual characteristics of the email.

Apart from generating donations, campaign emails can also help increase turnout and mobilize volunteers (Nickerson 2007; Vaccari 2008). The 2020 Democratic party primaries and caucuses took place between February 2020 and August 2020, and the emails used in our analysis are from May 2019 to June 2019; this suggests that increasing turnout and mobilizing election-day volunteers were distant concerns during the email period we examine. However, our analysis does not require us to assume that candidates are sending emails solely to maximize donations. Even if candidates have goals other than generating donations, our modeling framework and our results can still provide guidance regarding how donations will be affected by the textual characteristics of the emails that they choose to send out.

Independent Variables

Our main set of independent variables is derived from the LIWC dictionary (Pennebaker, Francis, and Booth 2001; Tausczik and Pennebaker 2010). For each email in our data, we first analyze the text in the body of the email with LIWC. Specifically, we use 42 variables of psychological processes in LIWC, including keywords related to cognitive processes, perceptions, drives, time orientations, and personal concerns (see the Web Appendix for a full explanation of our text analysis framework). Given this large number of LIWC variables, we then reduce the data set's dimensionality by performing PCA. The PCA analysis suggests that to explain 70% of variation in the data, 14 components are required, as shown in Figure W1 of the Web Appendix. The scree plot, shown in Figure W2 of the Web Appendix, also reveals that there are 14 components that have an eigenvalue greater than or equal to one (Cattell 1966).

In essence, PCA decomposes a given data set into loadings and scores. The components and their LIWC dimensional loadings are summarized in Table W1 of the Web Appendix. Based on the most contributive LIWC variables to each component, we give it a label. For instance, for the first component (or textual characteristic), which we label empathy, some of its most contributive LIWC variables are affect, negative emotions, hear, and perception. To illustrate with a concrete example, the following is an excerpt of an email from Elizabeth Warren that has a high empathy score: "We've got more work to do to make sure everyone in the LGBTQ community is safe to be who they are and love who they love without having to face the fear of violence. … We remember the victims, we grieve with their loved ones, and we honor them with action."

Besides empathy, we label the remaining components (or textual characteristics) as follows: conversation, reward, reflection, family, solicitation, mobilization, risk, masculinity, femininity, contrast, vulnerability, localization, and mortality. Table W1 of the Web Appendix shows the resulting 14 components (or textual characteristics), their contributing LIWC variables, and example keywords. To provide additional discussion and contextual insights, we also include examples of emails that reflect each component in the Web Appendix.

To account for the fact that a candidate might have sent multiple emails in a short period of time, we operationalize a candidate's email communication by calculating the three-day rolling averages of each of the PCA components. As we mention previously, our dependent variables, total donation amount and small donation amount, are aggregated to the candidate-state-day level. However, emails are sent at a less predictable schedule:

[6] Another challenge with focusing on individuals' behavior is that we do not observe the recipients for each email. This means that we cannot describe individual-level responses or any intermediate behaviors like opening emails, clicking on links, and so forth. However, this data limitation does not affect the main analysis of the article, which focuses on measuring how different textual characteristics affect overall donations.

candidates sometimes send multiple emails per day, and sometimes they go multiple days without sending emails. In order for the email data to comport with the dependent variables, we aggregate the email characteristics to the candidate-day level. We do so by assuming that there is a three-day window in which emails have an effect on donations: donations made on day t are attributed (in part) to emails that were sent on day $t-2$, day $t-1$, and day t. Therefore, for day t, we calculate the averages of the PCA components for all the emails sent by that candidate on day $t-2$, day $t-1$, and day t. These average PCA components are then used as independent variables in our regression.

In addition to these average PCA components, we also include a parallel version of them that normalizes the data and focuses on within-candidate deviations from each candidate's usual email behavior. To construct these parallel components, we take z-scores (relative to the candidate average across all emails) for each of the PCA components. Therefore, the value of these z-scored components represents how much more or less the day's emails emphasize each particular PCA component, relative to the average email behavior of that candidate overall.

Data Descriptives

Candidates in our data vary in terms of how often they send out emails; these patterns are displayed in Figure 1. The corresponding donation amounts for each candidate are displayed in Figure 2. Note that there does not appear to be a strong relationship between the candidates' email frequency and their donations; that is, candidates who send out more emails (in Figure 1) do not necessarily receive greater donations (in Figure 2). Based on the patterns in Figure 2, there are four "major" candidates who received the majority of donation dollars: Joe Biden, Pete Buttigieg, Bernie Sanders, and Elizabeth Warren.

In addition to having different email frequency patterns, candidates also differ in terms of their textual characteristics. Figure 3 provides averages of each of the 14 PCA components, by candidate. Candidates whose average values are near zero can be seen as having email patterns that are broadly in line with their competitors, while candidates whose average values are further away from zero have email patterns that deviate more strongly from the herd: positive values mean that they use a particular textual characteristic more than others, whereas negative values imply that they use it less. The single highest value is Joe Biden's emphasis on using empathetic language (empathy), while the lowest values represent Cory Booker rarely focusing on directly soliciting money from email recipients (solicitation) and Kamala Harris refraining from using reward-oriented language (reward). Candidates like Beto O'Rourke, Michael Bennet, and Pete Buttigieg have the least distinguishable emails: they do not have very low or very high values on any of the 14 textual characteristics we examine.

Our final data are aggregated to the candidate-state-day level. To estimate the relevant parameters in our model, we need substantial variation in how much gets donated from day to day. Figure 4 shows that the daily donation amount is highly variable, both across days and states. Note that we have 52 "states" in our data: this corresponds to the 50 U.S. states, plus Washington, DC, and an "other" category that contains donations from U.S. citizens living abroad.

Baseline Model

Our main regression model focuses on total donation amount (in dollars) as the dependent variable. The key independent variables in this regression are the 14 average PCA components representing the textual characteristics of the emails sent in the past three days. We allow for heterogeneity across the different candidates in terms of how effective the different email characteristics will be; to account for this, we estimate candidate-specific coefficients for each of these PCA components. As a control, we also account for the number of emails sent in the past three days. Finally, our model includes fixed effects by day as well as fixed effects by candidate-state.

Regression Specification

Recall that each observation in our data is a unique candidate c, state s, and day t. Formally, we estimate the following regression:

$$\begin{aligned}\log(\mathrm{DonationAmount}_{cst}+1) = {} & \alpha_{1sc} + \alpha_{2t} \\ & + \beta_1 \log(\mathrm{NumEmails}_{ct}) \\ & + \beta_{2c}\mathrm{Contrast}_{ct} \\ & + \beta_{3c}\mathrm{Conversation}_{ct} \\ & + \beta_{4c}\mathrm{Empathy}_{ct} \\ & + \beta_{5c}\mathrm{Family}_{ct} \\ & + \beta_{6c}\mathrm{Femininity}_{ct} \\ & + \beta_{7c}\mathrm{Localization}_{ct} \\ & + \beta_{8c}\mathrm{Masculinity}_{ct} \\ & + \beta_{9c}\mathrm{Mobilization}_{ct} \\ & + \beta_{10c}\mathrm{Mortality}_{ct} \\ & + \beta_{11c}\mathrm{Reflection}_{ct} \\ & + \beta_{12c}\mathrm{Reward}_{ct} \\ & + \beta_{13c}\mathrm{Risk}_{ct} \\ & + \beta_{14c}\mathrm{Solicitation}_{ct} \\ & + \beta_{15c}\mathrm{Vulnerability}_{ct} + \varepsilon_{cst}.\end{aligned} \tag{1}$$

Model Identification

Average donation amounts vary both across candidates and across states. The inclusion of the state-candidate fixed effect α_{1sc} accounts for this fact, and it ensures that we are only comparing outcomes across time within a specific state-candidate

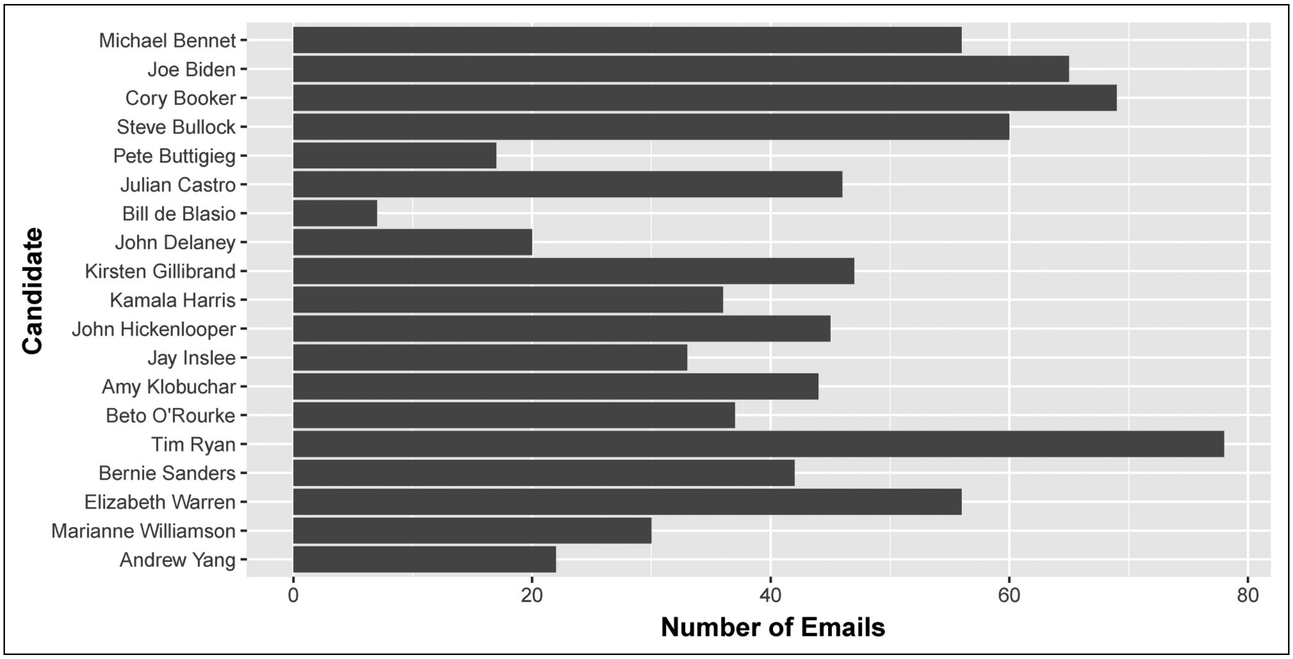

Figure 1. Number of Emails Sent by Each Candidate.

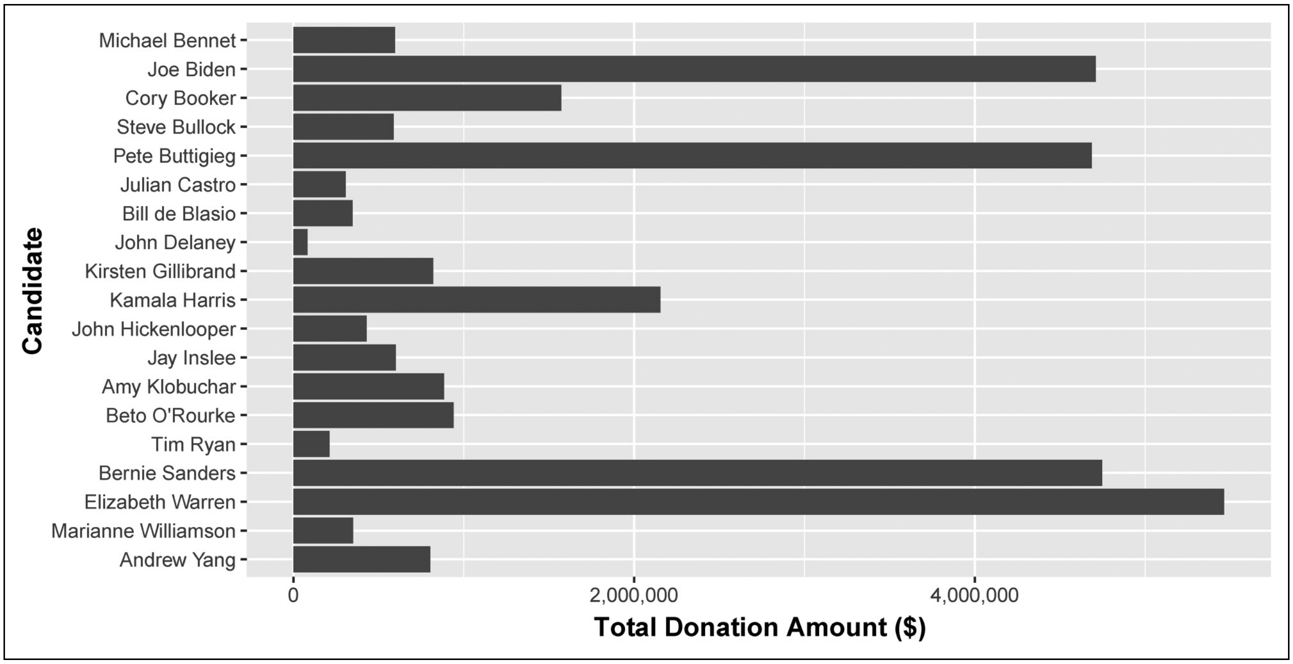

Figure 2. Total Donation Amount by Candidate.

dyad. Similarly, the day fixed effect α_{2t} accounts for the fact that some days may bring in more donations in general.

The key parameters of interest are the parameters $\beta_{2c} \ldots \beta_{15c}$, which represent the candidate-specific coefficients for the 14 textual characteristic variables. These coefficients are identified through linguistic variation in each candidate's emails over time. Our identifying assumption is that these textual characteristic variables are not correlated with the unobservable term ϵ_{cst}, conditional on the fixed effects α_{1sc} and α_{2t} that are also included in our model. This identification strategy is based on Rossi (2014), which argues that high-dimensional fixed effects can yield valid results because they minimize the unobserved error and therefore significantly reduce the possibility that the error term might be correlated with one of the independent variables. If this assumption holds true, then the residual variation in textual characteristics (within candidate, across time) is plausibly exogenous and our model estimates

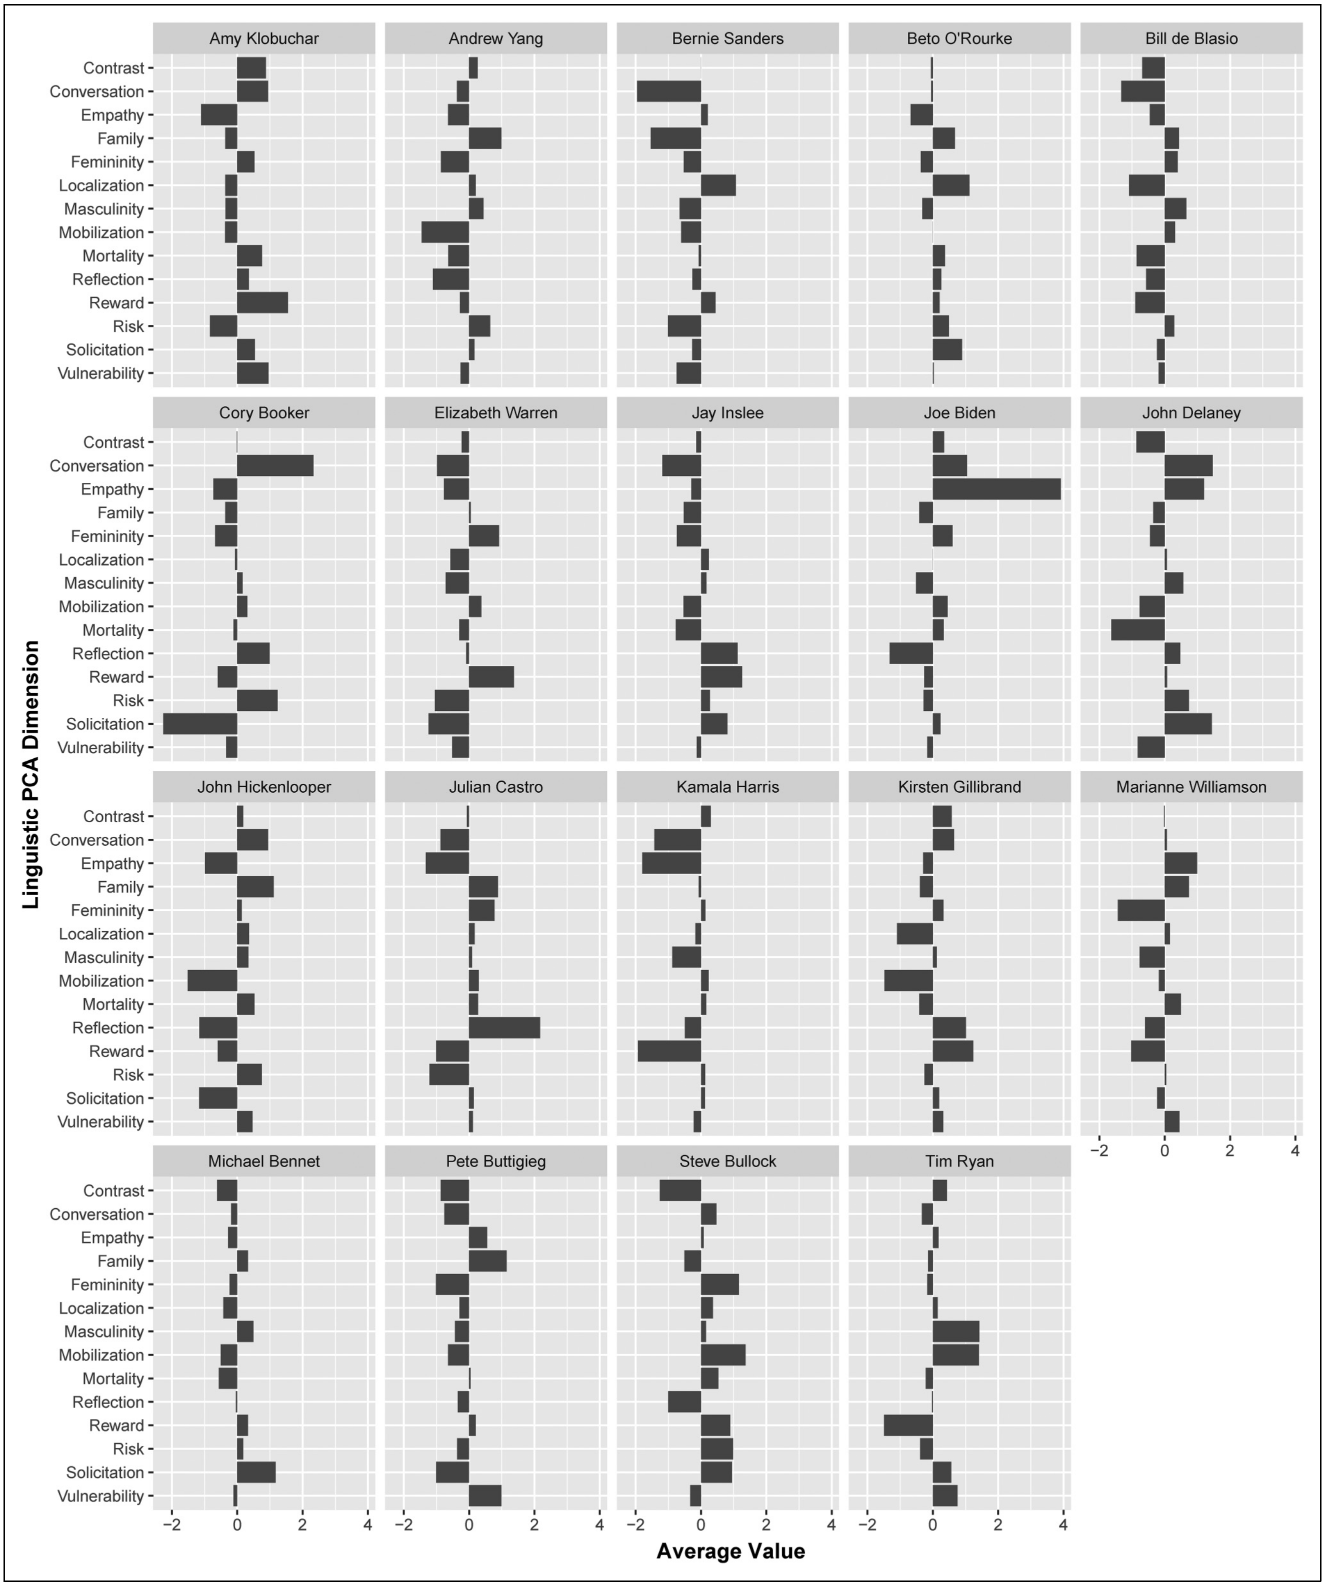

Figure 3. Average Textual Characteristic Patterns by Candidate.

can be interpreted causally. This assumption means that variations in textual characteristics across emails are strictly exogenous, conditional on the fixed effects that are incorporated in the model. The fixed effect α_{1sc} is particularly crucial here because it means that we are not leveraging across-candidate variation in email content to identify any of the parameters $\beta_{2c} \ldots \beta_{15}$. Instead, those parameters

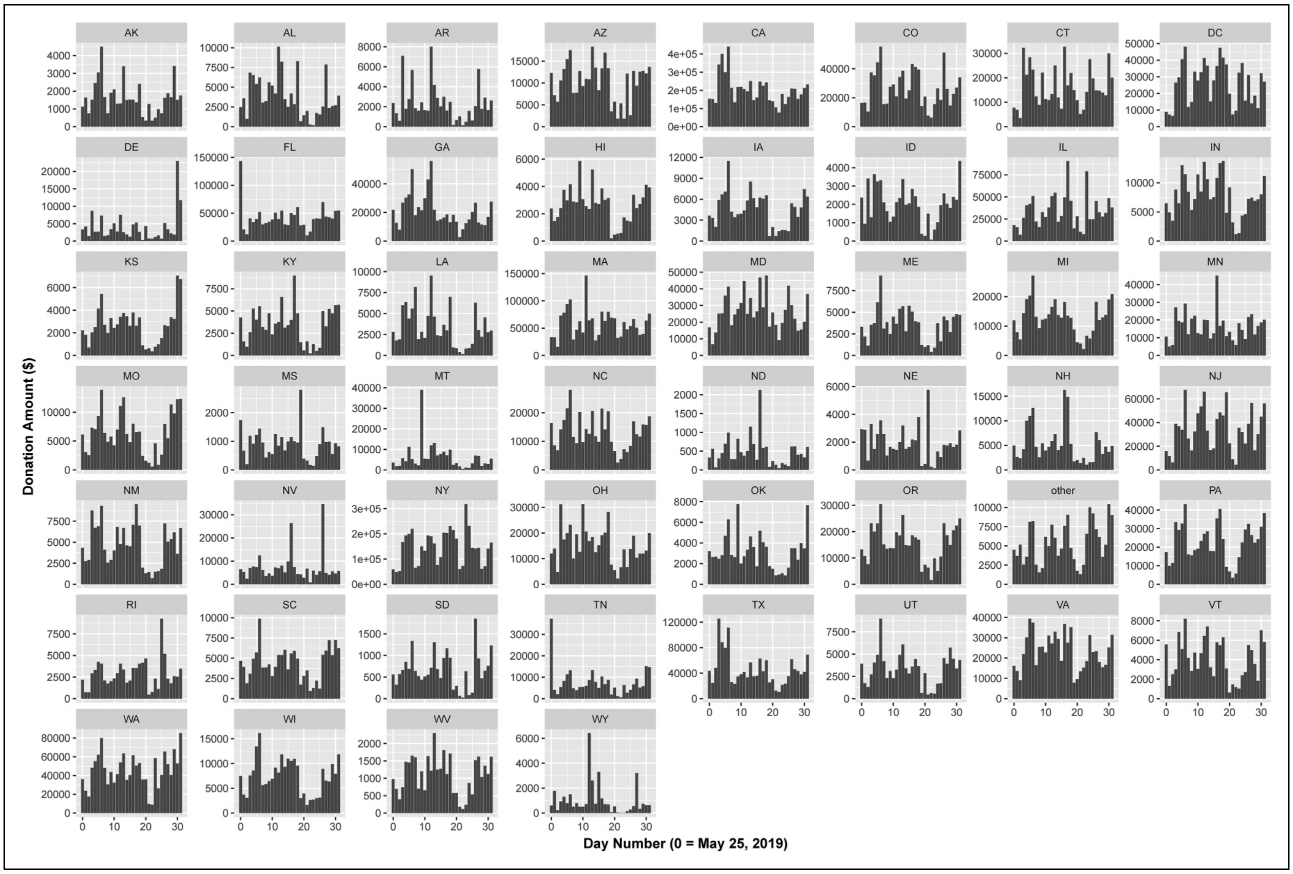

Figure 4. Donation Patterns by State and Date.

are identified solely through within-candidate variation in email content over time.

Note that our plausible exogeneity assumption does not imply that campaigns are sending emails at random. Campaigns may have a "house style" in which they tend to emphasize a particular set of textual characteristics. These campaign-level differences are summarized in Figure 3, and they do not violate our plausible exogeneity assumptions because they are subsumed by the fixed effect α_{1sc} in our model. One way of interpreting our identification strategy is that we are leveraging within-campaign deviations from their usual house style: as long as they occasionally vary their emails by emphasizing some textual characteristics or de-emphasizing others, this generates the variation that is necessary for our model to be identified.

A key source for these within-campaign deviations across emails is the identity of the email writer. Modern presidential campaigns typically have over 20 staff members who write campaign emails (Sutton 2013). In our data, we do not observe the identity of the staff members who write each email.[7] However, this institutional detail explains why we have a key source of variation for our analysis. If different staff members tend to write emails in slightly different ways, this would provide a plausibly exogenous source of variation for each email's textual characteristics.

Our plausible exogeneity assumption also does not imply that campaigns are unresponsive to current events or other temporal shocks. For instance, if there were an important event in the news, multiple campaigns may choose to respond to this event and discuss it in their emails. This type of behavior would not violate our plausible exogeneity assumption, since the day fixed effect α_{2t} controls for any temporal factors that affect all campaigns. These current events would in fact aid in our identification if campaigns responded to the same event in different ways: for instance, one candidate may choose to discuss a tragic news story by being empathetic, a second candidate by being reflective, a third candidate by being vulnerable, and so forth.

In practice, our plausible exogeneity assumption would be violated if candidates are able to perfectly optimize their email content each day because they know ex ante that certain kinds of textual characteristics will be more effective for them on some days than on others. This kind of optimization is unlikely in our setting because it requires the following unlikely conditions to hold: (1) donors' responsiveness to certain textual characteristics varies over time, (2) this donor responsiveness to certain textual characteristics is candidate-specific and therefore

[7] The identity of the staff member who wrote the email is usually not reflected in the "from" field of the email. Typically, the emails are nominally from a prominent figure like the candidate, their campaign manager, or the candidate's spouse, even though these individuals are not the ones actually writing the email content.

not subsumed by the fixed effect α_{2t}, and (3) the campaigns know how donors' responsiveness to certain textual characteristics is going to vary ahead of time, and they can use this information to craft their emails. As long as the email content is not optimized in such a fashion, our model estimates can be interpreted causally.

At a high level, our approach of estimating the effect of email textual characteristics through their temporal variation is similar to how researchers have also estimated the effect of TV advertising (see Shapiro, Hitsch, and Tuchman [2021] for a recent example). In our setting, we observe changes in the intensity of textual characteristics and connect these to changes in candidates' donation amounts. In the Shapiro, Hitsch, and Tuchman (2021) context, they observe changes in the intensity of advertising exposure and connect these to changes in brands' retail sales.

Our model (Equation 1) is parsimonious and does not control for a host of additional variables beyond those that summarize the textual characteristics of candidates' emails. Omitted variable bias could be a threat to causal interpretation of our results, but only if the omitted variables were correlated with the 14 textual characteristic variables that we use in our model; for example, if campaign A emphasized a certain textual characteristic only on days when they received positive press. In our setting, these kinds of patterns do not appear to be present. Instead, campaigns' emphasis on different textual characteristics seems to vary more idiosyncratically, perhaps due to the fact that the emails are written by different staff members who have their own stylistic preferences. Any omitted variables that are not correlated with usage of the 14 textual characteristics would simply mean that there will be more unexplained variation ϵ_{cst}; however, our coefficient estimates $\beta_{2c} \ldots \beta_{15c}$ will remain unbiased. This means that although we would be unable to comment on the effect of other marketing levers on campaign donations, our measurement of the effect of email textual characteristics would be accurate.[8]

A broader concern is the presence of confounds that might cause us to misattribute fluctuations in donation amounts. For instance, our estimates could be misleading if a candidate ran a widely watched TV advertisement emphasizing their empathy and then also concurrently sent out an email along the same theme—our regression would attribute jumps in the donation amount to the textual characteristics of the email, when in reality the TV advertisement might have been the main catalyst. As with any research relying on observational data, we cannot rule out all potential confounds in our data; however, we are able to mitigate some of the more severe concerns as described subsequently.

[8] Similarly, our analysis also does not model whether people donate to a specific focal candidate because of an email from a rival candidate. The textual characteristics of emails from these rival candidates are unlikely to be correlated with the textual characteristics of emails from the focal candidate (above and beyond the fixed effects already in the model), so this omission should not affect our coefficient results.

Television advertising. Although TV advertising is heavily used by campaigns later in the election cycle, usage is near zero during the early part of the campaign that we are studying here. See Bycoffe (2020) for a summary of how presidential campaigns used TV advertising during the 2020 primary election. Furthermore, TV advertising is typically used to convince people to vote for their focal candidate, rather than to convince them to donate. In fact, most TV advertising by political campaigns is feasible only because they have already raised substantial funds through email campaigns (Doubek 2015).

Online advertising. Advertising through Google Ads or Facebook is minimal during the early part of the primary election. For example, the Biden campaign and its associated political action committees spent about .6% as much money on Google Ads during the first week of data we analyze (ending June 2, 2019), compared with the same time period one year later (OpenSecrets 2021).

Retail politics. The phrase "retail politics" refers to politicians engaging in direct contact with citizens, typically by attending local events or rallies, or by visiting public settings like restaurants or county fairs. These kinds of events may influence citizens to vote for specific candidates, and they may also lead to additional campaign donations from people living in a specific geographic area. In our context, these events are unlikely to have a major effect on our results for two reasons: (1) it is unlikely that the success of particular campaign events will be systematically correlated with the textual characteristics of emails that get sent on the same day, and (2) the effects of retail politics are likely to be limited to donors who live nearby, and these effects would therefore be minimized when considering donations from the entire country as we do in this analysis.

Social media. Although social media is popular among political campaigns because it is an inexpensive way to contact a broad group of people, it is not a major source of political donations. Fundraising leaders for presidential campaigns estimate that "our campaigns will do 70 percent plus of their fundraising through email" (Doubek 2015). Furthermore, our estimates for the effects of textual characteristics would still be valid unless campaigns were specifically using different text strategies across these two tools (i.e., if there was a negative correlation in usage of textual characteristics between emails vs. their social media). As long as that is not the case, our approach can still yield unbiased estimates for the effects of different textual characteristics.

Fundraisers. Campaigns occasionally host exclusive events that are only open to people that are willing to make a high-value campaign donation. In 2019, Pete Buttigieg was criticized by his competitors for hosting an event at a Napa Valley "wine cave" where supporters donated thousands of dollars to attend (Higgins 2019). Although these events can lead to a substantial spike in donation amounts (Goldmacher 2020), we show that our results still hold when we focus on small donors (i.e., those that are not driven by these fundraisers or other high-value outlier donations). Details from this robustness check are provided in the Web Appendix.

Public statements by President Trump. In 2019, public statements and Twitter posts from Donald Trump were a frequent topic of discussion for the Democratic presidential candidates. Many candidates criticized controversial statements from President Trump in their emails as a way to motivate their potential donors and raise money from them. Since President Trump's actions and statements would affect all candidates, their effect on candidates' donations would be absorbed by the day fixed effect α_{2t}.

Holidays and major events. People may be more likely to donate on specific days such as Memorial Day or Independence Day that emphasize patriotic values. Furthermore, campaigns may be more likely to use specific textual characteristics (e.g., focusing on reflection and empathy) in their email communications around these same holidays. The simultaneous combination of these two patterns would typically cause a confound in our analysis, but we adjust for this through the inclusion of the day fixed effect α_{2t}. This enables us to control for changes in donation behavior caused by holidays or other major events, rather than the email content specifically.

Results

The key coefficients of interest are the 256 candidate-specific textual characteristic coefficients $\beta_{2c} \ldots \beta_{15c}$.[9] These coefficient estimates are displayed in Figure 5; the coefficient point estimates are represented by dots, and the 95% confidence intervals for each estimate are represented by the horizontal bars. Out of these 256 coefficients, 143 of them are statistically significant. The full regression output from this model appears in the Web Appendix.

We next examine the coefficients for the four "major" individual candidates in detail. The coefficients for Joe Biden are provided in Figure W3 of the Web Appendix. There are five textual characteristics that have a positive and significant effect on Biden's donations: empathy, masculinity, mobilization, reflection, and solicitation. The Biden campaign would benefit financially from increasing their use of these language components. Meanwhile, there are two textual characteristics that have a negative and significant effect: family and vulnerability. Campaign donations would drop if the Biden campaign were to increase these values, so instead they would benefit from reducing the frequency with which they use these textual characteristics.

The corresponding coefficient plots for Pete Buttigieg, Bernie Sanders, and Elizabeth Warren are in Figures W4–W6 of the Web Appendix. Buttigieg's coefficient estimates are overall larger in magnitude than those of the other candidates, thereby suggesting that the Buttigieg campaign's email recipients tend to be more affected by changes in textual characteristics. Sanders benefits the most from mobilization and femininity, and he loses donations with more masculinity. Finally, Warren receives the biggest jump in donations when she stresses localization and vulnerability, and she suffers the largest drop with mortality.

For these four major candidates, we summarize the significant coefficient results in Table 1. Even just among these four candidates, there is no single textual characteristic that has positive effects for all four candidates, nor is there any single textual characteristic that has negative effects for all four candidates. This indicates that accounting for heterogeneity across candidates in our model is vital; our primary substantive finding is that no two candidates should be focusing on the same textual characteristic.

Robustness Checks and Extensions

The main results from our baseline regression are twofold: (1) the textual characteristics of candidates' emails do have a significant effect on the campaign donations that they subsequently receive, and (2) the effects of different textual characteristics vary dramatically across candidates. The combination of these two results indicate that presidential candidates would benefit financially from carefully crafting their emails, but that the optimal email looks quite different for each candidate. We next examine whether these results are robust to alternative model specifications and extensions. Full details for these analyses are provided in the Web Appendix, but we summarize the key findings as follows:

1. Candidates are not already optimizing their emails' textual characteristics, and they do not seem to know which characteristics work best for them. We show this by using z-scored versions of the textual characteristic variables, and we find that deviations from the candidate's baseline usage patterns often lead to improved outcomes.
2. Our main results are not driven by high-value donations. We show this by restricting our analysis only to small donations (donations below \$200), and we find that the main results are consistent.
3. The point estimates of our coefficient results remain consistent when we include candidate-state-week fixed effects in our model rather than candidate-state fixed effects. However, the standard errors become wider, and more of the coefficients are statistically insignificant. Because we have only one month of data, we are

[9] With perfect variation in the variables, there would be a total of 266 estimated coefficients. However, 10 coefficients were unable to be estimated because there was not sufficient variation in that particular variable for that specific candidate over time.

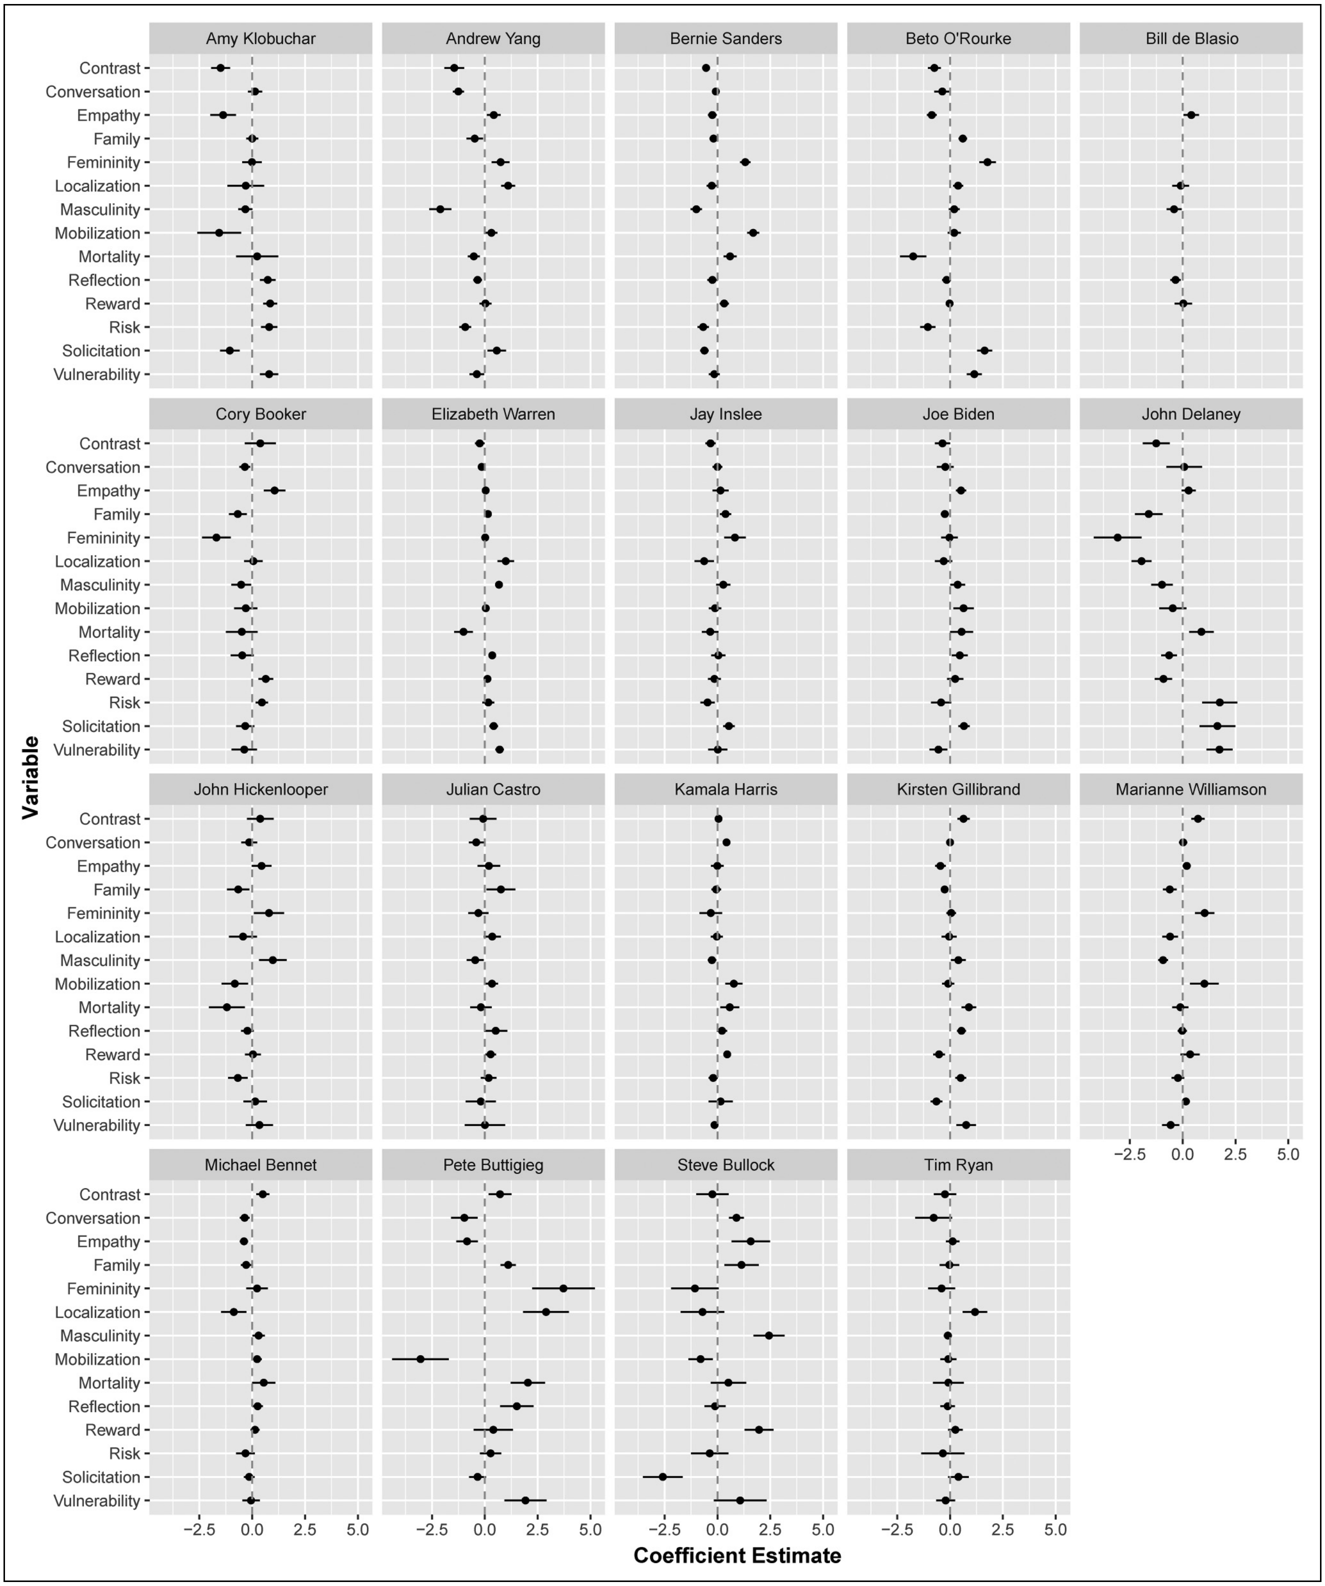

Figure 5. Coefficient Estimates from the Baseline Regression Model.

unable to examine more granular time dynamics in terms of candidates' strategies and/or which textual characteristics are effective.

4. Our coefficient results remain consistent when we include topics, first derived from the latent Dirichlet allocation model, as control variables.

Table 1. Significant Effects for the Four Major Candidates (Baseline Regression Model).

	Positive Significant Coefficients		Negative Significant Coefficients	
Candidate	**Variable**	**Coef. Estimate**	**Variable**	**Coef. Estimate**
Joe Biden	Empathy	.52	Family	−.25
	Masculinity	.36	Vulnerability	−.55
	Mobilization	.64		
	Reflection	.46		
	Solicitation	.66		
Pete Buttigieg	Contrast	.72	Conversation	−.97
	Family	1.11	Empathy	−.84
	Femininity	3.72	Mobilization	−3.04
	Localization	2.90		
	Mortality	2.04		
	Reflection	1.51		
	Vulnerability	1.92		
Bernie Sanders	Femininity	1.32	Contrast	−.54
	Mobilization	1.69	Empathy	−.24
	Mortality	.60	Localization	−.27
	Reward	.32	Masculinity	−1.00
			Risk	−.67
			Solicitation	−.61
Elizabeth Warren	Family	.15	Contrast	−.24
	Localization	.99	Conversation	−.15
	Masculinity	.67	Mortality	−1.01
	Reflection	.35		
	Solicitation	.42		
	Vulnerability	.70		

5. Our coefficient results remain consistent when we estimate a Bayesian hierarchical linear model rather than our standard linear regression. In the Bayesian model, each textual characteristic has a common effect on donation outcomes and there are also candidate-specific deviations from that common effect. We find that the common effects are nearly all statistically insignificant, and the "total effects" of the common effect and the candidate-specific deviations are very consistent with the main results we show previously.
6. We are unable to quantify whether the textual characteristics have a bigger effect on people's decision of whether to donate versus their decision of how much to donate. We examine this issue by estimating two models: one for donation count and one for average donation amount. We find that the standard errors are quite large for most of the coefficients, which means that we are unable to draw strong conclusions about which textual characteristics affect each of these outcomes separately.
7. To demonstrate the practical benefits of our approach, we estimate back-of-the-envelope predictions for candidates' donation outcomes if they were to make small changes to their email communications, either by slightly increasing their usage of textual characteristics that are most effective for them or by slightly decreasing their usage of textual characteristics that are most damaging for them. On average, candidates can increase their campaign donations by about 9% through these relatively minor interventions, but there is substantial heterogeneity across candidates.

General Discussion

American political campaigns are expensive: recent presidential campaigns have spent over a billion dollars on expenditures such as advertising, transportation, consulting fees, and campaign staff salaries. Campaigns rely on campaign donations from individual donors to pay for these expenditures, and the majority of these campaign donations can be attributed to fundraising emails sent by campaigns to their supporters. The 2020 presidential campaigns raised nearly $4 billion in total, so this is a context that is consequential in terms of both its financial size and its societal importance (OpenSecrets 2020).

This article provides a flexible methodological framework that can aid political managers by linking the textual characteristics of campaign emails with the donations that their campaigns subsequently receive. We demonstrate the potential value of this framework by analyzing emails sent by Democratic presidential candidates in 2019. First, we use a dictionary-based method (LIWC) to quantify the textual content of candidates' emails. We then reduce the dimensionality of the LIWC output with PCA, which yields a more manageable set of textual characteristics. Then, by merging the quantified textual characteristics with publicly available donation data, we investigate the relationship between these textual characteristics and subsequent donations using an econometric model.

The application of our framework yields three important takeaways for campaign managers: which textual characteristics are effective for their candidate, which specific words they should be using more or less, and how much their donations would increase if they implemented small improvements to their email copy. To the best of our knowledge, this work is the first to combine automated text analysis with an econometric model to develop a unified framework for analyzing political emails and the campaign contributions that result from them.

In essence, our framework gives political campaigns a roadmap that they can implement when trying to determine how best to construct their campaign emails to maximize donations. One benefit of our approach is that it does not require campaigns to run an expensive field experiment or to collect any additional data beyond what they typically have: the specific wording of their emails and subsequent daily donation amounts. As a result, it provides a convenient, easy-to-implement way for campaigns to increase their donation intake without much investment or disruption.

Besides demonstrating the merits of our framework, our analysis shows that there is not a one-size-fits-all strategy for writing campaign emails. Instead, we find that the effects of textual characteristics differ significantly across candidates; each candidate has their own unique strategies for increasing donations. There are two explanations that likely contribute to this pattern: heterogeneity across candidates, and heterogeneity across audiences. Heterogeneity across candidates means that candidates have different rhetorical strengths and weaknesses, different policy focuses, different prior experiences, and different public images—all of this affects their ability to be compelling when making a particular written appeal. Heterogeneity across audiences means that the people that constitute different candidates' email lists are not the same, and therefore these email recipients may differ in terms of their responsiveness to different kinds of emails. Since we do not observe exactly who receives emails from each candidate, we are unable to separate these two factors or quantify their relative importance. The most likely answer is that both contribute to the final outcome: for instance, the positive effect for Biden's mobilization is likely both because he is particularly effective at using mobilization-related language and also because his audience is particularly receptive to it.

In addition to generating insights related to LIWC variables, our approach enables us to make candidate-specific email communication recommendations by determining which textual characteristics have the greatest positive versus negative effects on donation behavior. Our what-if analyses enable us to determine the "most effective" words for each candidate based on the overlap between their usual email word choice and textual characteristics that are associated with greater donation amounts. Thus, we are able to provide concrete, candidate-specific recommendations on which words to use and which ones to avoid. Our recommended changes are small yet capable of generating a sizable increase in donations (especially for less popular candidates). Importantly, campaigns do not need to make drastic changes to their communication or their platforms and need only increase or decrease the frequency of certain words that they are already using in their emails to boost donations. The advantage of our framework is that it can be applied to any data set to yield immediate insights and increase donations without necessitating wholesale changes to a campaign's existing communication strategy that may be difficult to implement.

Our data also have some limitations that may provide fruitful avenues for future research. First, the relatively short data sample period means that we are unable to measure changes in candidates' usage of different textual characteristics over time or changes in the efficacy of those textual characteristics over time, but these limitations could be addressed by future researchers with access to a longer panel of data. Second, our data set relies on observational data, and therefore we cannot rule out all possible endogeneity concerns, but researchers would be able to avoid these issues if they were able to work hand in hand with a political campaign and run a field experiment. Third, we are unable to model individual-level heterogeneity because our data do not show the individual recipients of each email, but this could also be addressed if researchers were able to receive data directly from the political campaign. Nevertheless, we hope that the methodological approach illustrated in this article provides a flexible framework that can serve both academics and practitioners in the future, as they investigate the relationships between textual characteristics in communication and downstream behaviors such as donations.

Acknowledgments

The authors thank Ashlee Humphreys, Sarah Moore, Grant Packard, Hema Yoganarasimhan, Jonathan Zhang, and participants in the Summer Language Lab seminar for helpful comments and suggestions. The authors also thank Tobin Van Ostern, Brandon English, and Lauren Miller for discussing institutional details regarding presidential campaigns and their email strategy.

Editor

Arvind Rangaswamy

Declaration of Conflicting Interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iDs

Simha Mummalaneni https://orcid.org/0000-0003-3364-2088
Mathew S. Isaac https://orcid.org/0000-0002-3182-1341

References

Alvarez, R. Michael, Jonathan N. Katz, and Seo-young Silvia Kim (2020), "Hidden Donors: The Censoring Problem in U.S. Federal Campaign Finance Data," *Election Law Journal: Rules, Politics, and Policy*, 19 (1), 1–18.

Barber, Michael (2016), "Donation Motivations: Testing Theories of Access and Ideology," *Political Research Quarterly*, 69 (1), 148–59.

Bycoffe, Aaron (2020), "Tracking Every Presidential Candidate's TV Ad Buys," FiveThirtyEight (April 8), https://projects.fivethirtyeight.com/2020-campaign-ads/.

Cattell, Raymond B. (1966), "The Scree Test for the Number of Factors," *Multivariate Behavioral Research*, 1 (2), 245–76.

Doubek, James (2015), "Political Campaigns Go Social, But Email Is Still King," NPR (July 28), https://www.npr.org/sections/itsallpolitics/2015/07/28/426022093/as-political-campaigns-go-digital-and-social-email-is-still-king.

Epstein, Ben and Jeffrey D. Broxmeyer (2020), "The (Surprisingly Interesting) Story of E-Mail in the 2016 Presidential Election," *Journal of Information Technology & Politics*, 17 (3), 232–48.

Evers-Hillstrom, Karl and Camille Erickson (2019), "Your Email Is For Sale — and 2020 Candidates Are Paying Up," OpenSecrets Center for Responsive Politics (June 13), https://www.opensecrets.org/news/2019/06/email-list-for-sale-2020-candidates-are-paying/.

Gallo, Amy (2017), "A Refresher on A/B Testing," Harvard Business Review (June 28), https://hbr.org/2017/06/a-refresher-on-ab-testing.

Gaynor, SoRelle Wyckoff and James G. Gimpel (2021), "Small Donor Contributions in Response to Email Outreach by a Political Campaign," *Journal of Political Marketing*, 23 (1), 51–73.

Goldmacher, Shane (2020), "Pete Buttigieg's Dash for Cash: 10 Fund-Raisers in Two Weeks," *New York Times* (February 15), https://www.nytimes.com/2020/02/15/us/politics/pete-buttigieg-fundraising.html.

Green, Joshua (2012), "The Science Behind Those Obama Campaign E-Mails," *Bloomberg Businessweek* (November 29), https://www.bloomberg.com/news/articles/2012-11-29/the-science-behind-those-obama-campaign-e-mails.

Han, Hahrie C. (2009), "Does the Content of Political Appeals Matter in Motivating Participation? A Field Experiment on Self-Disclosure in Political Appeals," *Political Behavior*, 31 (1), 103–16.

Hassell, Hans J.G. (2011), "Looking Beyond the Voting Constituency: A Study of Campaign Donation Solicitations in the 2008 Presidential Primary and General Election," *Journal of Political Marketing*, 10 (1–2), 27–42.

Heerwig, Jennifer A. (2016), "Donations and Dependence: Individual Contributor Strategies in House Elections," *Social Science Research*, 60, 181–98.

Heng, Yu Tse, David T. Wagner, Christopher M. Barnes, and Cristiano L. Guarana (2018), "Archival Research: Expanding the Methodological Toolkit in Social Psychology," *Journal of Experimental Social Psychology*, 78, 14–22.

Hersh, Eitan D. (2018), Responses to Questions for the Record: Cambridge Analytica and the Future of Data Privacy, U.S. Senate Committee on the Judiciary, https://www.judiciary.senate.gov/imo/media/doc/Hersh%20Responses%20to%20QFRs.pdf.

Hersh, Eitan D. and Brian F. Schaffner (2013), "Targeted Campaign Appeals and the Value of Ambiguity," *Journal of Politics*, 75 (2), 520–34.

Higgins, Tucker (2019), "'Pete's Wine Cave': Why Buttigieg's Rivals Are Slamming His Napa Fundraiser," CNBC (December 20), https://www.cnbc.com/2019/12/20/pete-buttigiegs-wine-cave-what-to-know-about-the-napa-fundraiser.html.

Hughes, Adam (2017), "5 Facts About U.S. Political Donations," Pew Research Center (May 17), http://pewrsr.ch/2qvvktm.

Larson, Jeff (2012a), "Buying Your Vote: Explore Hundreds of Campaign Emails in the Message Machine," ProPublica (July 17), https://www.propublica.org/article/explore-hundreds-of-campaign-emails-in-the-message-machine.

Larson, Jeff (2012b), "Buying Your Vote: Message Machine Starts Providing Answers," ProPublica (October 18), https://www.propublica.org/article/message-machine-starts-providing-answers.

Madrigal, Alexis C. (2012), "Hey, I Need to Talk to You About This Brilliant Obama Email Scheme," *The Atlantic* (November 29), https://www.theatlantic.com/technology/archive/2012/11/hey-i-need-to-talk-to-you-about-this-brilliant-obama-email-scheme/265725/.

Mehta, Dhrumil (2019), "What Our Inbox Tells Us About How Democrats Are Tackling Trump," FiveThirtyEight (July 11), https://fivethirtyeight.com/features/which-democrats-are-campaigning-on-trump/.

Moth, David (2013), "Seven Lessons Obama's Digital Team Learned from A/B Testing Emails," Econsultancy (October 28), https://econsultancy.com/seven-lessons-obama-s-digital-team-learned-from-a-b-testing-emails/.

Neslin, Scott A., Gail Ayala Taylor, Kimberly D. Grantham, and Kimberly R. McNeil (2013), "Overcoming the 'Recency Trap' in Customer Relationship Management," *Journal of the Academy of Marketing Science*, 41 (3), 320–37.

Nickerson, David W. (2007), "Does Email Boost Turnout?" *Quarterly Journal of Political Science*, 2 (4), 369–79.

OpenSecrets (2020), "2020 Presidential Race: Campaign Funding," OpenSecrets Center for Responsible Politics (October 23; accessed November 27, 2020), https://www.opensecrets.org/2020-presidential-race.

OpenSecrets (2021), "Online Ad Spending for Joe Biden," OpenSecrets Center for Responsible Politics (accessed December 29, 2021), https://www.opensecrets.org/2020-presidential-race/joe-biden/online-ad-spending?id=N00001669.

Pennebaker, James W., Martha E. Francis, and Roger J. Booth (2001), *Linguistic Inquiry and Word Count: LIWC 2001*. Lawrence Erlbaum Associates.

Rossi, Peter E. (2014), "Even the Rich Can Make Themselves Poor: A Critical Examination of IV Methods in Marketing Applications," *Marketing Science*, 33 (5), 655–72.

Sabato, Larry J. (1981), *The Rise of Political Consultants: New Ways of Winning Elections*. Basic Books.

Shapiro, Bradley T., Günter J. Hitsch, and Anna E. Tuchman (2021), "TV Advertising Effectiveness and Profitability: Generalizable Results from 288 Brands," *Econometrica*, 89 (4), 1855–79.

Shea, Daniel M. and Michael J. Burton (2006), *Campaign Craft: The Strategies, Tactics, and Art of Political Campaign Management*. Praeger.

Sutton, Adam (2013), "Email Testing: How the Obama Campaign Generated Approximately $500 Million in Donations from Email Marketing," Marketing Sherpa (May 7), https://www.marketingsherpa.com/article/case-study/obama-email-campaign-testing.

Tausczik, Yla R. and James W. Pennebaker (2010), "The Psychological Meaning of Words: LIWC and Computerized Text Analysis Methods," *Journal of Language and Social Psychology*, 29 (1), 24–54.

Vaccari, Cristian (2008), "From the Air to the Ground: The Internet in the 2004 Presidential Campaign," *New Media & Society*, 10 (4), 647–65.

Verba, Sidney, Kay Lehman Schlozman, and Henry E. Brady (1995), *Voice and Equality: Civic Voluntarism in American Politics*. Harvard University Press.

Viveiros, Beth Negus (2019), "Email Personalization Boosts Bernie 2020 Donations," Chief Marketer (October 18), https://www.chiefmarketer.com/email-personalization-boosts-bernie-2020-donations/.

Williams, Andrew Paul and Kaye D. Trammell (2005), "Candidate Campaign E-Mail Messages in the Presidential Election 2004," *American Behavioral Scientist*, 49 (4), 560–74.