Campaigns in control: Analyzing controlled interactivity and message discipline on Facebook

Deen Freelon

School of Communication, American University

dfreelon@gmail.com

Abstract: Most American political campaigns use social media as one component of a broader communication strategy. Campaign use of social media is typically governed by *controlled interactivity*, a philosophy that attempts to leverage citizens' online behavior toward the goal of electing the candidate. One key outcome of controlled interactivity is high levels of *message discipline*, the degree of correspondence between the campaign's and its audience's political speech. This study quantifies message discipline as it flows through two highly-visible controlled-interactive spaces—Barack Obama's and Mitt Romney's respective official campaign Facebook pages—during the 2012 US presidential campaign. The results of a lexicon analysis indicate that Romney's campaign controlled its audience's interactivity more effectively than the Obama campaign and that both audiences departed from message discipline most sharply on the issues of civil rights and religion.

Campaigns in control: Analyzing controlled interactivity and message discipline on Facebook

Political candidates have gradually come to embrace interactive digital media as integral elements of their election campaigns. Their ultimate goal—getting elected—is the same as it has always been, but campaigns now have many new channels through which to engage the electorate. In some ways, the rise of social media in campaigns mirrors that of previous "new media" such as print media, radio, and television in the 20th century. But unlike analog media, digital media allow citizens to play active roles in sharing, modifying, and commenting on campaign content.

Nearly every Democratic and Republican candidate for national office has one or more social media pages dedicated to their campaign (Gulati & Williams, 2013). Social media have been praised for enabling citizens to engage in "mass self-communication" on a grand scale (Castells, 2007), but politicians typically care little for such undirected conversation. To further the ultimate end of their campaigns, most politicians use social media mostly in top-down ways, distributing information to voters and the mass media while circumventing the latter's tendency to alter their messages. Campaigns also typically ask likely voters to take various actions on their behalf such as buying merchandise, contributing money, or simply spreading favorable messages. In this way, they attempt to control the interactivity of their social media audiences to facilitate the candidate's victory (Stromer-Galley, 2014).

A number of studies have analyzed campaign use of digital media from the perspective of *controlled interactivity*. These have mostly been explorations of the features of the technologies and the motivations of the campaign staff who created and implemented them. The current study analyzes controlled interactivity directly by quantifying flows of information from campaigns to their digital audiences. It incorporates the concept of *message discipline* as an important indicator of the success of controlled interactivity: the more the audience stays on-message, the greater the degree of control. It

uses Barack Obama and Mitt Romney's official campaign pages during the 2012 election as cases to

demonstrate the extent to which message discipline manifests in spaces of controlled interactivity.

Controlled interactivity

The concept of controlled interactivity (Stromer-Galley, 2014) is a useful one for examining

presidential candidates' and their constituents' uses of social media. It assumes that a candidate's sole

goal is to get elected (Doherty, 2012; Mayhew, 1974), not to educate or seek input from the public. To

whatever extent this conceptual simplification holds true across all candidate activities, it is undoubtedly

the case for campaigns. Thus, a communication technology's value to a campaign is measured solely in

terms of its ability to facilitate election/reelection. In practical terms, this means that campaigns will

typically use social media to enable their bases to support their candidates in preapproved ways

(Stromer-Galley, 2014, p. 12).

This represents something of a disappointment from the standpoint of deliberative democratic

theory, which considers such online environments "weak spaces" due to the absence of any meaningful

connection to decision-making processes (Janssen & Kies, 2005). In such spaces, control lies

overwhelmingly, but not completely, with the campaign's decision-makers. This manifests in technical

and content limitations on how users can express themselves. But although interactivity in campaign

spaces may be controlled to some extent, what differentiates them from top-down media such as

television is users' capacity to innovate and communicate within those constraints. The messages they

circulate to friends and family online may have greater impact than campaign ads due to the personal

connections between the senders and receivers (Allsop, Bassett, & Hoskins, 2007; Brown, Broderick, &

Lee, 2007). Thus, users in such spaces have some agency despite their general lack of control.

Indeed, as Stromer-Galley notes, "the challenge for campaigns is in determining how best to

engage citizens to help the campaign win, while not getting drowned out or sidetracked by those

citizens" (Stromer-Galley, 2014, p. 14). Campaigns aim to get their candidate elected, but citizens have many different goals for their online participation, so an effective controlled-interactive space needs to balance these two interests. Campaigns may choose to do so in a variety of ways: for example, they can manage the types of content users can post (e.g. event planning, donation, offline group meetups), as the MyBO platform did for the Obama campaign in 2008 (Kreiss, 2012). They can discourage users from leaving the space by avoiding external hyperlinks, as was common in the early 2000s (Schneider & Foot, 2006). They can cut off potential communication pathways between citizens and candidates, as do weak digital spaces and email action alerts (Karpf, 2010). And they can control the visibility of citizen contributions by creating discussion platforms in which only the campaign can initiate conversations, as in the Facebook pages examined in this study.

Facebook may be the most widely-used controlled-interactivity platform in American politics today. As of 2012, the most recent year for which data are available, the president and 87.2% of members of Congress operated official Facebook pages (Greenberg, 2012; see also Gulati & Williams, 2013). (That percentage is likely higher today.) Compared to Twitter, which also hosts nearly all elected members of the legislative and executive branches, Facebook has a much larger user base and offers a more configurable interactive environment. Page creators can, for example, decide whether they want to allow page followers to post directly to it, or only reply to official posts. The latter is how both the Barack Obama and Mitt Romney campaigns configured their official Facebook pages in 2012, and although an extensive literature review search found no studies of how members of Congress configure their Facebook pages, it would be surprising if many allowed non-campaign users to post free-standing messages. This would amount to "uncontrolled interactivity" and would likely backfire very quickly on any politician who attempted it.

One of this paper's contributions is the idea that different types of interactive campaign environments and communication tools enact different degrees and genres of control. Facebook pages

are relatively lightly controlled compared to heavily managed tools like MyBO and email alerts, which permit only a narrow range of user behaviors. Although citizens can only reply to official posts on campaign Facebook pages, they are free to speak their minds within those confines and to address their messages to whomever they wish, though it seems unlikely that many such messages ever reach the candidates. However, citizen replies can be seen by, and thus may inform the political thinking of, the poster's Facebook friends and others browsing the page.

Existing research on controlled interactivity and related concepts largely focuses on the architecture and implementation of digital campaign communication tools as opposed to the interactive processes that occur within them. The earliest of this work examined web sites and email as campaigns slowly explored how the web could help them elect their candidates. Some studies sought theoretical ways to capture the technological features of which campaigns made most extensive use (Gulati, 2004; Lawson-Borders & Kirk, 2005; Schneider & Foot, 2006). Others looked directly at the handling of user data and other digital practices of campaign insiders (Howard, 2005; Kreiss, 2012, 2014). More recent studies of controlled interactivity have tended to focus on how the features of the tools themselves enable and constrain various forms of citizen participation (Lilleker, 2016; Schweitzer, 2011; Stromer-Galley, 2014). Few if any have systematically analyzed the substance of controlled interactivity to ascertain just how much control campaigns are able to exercise.

Research on the relationships between candidate and citizen communications also informs this study. However, they are seldom analyzed through the lens of controlled interactivity. Many such studies are exploratory, asking research questions that are variations of "how did candidates and/or citizens use this social media platform(s)?" (Bruns & Highfield, 2013; Christensen, 2013; Conway, Kenski, & Wang, 2013; Goodnow, 2013; Graham, Broersma, Hazelhoff, & van 't Haar, 2013; Skovsgaard & Van Dalen, 2013). More theoretical approaches have applied concepts such as the public sphere (Robertson,

Vatrapu, & Medina, 2010), diffusion of innovations (Gulati & Williams, 2013), and hypotheses drawn

from the online campaigning literature (Vaccari & Nielsen, 2013).

Message discipline: The goal of control

One of the most important outcomes for political campaigns is control of "the message"—the

thematically unified collection of issues, frames, talking points, concepts, and images that define the

candidate. This is why staying "on-message" has been a key goal for campaigns since long before the

dawn of the digital age (Benoit et al., 2011; Needham, 2005; Norris, Curtice, Sanders, Scammell, &

Semetko, 1999). In contexts of controlled interactivity generally, and in lightly-controlled conversation

platforms specifically, control prevails to the extent that citizens maintain message discipline. Message

discipline is defined here as a close correspondence between the political issues candidates and their

supporters discuss and how they discuss them. The process of promoting message discipline is similar

to, but distinct from, two well-known communication theories: the two-step flow and agenda-setting.

The two-step flow emphasizes the transmission of ideas from political elites and the media to citizens

through opinion leaders, and has found new relevance in the 21[st] century. However, empirical studies

thereof rarely examine how ideas change as they flow (cf. Choi, 2015; Lotan et al., 2011; Wu, Hofman,

Mason, & Watts, 2011), even though this is a key feature of both the original theory (Bennett &

Manheim, 2006) and the dynamics of information transmission online (Leskovec, Backstrom, &

Kleinberg, 2009). Agenda-setting research concerned with the degree of correspondence between the

salience of items on elite issue agendas and on citizen agendas (McCombs, 2005). Yet it is prohibitively

difficult to obtain a representative participant sample from a large online forum to ask about how they

assign importance to various issues.

Message discipline as enacted by a campaign's audience is the end state of a communication

flow, to be sure. But what matters is less the number of steps in it or the full scope of

citizens'/candidates' agendas and more the extent to which information mutates as it moves from the top down. A campaign's message will inevitably change as it moves through controlled spaces, but these changes cannot be blamed on low copying fidelity, as in the children's game "Telephone." Rather, they are more likely introduced by participants with political priorities of their own, which may or may not dovetail with the campaign's. This tension between what campaigns want participants to discuss and what the latter actually discuss lies at the heart of the current analysis. To the extent that the two match (i.e., that participants remain on-message), the campaign can claim to have successfully controlled participants' interactivity. This is an important step to "empower supporters to persuade others in their own social network through word of mouth online or offline"(Stromer-Galley, 2014, p. 15). On the other hand, low levels of message discipline would indicate that citizens had introduced major changes to the message: in other words, a loss of control. Depending on the nature of the changes, such a result might call into question the value of lightly-controlled spaces for the single-minded reelection seeker (though conventional wisdom and peer pressure might prevent candidates from abandoning them entirely).

Existing research offers several suggestions as to exactly how on-message we might expect participants in lightly-controlled spaces to be, although not quite as many as we might want. The current data allow comparisons between Barack Obama and Mitt Romney's Facebook pages, but prior research offers little basis for expecting either's participants to remain more on-message than the other's. Studies of Usenet, a pre-Web forum whose technical features are similar to that of campaign-controlled Facebook pages, suggest that threaded conversations can veer off-topic relatively easily (Herring, 1999; Hewitt, 2001; Lambiase, 2010). Presumably participants in controlled spaces like Facebook pages should stay somewhat more on-message, since their expressive latitude is less than in most Usenet newsgroups. Still, quantitative evidence either way is difficult to come by, with the notable exception of Hawthorne, Houston, and McKinney (2013).

Hawthorne et al. examined differences between candidate and citizen talk on Twitter during one of the 2012 GOP primary debates, in effect measuring message discipline. Using a lexicon-based approach similar to that applied here, they found no substantial differences between issues mentioned by political elites (defined as journalists, pundits, and elected officials) and nonelites. In other words, message discipline was extremely high. But the current context differs in several ways from Twitter during a debate: first, Facebook's audience is not only much larger than Twitter's, it differs in terms of its users' interest in and general attitude toward political content. For example, Facebook users are more likely to comment on political content than Twitter users, who are more likely to use that platform to follow breaking news as it occurs (Shearer, 2015). Second, the more than six months of data examined here may yield different results than data from a single evening. Third, Twitter hashtags are not controlled-interactive environments: there is nothing a campaign can do from a technical standpoint to control how participants use them.

Given the lack of clear predictions from the literature, this study's empirical work is driven by three related research questions:

- RQ1: How will Obama's and Romney's respective Facebook audiences differ in the extent to which they remain on-message?

- RQ2: When the candidates' audiences veer off-message, which political issues will they emphasize disproportionately more and less often?

- RQ3: When the candidates and their audiences discuss the same issues, how similarly will they discuss them?

Data and methods

Before describing this study's methods, it is important to explain how message discipline will be operationalized. Here the concept is divided into two components: 1) the broad political *issues*

discussed by the campaigns and their followers, and 2) the specific *keywords* that indicate the presence

of each issue. A high degree of issue correspondence between campaigns and participants suggests that

the latter are discussing the issues the former would prefer them to discuss. A high degree of keyword

correspondence within a specific issue indicates that participants are using the same words as the

campaign on that issue, a key aspect of message discipline. While this issue-based approach may not

capture all possible aspects of message discipline—in particular, it is likely to miss complex concepts that

span more than a few words—it offers a general, replicable impression of how similarly campaigns and

their followers discuss the issues. It is also consistent with previous studies of message discipline which

employ traditional content analysis (Benoit et al., 2011; Norris et al., 1999).

In a departure from these earlier studies, the main method used here is lexicon analysis, which

detects the presence of different topics in a corpus of texts by searching for keywords associated with

each topic. In this case, the topics are political issues and the keywords within each topic are all related

to that issue. For example, the keywords grouped under the topic "finance, economy, and labor" include

"job," "employ" (which captures related terms like "employment" and "employer") and "tax." This

method is an apt fit for measuring message discipline between multiple corpuses of text in large-scale

social media conversations: the datasets are too large for traditional content analysis and unsupervised

text classification often misses important nuances (Klebanov, Diermeier, & Beigman, 2008). Lexicon

analysis has been used effectively to categorize political speech in both online and offline contexts

(Hawthorne et al., 2013; Pennebaker, Slatcher, & Chung, 2005; Simon & Jerit, 2007).

The data analyzed in this study were collected from Facebook. They include all messages posted

to Barack Obama's and Mitt Romney's official Facebook pages between April 25, 2012 (the day the

Republican National Committee officially endorsed Romney) and November 5, 2012 (the day before the

election).[1] During this period the Obama campaign posted 268 messages while Romney's posted 584.

Although each campaign allowed only official messages to be posted to its Facebook wall, visitors could

append comments to these messages. All publicly-available comments appended to the official

candidate messages during the study period were collected on November 10, 2012 from the Facebook

Graph API using a custom PHP script developed by the author.[2] In total, 233,129 Obama comments and

624,326 Romney comments were collected. The combined N of posts (official messages and comments)

analyzed in this study is 858,307.

Some previous lexicon-based political communication research has used prepackaged lexicons

(Pennebaker et al., 2005; Simon & Jerit, 2007) or short, ad-hoc keyword lists (Neuman, Guggenheim,

Jang, & Bae, 2014), but such techniques are less than ideal. Because social media users may use a wide

array of slang, neologisms, and other distinctive language, fully-automated or "canned" approaches may

miss much of theoretical interest (Petchler & Gonzalez-Bailon, 2015). Therefore, this study's lexicon

analysis used as a starting point the Lexicoder Policy Agendas lexicon (LPA), an open-access list of 1,402

public policy-related terms grouped into 28 issue categories (Albaugh, Sevenans, & Soroka, 2013;

Albaugh, Sevenans, Soroka, & Loewen, 2013). Although the LPA contained many terms relevant to the

present case, a cursory comparison between it and the data revealed that many essential terms were

lacking. To remedy this, the author and a research assistant each read through 1,004 Facebook posts—

306 candidate messages and 698 user comments—in search of additional terms to add to the LPA.

These posts had been coded as containing issue-relevant content by other coders as part of an earlier,

unpublished project. Those coders analyzed all official candidate messages and 3,080 randomly-sampled

comments from the same larger dataset used in this study. Thus, the subset of messages re-examined

for issue terms in this study represents all those from the previous project's random sample that were

judged to mention one or more political issues.

Most of the new lexicon terms culled from the Facebook posts were added to the most

appropriate issue category. Terms that did not unambiguously denote a single category were discarded

(including a few from the original LPA), and terms sharing a single word stem were merged (e.g.

"taxpayer" and "taxes" were combined as "tax"). In addition, several of the original LPA categories were merged for the sake of parsimony, resulting in a final total of 1,476 terms across 18 categories (see the appendix for a complete list of categories and terms).[3] Counting all new term additions, mergers, truncations, and spelling corrections, a net total of 305 terms were added to the original LPA.

Once the lexicon categories were finalized, the data were preprocessed for analysis. Rigorous data preprocessing is extremely important in lexicon analysis, as its absence may result in unacceptable rates of Type I and/or II error (Petchler & Gonzalez-Bailon, 2015). First, all the multi-word lexicon terms were concatenated with underscores in the data so that they would not be broken apart by the tokenization process. Next, each post was tokenized, or separated into its component terms and stripped of punctuation (excluding underscores), leaving all relevant multi-word terms remaining intact. Most of the tokenized terms were then lemmatized (reduced to their canonical or dictionary-entry forms) using the NLTK package for the Python programming language (e.g. instances of "taxed" were changed to "tax") (McEnery, Xiao, & Tono, 2006; Perkins, 2010). A few terms were not lemmatized because their non-canonical forms or tenses had distinct political meanings. A representative example is the term "illegals": this pluralized form is often used by conservatives to refer to undocumented immigrants, but the lemmatized singular "illegal" does not share this meaning. In addition to the lemmatization process, certain lexicon terms were designated to be stemmed, or have their affixes such as "ing" or "pre" removed (Perkins, 2010; Stock & Stock, 2013). Stemming creates text strings (stems) that may or may not be actual words (such as "racis" which matched "racist," "racists," and "racism"). This list of words to stem was created manually because existing general-purpose stemmers did not consistently produce the appropriate stems.

After the data were preprocessed in the above manner, the lexicon analysis began. Upon completion, it yielded a set of overall counts for each issue category and term. That is, the lexicon output contains the number of posts in which each candidate and his commenters 1) mentioned each

individual term and 2) mentioned at least one term from each issue category. These counts form the

basis of the empirical analysis that follows.

Results

Before we address the research questions, let us first consider some of the more relevant descriptive

statistics. In total, 81 of Obama's 268 posts (30.2%) contained at least one issue term while 292 (50.0%)

of Romney's 584 did. Similarly, Obama's commenters discussed policy issues proportionally less often

(80,795 posts of 233,129; 34.6%) than Romney's did (266,335 posts of 624,326; 42.7%). Thus both

Romney and his followers were more focused on policy than their Democratic counterparts, at least on

Facebook. Of course, many posts mentioned more than one issue.

Figure 1 plots the number of times each issue's terms appear on the two candidates' pages (in

official posts and comments) against the number of terms in each issue. It shows that although the issue

categories contained differing numbers of terms, those numbers were only weakly correlated with issue

prevalence in the data. In fact, if the extreme outlier of FEL (finance/economics/labor) is removed, the

correlations for both Obama and Romney are non-significant.

[Figure 1 here]

Figures 2a and 2b reveal each subset's relative issue emphasis, with each chart's y-axis

indicating the proportion of all issue category mentions. Figure 1a, for example, shows that Obama's

two most-mentioned issues were civil rights and FEL, followed by healthcare, with no other issue

exceeding 10% of all issue mentions. In Figure 1b we see that Romney's messages were dominated by

economic concerns and that the next most popular issue, healthcare, accounted for only about 7% of

issue mentions. Unsurprisingly given their much larger post volumes, the commenters' issue mentions

were more evenly distributed than the candidates', although this was more the case for Obama's

commenters than Romney's. Both were most interested in economics with religion and civil rights

following close behind, as the remaining issues gently tapered off.

[Figures 2a and 2b here]

The first research question asks about how closely each candidate's Facebook audience

maintained message discipline. I explore this question in two ways. First, each group of commenters'

issue proportions were entered as a dependent variable into two OLS regressions: one using only the

corresponding candidates' issue proportions as the sole predictor, and the other adding the remaining

two issue proportion subsets as predictors. This hierarchical approach includes the focal test of each

candidate's agenda-setting power over his audience in two separate regressions to lower the likelihood

that the results are spurious. The unit of analysis in all four regressions is the issue category (n = 18).

[Tables 1a and 2a here]

Table 1a contains the output of the regressions of Obama's commenters, while Table 1b

contains those for Romney's commenters. Each table's Model 1 includes the candidate's issue

proportions as the sole predictor, while Model 2 adds the remaining two subsets as controls. For the

current purposes the standardized regression coefficients are the appropriate quantities to compare

between models and candidates; unstandardized coefficients are included for the sake of completeness.

In answer to RQ1, comparisons between Obama's Model 1 and 2 betas and the corresponding

coefficients in Romney's models reveal that the latter's audience was more consistently on-message.

Interestingly, the addition of the other subsets as predictors substantially increases the overall

predictive power of Model 2 in both cases. The Model 2s also show that the two audiences' agendas

were more similar to each other than either was to its corresponding candidate, although this

phenomenon is more pronounced in Obama's case than in Romney's.

This regression approach provides suggestive evidence in favor of the Romney campaign's ability to control its followers' interactivity, but its value is somewhat limited by its low statistical power as well as the fact that it does not directly analyze message discipline within comment threads. Even though candidates and commenters may discuss issues in similar proportions overall, they may fall out of sync within specific message threads. To investigate this possibility directly, for all posts in which each candidate mentioned at least one issue, the median proportion of comments in which his commenters mentioned that issue was calculated. For posts in which a candidate mentioned multiple issues, commenter mentions of any of those issues were included in the proportion's numerator. This resulted in two median proportions, one for Obama and one for Romney, higher magnitudes of which indicated greater message discipline on a per-thread basis.

This analysis corroborates the regressions above: the median percentage of comments mentioning the issue(s) Obama mentioned in those posts where he mentioned at least one issue was 20.7%, while the corresponding median for Romney was 29.5%. The results indicate that Romney's ability to control message discipline is nearly 50% greater than Obama's, though in neither case was the proportion anywhere close to a majority. The convergent evidence from these two distinct analytical procedures permits a much stronger conclusion in Romney's favor than either would have by itself.

[Figures 3a and 3b here]

While the regression analysis allows us to compare the two candidates' audiences in terms of message discipline, it is also important to understand the specific issues the candidates and their audiences discuss disproportionately more and less often (RQ2). To do so, the corresponding issue proportions between each candidate-audience pair were subtracted and graphed (Figures 3a and 3b). Both figures clearly show that the gaps in issue interest between each pair of subsets is heavily

concentrated among a few issues: for the vast majority, the disparity falls below five percentage points. But the exceptions are worth exploring further.

Obama talks about his most-emphasized issue, civil rights, proportionally much more than his audience does. As the RQ3 analysis shows, this is primarily due to his strong interest in women's issues. His commenters, on the other hand, are disproportionately more interested in matters of religion, which the president mentions only once in all of his 268 posts. While Romney discusses religion somewhat more than Obama, his commenters care about it even more than he does. And although both the governor and his audience mention economic terms more than any other category, the former's interest exceeds the latter's by 22 percentage points.

Message discipline is not only about which issues are discussed—it also encompasses how they are discussed (RQ3). Therefore, this study's final analysis takes a look at the extent to which the candidates and their respective audiences used specific issue keywords differentially. It does so by computing keyword usage differences within issue categories between the subsets. Because space limitations prohibit an exhaustive exploration of all issues, this section will examine the four most-frequently-mentioned issues across all four subsets: FEL, civil rights, healthcare, and religion. For each of these issues, Table 2 presents each subset's five most-used keywords. To the right of each keyword in parentheses is the percentage of total posts from each campaign or audience in which it appeared.

[Table 2 here]

Perhaps the most readily apparent pattern in this table is the strong similarity of terms within most issue categories. This is probably clearest for FEL and healthcare, within which many of the same keywords are repeated across subsets ("job," "tax," "obamacare," "medicare") in slightly differing proportions. Those terms that are not repeated do not seem to indicate a consistent focus on fundamentally different issue attributes compared to those that are.

Civil rights represents a partial departure from this pattern. While all four subsets are substantially focused on women's issues—a clear indication of message discipline—the two candidates are in near-total lexical lockstep with one another, with most terms appearing either verbatim or as close near-synonyms on both sides. But while the commenters share this concern, they also both veer off-message with discussions of abortion and racism, two topics neither candidate mentions at all. The current methods do not permit any claims about which sides of each are most popular (e.g. the pro- vs. the anti-abortion side), but it is interesting to see the commenters raising issues the candidates would prefer to avoid.

Nowhere is this departure from message discipline more evident than on the issue of religion. Obama makes only one passing reference to it, in an exhortation to his supporters to "keep the faith," which at best carries only a vague religious connotation. Romney mentions religion more often, but usually with nonspecific terms like "prayer" and "god," with the exception of "catholic" which he uses only once. Meanwhile, the commenters do not hesitate to discuss the specifics of religions that played controversial roles in the campaign, especially Islam. Romney's religion of Mormonism was a popular topic among his commenters, albeit less so than Islam. Overall, these differences in how and how much audiences and candidates discuss religion speak volumes about both campaigns' general lack of control on this particular issue.

Discussion

This study offers a novel conceptualization of controlled interactivity and also specifies empirical criteria that can be used to infer its efficacy. To wit, participants remaining on-message in an online, candidate-sponsored discussion space is evidence that the space's interactivity was effectively controlled. This study's results reveal a high degree of control in both Obama and Romney's official campaign Facebook pages, which is consistent with prior research (Hawthorne et al., 2013). But there were substantial

differences between the campaigns, with Romney's audience being much more on-message than Obama's. Also, despite the general trend toward message discipline, the audiences emphasized certain issues disproportionately more than the candidates, including religion and the economy. Similarly, the audiences tended to follow the campaigns' lead in the keywords they used to discuss the top issues, but deviated most on the issues of civil rights and religion.

This analysis provides an important complement to studies of the structure and features of controlled-interactive spaces (e.g. Lilleker, 2016; Schneider & Foot, 2006; Schweitzer, 2011; Stromer-Galley, 2014) by focusing on the extent to which they actually serve their intended purposes. From the perspective of both the theory and the campaigns, in this case controlled activity seems to be doing its job for the most part. The campaigns are discussing certain issues in certain ways, and the audiences are generally following their leads. But one of this study's main findings is that some issues are more difficult to control than others. In particular, it seems that candidates' ability to control the form and content of discussions on some social issues (civil rights and religion in particular) may be less than for other issues. Strategic vagueness, a time-honored strategy for avoiding key constituencies (Carey, 1997), is no defense in controlled interactive spaces: when citizens deem it appropriate, they do not hesitate to go off-message.

The extent to which these results will recur in subsequent studies is not clear. One possible conclusion is that they represent fundamental differences between incumbents and challengers, or between Republicans and Democrats. However, studies of campaign communication over multiple elections have yielded results that do not map consistently onto those boundaries (Benoit, 2007; Kaid & Johnston, 2001). Other factors that may be relevant in determining differences between presidential candidates in terms of message control include economic conditions, the most prominent issues that election year, recent external shocks like major terrorist attacks and natural disasters, and each candidate's popularity relative to the other. Still, the conceptualization of message discipline offered

here allows future researchers to explore the outcomes of controlled interactivity in future elections. The importance of such networked environments for campaigns only seems likely to increase as digital media continue to permeate our society.

While Facebook is an important case of controlled interactivity to study due to its near-universal uptake among American political candidates, it represents only one among a variety of types and degrees of such control. Compared to other environments, the campaigns chose to control their pages fairly lightly, largely taking a hands-off approach to the conversation (possible moderation for extremely offensive content notwithstanding). Other interactive tools with different intended purposes may require varying degrees of control. For example, some email campaigns enable only very specific actions such as donations, while others allow participants to modify boilerplate messages to send to legislators (Karpf, 2010). MyBO's group organizing tool from 2008 was also more controlled than Facebook in 2012, as its features are oriented specifically toward that task (Kreiss, 2012). Future research could seek ways to measure message discipline as it emanates outward from spaces like these.

Questions remain about the possible effects of controlled interactivity and message discipline. As noted earlier, the possibility that political messages may be especially effective coming from friends and family makes controlled-interactive spaces attractive for campaigns (Stromer-Galley, 2014). Research on offline political conversation has found persuasion effects (Huckfeldt, Johnson, & Sprague, 2004; Sinclair, 2012), and on the basis of this work it seems logical that digital messages could produce similar effects. Yet this prospect has received little if any attention in digital contexts, and until it does, the complete value of controlled interactivity for campaigns will remain unknown. If citizen-created social media messages are eventually shown to exercise minimal or no persuasive effects, campaigns might lose interest in attempting to control them. Linking specific messages to opinion changes represents a major methodological challenge, however, and so campaigns may continue to be guided by their intuitions in this matter for the foreseeable future.

In addition to the above theoretical contributions, this study also makes a small methodological contribution. With Hawthorne et al. (2013), it recommends against using only prefabricated dictionaries in lexicon analysis. This is particularly important for social media, where linguistic innovation moves very quickly (Eisenstein, O'Connor, Smith, & Xing, 2014) and static dictionaries will miss more and more keywords of interest over time. This study offers a systematic method of adding keywords to existing dictionaries from the text of the dataset being analyzed, reducing the likelihood of omitting relevant content. It also exemplifies one of the rarely-noted advantages of lexicon classification over machine learning: the former ensures that keywords of direct theoretical interest are included as indicators of the category labels, while the former does not. Machine learning algorithms rely on the statistical distribution of N-grams (single- or multi-word text strings) throughout a dataset to classify its elements, which means they may omit or deprecate certain N-grams simply because of how they are distributed. Lexicon methods include keywords based on their theoretical importance, ensuring that they can play an appropriate role in allocating category labels. (Of course, a key weakness is that they usually exclude terms that are statistically associated with a category label but semantically unrelated to it.)

Like all studies, this one has its limitations, some of which have already been mentioned above. Aside from those, the extent to which these findings might generalize beyond the American executive branch is not clear. There may be different message discipline dynamics at work in Congressional campaigns, and future research should address this possibility directly. But perhaps this study's most consequential limitation lies in its operational definition of message discipline. Political issues and the keywords used to discuss them certainly constitute an important component of the concept, but they do not exhaust it. The methods used here almost certainly excluded an unknown number of false negatives, including complex images, concepts, and perspectives that are difficult to reduce to specific words and phrases. Because these elements are difficult to capture at scale, followup work could take a mixed-

methods approach to analyze one or a small number of particularly important messages as they travel from campaigns through citizen networks.

Nevertheless, this study has yielded interesting findings about the consequences of controlled interactivity. True to scholarly expectations (and the hopes of campaigns), online participants usually stay on-message. But the instances where they deviate are instructive, and demonstrate that campaigns' control over their audiences' interactivity is not absolute. The relationships between the controllers and the controlled will no doubt continue to evolve as political campaigns do, and researchers should build on the work presented here as they study them.

References

Albaugh, Q., Sevenans, J., & Soroka, S. (2013). Lexicoder Topic Dictionaries. McGill University. Retrieved

from http://www.lexicoder.com/download.html

Albaugh, Q., Sevenans, J., Soroka, S., & Loewen, P. J. (2013). The Automated Coding of Policy Agendas: A

Dictionary-Based Approach. Presented at the 6th annual Comparative Agendas Project (CAP)

conference, Antwerp.

Allsop, D. T., Bassett, B. R., & Hoskins, J. A. (2007). Word-of-Mouth Research: Principles and

Applications. *Journal of Advertising Research*, *47*(4), 398–411.

http://doi.org/10.2501/S0021849907070419

Bennett, W. L., & Manheim, J. B. (2006). The One-Step Flow of Communication. *The ANNALS of the*

American Academy of Political and Social Science, *608*(1), 213–232.

http://doi.org/10.1177/0002716206292266

Benoit, W. L. (2007). *Communication in Political Campaigns*. Peter Lang.

Benoit, W. L., Glantz, M. J., Phillips, A. L., Rill, L. A., Davis, C. B., Henson, J. R., & Sudbrock, L. A. (2011).

Staying "On Message": Consistency in Content of Presidential Primary Campaign Messages

Across Media. *American Behavioral Scientist*, *55*(4), 457–468.

http://doi.org/10.1177/0002764211398072

Brown, J., Broderick, A. J., & Lee, N. (2007). Word of mouth communication within online communities:

Conceptualizing the online social network. *Journal of Interactive Marketing*, *21*(3), 2–20.

http://doi.org/10.1002/dir.20082

Bruns, A., & Highfield, T. (2013). Political Networks on Twitter. *Information, Communication & Society*,

16(5), 667–691. http://doi.org/10.1080/1369118X.2013.782328

Carey, J. W. (1997). The press, public opinion, and public discourse: On the edge of the postmodern. In

E. S. Munson & C. A. Warren (Eds.), *James Carey: A critical reader* (pp. 228–260). Minneapolis:

University of Minnesota Press.

Castells, M. (2007). Communication, power and counter-power in the network society. *International

Journal of Communication*, *1*(1), 238–266.

Choi, S. (2015). The Two-Step Flow of Communication in Twitter-Based Public Forums. *Social Science

Computer Review*, *33*(6), 696–711. http://doi.org/10.1177/0894439314556599

Christensen, C. (2013). Wave-Riding and Hashtag-Jumping. *Information, Communication & Society*,

16(5), 646–666. http://doi.org/10.1080/1369118X.2013.783609

Conway, B. A., Kenski, K., & Wang, D. (2013). Twitter Use by Presidential Primary Candidates During the

2012 Campaign. *American Behavioral Scientist*. http://doi.org/10.1177/0002764213489014

Doherty, B. J. (2012). *The Rise of the President's Permanent Campaign*. University Press of Kansas.

Eisenstein, J., O'Connor, B., Smith, N. A., & Xing, E. P. (2014). Diffusion of Lexical Change in Social Media.

PLoS ONE, *9*(11), e113114. http://doi.org/10.1371/journal.pone.0113114

Goodnow, T. (2013). Facing Off: A Comparative Analysis of Obama and Romney Facebook Timeline

Photographs. *American Behavioral Scientist*. http://doi.org/10.1177/0002764213489013

Graham, T., Broersma, M., Hazelhoff, K., & van 't Haar, G. (2013). Between Broadcasting Political

Messages and Interacting with Voters. *Information, Communication & Society*, *16*(5), 692–716.

http://doi.org/10.1080/1369118X.2013.785581

Greenberg, S. (2012). *Congress + Social Media*. Center for Politics and Governance, University of Texas at

Austin. Retrieved from http://sites.utexas.edu/cpg/files/2014/05/2012-Congress-Social-

Media.pdf

Gulati, G. J. (2004). Members of Congress and presentation of self on the World Wide Web. *The

International Journal of Press/Politics*, *9*(1), 22–40.

Gulati, G. J., & Williams, C. B. (2013). Social Media and Campaign 2012 Developments and Trends for Facebook Adoption. *Social Science Computer Review*, *31*(5), 577–588. http://doi.org/10.1177/0894439313489258

Hawthorne, J., Houston, J. B., & McKinney, M. S. (2013). Live-Tweeting a Presidential Primary Debate: Exploring New Political Conversations. *Social Science Computer Review*. http://doi.org/10.1177/0894439313490643

Herring, S. (1999). Interactional Coherence in CMC. *Journal of Computer-Mediated Communication*, *4*(4), 0–0. http://doi.org/10.1111/j.1083-6101.1999.tb00106.x

Hewitt, J. (2001). Beyond Threaded Discourse. *International Journal of Educational Telecommunications*, *7*(3), 207–222.

Howard, P. N. (2005). Deep Democracy, Thin Citizenship: The Impact of Digital Media in Political Campaign Strategy. *The ANNALS of the American Academy of Political and Social Science*, *597*(1), 153–170. http://doi.org/10.1177/0002716204270139

Huckfeldt, R., Johnson, P. E., & Sprague, J. (2004). *Political Disagreement: The Survival of Diverse Opinions Within Communication Networks*. New York: Cambridge University Press.

Janssen, D., & Kies, R. (2005). Online forums and deliberative democracy. *Acta Politica*, *40*(3), 317–335.

Kaid, L. L., & Johnston, A. (2001). *Videostyle in Presidential Campaigns: Style and Content of Televised Political Advertising*. Greenwood Publishing Group.

Karpf, D. (2010). Online Political Mobilization from the Advocacy Group's Perspective: Looking Beyond Clicktivism. *Policy & Internet*, *2*(4), 7–41. http://doi.org/10.2202/1944-2866.1098

Klebanov, B. B., Diermeier, D., & Beigman, E. (2008). Automatic Annotation of Semantic Fields for Political Science Research. *Journal of Information Technology & Politics*, *5*(1), 95–120. http://doi.org/10.1080/19331680802149640

Kreiss, D. (2012). *Taking Our Country Back: The Crafting of Networked Politics from Howard Dean to Barack Obama*. New York: Oxford University Press.

Kreiss, D. (2014). Seizing the moment: The presidential campaigns' use of Twitter during the 2012 electoral cycle. *New Media & Society*, 1461444814562445. http://doi.org/10.1177/1461444814562445

Lambiase, J. J. (2010). Hanging by a thread: Topic development and death in an online discussion of breaking news. *Language@Internet*, *7*(9). Retrieved from http://www.languageatinternet.org/articles/2010/2814

Lawson-Borders, G., & Kirk, R. (2005). Blogs in Campaign Communication. *American Behavioral Scientist*, *49*(4), 548–559. http://doi.org/10.1177/0002764205279425

Leskovec, J., Backstrom, L., & Kleinberg, J. (2009). Meme-tracking and the Dynamics of the News Cycle. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 497–506). New York, NY, USA: ACM. http://doi.org/10.1145/1557019.1557077

Lilleker, D. G. (2016). Comparing online campaigning: The evolution of interactive campaigning from Royal to Obama to Hollande. *French Politics*, *14*(2), 234–253. http://doi.org/10.1057/fp.2016.5

Lotan, G., Graeff, E., Ananny, M., Gaffney, D., Pearce, I., & boyd, d. (2011). The Revolutions Were Tweeted: Information Flows During the 2011 Tunisian and Egyptian Revolutions. *International Journal of Communication*, *5*.

Mayhew, D. R. (1974). *Congress: The Electoral Connection*. New Haven, CT: Yale University Press.

McCombs, M. (2005). A Look at Agenda-setting: past, present and future. *Journalism Studies*, *6*(4), 543–557. http://doi.org/10.1080/14616700500250438

McEnery, T., Xiao, R., & Tono, Y. (2006). *Corpus-based Language Studies: An Advanced Resource Book*. New York: Routledge.

Needham, C. (2005). Brand Leaders: Clinton, Blair and the Limitations of the Permanent Campaign.

    Political Studies, *53*(2), 343–361. http://doi.org/10.1111/j.1467-9248.2005.00532.x

Neuman, W. R., Guggenheim, L., Jang, S. M., & Bae, S. Y. (2014). The Dynamics of Public Attention:

    Agenda-Setting Theory Meets Big Data. *Journal of Communication*, *64*(2), 193–214.

    http://doi.org/10.1111/jcom.12088

Norris, P., Curtice, J., Sanders, D., Scammell, M., & Semetko, H. A. (1999). *On Message: Communicating

    the Campaign*. Thousand Oaks, CA: SAGE.

Pennebaker, J. W., Slatcher, R. B., & Chung, C. K. (2005). Linguistic Markers of Psychological State

    through Media Interviews: John Kerry and John Edwards in 2004, Al Gore in 2000. *Analyses of

    Social Issues and Public Policy*, *5*(1), 197–204. http://doi.org/10.1111/j.1530-2415.2005.00065.x

Perkins, J. (2010). *Python Text Processing with NLTK 2.0 Cookbook*. Birmingham, UK: Packt Publishing.

Petchler, R., & Gonzalez-Bailon, S. (2015). Automated content analysis of online political

    communication. In S. Coleman & D. Freelon (Eds.), *Handbook of Digital Politics* (pp. 433–450).

    Northampton, MA: Edward Elgar.

Robertson, S. P., Vatrapu, R. K., & Medina, R. (2010). Off the wall political discourse: Facebook use in the

    2008 U.S. presidential election. *Information Polity*, *15*(1), 11–31. http://doi.org/10.3233/IP-

    2010-0196

Schneider, S. M., & Foot, K. A. (2006). Web campaigning by US presidential primary candidates in 2000

    and 2004. In A. P. Williams & J. C. Tedesco (Eds.), *The Internet election: Perspectives on the Web

    in campaign 2004* (pp. 21–36). Lanham, MD: Rowman & Littlefield. Retrieved from

    https://books.google.com/books?hl=en&lr=&id=GYZjlgs0wOIC&oi=fnd&pg=PA21&dq=%22camp

    aigns%22+%22kirsten+foot%22&ots=YsMs6m9Tur&sig=QxIs-vNp_MKk5vqey_0bHJOLens

Schweitzer, E. J. (2011). Controlled interactivity: Parties' online campaigns in the 2009 German national

    elections. In D. G. Lilleker & N. A. Jackson (Eds.), *Political campaigning, elections, and the*

    internet: Comparing the US, UK, France, and Germany (pp. 99–120). New York: Routledge.

Shearer, E. (2015, July 14). 5 key takeaways about Twitter, Facebook and news use. Retrieved from

    http://www.pewresearch.org/fact-tank/2015/07/14/5-key-takeaways-about-twitter-facebook-

    and-news-use/

Simon, A. F., & Jerit, J. (2007). Toward a Theory Relating Political Discourse, Media, and Public Opinion.

    Journal of Communication, *57*(2), 254–271. http://doi.org/10.1111/j.1460-2466.2007.00342.x

Sinclair, B. (2012). *The Social Citizen: Peer Networks and Political Behavior*. Chicago: University of

    Chicago Press.

Skovsgaard, M., & Van Dalen, A. (2013). Dodging the Gatekeepers? *Information, Communication &*

    Society, *16*(5), 737–756. http://doi.org/10.1080/1369118X.2013.783876

Stock, W. G., & Stock, M. (2013). *Handbook of Information Science*. Boston, MA: de Gruyter.

Stromer-Galley, J. (2014). *Presidential campaigning in the Internet age*. New York: Oxford University

    Press.

Vaccari, C., & Nielsen, R. K. (2013). What Drives Politicians' Online Popularity? An Analysis of the 2010

    U.S. Midterm Elections. *Journal of Information Technology & Politics*, *10*(2), 208–222.

    http://doi.org/10.1080/19331681.2012.758072

Wu, S., Hofman, J. M., Mason, W. A., & Watts, D. J. (2011). Who says what to whom on twitter. In

    Proceedings of the 20th international conference on World wide web (pp. 705–714).

Table 1a: OLS regressions of Obama commenters' issue proportions.

Predictor	Model 1			Model 2		
	B	Beta	SE	B	Beta	SE
Obama issues	0.515***	0.762	0.109	0.12**	0.178	0.039
Romney issues	.	.	.	-0.303***	-0.708	0.047
Romney commenters' issues	.	.	.	1.161***	1.474	0.094
Constant	0.026	.	0.01	0.001	.	0.003
R^2	0.581***			0.976***		
ΔR^2	.			0.395***		

* = $p < 0.05$, ** = $p < 0.01$, *** = $p < 0.001$

Table 1b: OLS regressions of Romney commenters' issue proportions.

	Model 1			Model 2		
Predictor	B	Beta	SE	B	Beta	SE
Romney issues	0.505***	0.928	0.051	0.276***	0.508	0.023
Obama issues	.	.	.	-0.082*	-0.095	0.036
Obama commenters' issues	.	.	.	0.790***	0.622	0.064
Constant	0.0276	.	0.006	0.001	.	0.003
R^2	0.861***			0.99***		
ΔR^2	.			0.129***		

* = $p < 0.05$, ** = $p < 0.01$, *** = $p < 0.001$

Table 2: Ranked terms for four most common issue topics across candidate and audience data subsets

| Issue	Rank	Obama	O comments	Romney	R comments
Civil rights	1	woman (5.97)	woman (3.39)	woman (1.54)	woman (2.05)
	2	daughter (1.12)	rights (1.37)	mother (0.34)	rights (1.04)
	3	mother (0.75)	daughter (0.93)	rights (0.17)	racis (0.79)
	4	latino (0.75)	abortion (0.68)	latino (0.17)	daughter (0.57)
	5	voter_registration (0.75)	racis (0.68)	voting_right (0.17)	abortion (0.48)
Finance/ economics/ labor (FEL)	1	job (4.48)	job (4.36)	job (10.45)	job (7.03)
	2	econom (2.61)	tax (3.27)	economy (9.25)	tax (5.75)
	3	tax (1.87)	econom (2.22)	middle_class (5.82)	econom (3.01)
	4	middle_class (1.49)	dollar (1.25)	tax (3.25)	middle_class (1.4)
	5	private_sector (1.12)	debt (1.13)	unemploy (2.74)	dollar (1.4)
Healthcare	1	health (2.99)	health (1.54)	obamacare (3.42)	health (1.31)
	2	obamacare (1.12)	obamacare (0.76)	medicare (1.88)	obamacare (1.13)
	3	medicare (0.75)	medicare (0.75)	health (0.51)	medicare (0.98)
	4	copay (0.37)	medical (0.43)	nurs (0.17)	medical (0.4)
	5	preventive (0.37)	doctor (0.28)	medicin (0.17)	obama_care (0.29)
Religion	1	faith (0.37)	god (3.95)	prayer (0.86)	god (3.92)
	2	-	muslim (2.31)	faith (0.68)	muslim (1.61)
	3	-	islam (1.46)	religio (0.17)	religio (1.02)
	4	-	religio (0.89)	god (0.17)	islam (0.8)
	5	-	christian (0.58)	catholic (0.17)	mormon (0.75)

Numbers in parentheses represent the percentages of all posts by the candidate or audience in which each term appears.

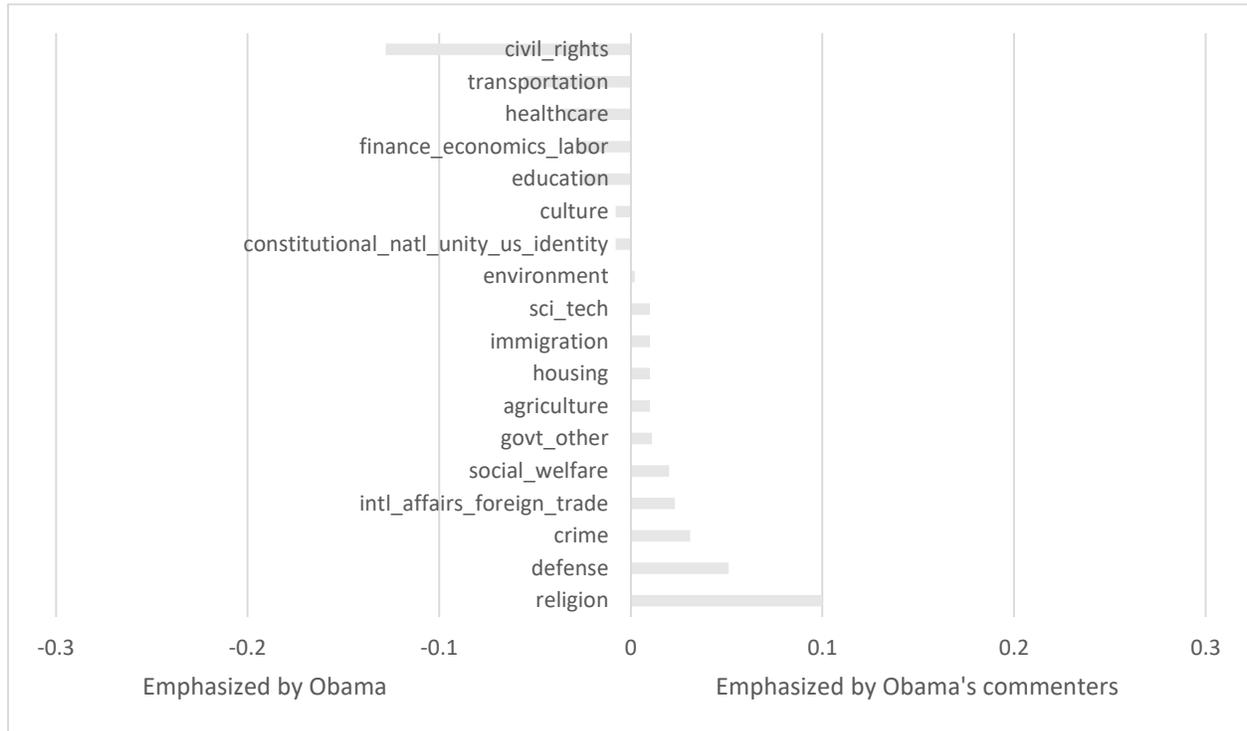**Figure 1:** N of posts in which issue terms appear as a function of the N of issue terms

Figure 2a: Obama's and his commenters' respective issue proportions
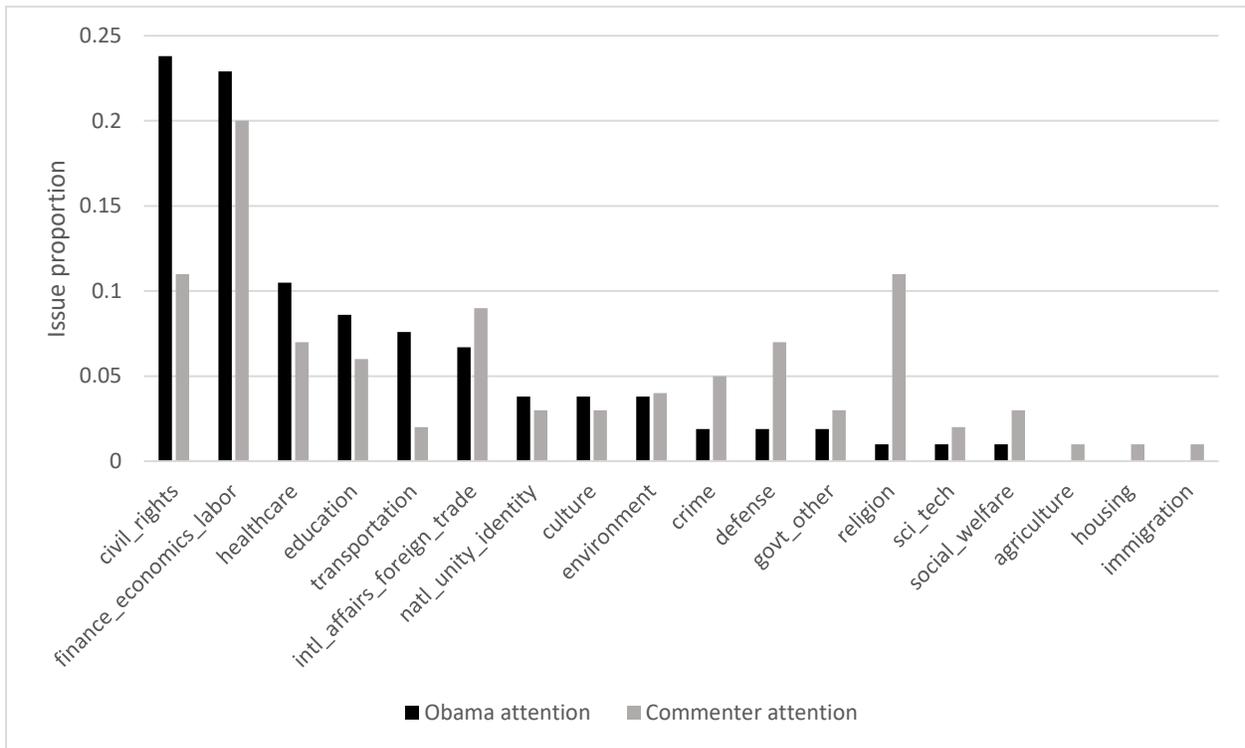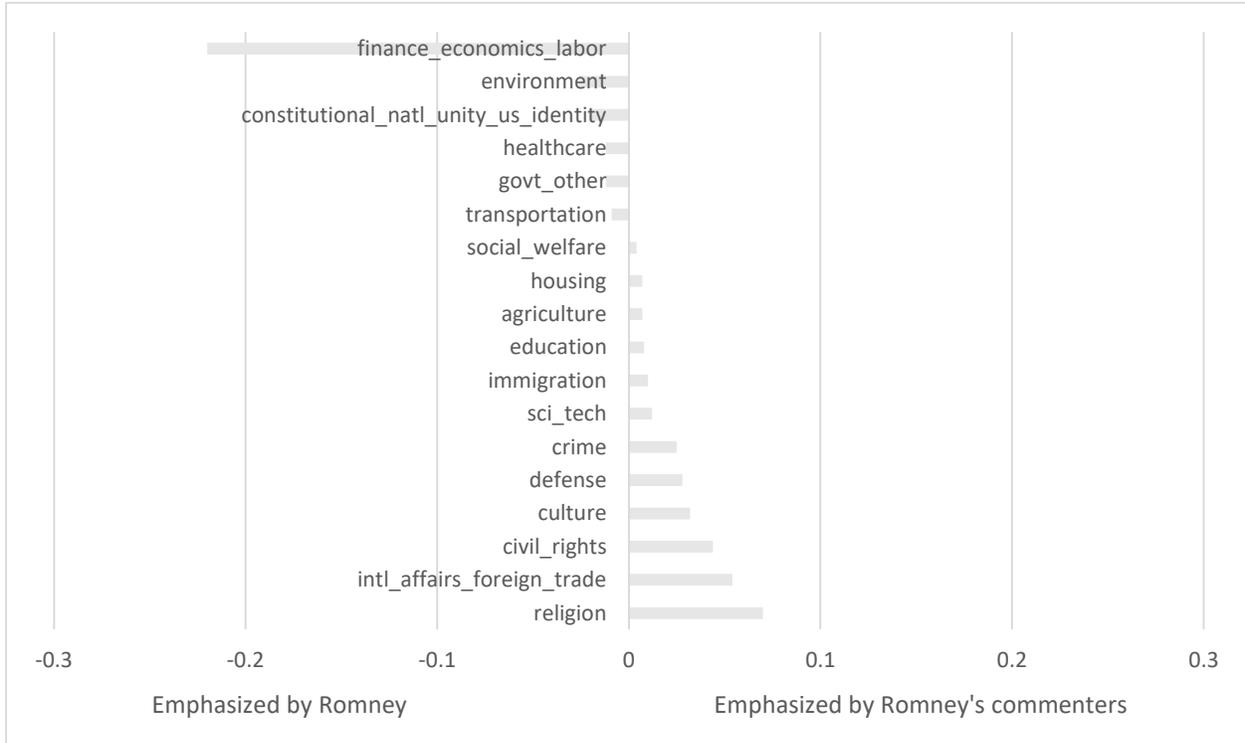
Figure 2b: Romney and his commenters' respective issue proportions
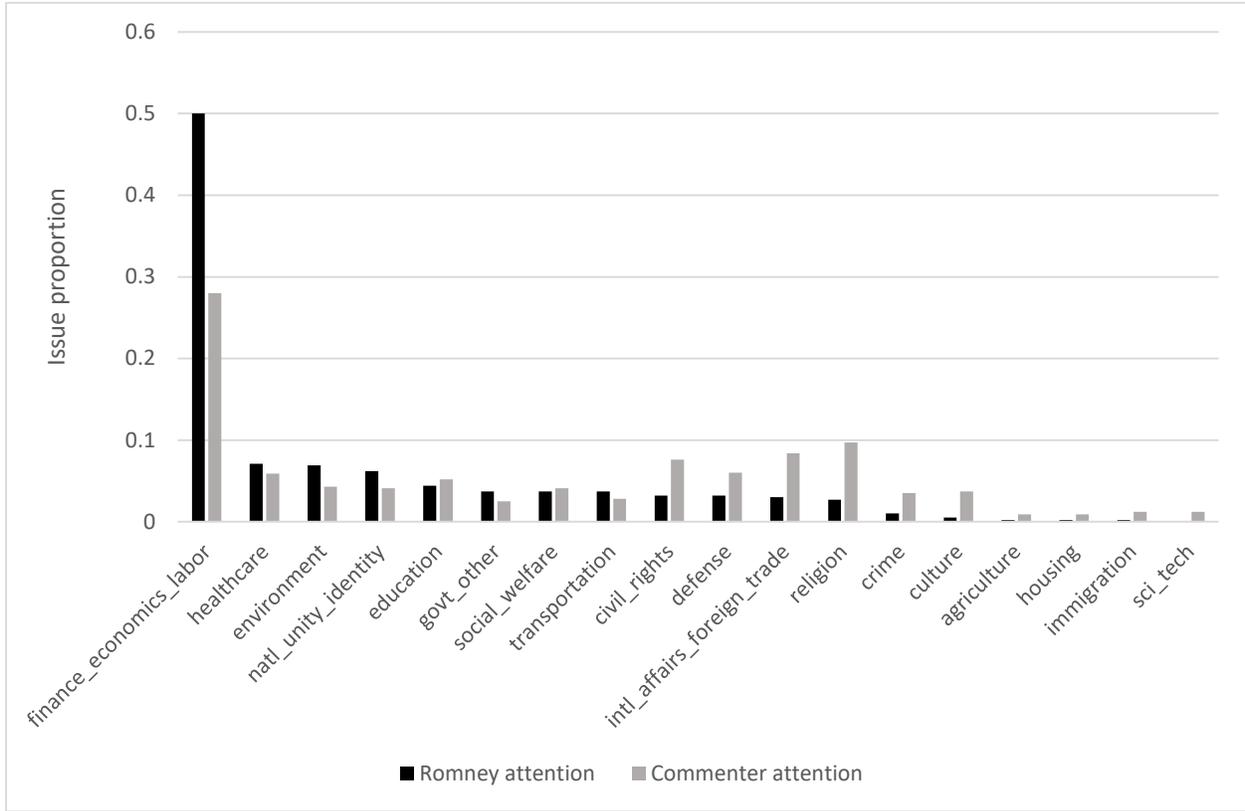
Figure 3a: Issue gaps between Obama and his commenters

Figure 3b: Issue gaps between Romney and his commenters

Notes

---

[1] http://facebook.com/barackobama and http://facebook.com/mittromney, respectively.

[2] An unknown number of comments were omitted from the dataset because the commenters' privacy settings prevented them from being collected.

[3] An anonymized online appendix containing all lexicon categories and terms is available here: http://bit.do/disintermediation