Does negative campaign advertising stimulate uncivil communication on social media? Measuring audience response using big data

A pre-print of the article that appeared in *Computers in Human Behavior, 68*, 368–377. https://doi.org/10.1016/j.chb.2016.11.034

Hopp, T., & **Vargo, C.** (2017).

Abstract: Using the 2012 presidential election as a case study, this work set out to understand the relationship between negative political advertising and political incivility on Twitter. Drawing on the stimulation hypothesis and the notion that communication with dissimilar others can encourage incivility, it was predicted that (1) heightened levels of negative campaign advertising would be associated with increased citizen activity on Twitter, (2) increased citizen activity would predict online incivility, and (3) that increases in citizen activity would facilitate a positive indirect relationship between negative advertising volume and citizen incivility. This theoretical model was tested using data collected from over 140,000 individual Twitter users located in 206 Designated Market Areas. The results supported the proposed model. Additional analyses further suggested that the relationship between negative political advertising and citizen incivility was conditioned by contextual levels of economic status. These results are discussed in the context of political advertising and democratic deliberation.

1.0 Introduction

Social media platforms play an increasingly prominent role in how Americans engage with politics. However, this increased engagement comes with a variety of concerns related to the functional *quality* of such engagement. Mirroring broader concerns over civility in the modern age, Forni (2011) has observed that "In today's America, incivility is on prominent display: in the schools… in the workplace…in politics, where strident intolerance takes the place of earnest dialogue; and on the Web, where many check their inhibitions at the digital door" (par. 1). Responding to these concerns, scholars have explored incivility from a number of angles, including the relationship between news media and citizen discussion quality (Borah, 2014), the relationship between individual level traits and uncivil discourse online (Hmielowski, Hutchens, & Cicchirillo, 2014), and the effects of uncivil commentary on people's understanding of information (Anderson et al, 2014). Despite the obvious potential of social media platforms to extend the public sphere, to-date research tends to treat online democratic communication processes in prospective, rather than observed, terms.

Additionally, there exist a wide array of questions relative to the degree that democratic discourse on social media is responsive to (and/or a product of) offline contextual realities. Research emanating from the political science, sociology, and mass communication fields suggests that offline civic involvement has been previously associated with a host of contextual factors, including socio-economic conditions, community heterogeneity, ideological polarization, and media environment (e.g., Costa & Kahn, 2003; Fiorina & Abrams, 2008; Goldstein & Freedman, 2002). However, such models of offline behavior have not been applied to online behavior.

In light of the foregoing, the goal of this study was to better understand the relationship between contextual political advertising factors and observed online behavior in terms of both volume and quality (i.e., *civility*) during the 2012 presidential election. To do so, we draw upon literature describing the so-called stimulation effect (the tendency for negative political advertising to stimulate voter engagement with the election) to develop and test a model that suggests (*a*) increased levels of negative campaign advertising will predict increased levels of campaign-relevant participation on social media and (*b*) increased participation on social media will, itself, be associated with higher levels of citizen incivility online. To test the above-described propositions, we capitalize on recent advances in computational social science to use novel big data and computer-assisted textual analysis techniques to analyze the campaign-relevant behavior of approximately 140,000 Twitter users located in 206 designated market areas (DMAs).

2.0 Literature Review

2.1 Negative Advertising and Civic Behaviors

During the past three decades there has been a fairly stiff debate over the degree to which negative, or *attack*, advertising influences citizen-based outcomes such as political efficacy and voting behavior (e.g., Ansolabhere & Iyengar, 1995; Lau, Sigelman, Heldman & Babbit, 1999). Perspectives framing the debate can be segmented into two camps. The first perspective holds that as the tone of political discourse becomes increasingly uncivil, Americans are more likely to become dissatisfied with the state of politics and discontinue their involvement in the political process (e.g., Ansolabhere & Iyengar, 1995). Conversely, the alternate perspective suggests just the opposite, stipulating that as negative, attack-oriented advertisements increase, so does citizen engagement with the political process (e.g., Brooks, 2006; Goldstein & Freeman, 2002; Lau, Sigelman, Heldman & Babbit, 1999; Lau, Sigelman, & Rovner, 2007; Stevens, Sullivan, Allen, & Alger, 2008).

Recent evidence has increasingly supported the latter proposition that negative campaign advertising stimulates – often on a conditional basis – citizen involvement with the election. A number of presumed mechanisms underlie this effect. First, negative campaign messages communicate "that something important is at stake in the outcome of the election" (Goldstein & Freeman, 2002, p. 735). This perceived importance effectively serves as an arousing agent (Martin, 2004), driving citizen involvement. Moreover, negative advertisements are most commonly used in competitive electoral races (Kahn & Kenny, 1999). A high volume of negative political advertising may therefore signal to the voter his or her individual vote may well play a role in determining important future outcomes.

Secondly, Martin (2004) points out that negative campaign advertising may incur a sense of republican duty. Here it is again assumed that "American citizens share some deep concern over the future of the country and that this concern can be stimulated to encourage participation" (p. 549). Attack advertisements suggest that maintenance of the status quo will result in negative future outcomes for the country. Attempts to disrupt citizens' faith in the status quo therefore exert a mobilizing effect on the citizenry (Berelson, Lazarsfeld, & McPhee, 1954; Martin, 2004).

Third, scholars (e.g., Derryberry, 1991; Pratto & John 1991; Taylor, 1991) have shown that people pay more attention to negative information (when compared to positive and neutral information). This effect is thought to be automatic in nature. As illustrated by Pratto and John (1991), "Events that may negatively affect the individual are typically of greater time urgency than are events that lead to desirable consequences. Averting danger to one's well being, such as preventing loss of life or limb, often requires an immediate response" (p. 380). In the context of political advertising, attack ads may be attended to more than non-attack advertisements (Marcus, 2003), resulting in greater engagement and, ultimately, heightened levels of voter turnout (Martin, 2004).

Fourth, attack advertisements facilitate the formation of negative emotional states among message receivers (e.g., Brader, 2005; Fridkin & Kenney, 2008; Valentino, Hutchings, Banks, & Davis, 2008). These negative emotional states, in turn, lubricate voter mobilization. Negative emotional states are particularly effective at altering behavior because of the heightened state of arousal that is associated with them. In his exploration of the relationship between negative campaigning and voter turnout, Martin (2004) cautioned that the “power of emotions such as anxiety to motivate participation should not be underestimated” (p. 550).

The literature reviewed above refers primarily to the relationship between negative campaign advertising and voting behaviors. Research interrelating negative advertising and citizens’ communication behaviors (in terms of both information seeking and self-expression) is both less common and less definitive. For example, a 2007 study by Shah and colleagues indicated that consumption of both traditional and online news was positively related to political advertising exposure. However, there was a *negative* relationship between exposure to negative advertising (calculated as the estimated proportion of negative to positive advertising experienced) and exposure to traditional news sources. Alternately, a separate study by Wang, Gabay, and Shah (2012) found that exposure to negative political advertisements among adolescents was associated with human-interest candidate knowledge, potentially suggesting that negative advertising may play an educational role among adolescents. As it specifically relates to social media, a recent study by Settle and colleagues (2015) indicated that Facebook users who reside in heavily contested “battleground” states were more likely to post election-relevant content than those living in less competitive "blackout" states. Although the authors did not explicitly measure the effects associated with advertising exposure, they contextualized their results by surmising that the competitive nature of campaigns, communicated in part through the heavy use of attack advertising, may drive user engagement online.

As it relates to the current study, we believe that Settle et al.'s (2015) findings coupled with research on negative advertising's so-called stimulation effect suggest that negative political advertising may stimulate broader levels of participation on Twitter. Building upon the discussed studies of voter behavior both on and offline, we specifically suggest that contextual environments featuring high levels of negative advertising may serve to prime users to believe that the election is a high stakes competition that is worthy of their attention. Thus, the following hypothesis is posited:

H1: DMAs with high levels of negative advertising will be associated with broader political participation on Twitter

2.2 Incivility

Incivility describes a wide array of communicative and non-communicative behaviors that range from rudeness and name-calling to vandalism and theft. Civil political discourse is central to a well-functioning democracy. As pointed out by Coe, Kenski, and Rains (2014), "commitment to civil discourse—the free and respectful exchange of ideas—has been viewed as a democratic ideal from the ancient Athenian forums to the mediated political debates of modern times" (p. 658). Similarly, Boyd (2008) posited that civil communication helps individual citizens "communicate respect for others and generate habits of moral equality in the everyday of life of a democracy" (p. 863).

What does uncivil online discourse look like? In Coe, Kenski, and Rains's (2014) study of uncivil discussion on online news comment forums, the authors defined incivility as "features of discussion that convey an unnecessarily disrespectful tone toward the discussion forum, its participants, or its topics" (pg. 660). Similarly, Santana (2014), defined incivility as possessing at least one of the following nine characteristics: (1) name calling; (2) threats; (3) vulgarities; (4) abusive or foul language; (5) xenophobia; (6) hateful language, epithets or slurs; (7) racist or bigoted sentiments; (8) disparaging comments on the basis of race/ethnicity; and (9) use of stereotypes. Speaking directly on the deliberativeness of conversational incivility, Papacharissi (2004) differentiated between *conversational impoliteness* and *collective impoliteness,* noting that conversational impoliteness is often a spontaneous emotional reaction while collective impoliteness is characterized by a purposeful disregard for democratic consequences.

2.3 Political Participation and Incivility

There are several reasons to believe discursive incivility increases as political participation broadens. First, those with high prior levels of political interest tend to engage with the election irrespective of the surrounding media atmosphere (e.g., Boulianne, 2011). Thus, embedded within the logic of the above-described stimulation hypothesis is the idea that the turnout-inducing effects of negative advertising are especially pronounced among those with moderate to low levels of "normal" political interest. Such stimulated citizens may possess opinions that are either underdeveloped (Freedman, Franz, & Goldstein, 2004) or exist outside the political mainstream (Calabrese, 2015; Kingwell, 1995). Moreover they may lack the knowledge necessary to fruitfully contribute to ongoing conversations, or may lack a thorough understanding of the social norms that govern political discussion (i.e., that it *ought* to be civil). In each of these cases, the lack of congruity between those citizens who normally engage in political discussion and those whose participation has been "stimulated" by the surrounding media environment could result in incivility.

Second, given the diverse racial, ethnic, and ideological composition of modern day America, it stands to reason that increased levels of political participation will *necessarily* result in the inclusion of a comparatively diverse array of voices. A core goal of political advertising (particularly in presidential campaigns) is to stimulate participation and interest among those who may not regularly vote. Indeed, recent years have seen presidential candidates aim messaging attempts toward non-regular/infrequent voters on the basis of race (i.e., Barrack Obama in the 2008 and 2012 elections), religious identification (i.e., Ted Cruz in the 2016 Republican primary), and policy preferences (i.e., Bernie Sanders in the 2016 Democratic primary). The heterogeneity-inducing effects of such messaging attempts may be especially apparent on Twitter, which has seen widespread adoption across demographic categories (Pew Research, 2015a) and the ideological spectrum (Pew Research, 2015b).

For its part, heterogeneity has been previously linked to incivility (e.g., Huysman & Wulf, 2004; Newton, 2001). For instance, as Schudson (1997) remarked, "democratic talk is not necessarily egalitarian but it is essentially *public*, and if this means that democratic talk is talk among people of different backgrounds and values, it is also profoundly *uncomfortable*" (p. 299, emphasis from the author). In other words, as more people with a wider array of individual and ideological attributes participate, it stands to reason that incivility levels may increase (Sunstein, 2001). A number of theoretical approaches support this contention. For instance, social capital theory (e.g., Boudieu, 1986; Coleman, 1988; Putnam, 2000) asserts that when people interact with others who are similar, they share important social resources such as norms of reciprocity, social trust, and feelings of mutual obligation (Ellison, Steinfeld, & Lampe, 2007; Huysman & Wulf, 2004). The availability of these resources helps individuals avoid and/or mitigate potential disagreements (Peck, 1993). In addition to the social capital perspective, the tendency for heterogeneity to accompany incivility can be similarly explained via a number of inter-related socio-psychological theories (see Yuan & Gay, 2006 for review), including self-categorization theory (e.g., Turner, 1987), the similarity–attraction hypothesis (e.g., Byrne, 1971), and the theory of homophily (McPherson & Smith-Lovin, 1987). Broadly understood, these theories hold that ‘‘similarity breeds connections’’ (McPherson, Smith-Lovin, & Cook, 2001, p. 415) and that society has embedded within it a ‘‘natural aversion to heterogeneity’’ (Alesina & La Ferrara, 2000, p. 225). Communication among heterogeneous others, therefore, is increasingly prone to incivility because participants lack the fundamental social tools (i.e., trust, credibility, mutual feelings of obligation) necessary for both functionally civil conversation. Moreover, when communicative conflicts do arise, the relative unavailability of social tools for conflict resolution may foster increased incivility (Balliet & Van Lange, 2012).

In light of the literature, our expectation was that broader levels of within-DMA participation would associate with heightened contextual levels of discussion-based incivility:

H2: DMAs with broader levels of Twitter-based participation will be associated with higher levels of incivility.

Finally, Hypotheses 1 and 2, when taken together, suggest that increased Twitter participation may facilitate a positive, indirect relationship between negative political advertising and campaign-related incivility:

H3: DMA-level Twitter participation will facilitate a positive, indirect effect between contextual negative advertising levels and contextual levels of campaign-related incivility.

3.0 Method and Materials

In the current study, all analyses were conducted at the DMA level. To test our hypotheses of interest, we retrieved data from several sources, as discussed below.

3.1 Datasets

3.1.1 Twitter Data

The first iteration of Twitter's Application Programming Interface (version 1.0) was called to download relevant Tweets during the 2012 presidential election period. The GET statuses/filter call was used to download public Tweets that included the terms "Obama" or "Romney." In all, 70 million tweets were collected. On August 1st, 2012 the collection started. It ended on the Election Day, November 6th. In all, 465,582 Tweets were collected with GPS coordinates. The researchers then created a program that used topojson files to resolve GPS coordinates (latitude and longitude) to the correct DMA.[1] Tweets were resolved successfully in 206 DMAs. In the case of 4 DMAs, no topojson files existed and thus the researchers were unable identify the necessary GPS coordinates. In all, 400,976 messages from 142,312 Twitter users were resolved to the DMAs. The average number of users per DMA was 690.83 (SD = 1,151.94; Range = 2 – 9,658). The average number of individual Tweets analyzed in each DMA was 1,946.49 (SD = 3,266.47; Range = 6 – 28,280).

Users in our dataset allowed tweets to be geo-tagged. Twitter asks users to make a one-time decision on whether their Tweets will be geo-tagged. On a per tweet basis it can then be optionally toggled. In 2012, roughly 1% of all Tweets were geo-tagged (Mahmud, Nichols & Drews, 2014). Studies survive this limitation and are able to predict larger general phenomena. Geo-tagged Tweets have been used to detect infectious disease transmission and breaking news stories (e.g., Sadilek, Kautz, & Silenzio, 2012). Researchers note key advantages of geo-tagged data for public discourse on issues and topics. First, spam bots that artificially inflate Twitter do not "opt-in" to geo-tagging (Guo & Chen, 2014). If a bot were to do so, it might reveal its location in an emerging country where spam bots are prevalent (Thomas, Grier, & Paxson, 2012). There is no motivation for a bot to "spoof" a specific location inside of the US. It would require setting metadata that is optional and not required to send messages.[2] As such, an argument can be made that geo-tagged data is more genuine and less prone to spam. Nonetheless, the use of geo-tagged data is a clear limitation of this work.[3]

After resolving users to individual DMAs, Twitter data was aggregated by user. If a user sent Tweets from more than one DMA, the DMA in which the majority of Tweets originated from was considered that user's DMA. Average incivility was calculated for each user. Finally, all users for each DMA were averaged, and average incivility scores were created for each DMA.

3.1.2 Kantar/CMAG Advertising Spending Data

Advertising spending for the 2012 election was acquired from Kantar Media's Campaign Media Analysis Group (CMAG). CMAG counts all political advertising in the 210 DMAs. This includes ads from candidates, parties, Super Political Action Committees (PACs) and interest groups that ran in support or against Barack Obama and Mitt Romney. All national broadcast, local TV stations, Spanish language networks and cable networks were monitored. Local cable advertisements are excluded. Each spot was monitored for certain attributes (here positive or negative valence). In addition to the number of ads that ran, CMAG relies on rate cards from media outlets and applies those rates to commercials that air to derive estimated spending.

3.1.3 Kantar/SRDS 2013 Market Profiles

Demographic information for each DMA was taken from Kantar Media's Standard Rate & Data Service (SRDS) DMA Market Profiles. The reported SRDS profiles are estimates generated using Census 2010 counts. Each DMA has a demographic profile. It also includes sales rank by merchandise category and rankings for TV and cable in that area. For this analysis, the 2013 version of the profiles were used. These numbers most accurately reflect the demographics of the DMAs in late 2012 when the data was collected.

3.2 Measures

3.2.1 Primary Measures

Incivility. Santana's (2014) definition of incivility was adopted as it broadly encompasses the conceptual approaches used by Papacharissi (2004) and Coe, Kenski, and Rains (2014) while also including constructs potentially relevant to the 2012 presidential election. Given the size of the dataset, the corpus could not be manually annotated for incivility. Computer-assisted content analysis enabled the researchers to derive incivility scores. Specifically, in the current study, researchers adopted an existing Python script (<withheld for peer review>, 2015) that was designed to detect political incivility in Twitter messages from the 2012 campaign. Using a dictionary of over 650 words associated with campaign-relevant incivility, the currently employed Python script took stemmed words and looked for a combination of windowed and unwindowed matches. To determine the degree to which the script accurately measured incivility in the current data, a trained human coder read a randomly selected subsample of 800 tweets and (1) verified words that were flagged were indeed uncivil and (2) analyzed the message to see if any words in the tweet were uncivil, but not detected by the computer. In both cases, external validity percentages were in excess of 98%. Once a computer has been verified to be valid, it is also reliable. Computers are not prone to human inconsistencies and as such calculating reliability is not necessary (Riffe, Fico, & Lacy, 2014; Zamith & Lewis, 2015).

After establishing the accuracy of the coding tool, incivility scores were calculated for each tweet by summing the number of words that appeared in that Tweet. In addition, when the matching word was in uppercase, or when the message contained an exclamation point, tweet scores were boosted by an additional point, as most affective boosting tools similarly do (Thelwall, 2010). After coding each individual tweet for incivility, a single aggregate index for incivility was created for each user. Average incivility scores were derived for each user based upon their "normal" pattern of interaction. These scores ranged from 0 (completely civil) to 23.00 (highly uncivil). Finally, the averaged individual scores were subsequently aggregated at the DMA level (M = 0.29, SD = 0.09; Range = 0.04 – 0.94).

To ensure criterion validity, the degree to which the presently derived construct correlated with logically related constructs such as arousal and sentiment was explored. Given that arousal and sentiment are more commonly identified constructs, previously validated classification schemes were used. For instance, <withheld for peer review> (2014) previously compared ANEW, AFINN and Sentistrength for sentiment of Twitter messages. Results echoed the findings of the creators and showed that Sentistrength performed best. As such, sentiment was coded with the same parameters recommend by <withheld for peer review> (2014). The ANEW list was chosen for arousal because it was the only available automated measure. It too has been found to be externally valid for Tweets (<withheld for peer review>, 2014). The expectation was that incivility would be negatively associated with sentiment and positively correlated with arousal. As expected, bivariate correlations of the aggregated measures (on the DMA level) were present in the expected directions (sentiment, $r = -.22$, $p < .01$; and arousal, $r = .20$, $p < .01$).

Negative campaign advertising. As the current study was interested in the general advertising climate surrounding the 2012 presidential election, we employed a measure that described the number of negative advertising spots shown in each DMA. This data was taken directly from the DMA dataset, which was coded by Kantar for positive/negative advertising (M

= 5, 334.04, SD = 10, 551.33; Range = 0.00 – 49,075.00). Notably, of the 1,264,383 spots included in the CMAG dataset (M = 6, 137.78, SD = 12, 288.61; Range = 0 – 59,197), 86.91% (*n* = 1,098,813) were negative in nature.

Twitter participation. To measure the extent to which citizens in each DMA were engaged with the 2012 presidential election on Twitter, we divided the number of unique Twitter handles that supplied election-relevant content by the total DMA population. The resulting number represented the percent of individuals, in each DMA, that posted election-relevant content on Twitter during the period of data collection (M = 0.05%, SD = 0.02%; Range = 0.006% - 0.142%).

3.2.2 Control Measures

DMA population. The total population of each DMA was taken from the SRDS DMA population estimates (M = 1,540,247.09; SD = 2,396,797.18; Range = 10,900 – 21,426,700).

DMA age. Because available data suggests that Twitter is more widely used by those 35 and younger, we controlled for the average age of each DMA. The SRDS data classifies population age in each DMA as follows: < 18 years; 18 -24 years; 25 – 34 years; age by 35 – 44 years; 45 – 54 years; 54 – 64 years; and ≥ 65 years. Using this information, we summed each category and calculated the percentage of the total population that was 34 years or younger (M = 46.08%, SD = 3.81%; Range = 34.26% - 58.16%).

Unemployment rate. The unemployment rate was represented as a percentage of the adult workforce and taken directly from the SRDS data (M = 5.70%, SD = 1.95%; Range = 2.00% - 24.90%).

Household income. The SRDS data provides household income estimates in the following categories: number of households whose annual income ranges from $10,000 to

$19,999; number of households whose annual income ranges from $20,000 to $34,999; number of households whose annual income ranges from $35,000 to $49,999; number of households whose annual income ranges from $50,000 to $74,999; number of households whose annual income ranges from $75,000 to $99,999; number of households whose annual income ranges from $100,000 to $124,999; number of households whose annual income ranges from $125,000 to $149,000; and number of households whose annual income is $150,000 or greater. Notably, at the time of data collection, the annual household income was $51,371 (U.S. Census Bureau, 2013). As such, we summed the number of households in each DMA that had annual income levels of $49,999 or lower and used this number to calculate the percentage of households in each DMA that earned less than the (approximate) national median (M = 47.66%, SD = 6.07%; Range = 24.79% - 63.56%).

Total advertising expenditures. Total advertising expenditures, in dollars, were taken directly from the SRDS data (M = $4,135,137.62, SD = $10,977,999.93; Range = $0.00 - $78,120,170).

Average Tweets per User. To control for the effects of highly active Twitter users on the incivility measure, we calculated a measure that described the average number of tweets sent by users in each DMA (M = 2.90, SD = 1.76; Range = 1.20 – 22.00).

3.3 Analytic Approach

To test Hypotheses 1 – 3, a series of OLS regression models were estimated using Hayes' (2013) PROCESS macro. The first model assessed the effects of the number of negative DMA advertising spots on (election relevant) Twitter-based participation in each DMA. The second OLS model assessed the effects of DMA-wide Twitter participation on incivility (aggregated at the DMA level), controlling for the direct effect of negative campaign advertising. Using the information from models 1 and 2, the PROCESS tool provides estimates of the model-implied indirect effects as facilitated by the proposed intercessory variable (Twitter participation). Statistical significance of the identified indirect effect was assessed using 5,000 bias-corrected re-samples of the data. If these bootstrapped confidence intervals do not contain 0, evidence of statistical significance is obtained. All analyses controlled for the effects of the following variables: DMA population size, DMA age, total DMA ad spending (in dollars), average number of Tweets per user, unemployment rate in each DMA, and the percentage of households in each DMA that had an annual average income below $50,000 USD. Notably, given the very strong correlation between number of advertising spots and number of negative advertising spots (r = .99), we controlled for the effects of overall advertising magnitude using the total spending amount (rather than total spot count, which introduced unacceptable levels of multicollinearity into our models). Finally, given the large variation in variable metrics, all variables were standardized (z-scored) prior to conducting the analyses. [4]

4.0 Results

4.1 Hypothesis Testing

First, an OLS model indicated the existence of a significant and positive relationship between number of negative campaign advertisements and campaign relevant participation on Twitter after controlling for the effects of total DMA population, DMA economic factors, total DMA campaign advertising expenditure, DMA age, and average number of Tweets per user, β = .33, $p < .01$. Full results are reported in **Table 1**. This result supported Hypothesis 1.

[INSERT TABLE 1 ABOUT HERE]

In a second regression analysis, which controlled for effects of the control variables *and* the number of negative campaign advertisements, there was a positive and significant relationship between campaign oriented Twitter participation and incivility on Twitter, $\beta = .21$, $p < .01$. A full report of this model is provided as **Table 1**. This finding supported Hypothesis 2.

Hypothesis 3 suggested that participation on Twitter would facilitate a positive indirect effect between negative campaign advertising and mean DMA incivility scores. The data suggested that the hypothesized indirect was, indeed, positive and significantly different than 0, $\beta = .07$, $95\%_{CI} = .01, .19$. Moreover, the model failed to indicate the existence of a direct effect between negative campaign advertising and incivility on Twitter, $\beta = -.07$, $p > .57$. These results supported Hypothesis 3.

4.2 Exploratory Analyses

Next, a series of exploratory moderation analyses were conducted. The goal of these analyses was to better understand if DMA characteristics condition the relationship between Twitter participation and average levels of incivility. Here, we specifically focused on economic issues as such issues were the central theme of the 2012 election. The first of these supplementary analyses examined the degree to which DMA household income moderated the relationship between Twitter participation and incivility. The interaction effect was significant, $\beta = -.12$, $p < .05$ (see **Table 2**, Model 1). Examination of the simple slopes revealed a pattern similar to that observed for unemployment rate. Specifically, the relationship between participation and incivility was stronger in DMAs with a greater percentage of households earning less than the national median income (10^{th} percentile: $\beta = .04$, $p >.71$; 25^{th} percentile: $\beta = .13$, $p >.12$; 50^{th} percentile: $\beta = .22$, $p <. 01$; 75^{th} percentile: $\beta = .29$, $p <. 01$; 90^{th} percentile: $\beta = .34$, $p <. 01$; see **Figure 1** for graphical representation of this effect).

[INSERT TABLE 2 ABOUT HERE]
[INSERT FIGURE 1 ABOUT HERE]

To follow up on the result reported above, we conducted a moderation analysis using a variable representing the percentage of DMA households earning over \$100,000 (roughly double the national median). Building on the foregoing, the expectation was that the relationship between Twitter participation and incivility would be strongest in those DMAs with a comparatively small percentage of households earning over \$100,000 annually. This expectation was confirmed. Specifically, a significant coefficient for the interaction term was obtained, $\beta = -.13$, $p < .05$, and examination of the simple slopes suggested that the relationship between the first-order predictor and the criterion variable was stronger in DMAs with a smaller percentage of households earning over \$100,000 annually (10^{th} percentile: $\beta = .34$, $p < .01$; 25^{th} percentile: $\beta = .28$, $p < .01$; 50^{th} percentile: $\beta = .22$, $p < .01$; 75^{th} percentile: $\beta = .13$, $p > .10$; 90^{th} percentile: $\beta = .04$, $p > .68$). **Table 2** (Model 2) provides details related to the regression model and **Figure 1** provides graphical representation of the interaction effect.[5]

Finally, given the above-identified interaction effects, we further explored the data to see if a moderated mediation effect (Hayes, 2013) was present. Specifically, the pattern of observed results suggested the existence of a second stage moderated mediation interaction effect whereby the indirect effect of negative ad volume on user incivility is conditioned by DMA economic status. Explicit tests of moderated mediation using 5,000 random re-samples of the data suggested the existence of a statistically significant effect for models employing both the measure representing percent of households earning below the national median (point estimate = 0.04, $95\%_{CI} = 0.002, 0.14$) and the measure representing percentage of households earning greater than \$100,000 annually (point estimate = -0.04, $95\%_{CI} = -0.16, -0.003$). As seen in **Table 3**, the indirect influence of negative advertising volume on citizen incivility was strongest in DMAs with lower household incomes. In DMAs with very high household incomes (i.e., those at the 75th and 90th percentiles), the indirect effect was no longer statistically significant at $p < .05$ level.[6]

[INSERT TABLE 3 ABOUT HERE]

5.0 Discussion

Using the 2012 presidential election as a case study, this study utilized a large volume of Twitter data to explore the relationship between contextual advertising environments, contextual levels of campaign-relevant participation on Twitter, and contextual levels of online incivility. The results suggested that negative advertising exerted a stimulating effect on Twitter participation, and that DMAs that had higher levels of Twitter participation also tended to feature higher levels of political incivility on Twitter. Exploration of the model-implied indirect effects suggested the existence of a positive and significant indirect effect, facilitated by contextual Twitter participation, between negative advertising and Twitter incivility. These effects were further conditioned by the finding that the indirect effect interrelating negative advertising exposure and citizen incivility was strongest in DMAs possessing comparatively low economic status levels. The implications of these findings are discussed below.

Consistent with prior studies on the effects of negative political advertising (e.g., Goldstein & Freeman, 2002; Lau, Sigelman, & Rovner, 2007), the current findings suggest that attack advertisements stimulate broader engagement with the election. Notably, most prior research has focused directly on the relationship between advertising and self-reported voting behaviors. This study, to the best of our knowledge, is the first study to specifically explore the stimulating effects of negative political advertising relative to *observed* citizen behavior on social media. In so doing, our results provide an initial indication that negative political advertising may influence citizen behaviors in ways that extend beyond the act of voting. Therein, the current data suggested a negative relationship between the average number of Tweets per user (in each DMA) and overall DMA incivility levels (see Table 2). This finding is consistent with our assumption that negative advertising primarily stimulates those with low levels of normal political engagement and suggests that a bulk of the uncivil discourse appearing on Twitter emanates from relatively infrequent contributors.

The current results also provide some insights into the structure of the relationship between political advertising and citizen discussion. Our *a priori* expectation that Twitter participation would facilitate a positive indirect effect between negative advertising volume and citizen incivility on Twitter was conditioned by the finding that the existence of this effect is dependent upon contextual economic conditions. In DMAs with a large percentage of high earning households, participation on Twitter (and, therefore, the above-identified indirect effect) was not significantly associated with citizen incivility. However, in comparatively poorer DMAs, we saw a linear strengthening of the indirect relationship between negative advertising and citizen incivility. This finding is consistent with years of research on the relationship between economic conditions and citizen behavior (e.g., Brady, Verba, & Schlozman, 1995), which has shown that behavioral incivilities are higher in economically depressed areas. At the same time, these findings extend our current understanding as they provide evidence that contextual factors, such as economic status, may influence civility in the online sphere.

As a whole, our results pose potentially important questions about the democratic and social value of attack ads. Negative political advertising often, if not always, seeks to capitalize on cynicism within the electorate. Attack ads often seek to manufacture citizen anger for the benefit of political elites. And, the broad reliance on such advertising techniques signals the existence of a political establishment that is perhaps more interested in manipulating public opinion than offering fruitful solutions to enduring problems. At the same time, our results suggest that negative campaign advertising was associated with broader levels of engagement with the election online. Such engagement, on or offline, is crucial to a well-functioning democratic society. Therefore, we wonder: what is the *net impact* of negative advertising on society? Clearly, this is a complicated question, and one that extends beyond the purview of the current work.

The present findings present similar questions about online incivility. There currently exists little empirical evidence on either the degree to which contemporary society is infiltrated by communicative incivility or the degree to which such incivility hinders democratic functioning. Broadly speaking, the prevailing notion is that “civil discussion lies at the heart of democracy” (Rowe, 2015, p. 121). However, as Schudson (1997) points out, truly democratic talk must encompass oppositional voices and perspectives. Undoubtedly, this means that discussions about serious issues (such as the selection of a president) should be accompanied by talk that is, at times, “profoundly uncomfortable” (p. 299). One way to rectify these ostensibly competing lines of thought is to recognize incivility as a natural bi-product of heterogeneous political discussion and, in so doing, conceptualize civility and incivility in *relative* terms. In other words, what is the proportion of uncivil to civil discourse? And, are there situations under which incivility serves as a catalyst for eventual cross-ideological understanding? By taking a broader view of incivility – and, in so doing, abandoning the notion that incivility is a *necessarily* an anti-democratic outcome – it may be possible to arrive at a more robust understanding of the degree to which society is/is not meaningfully influenced by incivility within the electorate.

Several key factors limit the current findings. First, the nature of our data imposes limitations on the generalizability of the current findings. Twitter behavior is spontaneous and cannot be understood as a comprehensive indicator of citizen communication habits. And, of course, Twitter messages do not represent public opinion in general. Furthermore, the data used in the current study were limited to a specific event (the 2012 general election). The researchers here believe the findings still yield important conclusions. Still, the present data cannot be used to make generalizations about the state of American political discussion in its entirety. Second, while the use of big data can offer novel insights, it is not without its drawbacks. These limitations are perhaps most apparent as they relate to the precision of our measures. For instance, previous efforts exploring the effects of negative advertising has been able to use individual-level outcomes via survey data (e.g., Goldstein & Freedman, 2002; Shah et al, 2007), resulting in the use of comparatively sophisticated estimates of negative political advertising exposure. However, in the current study, information related to individual Twitter users' media consumption habits was not available. As such, we choose to conduct our analyses at the DMA level. Undoubtedly, this introduced a substantive amount of noise into our measures that could, potentially, bias our results in an unknown manner. Furthermore, we recognize that our mechanism of conceptualizing Twitter participation is similarly imprecise. However, as shown in Table 1, this variable was not related to standard demographic controls, suggesting that our results were not confounded by obvious DMA factors such as population size, population age, household income, or unemployment rate. Third, it is important to note that while the current study broadly accounted for factors traditionally associated with the digital divide (i.e., socio-economic factors tied to individual abilities to access information and communication technologies), the nature of current study left us blind to the so-called "second-level" digital divide factors related to user skill (e.g., Hargitti, 2002). Moreover, while the SRDS data is clearly the most reliable available indicator of DMA demographic factors, it does not contain information about potentially meaningful socioeconomic factors such as educational obtainment. The use of DMA-level data additionally limited our ability to model voting behaviors and outcomes (i.e., polarity, voter turnout, and so on).

Future research should continue to combine big data and conditional modeling approaches as a means of better understanding the potentially complex interplays between contextual factors and online behaviors. Further research can be taken in a number of directions. First, given the case-study nature of the current work, researchers should attempt to replicate these findings using data taken from other elections. It may be especially interesting to conduct cross-national comparisons of the relationship between negative political advertising and social media-based incivility. Second, future research could explore other potential mediators and moderators governing the relationship between negative political advertising and online citizen incivility. Third, our focus on the aggregate-level effectively left us blind to individual level effects. Future research could better account for these effects by acquiring individual-level data and using multilevel modeling techniques to better understand the relationship between contextual influences and individual behavior.

Notes

1. The topojson files are available at the following GitHub repository: https://github.com/simzou/nielsen-dma

2. Bots do use various techniques to mask their true locations. Still, this does not involve setting the metadata to match locations in which the intermediary servers are located. Locking a spam bot into set locations (e.g., where a VPN server was) would make them easier to detect. As such, bots are likely to avoid geo-tagging altogether.

3. Prior research has found that user decisions to employ geotagging is not meaningfully associated with education and SES factors (<withheld for peer review>, 2015), thereby suggesting that it is likely not the product of first or second-level digital divide factors. Nonetheless, in the current dataset, we addressed, to the degree possible, the degree to which the geotagged data differed from the non-geotagged data. As the current sample was a subsample of a larger body of Tweets on the 2012 election (over 70 million tweets from more than 11 million users), we were able to compare the characteristics of the geotagged subsample to the larger sample comprised of those users who opted out of geotagging. Comparative examination of the raw (i.e., non-aggregated) individual-level data contained in each dataset suggested that the distributional characteristics of the incivility variable were generally similar across samples (geotagged sample: M = 0.28, SD = 0.74, skewness = 5.79; non-geotagged sample: M = 0.24, SD = 0.66, skewness = 7.46). Furthermore, within the aggregated data, we failed to find evidence that the percentage of geotagged-enabled users, relative to the total DMA population was, was related to percentage of DMA households earning less than $50,000 USD ($r = -.10$, $p > .14$), the percentage of household earning more than $50,000 ($r = .06$, $p > .42$), or the DMA-wide unemployment rate ($r = .01$, $p > .91$). We did, however, identify a small, positive correlation between the percent of geotag-enabled user accounts and percent of across DMA households earning more than $100,000 per year, $r = .16$, $p < .05$. As a whole, these results failed to provide strong evidence that, in the current dataset, Twitter geotagging is substantially different in terms of observed incivility or driven by macro socioeconomic factors.

4. Notably, one of the assumptions underlying the currently specified model is that increased participation will result in participatory diversity. To empirically test this assumption, we used the method described by Costa and Kahn (2003) to calculate DMA-level heterogeneity measures for race, household income, and age. For each measure, higher scores were indicative of higher levels of contextual heterogeneity. As expected, bivariate correlations suggested that the Twitter participation measure was positively correlated with both racial heterogeneity, $r = .16$, $p < .05$, and age-related heterogeneity, $r = .31$, $p < .01$. Moreover, we observed a positive and marginally significant relationship between Twitter participation and income heterogeneity, $r = .12$, $p = .08$.

5. We also conducted a regression analysis that explored the effect of DMA unemployment rate on the relationship between Twitter participation and incivility. Here, the interaction effect was marginally significant, $\beta = .13$, $p > .07$ [.073]. Exploratory examination of the simple slopes suggested that high levels of unemployment strengthened the relationship between campaign-relevant participation and user incivility. Specifically, at the 10^{th} percentile of unemployment, the relationship between participation and incivility was $\beta = .07$, $p > .53$; at the 25^{th} percentile of

unemployment, this relationship was $\beta = .13$, $p > .13$; at the 50th percentile, this relationship was $\beta = .20$, $p < .01$; at the 75th percentile of unemployment, the relationship between participation and incivility was $\beta = .25$, $p < .01$; and at the 90th percentile of unemployment, the relationship was $\beta = .32$, $p < .01$.

6. Given our use of aggregation, skewness issues presented less concern than if the analyses were conducted at the individual level. Examination of the variables of primary interest (i.e., those constituting either Hypotheses 1 - 3 or the exploratory moderation models) suggested that skewness was unlikely to bias the current results. With the exception unemployment rate (skewness = 4.73), all variables had skewness values ≤ 2.25.

References

Alesina, A., & La Ferrara, E. (2002). Who trusts others? *Journal of Public Economics, 85,* 207–234.

Anderson, A.A., Brossard D., Scheufele, D.A., Xenos, M.A., & Ladwig, P. (2014). The 'Nasty Effect': Online incivility and risk perception of emerging technologies. *Journal of Computer-Meditated Communication, 19*(3), 373-387.

Ansolabehere, S. & Iyengar, S. (1996). *Going negative: How political ads shrink and polarize the electorate.* New York, NY: Free Press.

Balliet, D., & Van Lange, P. A. M. (2012). Trust, conflict, and cooperation: A meta-analysis. *Psychological Bulletin, 139,* 1090 – 1112.

Berelson, B.B., Lazarsfeld, P.F., & McPhee, W.N. (1954). *A study of the opinion formation in a presidential campaign.* Chicago, IL: University of Chicago Press.

Borah, P. (2014). Doe it matter where you read the news story? Interaction of incivility and news frames in the political blogosphere. *Communication Research, 41*(6), 809-827.

Bouilianne, S. (2011). Stimulating or reinforcing political interest: Using panel data to examine reciprocal effects between news media and political interest. *Political Communication, 28*(2), 147-162.

Bourdieu, P. (1986). The forms of capital. In J.G. Richardson (ed). *Handbook of theory and research for the sociology of Education*. (New York, NY: Greenwood).

Boyd, R. (2006). The value of civility. *Urban Studies, 43*(5-6), 863 – 878.

Brader, T. (2005). Striking a responsive chord: How political ads motivate and persuade voters by appealing to emotions. *American Journal of Political Science, 49*(2), 388 – 405.

Brady, Henry, E., Sidney Verba, and Kay L. Schlozman (1995), "Beyond SES: A Resource Model of Participation," American Political Science Review, 89(2), 271-294.

Brooks, D. J. (2006). "The resilient voter: Moving toward closure in the debate over negative Campaigning and turnout. *Journal of Politics, 68*(3), 684 – 696.

Byrne, D. E. (1971). *The attraction paradigm.* New York, NY: Academic Press

Calabrese, A. (2015). Liberalism's disease: Civility above justice. *European Journal of Communication.* Online before print. doi: 10.1177/0267323115595306

Coe, K., Kenski, K., & Rains, S.A. (2014). Online and uncivil? Patterns and determinants of incivility in newspaper website comments. *Journal of Communication, 64*(4), 658 – 679.

Coleman, J. S. (1988). Social capital in the creation of human capital. *American Journal of Sociology, 94*, S95- S120.

Costa, D. L., & Kahn, M. E. (2003). Civic engagement and community heterogeneity: An economist's perspective. *Perspectives on Politics, 1*(1), 103 – 111.

Dayberry, D. (1991). The immediate effects of positive and negative feedback signals. *Journal of Personality and Social Psychology, 61*(2), 267 – 278.

Ellison, N. B., Stienfeld C., & Lampe C. (2007). The benefits of Facebook 'Friends:' Social Capital and College Students' use of online social network sites. *Journal of Computer-Mediated Communication, 12*, 1143 – 1168.

Fiorina, M. P., Abrams S.J. (2008). Political polarization in the American public. *Annual Review of Political Science, 11*(1), 563 – 588.

Forni, P.M. (2011, July 23). Why civility is necessary for society's survival. *The Dallas Morning News.* from: http://www.dallasnews.com/opinion/sunday-commentary/20100723-p.m.-forni-why-civility-is-necessary-for-society_s-survival.ece

Freedman, P., Franz, M., & Goldstein, K. (2004). Campaign advertising and democratic citizenship. *American Journal of Political Science, 48,* 723 – 741.

Fridkin, K. L., & Kenney P.J. (2008). The dimensions of negative messages," *American Politics Research, 36*(5), 694-723.

Goldstein, K., & Freedman P. (2002). Campaign advertising and voter turnout: New evidence for a stimulation effect. *The Journal of Politics, 64*(3), 721 – 740.

Guo, D., & Chen C. (2014). Detecting non-personal and spam users on geo-tagged twitter network. *Transactions in GIS, 18*(3), 370-384.

Hargittai, E. (2002). Second-level digital divide: Differences in people's online skills. *First Monday, 7(*4).

Hayes, A. F. (2013). *Introduction to mediation, moderation, and conditional process analysis*. New York, NY: The Guilford Press.

Hayes, A. F. (2015). An index and test of linear moderated mediation. *Multivariate Behavioral Research, 50*(1), 1-22.

Hmielowski, J. D., Hutchens M. J., & Cicchirillo, V. J. (2014). Living in an age of online incivility: Examining the conditional indirect effects of online discussion on political flaming. *Information, Communication, & Society, 17*(10), 1196 – 1211.

Huysman, M., & Wulf V. (2004). *Social capital and information technology*. Cambridge, MA: MIT Press.

Kahn, K. F., & Kenney, P.J. (1999). Do negative campaigns mobilize or suppress turnout? Clarifying the relationship between negativity and participation. *American Political Science Review, 93*(4), 877 – 889.

Kingwell, M. (1995). A civil tongue: Justice, dialogue, and the politics of pluralism. University Park, PA: Pennsylvania State University Press.

Lau, R., Sigelman, L., Heldman, C., & Babbit P. (1999). The effects of negative political advertisements: A meta-analytical assessment. *American Political Science Review, 93*(4), 851-76.

Lau, R., Sigelman, L.. & Rovner, I.B. (2007). The effects of negative campaigns: A met-analytic reassessment. *The Journal of Politics, 69*(4), 1176 – 1209.

Mahmud, J., Nichols, J., & Drews, C. (2012). Where is this tweet from? *ICWSM, 12*(1), 511 – 514.

Marcus, G. E. (2000) Emotions in politics. *Annual Review of Political Science, 3*, 221 – 250.

Martin, P. S. (2004). Inside the black box of negative campaign effects: Three reasons why negative campaigns mobilize. *Political Psychology, 25*(4), 545 – 562.

McPherson, M., & Smith-Lovin, L. (1987). Homophily in voluntary organizations: Status distance and the composition of face-to-face groups. *American Sociological Review, 52,* 370–389.

Newton, K. (2001). Trust, social capital, civil society, and democracy. *International Political Science Review, 22*(2), 201 – 214.

Pew Research (2015a). Demographics of key social networking platforms. Retrieved from: http://www.pewinternet.org/2015/01/09/demographics-of-key-social-networking-platforms-2/

Pew Research (2015b) Politics fact sheet. Retrieved from: http://www.pewinternet.org/fact-sheets/politics-fact-sheet/

Papacharissi, Z. (2004). Democracy online: Civility, politeness, and the democratic potential of online political discussion groups. *New Media & Society, 6*(2), 259 – 283.

Peck, D. L. (2002). Civility: A contemporary context for a meaningful historical concept. *Sociological Inquiry, 72*(3), 358 – 375.

Pratto, F., & John, O. P. (1991). Automatic vigilance: The attention-grabbing power of negative social information. *Journal of Personality and Social Psychology, 61*(3), 380 – 391.

Putnam, R. (2000). *Bowling alone: The collapse and revival of the American community*. New York: Simon & Schuster.

Riffe, D., Lacy, S., & Fico, F. (2014). *Analyzing media messages: Using 1uantitative content analysis in research*. New York, NY: Routledge

Rowe, I. (2015). Civility 2.0: A comparative analysis of incivility in online political discussion. *Information, Communication, & Society, 18*(2), 121 – 138.

Sadilek, A, Kautz, H.,, & Silenzio, V. (2012). Predicting disease transmission from geo-tagged micro-blog data. *The Twenty-Sixth AAAI Conference on Artificial Intelligence*. Retrieved from http://www.aaai.org/ocs/index.php/AAAI/AAAI12/paper/download/4844/5130

Santana, A. D. (2014). Virtuous or vitriolic: The effect of anonymity on civility in online newspaper reader comment boards. *Journalism Practice, 8*(1), 18 – 33.

Schudson, M. (1997). Why conversation is not the soul of democracy. *Critical Studies in Mass Communication, 14*(1), 1–13.

Settle, J. E., Bond, R. M., Coviello, L.,Fariss, C.J., Fowler, H.J., & Jones, J.J. (2015). From posting to voting: The effects of political competition on online political engagement. *Political Science Research and Methods.* Retrieved from http://dx.doi.org/10.1017/psrm.2015.1

Shah, D. V., Cho, J., Nah, S., Gotlieb, M R., Hwang, H., Lee, N., Scholl, R. M., &. McLeod, D. M. (2007). Campaign ads, online messaging, and participation: Extending the communication mediation model. *Journal of Communication, 57*(4), 676 – 703.

Stevens, D., Sullivan, J., Allen, B., & Alger D. (2008). What's good for the goose is bad for the gander: Negative political advertising, partisanship, and turnout. *The Journal of Politics, 70*(2), 527-541.

Sunstein, C. (2001). *Republic.com*. Princeton, NJ: Princeton University Press.

Taylor, S. E. (1991). Asymmetrical effects of positive and negative events: The mobilization-minimization hypothesis. *Psychological Bulletin, 110*(1), 67 – 85.

Thomas, K., Grier C., & Paxson V. (2012). Adopting social spam infrastructure for political censorship. *Proc. USENIX Security Symposium*. Retrieved from http://www.icir.org/vern/papers/kremlin-bots.leet11.pdf

Turner, J. C. (1987). *Rediscovering the social group: A self-categorization theory.* Oxford, England: Basil Blackwell.

U.S. Census Bureau. (2013), *Household income: 2012.* Retrieved From https://www.census.gov/prod/2013pubs/acsbr12-02.pdf

Valentino, N. A., Hutchings, V. L., Banks, A. J., & Davis, A. K. (2008). Is a worried citizen a good citizen? Emotions, political information seeking, and learning via the internet. *Political Psychology, 29*(2), 247 – 273.

Wang, M., Gabay, I., & Shah, D.D. (2012). The civic consequences of ‘Going Negative’: Attack ads and adolescents’ knowledge, consumption, and participation. *The ANNALS of the American Academy of Political and Social Science, 644* (1), 256 – 271.

Yuan, Y. C., & Gay, G. (2006). Homophily of network ties and bonding and bridging of social capital in computer-mediated distributed teams. *Journal of Computer-Mediated Communication, 11,* 1062–1084.

Zamith, R., & Lewis, S.C. (2015). Content analysis and the algorithmic coder: What computational social science means for traditional modes of media analysis. *The ANNALS of the American Academy of Political and Social Science, 659*(1), 307 – 318.

Table 1

Regression results for models predicting Twitter participation and Twitter Incivility

Variable	**Twitter Participation**	**Twitter Incivility**
DMA Population	.10	-.09
DMA Age $\leq$ 34 (%)	.13^t	.07
Total Ad Spending ($)	.08	-.05
Average Tweets Per User	-.06	-.16*
Unemployment Rate (%)	.01	.25**
Household Income < 50,000 (%)	.04	-.04
Negative Advertising Spots	.33**	-.07
Twitter Participation (%)		.21**
F	(7, 198) = 6.33, $p < .01$	$F(8, 197) = 4.56$, $p < .01$
R^2	.18	.16

Note: All estimates standardized; ** $p < .01$, * $p < .05$, ^{t} $p < .10$

Table 2

Regression results for models illustrating moderating effect of household income levels on the relationship between Twitter participation and Twitter incivility

Variable	**Twitter Incivility (Household Income Model 1)**	**Twitter Incivility (Household Income Model 2)**
DMA Population	-.11	-.13
DMA Age $\leq$ 34 (%)	.06	.05
Total Ad Spending ($)	.02	.05
Average Tweets Per User	-.17*	-.16*
Unemployment Rate (%)	.27**	.27**
Household Income < $50,000 (%)	-.07	
Household Income > $100,000 (%)		.11
Negative Advertising Spots	-.11	-.13
Twitter Participation (%)	.20**	.20**
Twitter Participation x Income (<50,000)	.12*	
Twitter Participation x Income (>100,000)		-.13*
F	(9, 196) = 4.54, $p < .01$	(9, 196) = 4.70, $p < .01$
R^2	.17	.18

Note: All estimates standardized; ** $p < .01$, * $p < .05$

Table 3

Strength of the indirect effect of negative advertising volume on incivility at various levels of DMA household income

	% of DMA Households Earning <$50,000	**% of DMA Households Earning >$100,000**
10th Percentile	.01 (-.05, .09)	.11 (.02, .31)
25th Percentile	.04 (-.003, .13)	.09 (.02, .25)
50th Percentile	.07 (.01, .19)	.08 (.02, .21)
75th Percentile	.10 (.02, .26)	.04 (-.001, .14)
90th Percentile	.11 (.02, .31)	.01 (-.06, .09)

Note: Reported estimates are standardized. Bootstrapped confidence intervals in parentheses

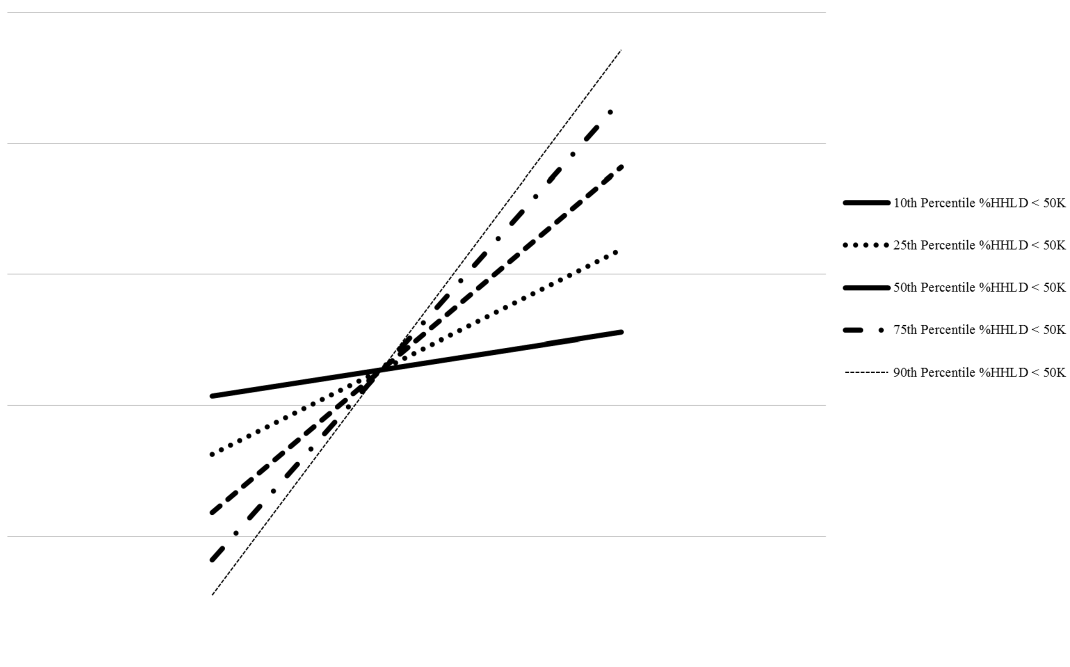

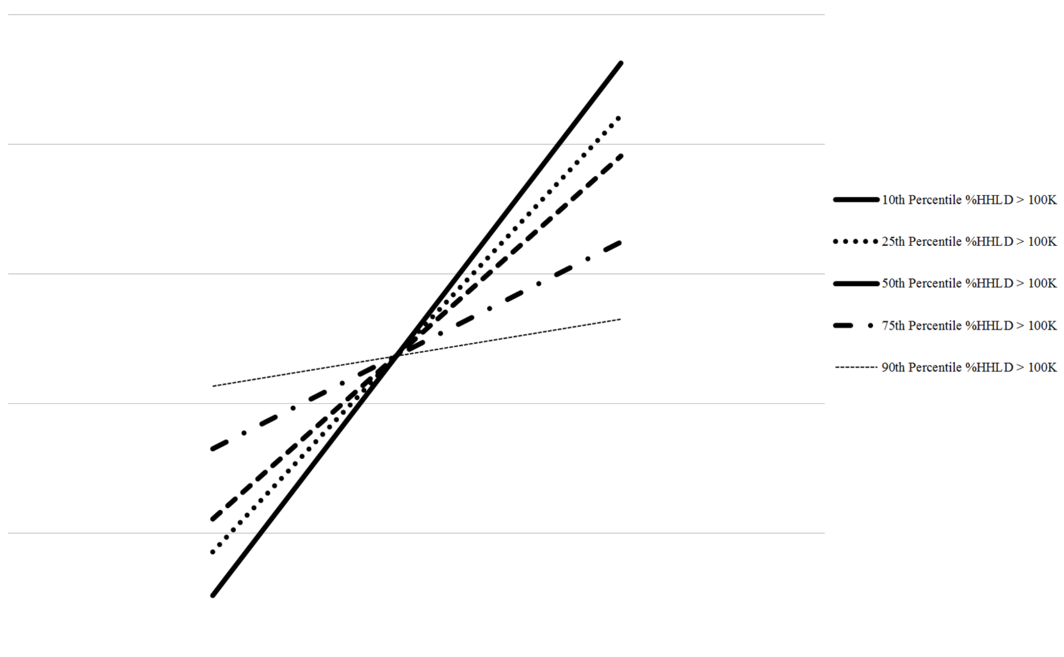

Figure 1. The relationship between DMA Twitter participation and DMA Twitter incivility at various levels of DMA household income (HHLD). In the top pane, the household income measure is the percentage of total DMA households with annual income levels below the approximate national median ($50,000 USD). In the bottom pane, the household income measure is the percentage of DMA households with annual income levels above $100,000. In both panes, Twitter participation is on the x-axis while incivility is on the y-axis