

Quantifying the power and consequences of social media protest

Deen Freelon

School of Communication, American University

Charlton McIlwain

Department of Media, Culture, and Communication, New York University

Meredith Clark

Mayborn School of Journalism, University of North Texas

Abstract: The exercise of power has been an implicit theme in research on the use of social media for political protest, but few studies have attempted to measure social media power and its consequences directly. This study develops and measures three theoretically-grounded metrics of social media power—unity, numbers, and commitment—as wielded on Twitter by a social movement (Black Lives Matter/BLM), a counter-movement (political conservatives), and an unaligned party (mainstream news outlets) over nearly ten months. We find evidence of a model of social media efficacy in which BLM predicts mainstream news coverage of police brutality, which in turn is the strongest driver of attention to the issue from political elites. Critically, the metric that best predicts elite response across all parties is commitment.

Keywords: Black Lives Matter, Twitter, protest, social movements, connective action, computational methods

Corresponding author:

Deen Freelon

School of Communication

American University

4400 Massachusetts Ave NW

Washington, DC 20016

freelon@american.edu

202-885-6462

Quantifying the power and consequences of social media protest

Social media (broadly defined) have become essential tools for 21st-century social movements. Accordingly, use of social media for political protest is a thriving research area, with studies applying both qualitative and quantitative methods to understand the nature and magnitude of the phenomenon. Most researchers in this area agree that social media can be consequential for social movements and their protests in at least some contexts (Bennett and Segerberg, 2012; Earl and Kimport, 2011; Shirky, 2011).

All successful social movements must exercise power to help bring about their chosen social goals. Movements have traditionally done so by a number of means, including protests, petitions, and directly lobbying politicians. Contemporary social movements such as Black Lives Matter (BLM), which we examine here, consider social media an important component of their overall strategies. But existing studies have not fully explored how movements harness power through social media. In particular, they have not adequately accounted for the fact that social movements are not alone in social media: other parties interested in the same topic almost always emerge to wield their own power alongside, against, or orthogonally with respect to the movement.

This paper introduces a new methodology to address this reality. It defines several forms of social media power that are particularly relevant to social movements, proposes accompanying techniques to measure them, and tests the extent to which they predict a key movement outcome—elite responses. Critically, non-movement parties may also wield these forms of power, which are rooted in Charles Tilly's concept of WUNC (worthiness, unity, numbers, commitment) (1999; Tilly and Wood, 2013). Using 40.8 million tweets about police

shootings of unarmed Black people in 2014 and 2015, we demonstrate that the digital manifestations of three of WUNC's four components can be measured quantitatively for both movement and non-movement constituencies. Our analysis of the relationships between these metrics and elite response suggests that certain of the former probabilistically cause the latter.

Social movement power

In an influential article, Diani (1992) outlines four essential components of social movements: “a) networks of informal interaction; b) shared beliefs and solidarity; c) collective action on conflictual issues; d) action which displays largely outside the institutional sphere and the routine procedures of social life” (1992, p. 7). This definition would seem to admit a wide range of structures and tactics, but from the development of resource mobilization theory in the late 1970s until very recently, formal social movement organizations (SMOs) have been considered all but essential for social movements (Earl and Kimport, 2011; McCarthy and Zald, 1977). The current study draws on contemporary frameworks that take digitally-enabled collective action seriously, in particular Bennett and Segerberg's connective action typology (2013; Earl and Kimport, 2011). We contend that connective movements are social movements, sharing all of Diani's definitional characteristics despite differing tactics and hierarchies.

It perhaps goes without saying that power is an indispensable resource for social movements, but as is often the case, the obvious warrants clarification. We define “power” for the purposes of this article as the capacity to bring about desired changes in society. This is consistent with the views of a broad range of scholars who view power as fundamental to all social systems (Bennett, 2003; Castells, 2012; Couldry and Curran, 2003; Giddens, 1987). Giddens, for example, refers to power as “the capability to intervene in a given set of events so as in some way to alter them” (1987: 7). We are concerned primarily with what is often labeled

“media power” (Couldry and Curran, 2003); that is, non-coercive power that flows through various forms of media.

Media power is especially important for connective movements. The assumption that shifts in discourse may eventually lead to broader social changes underlies every social movement’s communication efforts. In some cases, changing the conversation about the issue in question is the ultimate goal. In others, movement-led discussions of social issues on social media are not ends in and of themselves, but rather one means of addressing a larger problem. This is particularly true of movements like BLM whose goals involve institutional policy change (see Movement for Black Lives, 2016). Among other uses, social media allow activists to interact with lawmakers directly, given that many if not most of the latter have Twitter and Facebook accounts (at least in the US). While some recent connective movements, most notably Occupy and the Egyptian and Tunisian revolutions, have explicitly avoided engaging politicians directly (Castells, 2012), doing so is essential to fulfill policy-related goals. Movements pushing for institutional changes must seek the attention of those in charge, the same as any formal interest-based organization (Button, 1989: 6; Tarrow, 1998: 34). Elite attention is a key outcome of power in such cases.

Anyone who has ever observed or participated in a connective movement as it has pressed its case online knows that it does not operate in a vacuum. Movements fortunate enough to attract substantial public attention online quickly find themselves among allies, opponents, journalists, celebrities, curious onlookers, and would-be entertainers seeking to capitalize on the latest trend. Almost invariably, similar groups of individuals tend to cluster together in social media, communicating about the topic at hand mostly within like-minded communities (Adamic and Glance, 2005; Conover et al., 2011; Hargittai et al., 2008). Each of these communities is

involved in a power competition with the others, whether its participants are aware of it or not (Aouragh, 2012; Kahn and Kellner, 2004). The simple act of sharing one side's message rather than another's is a key component in this process.

This suggests that when researchers analyze social movements' power online, they should not focus solely on the movement. Instead, they should include other collective interests so that they may be compared. Aside from BLM, two additional interests will be analyzed here. First, movements with controversial or radical aims often attract counter-movements dedicated to thwarting them.¹ Though social media make confronting one's ideological adversaries easier than ever before, few studies have examined online counter-movements directly (exceptions include Croeser and Highfield, 2014; and Jensen and Bang, 2013). Second, the mass media typically cover movements that achieve a certain threshold of popularity. True to their ostensibly objective principles, they usually align neither with movements nor counter-movements consistently and are best considered "unaligned," for lack of a better term. **While US mainstream news outlets exhibit their own distinct ideology** (see e.g. Barnhurst, 2005; Reese, 1990), **their coverage does not consistently favor the left or the right**. On social media, most high-visibility movements will likely attract both counter-movements and unaligned observers. Because all these communities are embedded with one another in a system of digitally-mediated power relations, they can all potentially command the same forms of power.

Measuring social media power

One of this article's central claims is that movement power as exerted through social media can be quantified. Previous studies have found social media to be important for information sharing, frame building, and/or offline protest facilitation (Bastos et al., 2015; De Choudhury et al., 2016; González-Bailón et al., 2013; Theocharis, 2013; Tufekci and Wilson, 2012; Valenzuela et al.,

2014). The current study differs from this work in two important ways. First, it demonstrates how abstract concepts of power developed for offline social movements manifest and change over time in social media. Second, it presents evidence that these forms of power can help further movements' policy goals directly, as opposed to solely facilitating communication among activists.

Our power metrics are based on digital traces of social media activity such as retweets, hashtags, and screen names. But we will not simply assume without justification that particular traces signify particular theoretical constructs, as some studies have done (Author, Date). Instead, we will argue that certain trace-based metrics can be considered indicators of Tilly's concept of WUNC (Tilly, 1999; Tilly and Wood, 2013). WUNC is an acronym whose letters signify *worthiness*, *unity*, *numbers*, and *commitment*, all essential elements for social movements to wield adeptly. Tilly describes WUNC as both a set of defining characteristics of social movements and as a source or index of social movement power. He associates its elements with movement "strength" and notes that they "increase the plausibility of the implied threat that the claimant will use its weight to enter, realign, or disrupt the existing polity" (Tilly, 1999: 262; Vliegenthart and Walgrave, 2012). Thus it is no major conceptual leap to consider WUNC as power by our definition.

Tilly conceives of WUNC as a measurable set of properties. He writes of "high" and "low" values of its four components (Tilly, 1999), which clearly imply possibilities for quantification. Yet most empirical applications have been qualitative, with authors describing how specific social movements' characteristics fit the WUNC framework (Agbaria and Mustafa, 2012; Bennett and Segerberg, 2012). For example, in developing their theory of connective action, Bennett and Segerberg write that "digitally mediated action networks often seem to be

accorded higher levels of WUNC than their more conventional social movement counterparts” (2012: 742). Again we see a clear suggestion that WUNC can be measured—and in digital contexts no less—but it is followed by no methodological suggestions as to how.

Of WUNC’s four components, we propose to measure only the latter three. While quantifying worthiness may be possible, it seems to us prohibitively difficult compared to unity, numbers, and commitment. Tilly and Wood give the following offline examples of worthiness: “sober demeanor; neat clothing; presence of clergy, dignitaries, and mothers with children” (Tilly and Wood, 2013: 5). Demeanor, clothing, and religious identity on social media could perhaps be judged by human coders, but not at scale. And because not all mothers identify themselves as such online, it would likely be impossible to reliably code social media profiles for motherhood. Fortunately, the remaining three components of WUNC are much more empirically tractable.²

Unity

As a theoretical construct, unity makes a much smoother transition to social media contexts than worthiness. Tilly cites the “wearing or bearing of common symbols [and] direct affirmation of a common program or identity” (1999: 261) as signifiers of unity, among others. For movements that use social media extensively, few common symbols are as emblematic as their best-known hashtags. They are the digital analogues of hand-held signs at street protests. #Jan25, #Occupywallstreet, and #Blacklivesmatter are three iconic examples that instantly identify their corresponding movements. Creating hashtags based on victims’ names after police killings is a common practice within BLM, so much so that participants sometimes speak of their fear of “becoming a hashtag” (Moodie-Mills, 2015). The names of the most famous victims become metonyms for the everyday fears of many Black Americans.

Empirically, unity can be expressed through social media as a tendency for a given community to use a small number of movement-related hashtags disproportionately more often than others. This indicates that participants are conveying a unified message, particularly when the hashtag in question expresses a normative claim (e.g. #Blacklivesmatter). A lack of consensus in hashtag use suggests at a minimum a corresponding lack of unity in social media messaging, and perhaps also in deeper tactical or philosophical viewpoints. Like the other two metrics, hashtag inequality can be measured at the community level, thus permitting quantitative comparisons.

Numbers

Of WUNC's four elements, numbers is probably the most straightforward to measure in social media. Doing so is much easier than in offline protests, where journalistic and activist estimates of attendance frequently diverge (Mann, 1974). While overall counts of social media users over time are important, we are more concerned with the specific numbers of users associated with movements, counter-movements, and unaligned parties. We describe and implement a novel method of doing so in the "Data and Methods" section below. This method relies on a network analysis technique known as *community detection* to categorize users based on their retweeting behavior. We use network communities as the main unit of analysis throughout this paper because they intuitively approximate participants' tendency to congregate with ideological allies.

Once a set of communities has been identified and labeled, the participants in each can be counted just as easily as for the entire dataset. Importantly, our method allows us to aggregate community user counts per day so that longitudinal changes may be observed. It is perhaps self-evident that, **other things being equal** and barring purchased followers, "bots," and other obfuscatory shenanigans, numbers signify power.

Commitment

Tilly defines commitment as, among other things, “declarations of readiness to persevere” (1999: 261).³ The longitudinal nature of social media data allows us to improve upon this operational definition and directly observe perseverance itself. Having first disaggregated a social media conversation into multiple communities, and then reconstituted those communities on each individual day, it becomes possible to measure how committed each community’s participants are. We propose a simple method of doing so: computing the proportion of participants in a given community on any given day who tweet at least once during the following three days.⁴ Note that participants do not need to appear in the same community on the first day as in the next three—they simply need to post at least one relevant message in the latter.

Comparing this *repeat participation rate* between communities allows us to determine which are most and least committed. High proportions indicate that many participants from a given community are returning to continue promulgating its point of view. Low proportions, in contrast, indicate a high turnover rate and therefore a less committed and less stable community. Commitment as expressed in this way sends the message that movements and their interactants will not disperse (digitally speaking) when the next trending topic emerges.

Black Lives Matter

We apply these three power metrics to nearly ten months of Twitter conversations started by the Black Lives Matter movement. Rising to prominence in late 2014, BLM is a loosely-coordinated, nationwide movement dedicated to ending police brutality. It takes its name from a hashtag started by three Black feminist activists--Patrisse Cullors, Alicia Garza, and Opal Tometi—but the movement and the hashtag are not synonymous. BLM has achieved national prominence

through their online and offline organizing, obtaining extensive news media coverage and widespread public recognition (Pew Research Center, 2016). Participants have cited the importance of social media in helping them pursue their goals (Jackson and Welles, 2016; Stephen, 2015).

BLM is important to study for several reasons. First, it qualifies as an “organizationally enabled network” in Bennett and Segerberg’s (2013) typology of connective action. It operates both online and in the streets, with much of the coordination being handled by formal organizations such as Million Hoodies, the Black Youth Project, and Ferguson Action. But these organizations **do** not directly control the movement—rather, they **are** among many groups and individuals that **help** plan and organize protests and activist messaging. Second, the movement has succeeded in shifting police brutality from the margins of American politics to a much more prominent position. Our analysis strongly suggests that the movement and the news media, rather than the elites who usually control the political agenda, drove this shift. Third, BLM serves as an apt case to test the influence of social media activism on policy goals. Unlike the Arab Spring uprisings and Occupy, which were short on policy demands, BLM’s core demand is simple: “stop killing us” (Kang, 2015). And while other policy-oriented movements such as the anti-SOPA/PIPA campaign have used social media heavily (Benkler et al., 2015), many of these are relatively short-term affairs. Finally, this study adds to a small but growing collection of studies analyzing BLM and recent anti-police brutality protests in the US (Anderson and Hitlin, 2016; Bonilla and Rosa, 2015; De Choudhury et al., 2016; Gallagher et al., 2016; Jackson and Welles, 2015, 2016; Kelley, 2015; LeFebvre and Armstrong, 2016; Olteanu et al., 2015).

Research questions

This article will undertake two empirical tasks: 1) measuring social media power using the metrics described above, and 2) testing for associations between them and elite attention to police killings of unarmed Black citizens. Our dataset features three communities: one connective social movement (BLM), one counter-movement (Political Conservatives, or PC), and one unaligned community (Mainstream News, or MN). There is little theoretical basis for predicting how these communities are likely to differ from one another on each individual metric, or which metrics are likely to best predict elite response. If we consider social movements as issue publics strongly interested in their defining issue (Krosnick, 1990), we might conjecture that they would exercise the most power in conversations on that issue. However, strong interest does not guarantee strength—if movement opponents have greater access to the mass media or politicians, for example, they may be able to overwhelm even highly enthusiastic activists. It is also conceivable that mainstream news outlets could draw large numbers of united onlookers at times when major stories break. The phrasing of the following research questions reflects these uncertainties:

- RQ1: How do the three communities compare on each of the three power metrics, and how do these comparisons change over time?
- RQ2: How well does each community predict elite response?
- RQ3: How well does each metric predict elite response?
- RQ4: How often does each community's distinct users and hashtags appear in elite tweets?

Data and methods

This study analyzes Twitter data pertaining to police brutality. We purchased from Twitter all public tweets posted during the yearlong period between June 1, 2014 and May 31, 2015 containing at least one of 45 keywords related to BLM and police killings of Black people **that some perceived as unjustified** (see Table 1).⁵ The keywords consist mostly of the full and hashtagged names of 20 Black individuals killed by police in 2014 and 2015. The resulting dataset contains 40,815,975 tweets posted by 4,435,217 unique users.

[Table 1 here]

The names in Table 1 were collated from two sources: a series of tweets posted by the NAACP Legal Defense Fund's Twitter account (@naacp_ldf) on December 3, 2014 containing the names of unarmed Black people killed by police between 1999 and 2014; and a May 1, 2015 BuzzFeed article listing a number of unarmed Black males killed by police in 2014 and 2015 (Quah and David, 2015).⁶ Neither of these lists is necessarily complete, but they were the most comprehensive we could find. From the NAACP list we extracted all of the 2014 names, and from the BuzzFeed list we extracted all names except two, resulting in a combined total of 20 names.⁷ To these keywords we added the hashtags #blacklivesmatter and #ferguson (the birthplace of the movement) and the phrase "black lives matter."

We analyzed these tweets using Python and R scripts written by the first author. We included only tweets posted between August 8, 2014 (the day before Michael Brown was killed) and May 31, 2015 (the end of our data collection period) because many of the tweets posted before this period were false positives (e.g. referencing other individuals named Michael Brown). This 297-day period includes 99.4% of all tweets and 99.1% of all unique users in the full dataset (40,563,224 tweets; 4,393,926 users).

Next came the task of identifying the like-minded communities on which our analysis is based. While small-scale studies have identified social media communities manually (Adamic and Glance, 2005; Hargittai et al., 2008), unsupervised network community detection algorithms are more effective for larger datasets (Aragón et al., 2013; Conover et al., 2011). However, most of these methods only generate cross-sectional communities. Author (Date) describes a method of tracking network communities over a period of months, but it is not effective for smaller time units. Hence we introduce a novel method of identifying and tracking social media communities that is equally effective for all time units.

We began by creating a distinct retweet-based network for each of the 42 weeks of our dataset, as retweets have been observed to signify ideological affinity in politically-oriented Twitter networks (Aragón et al., 2013; Bode et al., 2015; Conover et al., 2011). Next, we used the Louvain community-detection algorithm (Blondel et al., 2008) to separate each network into a set of communities characterized by dense retweeting patterns. When applied to large Twitter networks, Louvain creates a small number of large communities and a large number of small communities, many of which consist of one or two users retweeting one another. Within each week, we retained the ten largest communities, as these are the ones most likely to represent politically consequential constituencies. This resulted in 420 retweet-based network communities, ten for each week.

The next major step was to separate the communities into categories based on membership similarity. To do this, we used Latent Dirichlet allocation (LDA), a popular method of unsupervised machine classification (also known as topic modeling). We created a document-by-term matrix in which the documents were communities and the terms were users to serve as the input. Each user was weighted by network in-degree so that users who were retweeted more

often were considered proportionally more important in the topic-generation process. Based on this input data, LDA created a series of *topics*, or collections of network communities with similar memberships. Because LDA requires researchers to set the number of topics (k) manually, and because there are no universal rules for choosing k , we ran LDA on our data ten separate times using k values ranging from four to 13. Next, we qualitatively identified three *clusters* of topics with similar sets of prominent participants across the ten LDA runs: one representing BLM (present in all ten runs), one representing mainstream news (“MN,” present in eight runs), and one representing political conservatives (“PC,” present in seven runs). These were by far the most frequently-recurring topics we could identify.

These clusters still needed some winnowing down, in part because certain participants appeared in more than one cluster. We discarded all communities that appeared in fewer than half of each cluster’s topics to ensure that only communities that were consistently placed together in the same topic were retained. We then placed participants appearing in multiple clusters into the cluster in which they appeared most often, discarding all those that appeared in at least two clusters equally often.

This entire process yielded three *persistent communities*—one representing BLM, one representing MN, and one representing PC—whose participants were consistently grouped together. As Table 2 shows, the PC community is by far the smallest, while BLM is only slightly larger than MN. The ten most-retweeted users in each community demonstrate the face validity of our method: all those in MN are institutional accounts for mainstream news outlets (including @blackvoices, which is operated by the Huffington Post). Most of the top PC users are conservative journalists and pundits, while BLM is dominated by anti-brutality activists, most of whom are Black. The three communities overlap a great deal in terms of hashtag use; #ferguson

is the most-commonly used hashtag across all three, and two other hashtags are also present in each community's top five (*#blacklivesmatter* and *#mikebrown*).

[Table 2 here]

Our main predictor variables are unity, numbers, and commitment measured on a per-day, per-community basis. To measure our main outcome variable, elite response, we manually compiled a list of the Twitter screen names (where available) of the following elected and appointed government officials:

- The US President, the First Lady, and all Cabinet members;
- The official accounts of the Cabinet-level federal agencies (Justice, Labor, State, etc.);
- All House and Senate members of the 113th and 114th Congresses;
- All US governors in office during the dataset timespan;
- The lieutenant governors, attorneys general, and members of the state legislatures of Maryland, Missouri, New York, and Ohio, the states in which five of the six most-discussed incidents occurred;⁸
- The mayors and top local prosecutors for the cities in which the above five incidents occurred.

This list contains 1,498 screen names, of which 298 (20%) tweeted at least once in the data. These 298 users contributed 2,524 total tweets. 169 names appeared in one or another of the three persistent communities; these were removed from the communities prior to analysis.

Results

RQ1 calls for a comparison between the three communities in terms of the three metrics. We begin with numbers, the most easily-interpreted metric. Figure 1 displays the number of unique

users from each community per day. While participation from each community spikes at the same times, BLM is nearly always the largest. Interestingly, PC is usually more active than MN on non-peak days, but MN tends to surpass it when attention focuses on a major event such as a killing or a major legal decision. Participation from all three communities spike around peak periods, but this effect is stronger proportionally for MN than for PC.

[Figure 1 here]

For the unity metric, which we operationalize as the Gini coefficient of hashtags used by each community, Figure 2 reveals substantial differences. BLM is consistently more unified than PC, which is more unified than MN. In other words, BLM's hashtag use was more concentrated among a smaller number of hashtags than were the other two communities. MN's unity values also fluctuate far more than PC's or BLM's: the variance of its daily Ginis is 0.0072, while PC's variance is 0.0021 and BLM's is 0.0014. These numbers corroborate the disparities in longitudinal variation that can be seen clearly in Figure 2.

[Figure 2 here]

Figure 3 shows the longitudinal changes in each community's repeat participation rates (i.e. commitment), which are simply the proportions of unique users on any given day that post at least once during the following three days. Those paying attention primarily to mainstream news are the least committed, with RPRs that usually fall below 0.25. BLM and PC are both higher than MN during non-peak periods, with BLM usually slightly higher than PC.

[Figure 3 here]

To summarize briefly before proceeding, BLM definitively exceeds the other two communities on all three power metrics most of the time. PC generally comes in second and MN

third. Spikes in attention seem to result in sharp increases of all three metrics for all three communities.

To answer RQ2 and RQ3, we estimate Granger causalities between the nine community/metric variables and the daily number of elite tweets. Extended discussions of this technique's logic and value for communication research are available elsewhere (Bastos et al., 2015; Neuman et al., 2014), so we will not repeat them here. Instead, we offer a highly condensed description: variable X Granger-causes variable Y if past values of X enable more accurate predictions of Y than past values of Y alone. (This should not be confused with commonsense notions of causality.) Granger causality can be estimated by computing one vector autoregression model (VAR) in which prior values of outcome variable Y are the sole predictors and a second model in which prior values of an independent predictor X are added to the first model. If the ratio of the variance of the first model's error term to that of the second is sufficiently greater than one, we conclude that X Granger-causes Y. After examining models with lags of one through five days using Breusch-Godfrey tests, we chose a four-day lag for all models because it yielded the lowest levels of autocorrelation. Despite this, we were unable to completely eliminate autocorrelation in some models.

Our Granger analysis proceeds in two steps. First, we estimate bidirectional Granger causalities between daily elite tweet counts (DET) and each of the nine community/metric variables.⁹ We call these *direct* Granger causes because there are no intermediate variables between them and the outcome. Second, we examine the extent to which each of these nine variables Granger-causes one another. We call these *indirect* Granger causes. Our results are summarized in Table 3, which requires some explanation. The coefficients in the second column from the right and the second column from the left are F-statistics giving the ratio described in

the preceding paragraph, which indicate the magnitude of the reduction in error term variance occasioned by the corresponding variable. Each F-statistic is one of a pair: for the second column from the right, arrows pointing right indicate metric-to-DET Granger causality, while those pointing left indicate DET-to-metric Granger causality (i.e. reverse Granger causality). The F-statistic with the greater value in each pair is indicated in bold. The variables in the middle column are the direct Granger causes of DET, while the leftmost column contains all statistically significant indirect Granger causes and -effects.

[Table 3 here]

The first important finding Table 3 reveals is that elites are clearly following the cues of the communities, as opposed to the reverse. The magnitudes of the metric-to-DET F-statistics for all nine direct Granger causes are much greater than those of their DET-to-metric counterparts. This is clear evidence that direct Granger causality overwhelmingly runs in one direction. In answer to RQ2, comparing the direct variables, MN is the clear leader in eliciting elite responses as well as the direct cause least affected by autocorrelation. MN commitment metrics are the first-, fourth-, and sixth-strongest direct Granger causes of DET; while BLM's are the second-, fifth-, and eighth-strongest. PC exerts the weakest direct influence, coming in at third, seventh, and ninth place.

Examining the significant forward and reverse indirect Granger causes generally supports this story. One concurring finding is that the significant indirect causes of MN commitment are fairly modest, with BLM clearly stronger than PC. But MN commitment is a much stronger cause of BLM numbers than the opposite, which may indicate the power of the media to draw users to BLM during periods of high attention. MN commitment is a weaker, though still

relatively strong, cause of PC numbers. BLM commitment seems to exert some influence on MN numbers, but the equation is autocorrelated, reducing confidence in its coefficient.

Turning to RQ3, which concerns the metrics of the greatest predictive capacity, commitment emerges as the leader. The commitment metrics are the top three direct Granger causes and 14 of the 23 significant indirect ones. Also, the commitment metrics are some of the strongest indirect Granger causes by F-statistic magnitude. However, the second and third-ranked direct Granger causes were produced from autocorrelated VARs, which is also the case for two of the 14 indirect causes. As for the other power metrics, numbers is a stronger direct Granger cause than unity, as it appears twice (in the fourth and fifth spots) before unity appears once. Numbers surpasses unity as an indirect cause, appearing twice as often and generally with slightly higher F-statistics.

RQ4, which concerns how often each community's users and hashtags appear in elite tweets, can be answered using basic computational techniques. For users, we simply counted the numbers of unique and total screen names mentioned by elites that belonged to each persistent community. Figure 4 shows that BLM users are mentioned more often by elites whether uniques or totals are considered. MN users are mentioned slightly more often than PC in each case.

[Figure 4 here]

The hashtag analysis is more complicated because the most popular hashtags were used extensively by all three communities (see Table 2). Therefore, we created a list of hashtags used disproportionately more often by each community compared to the other two. We call these each community's *distinctive hashtags*. To ensure that our results were not specific to the choice of a single disproportion constant, we used two, examining hashtags used by a given community's

participants in proportions at least 1.5 times and two times greater than the other two. Figure 5 shows how often elites used each community's distinctive hashtags at both the x1.5 and x2 levels. In both cases, BLM achieves only a slight advantage over the next-ranked community. But at the x1.5 level, PC comes in second, while MN occupies that position at the x2 level. Figure 5 suggests that although all three communities are sensitive to the choice of disproportion constant, BLM's presence is felt most consistently.

[Figure 5 here]

Discussion

This study presents convergent, highly suggestive evidence of power as projected through social media by a connective social movement and two competing communities. It introduces three theoretically-derived, movement-relevant metrics of social media power, measures them longitudinally over the course of nearly 300 days, and estimates the extent to which they Granger-cause elite responses. Our results indicate that unaligned news outlets and their audiences are more successful than the other two communities in provoking elite responses. We also find modest but convergent evidence that BLM helped to generate the media attention in the first place (we further substantiate this claim in Freelon et al., 2016).

These results contribute a novel answer to a central question in the literature on digitally-enabled social movements: how, if at all, does social media use contribute to movement goals? We demonstrate for the first time that social movement activity through social media can help attract elite attention as their concerns are broadcast through mainstream news outlets. The effect is stronger than that of the total number of individuals engaged on the issue at any given time and of unity as measured through hashtag use. The finding is consistent with evidence that offline

activism can influence elites through media coverage (Vliegenthart and Walgrave, 2012; Walgrave and Vliegenthart, 2012).

While Granger causality is not “true” causality, our method definitively fulfills two of the three criteria for causal inference and substantially, although incompletely, addresses the third. Causal inference is widely considered to be valid when three criteria obtain: correlation, time precedence, and the elimination of alternative explanations (Babbie, 2012: 93–94; Vogt and Johnson, 2015: 55). Granger causality demonstrates correlation through the vector autoregression models on which it is based, and time precedence through its use of lagged predictors. And while it cannot eliminate all potential rival explanations, it can account for some of the most obvious ones. First, reversing the Granger-causal order of each pair of variables tests for the presence of reverse and bidirectional causation. While this occurs to some degree in our results, in most instances Granger causality is much stronger in one direction than in the other. Second, we test the possibility that nine variables may directly Granger-cause DET, some of which turn out to be much more consequential than others. Third, we examine indirect causes to account for the possibility of a multistep causal process. These measures add additional support, though not definitive proof, of a probabilistic causal interpretation.

That said, the absence of non-Twitter variables is this article’s chief limitation, and may have caused some of the autocorrelation in the VAR models. It is likely that the political elites were motivated to speak out on this issue through a number of channels, with Twitter being only one. Other media channels, letters and phone calls from constituents, conversations with colleagues, and events occurring in one’s district are a few of the plausible possibilities. But the fact that mainstream news outlets sourced much of their reporting on police killings in 2014 and 2015 from social media (Author, Date) supports our multistep model of online protest power. It

is also impossible to completely separate the influence of offline protests from protest tweets given that they spiked around the same times. However, elites' extensive use of relevant hashtags and mentions of movement-associated participants and media outlets support the notion that the tweets had some impact. Further research that includes additional variables may well discover new causes.

The metrics of social media power we have introduced may exhibit predictive power in other studies, but they are interesting in and of themselves. Although unity (as operationalized through hashtags) proved to be the least powerful metric in our Granger analysis, it has the potential to contribute to the voluminous literature on collective action frames (Corrigan-Brown and Wilkes, 2011; Sanfilippo et al., 2008; Snow et al., 1986). Given its long-standing status as a key concern in studies of offline protest (McCarthy et al., 1996; Soule and Earl, 2005), numbers will likely remain so in social media contexts. And among its other potential uses, commitment in the form of repeat participation rates is a new method of examining "serial activism" in social media (Bastos et al., 2013; Bastos and Mercea, 2015).

We also contribute a computationally tractable method of identifying and tracking distinct Twitter communities over time. While community detection is relatively straightforward for cross-sectional research, it is far less so for longitudinal studies. As a result, the predominant cross-sectional approaches typically used in network studies have been unable to analyze much of theoretical interest in social media, which generate inherently longitudinal data. Our method creates persistent communities drawn from the entire dataset whose variables (unity, numbers, commitment, etc.) can be measured at any desired level of time granularity. Its utility is not limited to the study of social media power: it can be applied to any large-N Twitter conversation in which multiple distinct communities participate.

This study's two main contributions go hand in hand: a falsifiable model of social media power as exercised by social movements and others interested in a given issue; and an innovative methodology for measuring it. Future studies may use our methods to investigate the extent to which the model applies to other social movements. We might expect that connective movements with similar characteristics to BLM—situated within an advanced democracy, led by marginalized but tech-savvy youth, and eager for policy change—may use Twitter to similar effect. But it may or may not apply equally well to other platforms or types of movements. Nevertheless, the finding that social movements can, under certain circumstances, further policy-relevant goals directly through tweeting is one with powerful theoretical and practical implications.

References

- Adamic LA and Glance N (2005) The political blogosphere and the 2004 US election: divided they blog. In: *Proceedings of the 3rd international workshop on Link discovery*, New York, NY: ACM, pp. 36–43.
- Agbaria AK and Mustafa M (2012) Two states for three peoples: the ‘Palestinian-Israeli’ in the Future Vision Documents of the Palestinians in Israel. *Ethnic and Racial Studies* 35(4): 718–736.
- Anderson M and Hitlin P (2016) *Social media conversations about race: How social media users see, share and discuss race and the rise of hashtags like #BlackLivesMatter*. Pew Research Center. Available from: http://www.pewinternet.org/files/2016/08/PI_2016.08.15_Race-and-Social-Media_FINAL.pdf.
- Aouragh M (2012) Social Media, Mediation and the Arab Revolutions. *tripleC: Communication, Capitalism & Critique. Open Access Journal for a Global Sustainable Information Society* 10(2): 518–536.
- Aragón P, Kappler KE, Kaltenbrunner A, et al. (2013) Communication dynamics in twitter during political campaigns: The case of the 2011 Spanish national election. *Policy & Internet* 5(2): 183–206.
- Babbie E (2012) *The Practice of Social Research*. Belmont, CA: Cengage Learning.
- Barnhurst KG (2005) News ideology in the twentieth century. In: Horst Pöttker and Svennik Høyer (eds), *Diffusion of the News Paradigm 1850-2000*, Gothenburg, Sweden: Nordicom, pp. 239–262.
- Bastos MT and Mercea D (2015) Serial activists: Political Twitter beyond influentials and the twittertariat. *New Media & Society*: 1461444815584764.
- Bastos MT, Puschmann C and Travitzki R (2013) Tweeting Across Hashtags: Overlapping Users and the Importance of Language, Topics, and Politics. In: *Proceedings of the 24th ACM Conference on Hypertext and Social Media*, HT ’13, New York, NY, USA: ACM, pp. 164–168. Available from: <http://doi.acm.org/10.1145/2481492.2481510> (accessed 25 June 2015).
- Bastos MT, Mercea D and Charpentier A (2015) Tents, Tweets, and Events: The Interplay Between Ongoing Protests and Social Media. *Journal of Communication* 65(2): 320–350.
- Benkler Y, Roberts H, Faris R, et al. (2015) Social Mobilization and the Networked Public Sphere: Mapping the SOPA-PIPA Debate. *Political Communication* 32(4): 594–624.
- Bennett WL (2003) New media power. In: Couldry N and Curran J (eds), *Contesting media power*, Lanham, MD: Rowman & Littlefield, pp. 17–37.
- Bennett WL and Segerberg A (2012) The Logic of Connective Action: Digital media and the personalization of contentious politics. *Information, Communication & Society* 15(5): 739–768.

- Bennett WL and Segerberg A (2013) *The Logic of Connective Action: Digital Media and the Personalization of Contentious Politics*. Cambridge University Press.
- Blondel VD, Guillaume J-L, Lambiotte R, et al. (2008) Fast unfolding of communities in large networks. *Journal of Statistical Mechanics: Theory and Experiment* 2008(10): P10008.
- Bode L, Hanna A, Yang J, et al. (2015) Candidate Networks, Citizen Clusters, and Political Expression Strategic Hashtag Use in the 2010 Midterms. *The ANNALS of the American Academy of Political and Social Science* 659(1): 149–165.
- Bonilla Y and Rosa J (2015) #Ferguson: Digital protest, hashtag ethnography, and the racial politics of social media in the United States. *American Ethnologist* 42(1): 4–17.
- Brulle RJ (2014) Institutionalizing delay: foundation funding and the creation of U.S. climate change counter-movement organizations. *Climatic Change* 122(4): 681–694.
- Burbidge JB, Magee L and Robb AL (1988) Alternative Transformations to Handle Extreme Values of the Dependent Variable. *Journal of the American Statistical Association* 83(401): 123–127.
- Button JW (1989) *Blacks and Social Change: Impact of the Civil Rights Movement in Southern Communities*. Princeton, NJ: Princeton University Press.
- Castells M (2012) *Networks of outrage and hope: Social movements in the Internet age*. Malden, MA: Polity.
- Conover MD, Ratkiewicz J, Francisco M, et al. (2011) Political polarization on twitter. In: *Proc. 5th Intl. Conference on Weblogs and Social Media*, Barcelona, Spain: AAAI, pp. 89–96.
- Corrigall-Brown C and Wilkes R (2011) Picturing Protest: The Visual Framing of Collective Action by First Nations in Canada. *American Behavioral Scientist*: 2764211419357.
- Couldry N and Curran J (2003) The paradox of media power. In: Couldry N and Curran J (eds), *Contesting media power: Alternative media in a networked world*, Lanham, MD: Rowman & Littlefield, pp. 3–15.
- Croeser S and Highfield T (2014) Occupy Oakland and #oo: Uses of Twitter within the Occupy movement. *First Monday* 19(3). Available from: <http://ojphi.org/ojs/index.php/fm/article/view/4827> (accessed 16 June 2016).
- De Choudhury M, Jhaver S, Sugar B, et al. (2016) Social Media Participation in an Activist Movement for Racial Equality. In: *Tenth International AAAI Conference on Web and Social Media*. Available from: http://www.munmund.net/pubs/BLM_ICWSM16.pdf (accessed 15 June 2016).
- Diani M (1992) The concept of social movement. *The Sociological Review* 40(1): 1–25.
- Earl J and Kimport K (2011) *Digitally Enabled Social Change: Activism in the Internet Age*. Cambridge, MA: MIT Press.

- Freelon D, McIlwain C and Clark MD (2016) *Beyond the hashtags: #Blacklivesmatter, #Ferguson, and the online struggle for offline justice*. Washington, DC: Center for Media and Social Impact, American University. Available from: <http://cmsimpact.org/blmreport>.
- Gallagher RJ, Reagan AJ, Danforth CM, et al. (2016) Divergent discourse between protests and counter-protests: #BlackLivesMatter and #AllLivesMatter. *arXiv:1606.06820 [cs]*. Available from: <http://arxiv.org/abs/1606.06820> (accessed 18 September 2016).
- Giddens A (1987) *A Contemporary Critique of Historical Materialism: The nation-state and violence*. Berkeley: University of California Press.
- González-Bailón S, Borge-Holthoefer J and Moreno Y (2013) Broadcasters and Hidden Influentials in Online Protest Diffusion. *American Behavioral Scientist* 57(7): 943–965.
- Hargittai E, Gallo J and Kane M (2008) Cross-ideological discussions among conservative and liberal bloggers. *Public Choice* 134(1): 67–86.
- Jackson SJ and Welles BF (2015) Hijacking #myNYPD: Social Media Dissent and Networked Counterpublics. *Journal of Communication* 65(6): 932–952.
- Jackson SJ and Welles BF (2016) #Ferguson is everywhere: initiators in emerging counterpublic networks. *Information, Communication & Society* 19(3): 397–418.
- Jensen MJ and Bang HP (2013) Occupy Wall Street: A New Political Form of Movement and Community? *Journal of Information Technology & Politics* 10(4): 444–461.
- Kahn R and Kellner D (2004) New Media and Internet Activism: From the ‘Battle of Seattle’ to Blogging. *New Media & Society* 6(1): 87–95.
- Kang JC (2015) ‘Our Demand Is Simple: Stop Killing Us’. *The New York Times*, 4th May. Available from: <http://www.nytimes.com/2015/05/10/magazine/our-demand-is-simple-stop-killing-us.html> (accessed 22 January 2016).
- Kelley RDG (2015) Beyond Black Lives Matter. *Kalfou* 2(2). Available from: <https://tupjournals.temple.edu/index.php/kalfou/article/view/71> (accessed 4 February 2016).
- Klandermans B (1997) *The Social Psychology of Protest*. 1 edition. Oxford, UK ; Cambridge, Mass: Wiley-Blackwell.
- Krosnick JA (1990) Government policy and citizen passion: A study of issue publics in contemporary America. *Political Behavior* 12(1): 59–92.
- LeFebvre RK and Armstrong C (2016) Grievance-based social movement mobilization in the #Ferguson Twitter storm. *New Media & Society*: 1461444816644697.
- Mann L (1974) Counting the Crowd: Effects of Editorial Policy on Estimates. *Journalism & Mass Communication Quarterly* 51(2): 278–285.

- McCarthy JD and Zald MN (1977) Resource Mobilization and Social Movements: A Partial Theory. *American Journal of Sociology* 82(6): 1212–1241.
- McCarthy JD, McPhail C and Smith J (1996) Images of Protest: Dimensions of Selection Bias in Media Coverage of Washington Demonstrations, 1982 and 1991. *American Sociological Review* 61(3): 478–499.
- Meyer DS and Staggenborg S (1996) Movements, Countermovements, and the Structure of Political Opportunity. *American Journal of Sociology* 101(6): 1628–1660.
- Moodie-Mills D (2015) OpEd: Are We All Just One Bullet Away From Becoming a Hashtag? *NBC News*. Available from: <http://www.nbcnews.com/storyline/walter-scott-shooting/oped-are-we-all-just-one-bullet-away-becoming-hashtag-n338111> (accessed 27 February 2016).
- Movement for Black Lives (2016) A vision for black lives: Policy demands for black power, freedom & justice. Available from: <https://policy.m4bl.org/>.
- Neuman WR, Guggenheim L, Jang SM, et al. (2014) The Dynamics of Public Attention: Agenda-Setting Theory Meets Big Data. *Journal of Communication* 64(2): 193–214.
- Olteanu A, Weber I and Gatica-Perez D (2015) Characterizing the Demographics Behind the #BlackLivesMatter Movement. *arXiv:1512.05671 [cs]*. Available from: <http://arxiv.org/abs/1512.05671> (accessed 19 January 2016).
- Pew Research Center (2016) *On Views of Race and Inequality, Blacks and Whites Are Worlds Apart*. Pew Research Center. Available from: http://www.pewsocialtrends.org/files/2016/06/ST_2016.06.27_Race-Inequality-Final.pdf.
- Quah N and David LE (2015) Here's A Timeline Of Unarmed Black People Killed By Police Over Past Year. *BuzzFeed*. Available from: <http://www.buzzfeed.com/nicholasquah/heres-a-timeline-of-unarmed-black-men-killed-by-police-over> (accessed 22 January 2016).
- Reese SD (1990) The news paradigm and the ideology of objectivity: A socialist at the Wall Street Journal. *Critical Studies in Media Communication* 7(4): 390–409.
- Sanfilippo A, Franklin L, Tratz S, et al. (2008) Automating Frame Analysis. In: Liu H, Salerno JJ, and Young MJ (eds), *Social Computing, Behavioral Modeling, and Prediction*, Springer US, pp. 239–248.
- Shirky C (2011) The Political Power of Social Media: Technology, the Public Sphere Sphere, and Political Change. *Foreign Affairs* 90: 28.
- Smith M (2014) Affect and Respectability Politics. *Theory & Event* 17(3). Available from: <https://muse-jhu-edu.proxyau.wrlc.org/article/559376> (accessed 18 September 2016).
- Snow DA, Rochford Jr EB, Worden SK, et al. (1986) Frame alignment processes, micromobilization, and movement participation. *American sociological review*: 464–481.
- Soule S and Earl J (2005) A Movement Society Evaluated: Collective Protest in The United States, 1960–1986. *Mobilization: An International Quarterly* 10(3): 345–364.

- Stephen B (2015) How Black Lives Matter Uses Social Media to Fight the Power. *WIRED*. Available from: <http://www.wired.com/2015/10/how-black-lives-matter-uses-social-media-to-fight-the-power/> (accessed 27 February 2016).
- Tarrow S (1998) 'The Very Excess of Democracy': State building and contentious politics in America. In: Costain AN and McFarland AS (eds), *Social movements and American political institutions*, Lanham, MD: Rowman & Littlefield, pp. 20–38.
- Theocharis Y (2013) The Wealth of (Occupation) Networks? Communication Patterns and Information Distribution in a Twitter Protest Network. *Journal of Information Technology & Politics* 10(1): 35–56.
- Tilly C (1999) Conclusion: From Interactions to Outcomes in Social Movements. In: Giugni M, McAdam D, and Tilly C (eds), *How Social Movements Matter*, Minneapolis: University of Minnesota Press, pp. 253–270.
- Tilly C and Wood LJ (2013) *Social Movements 1768-2012*. Routledge.
- Tufekci Z and Wilson C (2012) Social Media and the Decision to Participate in Political Protest: Observations From Tahrir Square. *Journal of Communication* 62(2): 363–379.
- Valenzuela S, Arriagada A and Scherman A (2014) Facebook, Twitter, and Youth Engagement: A Quasi-experimental Study of Social Media Use and Protest Behavior Using Propensity Score Matching. *International Journal of Communication* 8(0): 25.
- Vliegenthart R and Walgrave S (2012) The interdependency of mass media and social movements. In: Semetko H and Scammell M (eds), *The sage handbook of political communication*, Thousand Oaks, CA: Sage, pp. 387–398.
- Vogt WP and Johnson RB (2015) *The SAGE Dictionary of Statistics & Methodology: A Nontechnical Guide for the Social Sciences*. SAGE Publications.
- Walgrave S and Vliegenthart R (2012) The Complex Agenda-Setting Power of Protest: Demonstrations, Media, Parliament, Government, and Legislation in Belgium, 1993-2000. *Mobilization: An International Quarterly* 17(2): 129–156.

Table 1: Twitter keywords

| Keyword |
|-------------------------------------|
| #ferguson |
| “michael brown”/“mike brown”/ |
| #michaelbrown/#mikebrown |
| #blacklivesmatter |
| “eric garner”/#ericgarner |
| “freddie gray”/#freddiegray |
| “walter scott”/#walterscott |
| “tamir rice”/#tamirrice |
| “black lives matter” |
| “john crawford”/#johncrawford |
| “tony robinson”/#tonyrobinson |
| “eric harris”/#ericharris |
| “ezell ford”/#ezellford |
| “akai gurley”/#akaigurley |
| “kajieme powell”/#kajiemepowell |
| “tanisha anderson”/#tanishaanderson |
| “victor white”/#victorwhite |
| “jordan baker”/#jordanbaker |
| “jerame reid”/#jeramereid |
| “yvette smith”/#yvettesmith |
| “phillip white”/#philipwhite |
| “dante parker”/#danteparker |
| “mckenzie cochran”/#mckenziecochran |
| “tyree woodson”/#tyreewoodson |

Table 2: Summary statistics for three persistent communities

| | | BLM (Black Lives Matter) | MN (Mainstream News) | PC (Political Conservatives) |
|-------------------------------|------|---------------------------------|-----------------------------|-------------------------------------|
| Total users | | 470,655 | 422,651 | 260,111 |
| | Rank | | | |
| Most-retweeted users | 1 | @deray | @cnnbrk | @chris_1791 |
| | 2 | @shaunking | @cnn | @waynedupreeshow |
| | 3 | @bipartisanism | @nbcnews | @foxnews |
| | 4 | @antoniofrench | @abc | @breaking911 |
| | 5 | @nettaaaaaaaa | @bbcbreaking | @larryelder |
| | 6 | @michaelskolnik | @nytimes | @fox2now |
| | 7 | @ryanjreilly | @rt_com | @ksdknews |
| | 8 | @talbertswan | @huffingtonpost | @amymek |
| | 9 | @bassem_masri | @ac360 | @patdollard |
| | 10 | @youranonglobal | @blackvoices | @michaelcalhoun |
| Most frequently-used hashtags | 1 | #ferguson | #ferguson | #ferguson |
| | 2 | #blacklivesmatter | #ericgarner | #tcot |
| | 3 | #mikebrown | #blacklivesmatter | #blacklivesmatter |
| | 4 | #ericgarner | #mikebrown | #mikebrown |
| | 5 | #freddiegray | #freddiegray | #michaelbrown |
| | 6 | #icantbreathe | #michaelbrown | #ericgarner |
| | 7 | #michaelbrown | #walterscott | #freddiegray |
| | 8 | #tamirrice | #icantbreathe | #darrenwilson |
| | 9 | #walterscott | #baltimore | #baltimore |
| | 10 | #baltimore | #fergusondecision | #pjnet |

Table 3: Direct and indirect Granger causes of daily elite tweets ($p=4$)

| Indirect Granger cause | Indirect F | Direct Granger cause | Direct F |
|-----------------------------|--------------------------|-----------------------------|--------------------------|
| BLM numbers | → 4.32 ← 71.1 | MN commitment | → 24.52 ← 0.98 |
| BLM commitment | → 3.79 ← 2.71 | | |
| PC numbers | → 3.02 ← 42.7 | | |
| MN numbers [†] | → 7.32 ← 83.9 | BLM commitment [†] | → 17.56 ← 0.69 |
| MN unity | → 7.34 ← 47.1 | | |
| MN commitment | → 2.71 ← 3.79 | | |
| MN numbers | → 2.87 ← 28.3 | PC commitment [†] | → 10.04 ← 1.04 |
| BLM commitment [†] | → 83.9 ← 7.33 | MN numbers | → 9.80 ← 1.02 |
| PC commitment | → 28.3 ← 2.87 | | |
| BLM numbers [†] | → 3.34 ← 1.46 | | |
| BLM unity [†] | → 2.41 ← 0.44 | | |
| MN commitment | → 71.1 ← 4.33 | BLM numbers | → 8.19 ← 0.57 |
| PC commitment | → 22.0 ← 1.90 | | |
| BLM commitment | → 47.14 ← 7.34 | MN unity | → 6.50 ← 0.67 |
| PC commitment | → 10.93 ← 1.10 | | |
| BLM unity [†] | → 3.14 ← 0.26 | | |
| BLM numbers | → 2.59 ← 0.32 | | |
| MN commitment | → 42.66 ← 3.02 | PC numbers | → 5.01 ← 1.35 |
| BLM commitment | → 32.76 ← 1.83 | | |
| MN commitment [†] | → 16.34 ← 1.51 | BLM unity | → 4.74 ← 0.89 |
| PC commitment | → 7.98 ← 2.09 | | |
| BLM commitment | → 15.04 ← 0.62 | PC unity [†] | → 2.97 ← 1.40 |
| MN commitment | → 10.75 ← 0.72 | | |

DET (Daily elite tweets)

F values above 2.4 = $p < 0.05$; above 3.35 = $p < 0.01$; above 4.73 = $p < 0.001$. Daggers indicate autocorrelated equations (Breusch-Godfrey $p < 0.05$). To reduce repetition, only the significant ($p < 0.05$) indirect causes of each direct cause are listed.

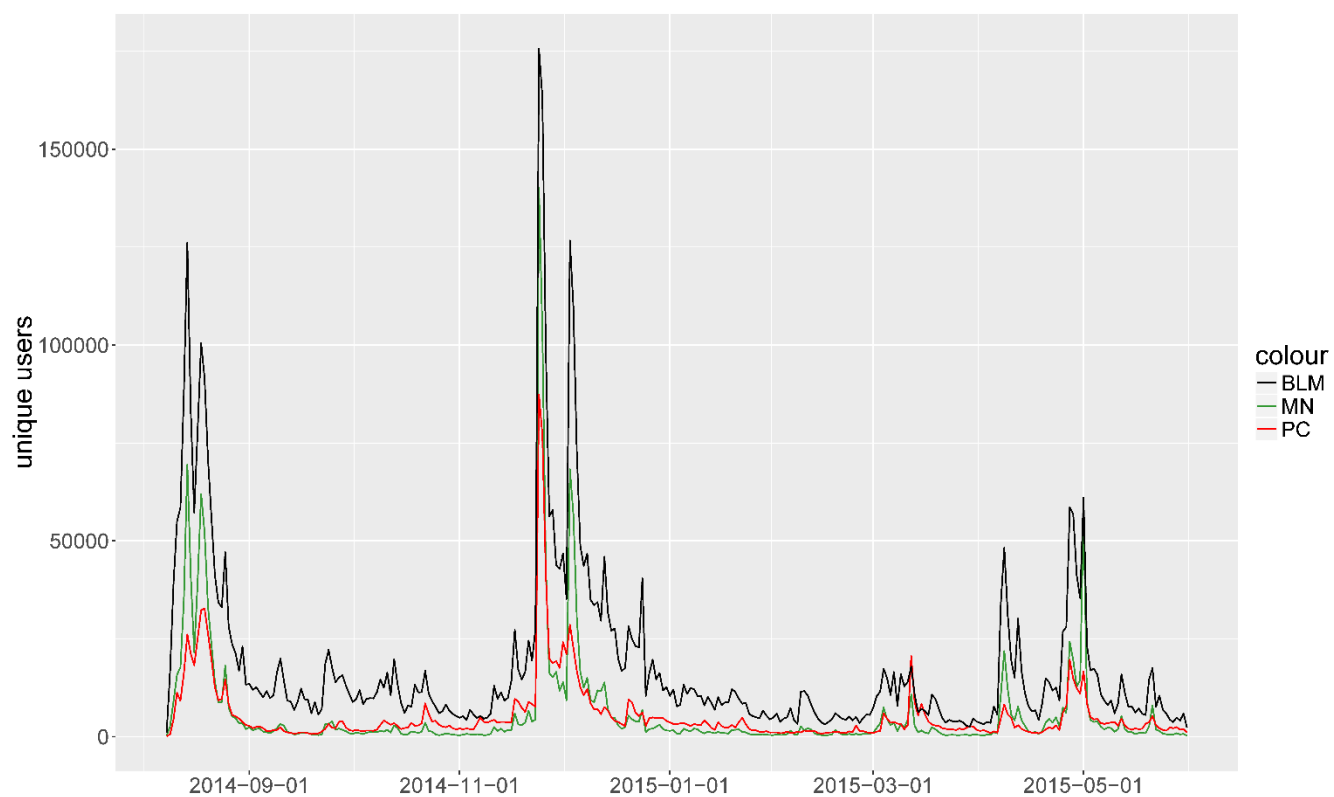


Figure 1: Numbers over time

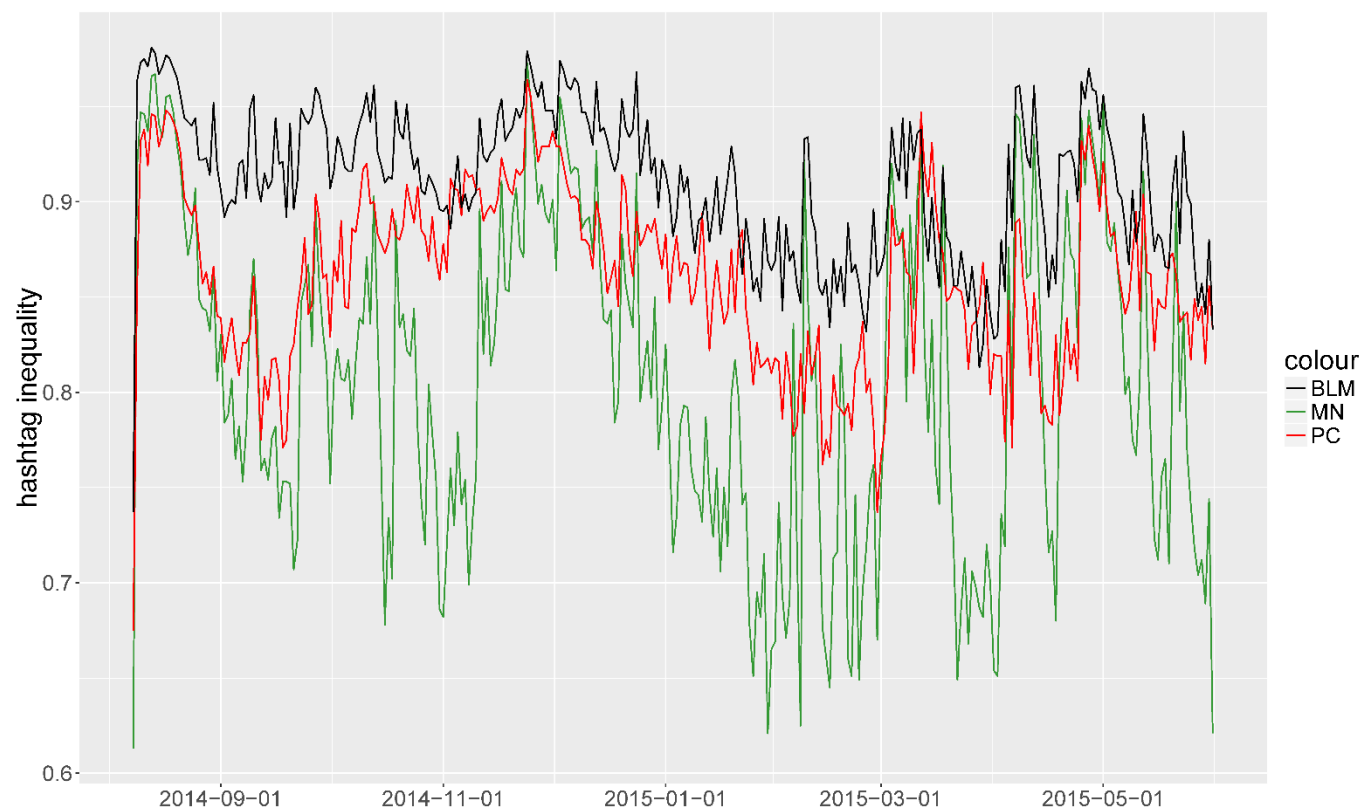


Figure 2: Unity over time

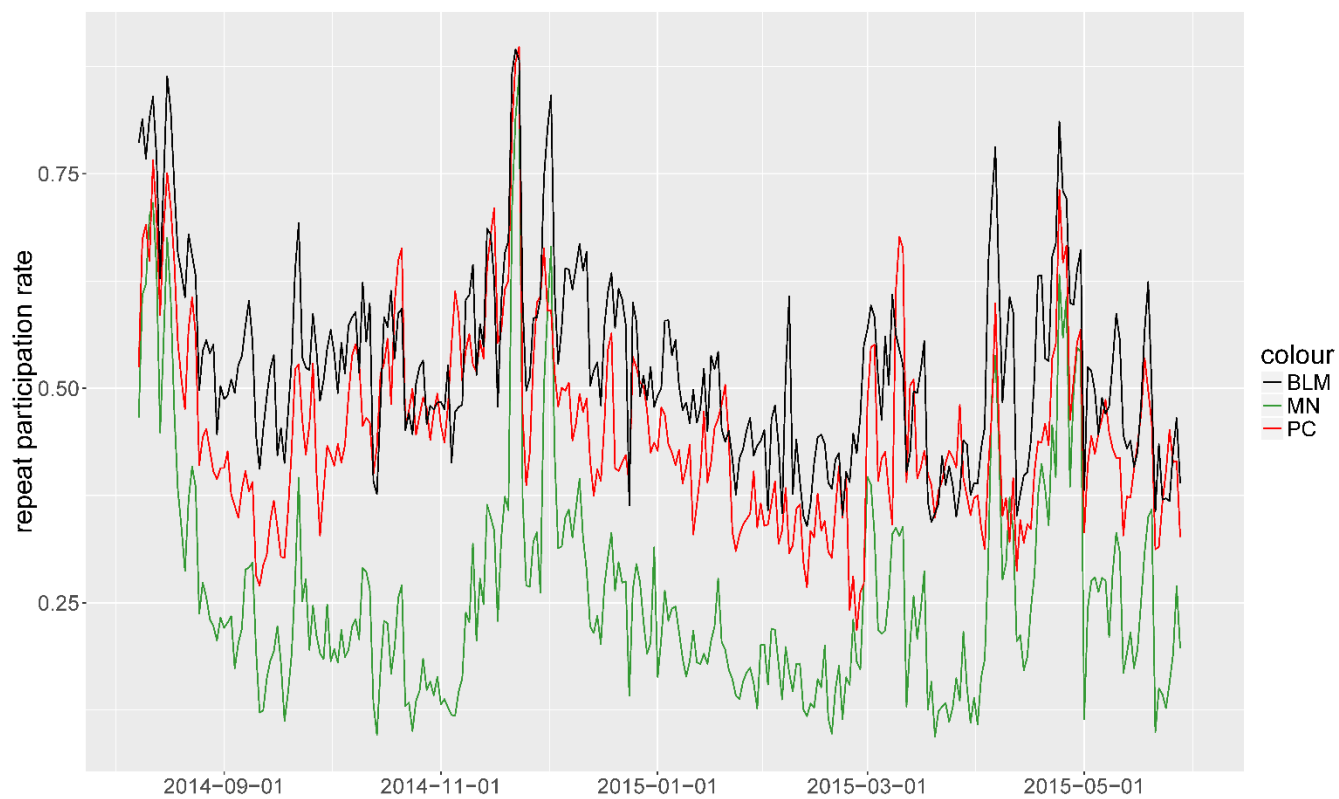


Figure 3: Commitment over time

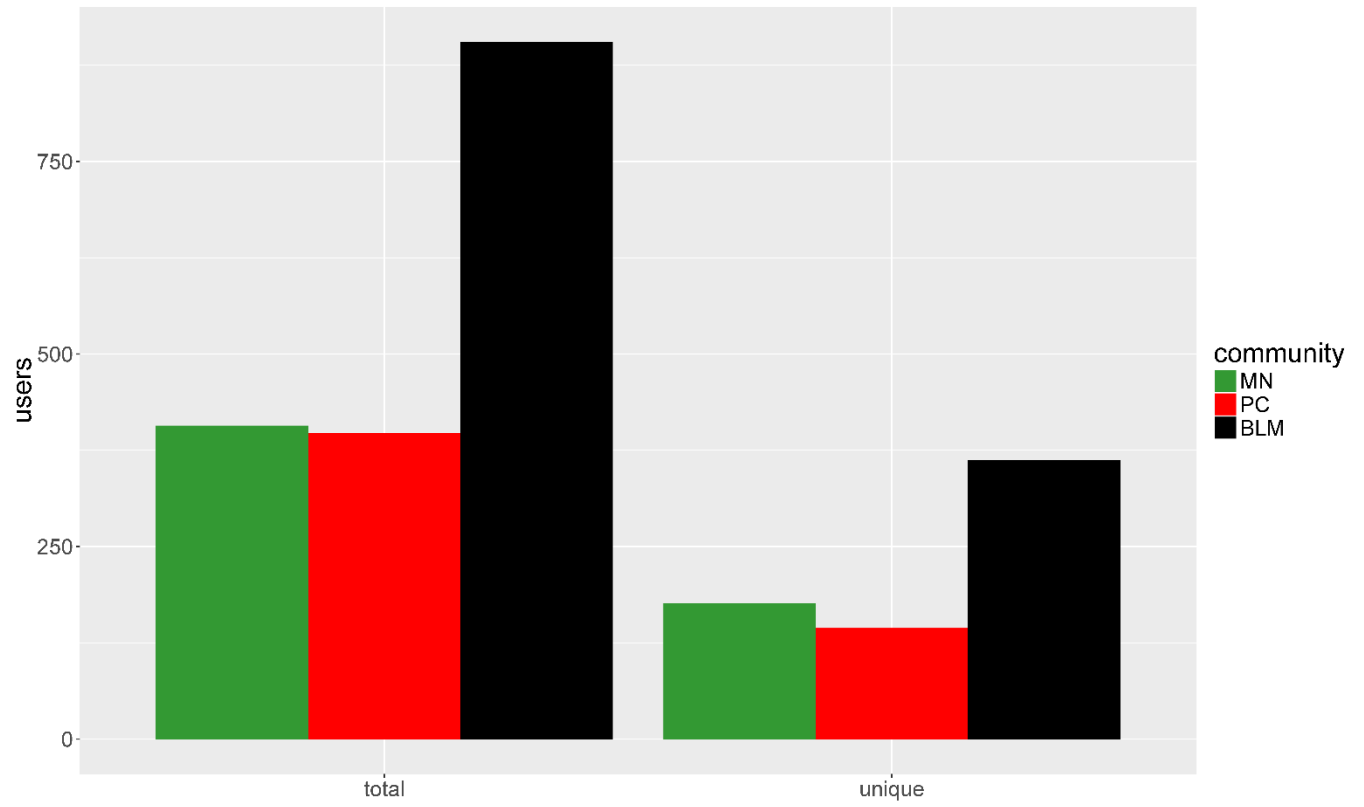


Figure 4: Total and unique users mentioned by political elites

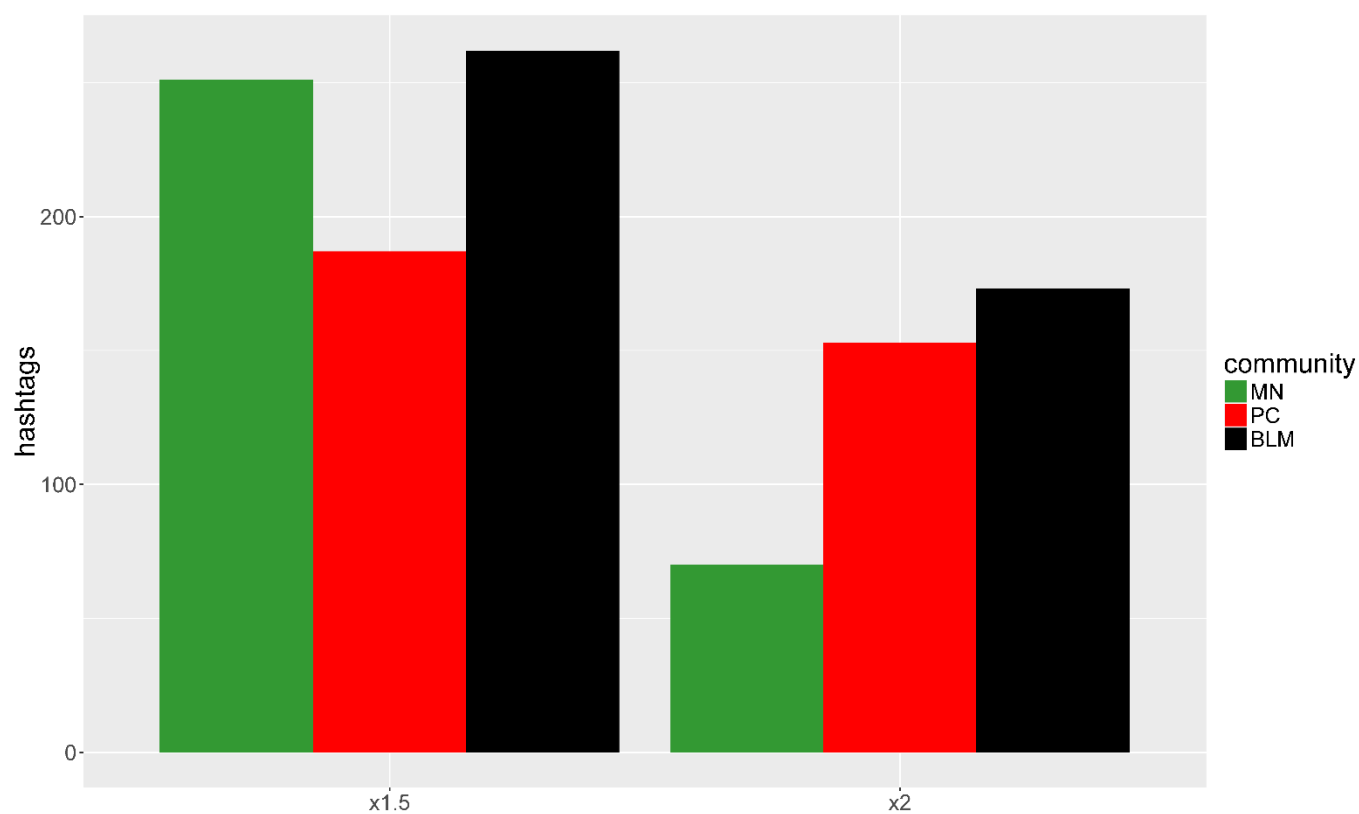


Figure 5: Distinct hashtags mentioned by political elites

¹ Counter-movements “make competing claims on the state on matters of policy and politics and vie for attention from the mass media and the broader public” (Meyer and Staggenborg, 1996: 1632) and, when opposing left-wing movements (as does the one analyzed here), “seek to maintain the currently dominant field frame and thus maintain the status quo by opposing, or countering, the efforts of movements seeking change” (Brulle, 2014: 683).

² In addition to these methodological considerations, BLM would likely condemn this conception of worthiness as counterproductive “respectability politics” (Smith, 2014).

³ See also Klandermans (1997), who concurs that “the more committed to a movement someone is, the more likely it is that he or she will continue to participate” (29).

⁴ We chose a three-day period to strike a balance between a week, which we felt would be too liberal, and one day, which would be too conservative.

⁵ Data purchased from Twitter includes all public tweets matching the buyer’s search criteria, which is not guaranteed when collecting data from the platform’s APIs (Jackson and Welles, 2016; LeFebvre and Armstrong, 2016).

⁶ The first tweet in this series is here: https://twitter.com/naacp_ldf/status/540250644658278401

⁷ Two names (Dontre Hamilton and Romain Brisbon) were omitted from our final list due to a clerical error.

⁸ We originally included the states where the five most-discussed incidents occurred, but since the fifth and sixth most-discussed incidents (Tamir Rice and John Crawford, respectively) both took place in Ohio, we decided to include it instead of South Carolina, where the fourth most-discussed incident (Walter Scott) occurred.

⁹ DET and the three numbers metrics were transformed prior to analysis using the inverse hyperbolic sine function (Burbidge et al., 1988) to satisfy the Granger method’s assumptions of normality and stationarity. The unity and commitment metrics were not transformed because they are already normalized.