

Online Issue Politicization:
**How the
Common Core and
Black Lives Matter
Discussions Evolved
on Social Media**

September 4, 2020

Authors:

*Sean Kates,
Zhanna Terechshenko,
Fridolin Linder,
Jonathan Nagler,
Richard Bonneau,
Mona Vakilifathi,
Joshua A. Tucker*



CSM_aP

CENTER FOR SOCIAL
MEDIA AND POLITICS



NYU

Online Issue Politicization: How the Common Core and Black Lives Matter Discussions Evolved on Social Media

Sean Kates^{1*}, Zhanna Terechshenko¹, Fridolin Linder², Jonathan Nagler^{1,3}, Richard Bonneau^{1,4}, Mona Vakilifathi⁵, Joshua A. Tucker^{1,3}

Abstract

Social media present an increasingly common path to issue politicization, as the distance between policy advocates and the masses is greatly reduced. In this Data Report, we analyze the discussions on Twitter of two issues (Black Lives Matter and Common Core State Standards) as they evolved over time. We show that politicization of the issues did not take the same path, and that different types of messages and senders were influential in expanding and shaping the discussions about the respective issues. For both issues, tweets by highly followed and verified users were widely shared, and contributed to a large downstream growth in the discussion. However, the substance of tweets mattered as well, with the use of angry language strongly correlated with measures of influence, alongside the important roles played by the use of hashtags. Finally, we find evidence that in the discussion around Common Core, some topics were far more important, including broaching issues of individual freedoms and personal values.

Keywords

Politicization — Social Media — Black Lives Matter — Common Core State Standards

¹Center for Social Media and Politics, New York University

²Siemens Mobility, Munich, Germany; fridolin.linder@siemens.com

³Department of Politics, New York University

⁴Departments of Biology and Computer Science, New York University

⁵Robert F. Wagner Graduate School of Public Service, New York University

*Corresponding author: sk5350@nyu.edu

Contents

	Introduction	1
1	Case Studies	2
2	Research Questions	3
3	Data and Methodology	4
3.1	Data	4
3.2	Methodological Framework	4
3.3	Timelines	4
4	Findings	5
4.1	Users	5
4.2	Retweets and Engagements	5
4.3	Hashtags	7
4.4	URLs	8
4.5	Sentiment and Tone	9
4.6	Substantive Topics	10
4.7	Influence and its Predictors	11
5	Conclusion	14
	Acknowledgments	14

References

15

Introduction

As social media have become an increasingly important source of news and entertainment for large swaths of the population, they also present an opportunity for the politicization of issues on a scale previously unimaginable. On the whole, social media make contact between activists and the masses far cheaper and easier.

While there has been a great deal of research on the effects of social media on political outcomes from voting to protesting, few researchers have been able to examine comprehensively how an issue becomes politicized online.¹ This is generally due to two related problems. First, and most important, is the difficulty of data collection. In order to properly study the full scope of issue politicization online, one must obtain a record of the discussion surrounding the issue that is plausibly complete or nearly so. Relying on smaller sam-

¹We use the term ‘politicized’ here to mean that the issue enters the political agenda: it becomes an issue over which political policy proposals are discussed by the general public.

ples of a larger discussion may yield inferences about what drives some forms of engagement, but make it challenging to say anything definitive about why individuals enter or leave a discussion.

Second, the process of politicization of an issue may vary depending on aspects of the issue itself. That is, some issues by their nature are more likely to be spread by and geared toward the masses, while others may remain largely elite-driven and -focused. Similarly, some issues may be more or less complex, and different entry points may be necessary depending on this complexity. Practically, this requires researchers to look at multiple issues (with nearly complete timelines for all) to speak more broadly about the dynamics of online issue politicization.

Our goal in this Data Report is to move carefully in the direction of explaining online issue politicization. We focus on two issues, the online Black Lives Matter campaign and the online discussion surrounding the adoption of the Common Core State Standards, for which we can reasonably argue we have nearly complete timelines of related posts on Twitter.

We chose these issues for both their similarities and differences. While we leave more extensive discussion of each issue to the section below, we note here that both topics' events occurred when social media, specifically Twitter, was available for elites and masses to directly engage with each other. Both topics initially lacked an existing partisan or policy issue framing, which allowed elites and masses to shape a topic's framing to influence social media users' responses online and offline. Each topic exhibits variation in social media user types and content, such that there is actual discussion, push back, and what one could broadly call "debate."

However, the issues also present clear differences. The size of discussion in Black Lives Matter is significantly larger both in total volume and in number of unique participants. Moreover, from a qualitative sense, we know that the discussions moved in opposite directions as regards consensus amid participants. While Black Lives Matter started as a tight-knit protest movement to engage the public in discussions of extralegal police violence, the discussion eventually split along racial and political lines, with counter-movements presenting with slogans such as "All/Blue Lives Matter" and constituting large parts of the online exchange. The discussion of Common Core evolved differently, with a steady focus on the perils and drawbacks of Common Core dominating all parts of the discussion. Indeed, there are very few users and tweets that might be considered "Common Core defenders".

Finally, the most obvious difference between the issues is the underlying subject matter. The discussion of Common Core is about education policy, and whether individual states and schools should be subject to national homogenized standards and testing. While there are very concrete experiences associated with it (new styles of homework and approaches to learning), it was a discussion about a specific set of policy proposals in a well defined domain (education). Black Lives Matter is far more existential in nature to a large portion of

the population. Fundamentally, the movement is about the killings of individuals belonging to a minority population by state-sanctioned law enforcement organizations. Our two discussions take place at very different pitches, with participants invested at very different fundamental levels.

In this Data Report, we cover how specific tweets were influential in shaping the discussion around both of these issues. This includes understanding the attributes of tweets that increase the politicization or volume of an online issue, including the substance, tone, and original author of the tweet.

We find that attributes of a tweet's creator are extremely powerful predictors of how influential a tweet is. In particular, verification (a "blue checkmark") is associated with large increases in immediate and long-term influence, as are increases in the numbers of followers. These results lend credence to claims that elites are more likely to be able to influence discussion on social media²

At the tweet level, we find that the substance of a tweet also matters. In both samples, we find evidence that anger in tone can be a powerful tool to increase influence, and the use of hashtags can also aid in increasing the spread of a message across the discussion. Neither the inclusion of links nor negativity in tone are associated with the same type of effects. Finally, we show that in both discussions, particular sub-topics *within* the broader discussion can wield greater or lesser influence on the entire conversation. On the whole, this Data Report allows us to better understand how discussions evolve over time, and the way they are shaped both by those participating in it, and the manner in which they are participating.

1. Case Studies

In this Data Report, we analyze two issues, and the way discussion around them evolved over time on social media. In this section, we briefly describe the issues, including the major events that shaped the movements and the most important participants.

Black Lives Matter

Black Lives Matter is a decentralized movement advocating primarily for the end of police violence against Black people. It first surfaced online with the use of the #BlackLivesMatter hashtag following the acquittal of George Zimmerman for the killing of Trayvon Martin, and rose to national prominence in the period immediately following the killing of Michael Brown and the ensuing protests in Ferguson, Missouri.

The movement as a whole has no centralized hierarchy, though the Black Lives Matter Network (BLMN) was founded in 2013 in hopes of creating a core group of principles that local chapters and proponents might use as guiding lights. These principles have developed over time, and the BLM movement has frequently partnered with other advocacy groups to lend

²In a second Data Report that we are releasing concurrently, we examine specifically the types of *individuals* who were most influential in the discussions around CCSS and BLM.

support for causes outside the primary focus on police violence against Black individuals.

From a policy perspective, the movement specifically calls for community control of the police, a reduction in the incarceration rate (especially for young, Black men), and a refocus on funding public education. The focus of the movement is on political expression through action, particularly in rallies and protests. This in-person agitation on the streets is combined with a massive social media presence in hopes of garnering attention and enacting change.

The immediate response to the movement was mixed, particularly along racial lines, and critics were quick to portray protests as riots harmful to the movement's goals. Counter-movements coined the terms "All Lives Matter" (to suggest that BLM was itself racist in its focus on Black people) and "Blue Lives Matter" (in support of police officers and their actions). However, the continuous displays of police brutality, captured on video and distributed widely through social media, spurred the movement and slowly won public opinion to its side. After George Floyd's killing in early 2020 and a new spark of the BLM campaign, polls showed that the majority of all races at least somewhat supported the movement (Pew 2020), although readers should note that 2020 is outside the time range of our current study.

Common Core State Standards

CCSS is the name given to a set of state K-12 English-language arts and mathematics standards to prepare students for college and career in the global workforce [1]. The Standards, adopted by the large majority of states, have become a source of contention inside and outside of the education policy world.

In the United States, state governments have the legal authority to regulate a state's K-12 education standards, which results in a large variation in the rigor of states' K-12 education standards and, ultimately, student achievement. The federal government has the legal authority to provide federal funding to state governments to incentivize state governments to adopt favorable policies, including K-12 education standards, upon receipt of funding.

In 2009, following the Great Recession, the federal government used the opportunities presented by crisis to tie stimulus and infrastructure funding to adoption of a core set of education policies. President Obama signed the American Recovery and Reinvestment Act of 2009 in February 2009, which created the Race to the Top Fund. This fund was set up to provide state and local governments with federal funding to invest in education reform relating to standards and assessments, data systems, teacher development, and school improvement. State governments were awarded grants on the basis of adopting a series of education reforms, including "developing and adopting common standards" and "developing and implementing common, high-quality assessments." State governments adopted suggested education reforms in order to apply for this grant and increase K-12 education revenue during the Great Recession.

The Standards were the natural extension of this plan. In order to have a consistent set of standards that states accepting the funds would adopt, the National Governors Association's Center for Best Practices and the Council of Chief State School Officers created the Standards, starting in April 2009, and finalizing them in June 2010 [2]. The Standards were touted as a product of collaboration among state government stakeholders, including governors, state education commissioners, researchers, curriculum experts, and educators [2].

In order to receive Race to the Top funds, 45 states and the District of Columbia initially adopted CCSS in 2010 and 2011 [3]. In addition, the Department of Education awarded grants to two state government coalitions to develop state assessment programs aligned with CCSS, the Partnership for Assessment of Readiness for College and Careers (PARCC) and the Smarter Balanced Assessment Consortium (SBAC), both in September 2010 [4]. 26 states initially adopted PARCC and 31 states adopted SBAC [4].

State governments implemented CCSS and PARCC or SBAC as early as the 2011-12 academic year [5]. Starting in 2010, Americans' support for CCSS decreased over time among all voters, Democrats and Republicans ([6]; [7]; [8]), with Republican voters' support declining at a larger rate than others. Democratic voters and teachers generally opposed CCSS because CCSS constrains a teacher's ability to tailor her curriculum to her students' needs, and incentivizes teachers to "teach to the test" rather than improve student achievement. In addition, CCSS represents an increase in the privatization of education or influence of philanthropic organizations in education. Republican voters, especially Tea Party Republican voters, generally opposed CCSS because CCSS represents an increase in the size of the federal government or the federal government's role in K-12 education and a threat to Americans' freedom and culture. In addition, the Standards were tied to President Obama, a Democrat widely distrusted by Republicans. This general dissatisfaction with the standards spread widely, and the Standards became a controversial topic in the 2016 presidential election, especially among Republican presidential candidates [9]. As a result of the public backlash, by the 2018-19 academic year, only 41 states and the District of Columbia used CCSS and only 15 states and the District of Columbia used PARCC or SBAC ([10]; [11]).

2. Research Questions

This Data Report focuses on the effect on politicization at the level of a single message, in this case "tweets" on the Twitter social media platform. We consider how influential particular tweets are in each of the two discussions, where influence is measured both in its rawest form (how many people interact with the tweet), and as the degree to which a tweet tends to increase participation in the discussion.

We identify the most influential tweets in each discussion, and attempt to answer some important questions about influence more generally. Primarily, we would like to know

how discussion online grows and is shaped by specific actors and their messages. What messages are most engaged with, and what types of attributes of tweets are likely to engender engagement?

Of particular interest are the ways in which new participants get involved in an online discussion. What types of tweets tend to bring in new discussants, and which tend to keep the new discussants involved? What are the roles played by tone, anger, and experience in politicizing an issue? Are there topics that tend to amplify the discussion compared to others?

While we expect that positivity and “civility” may be more likely to capture a previously non-committed individual’s attention, it may be that this is conditional on the type of issue. Anger and negativity may be far more powerful when the issue at hand is more basic and existential, as Blacks Live Matter is in comparison to Common Core.

We also expect some topics to resonate more deeply with potential discussants. For Common Core, we expect that topics appealing to personal freedoms and traditional values might be more broadly influential than those concerned with more arcane aspects of the Common Core curriculum and testing regime, despite the argument around the issue ostensibly focused on details of education policy.

Similarly, for Black Lives Matter, we expect that tweets focused on killings and violence against the black community will be far more influential in engendering discussion than more abstract calls for policy change or political considerations. Across both issues, we expect that direct appeals to tangible events and personal benefits will be more powerful than abstract arguments for progress or policy change.

3. Data and Methodology

In this section, we detail the unique data we bring to the questions above, as well as the methodological framework in which we use them.

3.1 Data

In order to analyze the politicization of our two topics as comprehensively as possible, we attempt to gather as close to the full population of tweets about each topic as is possible. From Twitter, we purchased all tweets meeting particular rules, with rules covering time and substance. We constructed a large set of keywords and hashtags about each topic as substance filters, and collected all tweets containing those words in the period from January 1, 2010 until December 31, 2018.³

The data set produced by this process contains 12,764,541 potentially CCSS-related tweets and 231,457,543 tweets with the potential to be BLM-related. However, keywords and (less likely) hashtags can be over-inclusive, as the inclusion of one does not necessarily mean the tweet is relevant to our topic. For this reason, we constructed a relevance classifier

³The full list of keywords and rules for data collection is available in Appendix A.

that weeded out tweets that met our conditions but that were ultimately unlikely to be about the topics we are focused on.⁴ From the original sample, we classified 6,331,690 tweets as relevant to the CCSS issue and 149,411,817 to the BLM topic. These tweets were created by 896,936 unique users in the CCSS data, and 13,363,384 users in the BLM discussion.

3.2 Methodological Framework

In section 4 of the Data Report, below, we first lay out a descriptive analysis of our two data sets. This includes analyzing the distributions of what will ultimately be our outcome variables, and any potential predictors we believe are associated with those variables. In the final subsection, we carry out a statistical analysis in an attempt to identify what attributes of a tweet are associated with greater influence in the politicization of an issue online.

As discussed in depth below, both data sets see most of our outcomes of interest (various ways of measuring *influence*) skewed heavily to the right, with a small minority of tweets representing the great majority of positive values, as well as a large majority of original tweets that have, for our purposes, effectively zero influence as we define it. This can create an analytical problem unless accounted for in the methodological approach.

Here, we address the large number of zeros in our data by modeling our outcomes as a negative binomial distribution, where our predictors are attributes of the tweet and the user who sends the tweet. This allows us to measure the association between our outcomes of interest and the substance and contexts of each tweet, while controlling for other aspects of the tweet. It also accounts for the large number of tweets with no discernible influence as we measure it.⁵

3.3 Timelines

The two issues were not entirely contemporary, as discussed above in the brief description of each movement. In Figures 1 and 2, we show timelines for each of the issues, designed to illustrate how our sample for each grows. These figures map the number of relevant tweets that enter our sample each day.

The timelines tell two different stories. As a topic of discussion, Black Lives Matter starts later, with the first real engagement taking place in mid- to late-2014. It continues to grow throughout our sample period, up to our final day at the end of 2018. Common Core, on the other hand, is steadily discussed throughout the early 2010s, and builds speed beginning in 2013. Discussion over Common Core

⁴In order to construct this classifier, we hand-coded 1000s of tweets from the sample as either “relevant” or “not relevant,” and then trained a simple classifying model to identify features that were strongly associated with relevance. We then applied this classification algorithm to the full body of tweets, marking each as likely “relevant” or “not relevant.” Our out-of-sample tests of this algorithm showed accuracy and precision metrics above 90%.

⁵We include the raw coefficients in our tables, though direct interpretation is not easy. Estimation of a negative binomial model actually predicts the logged count for our dependent variable, so effects need to be exponentiated to be better understood. For ease of understanding, then, we directly interpret our analyses in terms of their proportionate effects in subsection 4.7.

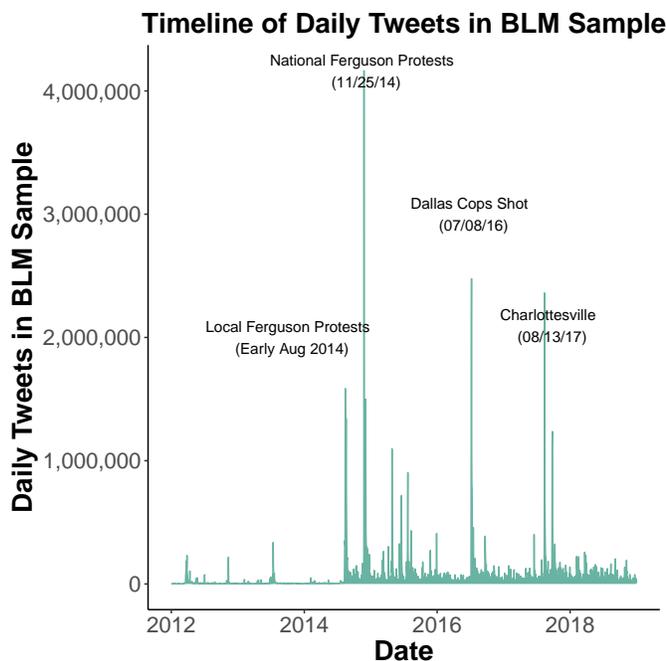


Figure 1. BLM Sample

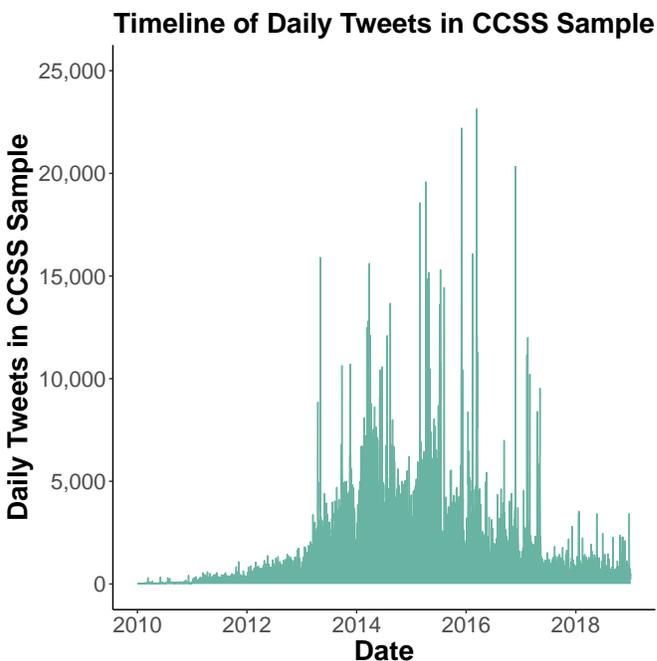


Figure 2. CCSS Sample

grows steadily throughout our sample period, but then levels off to some degree in 2016, and further in 2017.

The manner of growth also differs across the two issues. Common Core grows at an extremely steady pace, with few spikes that really dominate the collection. Even the heaviest days only see about 7-8 times the number of daily tweets as a normal day in the sample collection. Black Lives Matter, however, is far more responsive to events. There are at least four obvious spikes where discussion on the topic spiked in a single or multi-day period, with each of those cliffs identified in Figure 1. During each of these periods, somewhere on the order of 4-7% of our sampled tweets were sent.

4. Findings

We start by offering some descriptive statistics on each of our samples, broken down into various attributes of a tweet, before moving on to inferential analyses that attempt to identify the attributes of a tweet that determine its influence in engaging outsiders in the discussion.

4.1 Users

As discussed above, the samples for each issue differ greatly in size. After restricting the corpora to tweets we identified as relevant to each discussion, the Black Lives Matter data set is more than twenty (20) times larger by tweet volume, with approximately 15 times as many users, than the Common Core data set. In addition to this difference in size, there are differences in the distribution of the number of tweets across users.

As shown in Figure 3, there are far more returning users in the BLM discussion, as more than half of discussants show

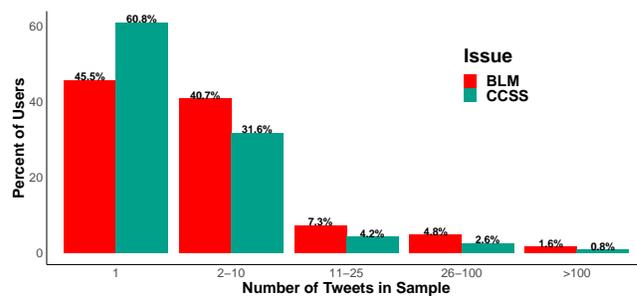


Figure 3. Distribution of Tweet Number

up in our sample more than once (i.e., they tweeted more than once about the topic). The same is not true of the Common Core data, where the modal user tweets a single time before disappearing from the discussion. We might gather from just this descriptive result that BLM is an issue that tends to keep individual users more interested and for longer periods of time. We also see a greater percentage of what might be termed "power discussants" in the BLM sample, with nearly 6.5% of users engaging in the discussion more than 25 times. Given the size of the BLM data, this sub-sample of frequent discussants is nearly as large as the entire number of unique users in the CCSS sample.

4.2 Retweets and Engagements

As a platform, Twitter is set up to encourage discussion, allowing for users to reply directly to tweets from others, as well as pass other users' tweets on to different networks via "retweeting." The act of retweeting serves to amplify the signal of the tweet, disseminating it to new networks that might not otherwise see it. For the majority of this Data Report, we are

using different retweeting behaviors as measures for influence of an original tweet. We discuss three specific measures.

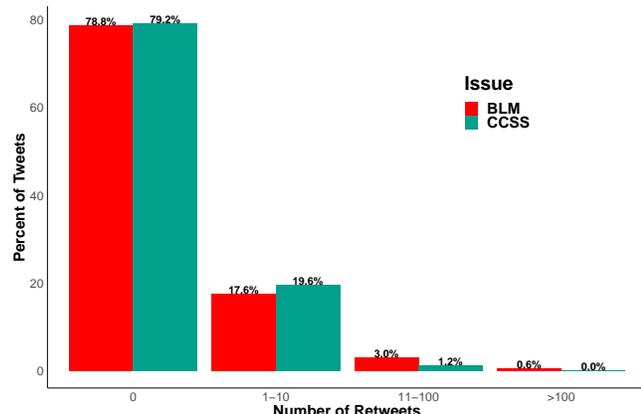


Figure 4. Distribution of Retweets of Each Tweet

First, we consider the *raw reach of a tweet*, using the number of retweets a tweet receives.⁶ Retweets are the most common way of proliferating information across the platform, and they represent the vast majority of engagement across our two issues. When a tweet is retweeted by a user, that user is consciously sharing and spreading the content of the tweet with everyone following the user, even if only *potentially*.⁷ While retweets do not necessarily mean endorsement, they do suggest that the retweeting user thought the content valuable enough to the discussion to pass on. Retweets, then, are a measure of influence, in that they reflect being considered worthwhile additions to the discussion. Both issues exhibit roughly the same pattern for retweets.

Figure 4 covers the distribution of retweets in our sample of original tweets. The vast majority of original tweets (in both cases, approximately 79%) receive no retweets, and those that *do* receive retweets mostly receive a very small number (i.e. less than 10 total retweets).

The BLM data does contain a higher percentage of tweets which were widely shared and approximately 6 of every 1000 original tweets (244,000 in total) were retweeted more than 100 times. CCSS had only 1682 total tweets that saw that level of engagement throughout our sample period. In another way of comparing the distribution of virality across topics, the top 1% of tweets in the BLM sample are responsible for an astounding 79.8% of total retweets, while the comparable number for CCSS is 44.1%.

However, our unique data gives us the ability to push further into retweeting behavior. Having a sample of tweets that closely approximates – or arguably matches – the full

⁶It should be noted that retweets are only recorded for “original” tweets - that is, tweets that are not themselves retweets or quotes of another tweet.

⁷We know that retweets *can* be seen by followers of the account that retweets, but we do not know if they were definitely seen because the tweet may have appeared when the follower was not logged on to Twitter. While Twitter collects these data (known as “impressions”), they are not provided with either the data we purchased from Twitter or the data we collected from the Twitter API.

population of tweets in the discussion allows us to judge what tweets engage new discussants, and bring them into the fold. We are able to measure what we call *new participant retweets* for each tweet. This is the number of times a tweet is retweeted, where that retweet is the first tweet by a user in our sample. This suggests that the original tweet “brought in” the new user, and thus enlarged the discussion.

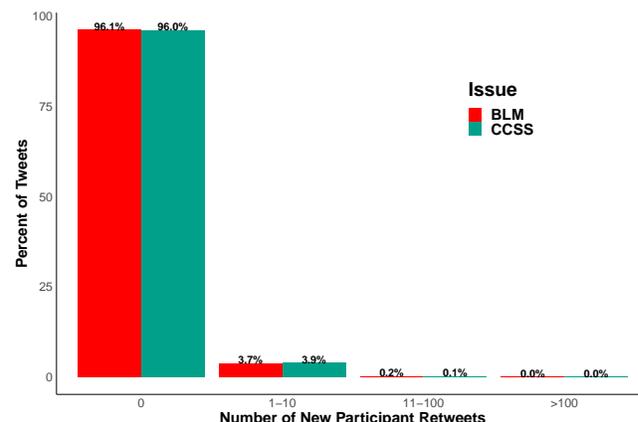


Figure 5. Distribution of New Participant Retweets

In Figure 5, we capture how many tweets receive retweets from users entering the discussion for the first time in our sample. This is an even more uncommon occurrence than merely being retweeted, and only approximately 4% of original tweets in either issue are ever retweeted by a user who was previously uninvolved in the discussion. The single most influential tweet by this measure was retweeted by new discussants 12,730 and 33,750 times in the CCSS and BLM samples, respectively.

Finally, we extend the idea behind “new participant retweets” and calculate a measure of *downstream engagement* that we can attribute to individual tweets. Here, we attribute to each individual tweet the total number of tweets created by users whose *first* tweet was a retweet of the tweet in question. This measure hopefully distinguishes between tweets that bring casual retweeters into the discussion for a short period of time, and those that engage longer-lasting, more involved discussants. We show the distribution for this measure in Figure 6.

Of those tweets that do bring in new users, most see very few of those new users stay active in the discussion. Only about 1 in every 300 original tweets in the BLM data brings in new users that total more than 100 future contributions to our sample. While small, this is more than three times the rate in the CCSS sample. As with many of our distributions, there is a long tail, and some tweets can be considered responsible for a very large number of tweets.

In the CCSS data, there are 8 tweets responsible for at least 10,000 future tweets, with the most influential tweet by this measure responsible for 42,856 knock-on engagements. Its message is straightforward and generally negative towards Common Core: “With #CommonCoreStandards I have com-

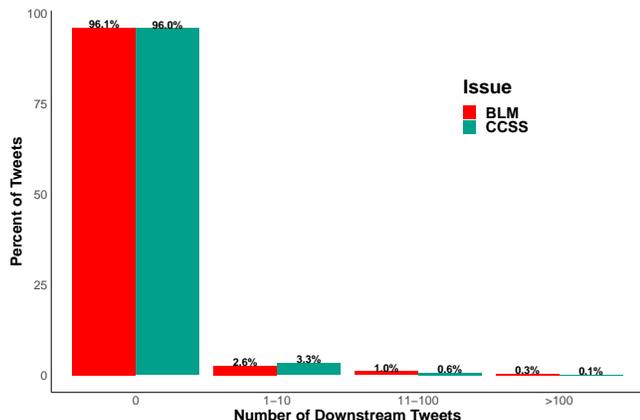


Figure 6. Distribution of Downstream Tweets

pletely lost faith in the public school system.” This tweet is emblematic of the rest of the most influential tweets by this measure. It is contributed by a non-verified account with a slightly above average number of followers and friends (approximately 2200 each at the time of the tweet).

Each of the other tweets most responsible for downstream behavior share this negativity, either discussing some flaw in Common Core or attempting to rally support against the standards. These tweets are just as likely to come from “regular” users as they are from verified accounts followed by large numbers of users. Interestingly, these top tweets are distributed relatively evenly across the period from early 2013 to early 2016, suggesting again the constant and slow burn of the CCSS issue across Twitter. In section 4.7, we more robustly analyze what attributes of tweets drive this measure of future engagement.

The BLM data is vastly larger in overall volume, as well as in tweets per user, so it is not surprising to find that the standout tweets in this area are responsible for many more downstream interactions. The top 10 tweets in this measure for BLM are responsible for at least 100,000 downstream tweets each, with the most influential tweet bringing in 331,501 future engagements. Again, the most influential tweet by this measure was sent by an unverified account in the 50th percentiles for both followers and friends. The text (“After Trayvon Martin was shot, I don’t remember the NRA saying that every black teenager should go out and get a gun for protection...”) is provocative in eliciting response in two possible online issues (Black Lives Matter *and* gun control), which helps to explain its downstream popularity. It should be noted that downstream tweets and new participant retweets are highly correlated, but not perfectly so. Here, the most influential tweet via the downstream measure actually has the second least new participant retweets of the top ten downstream creators.

These top ten tweets share rough similarities that help to explain their long-term influence. Most are both specific to a time and general in nature, allowing them to be retweeted not only at the time of creation, but much later in the movement’s

timeline.⁸ There is also a consistent focus on the failings of other institutions, like the press or police, to prevent and address the injustices.⁹ Finally, it should be noted that BLM tweets are far more likely to have long-lasting consequences if the original sender is at the top of the Twitter food chain, with all of the top 10 tweets by this measure, save the first, sent by someone in the 80th percentile or above in followers.

We use these three measures as proxies for tweet-level influence as we assess the tweet- and user-level characteristics that are associated with more influential tweets.

4.3 Hashtags

Hashtags (any term preceded by the number (#) sign) are an important way of centralizing information spread, and they are routinely used to organize different aspects of the online discussion. When hashtags become popular, they also become ways of signalling an author’s broader message without wasting the limited number of characters allowed in each tweet. Hashtags were used quite differently across our two issues.

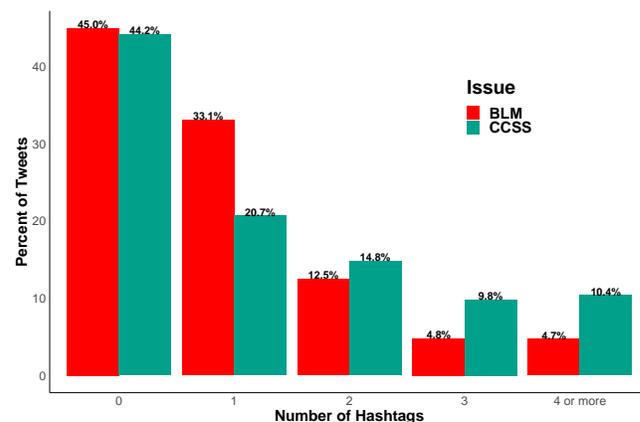


Figure 7. Distribution of Hashtags

First, in Figure 7, we see that while hashtags are used in roughly the same proportion of tweets in each discussion, there is a greater tendency in the CCSS discussion to use many hashtags, and, thus, a higher number of hashtags used per tweet.

In Tables 1 and 2, we see that the type of hashtag most common to each issue is also drastically different. In CCSS, the most popular hashtags are universally general in nature. They exhort readers to “stop” Common Core, or conduct “edreform.” The most specific that the hashtags ever get is referencing broad education platforms – such as ESEA (the forefather of the No Child Left Behind Act) and ECAA (an

⁸Two examples: “I’m a cop. I do not agree with or condone @PO-TUS remarks today on police brutality. Those that applauded and cheered should be ashamed.” and “The Civil Rights Act is 50 years old. These two pictures were taken 50 years apart. Behold our progress. #Ferguson <http://t.co/8PNn8eteO2>”

⁹“Hey Media? Maybe instead of sending cameras to Robin Williams’ house to be ghoulish, you could send cameras to #Ferguson to be journalists.” and “I interviewed the key witness to the Michael Brown shooting last night. The police haven’t. Think about that.” are good examples of this strain of tweet.

Table 1. Most Popular Hashtags, Common Core

Hashtag	Number of Uses
CommonCore	19.5%
StopCommonCore	9.9%
PJNET	7.6%
CCSS	4.6%
edreform	3.0%
edchat	2.6%
tcot	1.8%
StopESEA	1.8%
education	1.5%
ESEA	1.3%
StopECAA	1.2%
KeepYourPromise	1.2%

assessment program for young students) – that are impacted by the Common Core standards. This generality reflects an issue that was little impacted by specific events, but rather a slowly building discussion around certain policies and broader disenchantment with the education standards.

In contrast, the most popular Black Lives Matter hashtags present a wholly different picture. While there remain general slogans with great popularity (including both the main protest slogan #BlackLivesMatter and the counter-protests #BlueLivesMatter and #AllLivesMatter), many of the remaining popular hashtags reference specific events that drove the movement forward. These include both the individuals who lost their lives at the hands of police (Alton Sterling, Mike Brown, Philando Castille, etc.) and protest events that either carried the banner for the BLM movement (e.g. Ferguson) or seemed to serve as a reactionary response (Charlottesville).

This combination of general slogans and event-specific hashtags reflect an issue that was fueled by a broad overarching concern (systemic racism, particularly in law enforcement), but punctuated by individual events that provided evidence and support for the broader claims. The presence of hashtags such as #TakeAKnee and #ICantBreathe in this list also show how individual events (Colin Kaepernick taking a knee during the playing of the National Anthem before an NFL game to protest the killing of black individuals by police, and the words of Eric Garner as he choked to death) can move from the specific to the general by becoming rallying cries for the movement as a whole. There are few similarities to this pattern in the CCSS data, largely due to the absence of meaningful specific events on the same scale as those in the BLM timeline.

4.4 URLs

Links to other resources on the internet can serve a similar purpose to hashtags, particularly on a platform where com-

Table 2. Most Popular Hashtags, Black Lives Matter

Hashtag	Number of Uses
Ferguson	10.5%
BlackLivesMatter	9.3%
Charlottesville	3.0%
AltonSterling	2.0%
SandraBland	1.6%
TakeAKnee	1.4%
MikeBrown	1.3%
PhilandoCastile	1.3%
EricGarner	1.2%
BlueLivesMatter	1.2%
BLM	1.2%
FreddieGray	1.1%
CharlestonShooting	1.0%
AllLivesMatter	0.9%
TamirRice	0.7%
JusticeOrElse	0.7%
BaltimoreRiots	0.7%
ICantBreathe	0.7%

municative space is limited. The ability to lead interested individuals to more information can be a powerful resource, and it is frequently used in our sample.

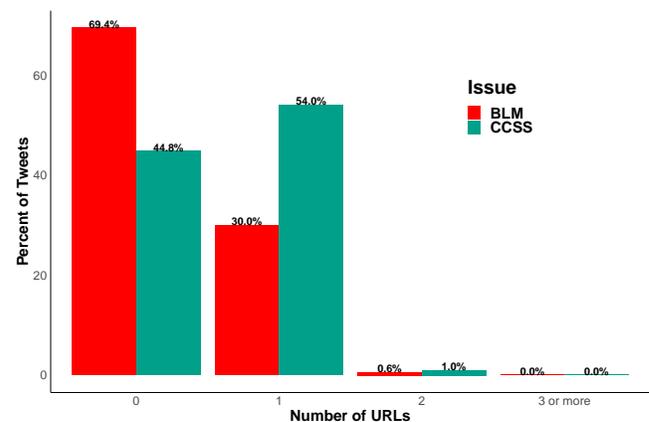


Figure 8. Distribution of URLs

In Figure 8, we see a similar pattern in the shared URLs that we observed in hashtags. Tweets in our CCSS sample are far more likely to contain links than those in the BLM sample. Additionally, the most popular links seem to have an altogether different purpose across the two topics.

In the Common Core data, the top URLs (all of which were shared between 5000 and 10000 times each, and many of which are now mothballed) tend to represent either extended arguments (delivered via a personal website) against Common

Core, or examples where it has harmed students. Consider, as one example, the video at <https://bit.ly/3i5LdRG>, which was linked 5521 times in our dataset.

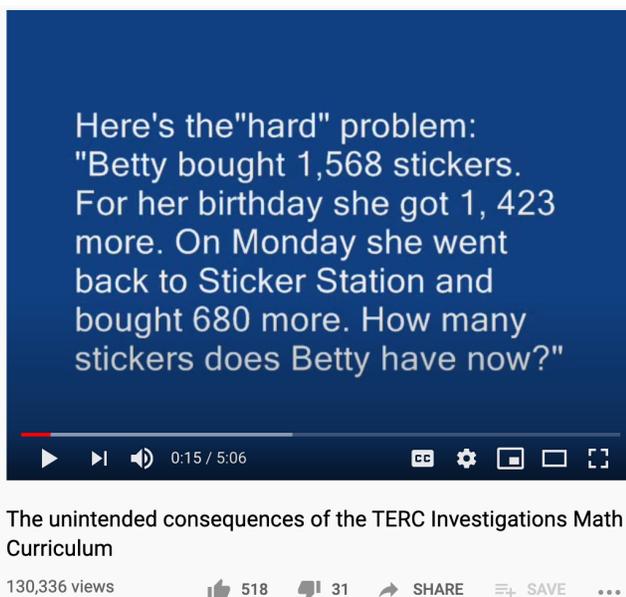


Figure 9. Most Shared URL, CCSS

The video (a screenshot of which can be seen in Figure 9) portends to show how Common Core has over-complicated simple addition problems for students and made them generally worse off. It shows a young girl attempting to explain a math problem to her mother, presumably in the way the girl has been instructed to solve it under Common Core standards. In general, the purpose of the URLs most shared in the CCSS data seems to be persuasion, rather than spreading information.¹⁰

Contrast this with the top URLs from the Black Lives Matter data, which tend to be to news stories or other Twitter accounts *about* news stories. Consider the link at <https://bit.ly/2CK1xtH>, shared over 90,000 times in our data. As can be seen in Figure 10, it shows a group dressed in unofficial militia gear heading towards a central gathering location on the day of the Charlottesville rallies. While there remain some attempts at persuasion (this link reports on the son of Alton Sterling in a sympathetic fashion), most of the sites in the top 15 most-shared urls in the Black Lives Matter data have an informational bent, spreading videos and pictures from the ground around protests.

4.5 Sentiment and Tone

As discussed previously, we might expect that the way in which a message is delivered will impact its influence, or capacity to engage new discussants. Specifically, we would like to know whether tweets that are roughly positive in tone are more likely to engender discussion than those more negative,

¹⁰Compare this video, for instance, with a hypothetical video that contrasted the CCSS-suggested protocol for math to the classic form of math pedagogy.



Figure 10. Most Shared URL, BLM

all else equal. We also want to gauge the efficacy of tweets containing more anger, particularly in the case of Black Lives Matter - arguably a more directly visceral and emotional issue than Common Core State Standards.

For these analyses, we subset the sample to only those original tweets in our data set; that is, we deal only with the first time a tweet shows up in our data, and exclude all retweeted versions. This leaves us with 3,486,608 tweets about Common Core, and 38,824,600 about Black Lives Matter.

We first classify tweets as to their tone - how positive or negative the message seems to be. We use three separate, common, and validated lexicons ([12], [13],[14]) to apply dictionary-based methods in classifying a tweet as largely positive, negative, or neutral in tone. We check that these measures are highly correlated, and produce a single index of negative sentiment using principle component analysis (PCA).¹¹ For purposes of analysis, we normalize this index so that we are capturing deviations from a mean negativity within the sample, rather than a wider baseline of completely neutral speech.¹²

It should be noted that positivity and negativity here do not measure stance toward Common Core or Black Lives Matter per se, but rather the tone in which whatever opinion the user has about the issue is expressed. Tweets expressing support for either issue can be quite negative in tone, and tweets communicating disagreement can potentially be very positive. In fact, these dynamics are very clearly present in each dataset, with many of the tweets in favor of the BLM message quite negative in their tone, and tweets opposed to Common Core quite positive in theirs. Negativity, then, is a measure of tone, rather than substance. We are interested

¹¹All three measures load heavily on a single primary dimension, which explains over 75% of the variance in the CCSS case, and nearly 74% in the case of BLM.

¹²More details about this process are available in Appendix B of this Report.

in the effect of how an idea is expressed (regardless of the content of the idea itself) on its capacity to engage future conversation.

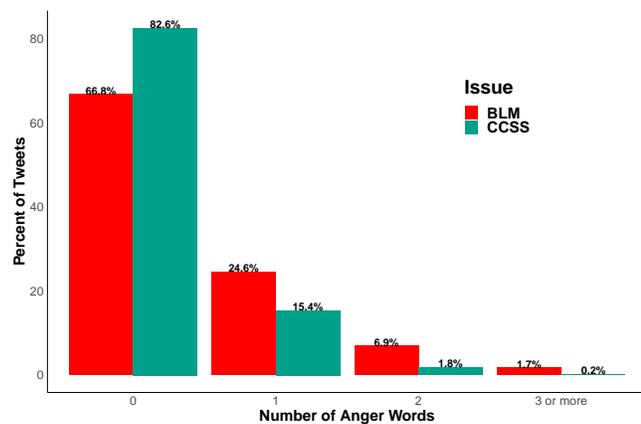


Figure 11. Distribution of “Anger” Scores

We also classify each tweet for its use of words associated with “anger” using the NRC Emotion Lexicon ([14]). The vast majority of tweets in each of the topics contain no words associated with anger, and those that do are still skewed toward the lowest levels of recorded anger.

In Figure 11, we plot the distribution of tweets in each topic by their anger score. We can immediately see the different way in which words of anger are present in the two discussions. Nearly a third of tweets in the Black Lives Matter data contain at least one word that could be construed as angry, while nearly 2% (or nearly 800,000 of the original tweets) include relatively high levels of anger. These numbers are substantial in comparison to CCSS, where only about 17% of tweets contain any anger, and fewer than 7,000 total tweets had even three instances of anger. It should not be surprising that anger plays a larger role in a discussion centered on violence against minority populations than in a discussion about appropriate education policy, though we reserve discussion of whether anger is more *influential* in engendering discussion for either issue until the final subsection.

4.6 Substantive Topics

There is also a natural assumption that specific sub-topics within each issue may resonate more with participants, both those new to the discussion and those already participating. In order to test this hypothesis, we designate six (6) different topics for each of the issues, and create a list of keywords associated with each of those topics. We then code tweets as pertaining to a topic based on whether the words used in each tweet match the keywords associated with the topic. This “dictionary” approach allows us to classify tweets in a simple, straightforward manner, and focus on the major topics of discussion in each issue.¹³

¹³The keywords for each topic can be found in Appendix C to this Data Report.

Table 3. Distribution of Topics in Tweets, CCSS and BLM

Topic	Issue	Percentage of Tweets
Protests	BLM	21.56%
Race and Racism	BLM	18.11%
Police and Law Enforcement	BLM	16.14%
Victims of Violence	BLM	9.06%
Calls for Reform or Justice	BLM	4.13%
Specific Politicians	BLM	3.11%
Education Policy	CCSS	19.73%
Individual Values	CCSS	10.94%
Legislative Policies	CCSS	4.28%
Academic Testing	CCSS	3.59%
Elite and Philanthropic Organizations	CCSS	0.83%
Specific Politicians	CCSS	0.47%

Table 3 lists the six non-mutually exclusive topics for each issue, and the percentage of original tweets in our sample that are classified under that topic. The topics were selected based on a review of the literature surrounding each issue, and specifically on how the movements were set up, and what the driving focus seemed to be in each. We discuss each more thoroughly below:

Common Core Topics

1. Education Policy - The Common Core standards were meant to fundamentally restructure K-12 education, and ensure that students across the country were given equal opportunities to learn about topics that would frame their future success. This topic deals with the discussions surrounding the education curriculum itself, including changes to best practices, homework, teacher responsibilities, etc.
2. Individual Values - The standards were also revolutionary in the extent to which they nationalized education policy, previously an area of policy that was highly localized. This topic deals with the pushback received from this angle, focusing on claims that the freedoms and rights of choice of local parents were being infringed upon.
3. Legislative Policies - The adoption of Common Core was necessarily a legislative endeavor, and as described above it involved at least 3-4 separate pieces of interacting legislation. Additionally, the calls for repeal of the standards were largely legislative in nature, with calls to action at both the state and local level. This topic contains the appropriate laws, as well as covering the calls for a reversal of the policies.
4. Academic Testing - One of the major areas that the Standards shifted policy on was academic testing of students. This shift toward focusing on meeting and exceeding benchmarks via academic testing was an easy target for critics, and this topic focuses on that part of the discussion.
5. Elite and Philanthropic Organizations - Critics frequently cast the standards as the product of elites and phil-

anthropic organizations disconnected from “everyday” people. In this topic, we cover the major organizations and institutions that were involved in the development and implementation of the Standards.

6. Specific Politicians - As the process of adoption and calls for reform were largely legislative in nature, specific politicians became associated with either support for or criticism of the standards. This topic focuses on those politicians who became associated with the debate surrounding the standards, including both supporters and critics.

Black Lives Matter Topics

1. Protests - The BLM movement was arranged around highly visible events of protest, including large gatherings across the country, and individual statements and actions like Colin Kaepernick’s kneeling before a National Football League game. This topic covers the most common of these protests, as well as more general protest-related terms.
2. Race and Racism - The movement’s major focus was how race impacted interactions between citizens and law enforcement. This topic covers topics related to race and racism at both the individual and institutional levels.
3. Police and Law Enforcement - The primary goal of the BLM movement was to stop the violence committed by police and members of law enforcement against black members of the community. This topic covers specific mentions of police, law enforcement, their counter-protest actions (tear gas, etc.), and the counter-movement focus on “blue lives matter” that sprung up as BLM gained speed.
4. Victims of Violence - One of the most powerful ways that BLM activists garnered attention for their cause was merely listing the victims of the type of violence they were fighting against. This topic compiles this list of names, and focuses on tweets that specifically mention the victims.
5. Calls for Reform or Justice - In addition to the end of police violence against black individuals, BLM activists frequently called for broader reforms, including but not limited to reparations, community policing, and removal of officials complicit with the violence. They also frequently called for justice in the form of charging and convicting police officers who committed violent acts. These prescriptive claims are catalogued under this topic.
6. Specific Politicians - Similarly to Common Core, particular politicians became associated with the movement, either as supporters or harsh critics. This topic captures those politicians.

There is substantial variation in the percentage of the sample each topic comprises, over both issues. Tweets about protests represent at least 21% of our original BLM tweets, while tweets covering Politicians or Reforms represent only about 3-4% of the sample each. In CCSS, the differences are more stark, with Education Policy being mentioned in nearly 20% of original tweets, and both Elites and Politicians being covered in less than 1% of same.

In both cases, the proportion of topics is roughly upheld whether we look at frequent participants or those individuals who contribute only once in our sample. In Common Core, education policy is slightly more prevalent for one-time discussants (~26.9% of their tweets), but the rest of the topics don’t see movement of more than a percentage point in either direction. Similarly, among Black Lives Matter topics, five of the six topics are mentioned at roughly the same rate by one-time discussants as by more frequent discussants, with only protests seeing a large difference. That topic is actually less pronounced amongst one-time participants, with only 16.8% of tweets associated.

In both cases, it is not clear that the existence of a topic in a particular number of tweets is the same as influence in the way discussed in this paper. While one topic may be far more common than another in our original tweets, it may be the latter type of tweets that attract new discussants, or engender longer term discussion. Thus, we turn now to an analysis of our chosen measures of influence, in order to explore out these relationships.

4.7 Influence and its Predictors

For the purposes of this paper, we consider “influence” in online politicization on multiple dimensions. We are primarily concerned with three questions:

1. What attributes of tweets and their creators are most likely to see those tweets disseminated widely?
2. What attributes are associated with attracting new discussants to an online issue?
3. What types of tweets bring in new discussants who remain engaged in the topic and provide additional long-term value?

In this part of the Data Report, we use the design elaborated in subsection 3.2 to analyze these questions, and provide insight into how different aspects of a tweet can affect its influence on an online discussion. Our three outcomes of interest for both topics are those described above in subsection 4.2, namely, the **number of retweets** received by a tweet, the number of **new participant retweets** it captures, and the number of **downstream tweets** for which it is responsible.

We model each of these outcomes as an additive function of both tweet- and user-specific attributes. In the former category, we use the scores of the tweet on both the anger and negativity indices, its substantive topic area(s), and both the existence and the number of hashtags and URLs in the

tweet.¹⁴ In the latter category, we use the verification status of the tweet's creator,¹⁵ the number of followers the creator had at the time of the tweet, the number of total tweets the account had sent at the time of the tweet, and the number of tweets the account had sent *about the issue* at the time of the tweet. These final three variables were cut into sample deciles (number of followers and number of lifetime posts) and quintiles (number of tweets sent to that point that had appeared in our sample).

Figures 12 and 13 present the results for CCSS and BLM, respectively. For ease of visualization, we have removed the results for the user attributes other than verification status, but still discuss those findings at length below.¹⁶ We start by interpreting the results for each issue separately, and then discussing differences across the issues.

In our Common Core data, we note a few broad trends across the metrics of tweet influence. First, a tweet by a verified user can be expected to be more influential than a tweet by a non-verified user, all else equal. The difference is quite substantial, with a verified user's tweet expected to receive about 5 times as many retweets and downstream tweets and nearly 7 times more new participant retweets.

The number of followers an user is an even stronger predictor of influence. The median user is expected to have approximately 7.5-12 times the influence as a user with no or very few followers, across our three measures, and moving from the median user to a user in our top decile of followers increases that advantage by an order of magnitude.

These findings are complicated by our other user-level measures of tweeting activity. The number of posts an individual user has sent in the past is a strongly negative predictor for all three outcomes, with those users in the median for previous posts receiving approximately 20% of the engagement as individuals who tweet very infrequently. Those at the very highest levels of tweet activity are even less likely to see their tweets picked up, with only about a quarter of the retweets of the median user, and one-sixth of their downstream tweets. Tweeting volume and numbers of followers are of course correlated, but taken together they suggest that more popular users do tend to see increased engagement, but tweeting more is not always an avenue to greater influence.

Finally on the user side, experience within the topic does not seem to increase influence on the whole. Individuals in the top quintile of on-subject tweets see about twice the number of retweets as those completely new to the subject, but see no significant gain in downstream tweets, and yield only 15%

¹⁴In this last set of predictors, we use a dummy for having hashtags (URLs) at all, as well as a predictor with the number of hashtags (URLs) the tweet contains. This allows us to say something both about the effect of hashtags generally, and whether their effect is declining or increasing in number.

¹⁵Twitter as a platform offers a "blue checkmark" to individuals who can reasonably claim that their account is of public interest. This "Verified" status is Twitter-granted and serves as an authenticating message that the account represents who it claims to represent. It can also serve as a proxy for expertise, though that is not Twitter's claimed focus.

¹⁶Full results, including the user-level results for each decile and quintile, are available in Appendix D.

as many new participant retweets. It's not clear that people entering the discussion of CCSS for a first time can clearly delineate between veterans of the topic and other novices.¹⁷

From a substantive tweet-level perspective, some interesting results pop out. The inclusion of hashtags tends to be associated with more retweets, and more downstream tweets, while URLs seem to have the opposite effect. This discrepancy may be reflective of the way the two tools are used. Hash-tags may make it easier to find (and then share) tweets, while URLs may provide information to the reader, but creating less necessity to share it.

In analyzing sentiment, both anger and negativity are positively correlated with retweets, but only anger keeps that positive relationship with downstream tweets. Negativity seems to find individuals willing to retweet, but not necessarily individuals that then contribute to the discussion long-term. Anger on the other hand, seems to resound with individuals, as each additional angry word corresponds with approximately 25% more downstream tweets.

Topics are perhaps the most intriguing category of tweet-level attributes to analyze, and they show that even rarely mentioned sub-topics can be quite powerful in politicizing a discussion. While elites and philanthropic organizations were mentioned in fewer than 1% of original tweets, they were associated with an out-sized increase in the number of new entrants to the discussion. So, too, were tweets dealing with academic testing and legislative policies like the abolition of the Common Core standards. However, the two most popular topics (education policy and individual values) were also associated with the largest predicted increases in downstream tweets, which suggests that that the scope of a topic may still play a vital role when it comes to influencing politicization.

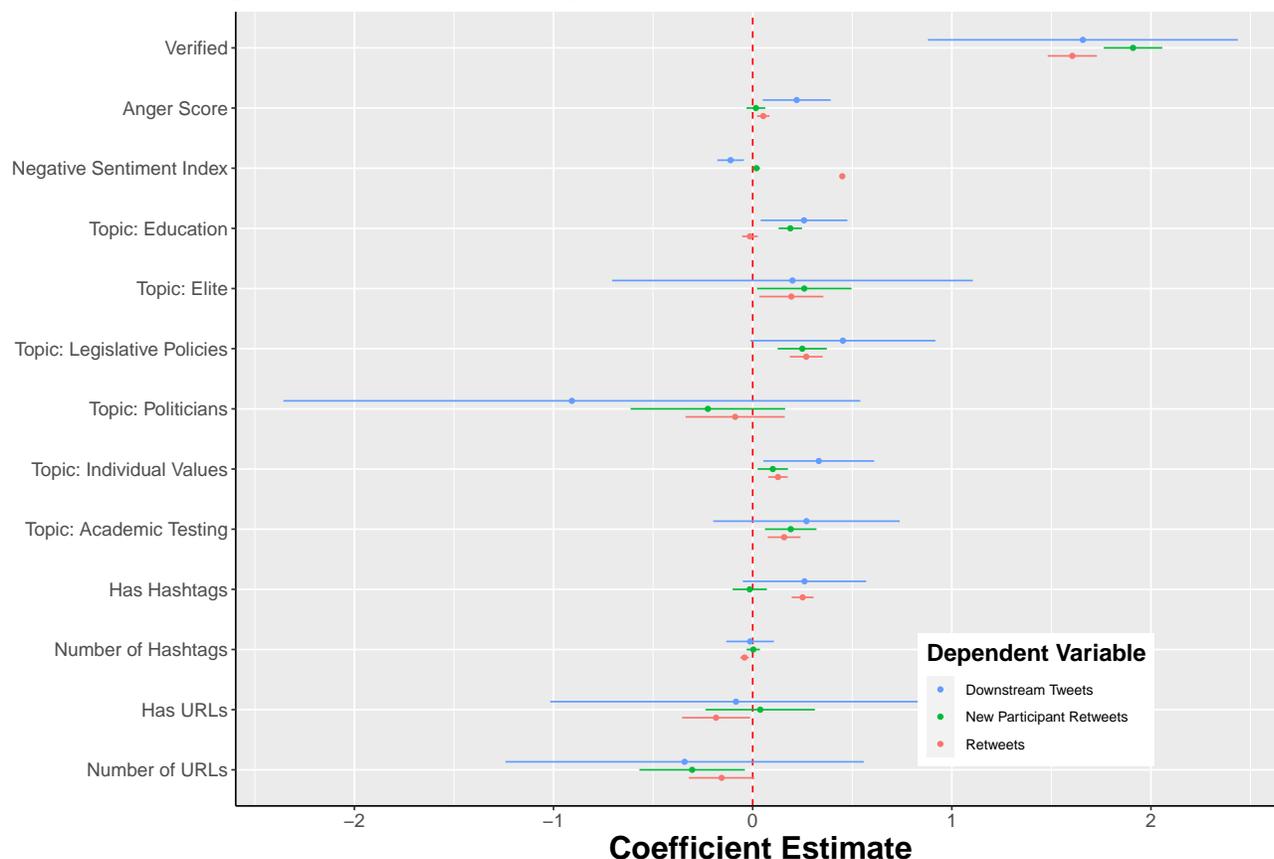
We see a near identical repeat for user-level attributes in our BLM data. Verification is associated with a large increase in all three measures of influence, as is the number of followers a particular user has at the time of tweet creation.

The number of posts in a user's lifetime up to that point is negatively associated with influence, while on-topic tweets have a mixed record, associated positively with retweets, and negatively with new participant retweets and downstream tweets. The total estimated effect sizes are roughly in line with those for the CCSS sample, which suggests a more general platform-wide relationship between user attributes and engagement.

At the tweet-level, hashtags and URLs remain positively and negatively correlated with engagement, respectively, though the effects here are much larger in total. While the addition of a single hashtag was associated with an increase of approximately 25% retweets and downstream tweets in the CCSS sample, here the relationship suggests closer to a 60% increase. The negative association for URLs is also stronger,

¹⁷As noted previously, in a companion Data Report being released concurrently as this one, we examine influence at the user-level, including a classification of users by their focus on these topics, and politics more generally. That Data Report can be found [HERE](#) (INSERT LINK).

Figure 12. Influence on Tweet Impact: Common Core State Standards Analysis



particularly in downstream tweets, where the expected loss moves from approximately one-third to over half.

At the level of sentiment, anger remains a steady and positive predictor of all three types of engagement, with every “anger word” associated with a 16% increase in downstream tweets. Negativity, however, seems to have no substantively interesting relationship with any measure of influence.

However, the discussion of influential topics is far muddier in the Black Lives Matter data. No single topic is positively or negatively associated with all three measures of influence at traditional levels of statistical significance or even substantive meaningfulness. Tweets about protests seem to suffer the most in comparison to other tweets, correlated with a decrease of approximately 20-30% in retweets and new participant retweets, while tweets surrounding race are positively associated with increased influence, albeit weakly. Only the topic centered on reform seems to affect the number of downstream tweets, though one should be reluctant to put too much weight on that single result in the context of the remaining findings.

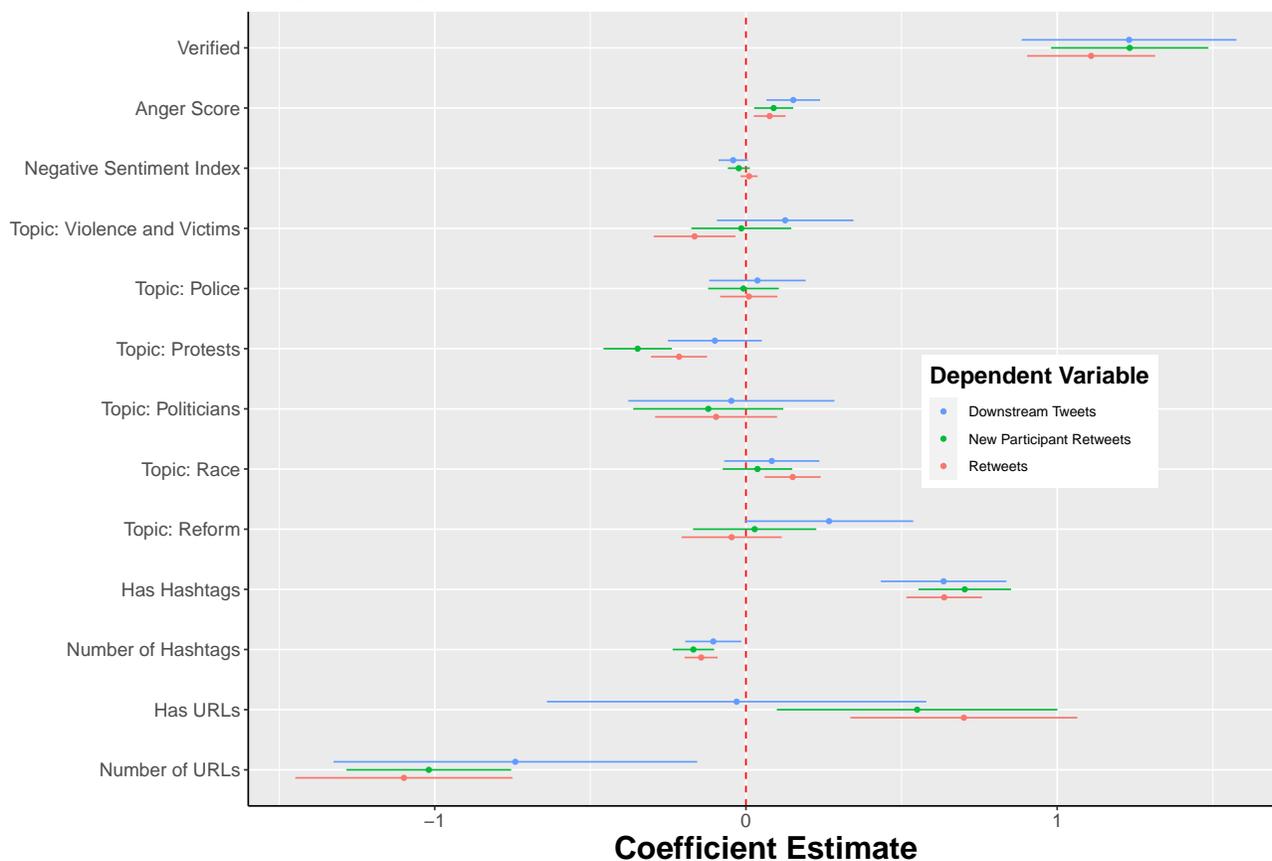
In comparing the two topics, one should first note that in both cases, the tweet-level effects are dwarfed by the user-level effects. What seems to matter most when discussing influence in politicizing an issue is the identity of the tweeter, more so than the substance of the tweet. While hashtags and

URLs have similar relationships with influence in both topics, their size is different and may be a function of the types of users in each sample and the information users are attempting to transmit.

In all, the analysis suggests three main takeaways:

1. The attributes of the accounts that send a tweet are, on the whole, more important than the substance of the tweet itself in determining the immediate and long-term influence of that tweet.
2. Still, substance matters. Some topics seem to resonate more strongly, bringing in more new discussants than others. For Common Core, tweets about individual values and legislative policies were most influential; for Black Lives Matter, discussing protests seemed to be counter-productive if the hope was to engender more discussion.
3. Tone may also play a role, as the existence of anger in a tweet was associated with greater influence. Negativity, however, was less important.

Figure 13. Impact on Tweet Influence: Black Lives Matter Analysis



5. Conclusion

In this Data Report, we examined how two different issues became politicized on social media. Using a comprehensive collection of Tweets, we analyzed how the discussion of Black Lives Matter and Common Core State Standards developed, and what types of engagements were likely to further shape and grow the conversation. We first show that the conversations developed differently, with the discussion around Black Lives Matter being far more responsive to particular events than the one around Common Core. We also show that the issues have different styles of interaction and participation, with Black Lives Matter seeing more returning discussants using a greater number of hashtags and anger words, though less direct links to other sources.

When we analyze the discussions for evidence of tweet-specific influence, we find that user attributes dominate. The popularity of an individual user, measured by their number of followers, is highly correlated with the number of retweets a tweet receives, as well as with the number of new discussants the tweet brings in, and the number of tweets those new discussants ultimately contribute to the conversation. Similarly, the “Verified” blue check mark bestowed upon particular users seems to serve as an imprimatur, and tweets by these users are generally more influential across all dimensions.

However, substance also plays a part in a tweet’s influence. In both discussions, tweets that utilize words expressing anger bring higher engagement, and particular topics of interest in each issue are found to be more important in encouraging new participants to join the conversation. The use of hashtags can increase a tweet’s spread throughout the discussion, though URL links do not seem to have the same effect.

This Report serves as a first step in theorizing and analyzing the ways in which issues become politicized online. We present new measures of influence that can be used at both the tweet-level (as they are here), or in the aggregate, as well as considerations of Future research should undoubtedly pursue this important topic in greater depth going forward.

Acknowledgments

We thank Cody Buntain, assistant professor in Informatics at New Jersey Institute of Technology, for his extensive support in data processing and measures, Megan A. Brown, research engineer and data scientist at the Center, for her assistance in data collection, Jen Rosiere Reynolds, the Center’s lab manager, for her edits, and the Bill and Melinda Gates Foundation for supporting this project. We also gratefully acknowledge that the Center for Social Media and Politics at New York University is supported by funding from the John S. and James

L. Knight Foundation, the Charles Koch Foundation, Craig Newmark Philanthropies, the William and Flora Hewlett Foundation, the Siegel Family Endowment, the Bill and Melinda Gates Foundation, and the National Science Foundation.

References

- [1] About the Standards | Common Core State Standards Initiative.
- [2] Development Process | Common Core State Standards Initiative.
- [3] Adopting Agency, March 2017.
- [4] U.S. Secretary of Education Duncan Announces Winners of Competition to Improve Student Assessments | U.S. Department of Education.
- [5] Catherine Gewertz. In National First, Kentucky Adopts Common Standards - Education Week. *Education Week*, February 2010.
- [6] Education Next. Ednext Poll Interactive Ten Year Trends Public Opinion, August 2016.
- [7] Anthony Cody, Network for Public Education co founder, and Former Teacher. Political Rivals Find Common Ground Over Common Core.
- [8] Who Is Fighting Against Common Core?
- [9] David Freedl and er On 09/28/15 at 10:54 AM EDT. How the GOP Turned on Common Core, September 2015.
- [10] Standards in Your State | Common Core State Standards Initiative.
- [11] Which States Are Using PARCC or Smarter Balanced? - Education Week. *Education Week*, February 2017.
- [12] F. Å. Nielsen. Afinn, mar 2011.
- [13] Bing Liu. *Sentiment analysis: Mining opinions, sentiments, and emotions*. Cambridge University Press, 2015.
- [14] Saif M. Mohammad and Peter D. Turney. Crowdsourcing a word-emotion association lexicon. 29(3):436–465, 2013.