1-11-2021

Let Them Tweet Cake: Estimating Public Dissent Using Twitter

Ben O. Smith
University of Nebraska at Omaha, bosmith@unomaha.edu

Ethan Spangler

Recommended Citation
Ethan Spangler & Ben Smith (2021) Let Them Tweet Cake: Estimating Public Dissent Using Twitter, Defence and Peace Economics, DOI: 10.1080/10242694.2020.1865042

Let Them Tweet Cake: Estimating Public Dissent using Twitter*

Ethan Spangler†
Ben Smith‡

November 21, 2020

Abstract

This paper establishes a new method of estimating public dissent, one that is both cost efficient and adaptable. Our approach utilizes the proliferation of social media and advances in text analysis. Twitter allows users to post short messages that can be viewed and shared by other users, creating a network of freely and easily observable information. Drawing data directly from Twitter, we collect tweets containing specified words and phrases from citizens voicing dissatisfaction with their government. The collected tweets are processed using a regular expressions based algorithm to estimate individual dissent; which is aggregated to an overall measure of public dissent. A comparative case study of Canada and Kenya during the summer of 2016 provides proof of concept. Controlling for user base differences, we find there is more public dissent in Kenya than Canada. This obvious, but necessary, result suggests that our measure of public dissent is a better representation of each country's internal dynamics than other more sporadic measures. As a robustness check, we test our estimates against real-world civil unrest events. Results shows our estimates of public dissent are significantly predictive of civil unrest events days before they occur in both countries.

Keywords: Public Dissent, Twitter, Civil Unrest, Protests

JEL Code: C80, O57, P16

Word Count: 8198

*For advice and comments, we thank Philip Wandschneider, Michael Brady, Thomas L. Marsh, Austin L. Wright, Zachary C. Steinert-Threlkeld, participants at the WEAI 2018 Annual Conference, as well as anonymous reviewers at Defence & Peace Economics.

†Corresponding Author, Independent Researcher, Email: eospangler@gmail.com

‡University of Nebraska-Omaha, Email: bosmith@unomaha.edu;

Introduction

An implicit rule of any government is that in order to rule they require most of the people, to follow must of the rules, most of the time. The issue is then how willing is the population to follow the rules of the government, which requires knowing the level of dissent amongst the populace. The importance of public dissent and collective action has been well theorized (Tullock 1971; Tullock 1974; Kuran 1989; Gurr 2015) as well as the empirical connection between public opinion and civil violence (Hirose et al. 2013; Blair et al. 2017). However, despite this significance, robust and efficient methods of obtaining measures of public dissent remain underdeveloped. Whereas previous estimates of public dissent relied on expert opinion, surveying, or other traditional methodology, this paper develops a measure of public dissent based on data collected from Twitter. Empirical testing shows that our estimates of public dissent are predictive of real world civil unrest events in both Canada and Kenya.

To illustrate why developing a better understanding measure of public dissent is important, let us examine what happens when its importance is ignored. Three commonly used measures of political stability are Political Risk Services (PRS), the Business Environment Risk Intelligence Index (BERI), and the Economist Intelligence Unit (EIU). Each of these indexes combine political, financial, and economic factors to assess a nation's political stability (Howell, 1998). Public dissent is not explicitly accounted for in these indexes as a precise quantitative value.

Instead political factors, such as public dissent, for each index are qualitatively determined and scored by panels of experts (Howell, 1998). These experts are usually former diplomats, scholars, and other suitably qualified individuals. While these teams of experts can be quite large and knowledgeable, it is still a relatively small group of people trying to assess an entire nation. These experts often make their decisions without a clear theoretical foundation. Since political factors compose 33-66% of each index (Howell, 1998), if these experts are somehow misinformed, the validity of the index could be greatly affected. In turn, this potential bias would affect all research based on these indexes. Additionally, Tetlock (2005) shows that political forecasts based on expert opinion are only marginally better than random chance.

The inherent issues of these measures can be seen in the country reports published prior to public uprisings. The PRS report on Ukraine published October 2012, stated that 'a repeat of the Orange Revolution...is unlikely.' and 'Ukrainians are disillusioned but in general they possess little appetite for protest.' (PRS 2013, 11). Mass protests began in November 2013 and by February 2014 the Yanukovych regime had fallen. The PRS report on Tunisia, published October 2010, called Tunisia an 'oasis of stability' (PRS 2011, 3) and postulated a 85% probability, the maximum probability PRS gives, that Tunisian dictator Ben Ali would retain power for the next 18 months. By January 2011, mass protests and revolt resulted in the dissolution of the ruling RCD party, the exile of Ben Ali to Saudi Arabia, and the establishment of an interim government.

While it is easy to critique these forecasts with the benefit of hindsight,[1] these examples demonstrate the inherent issues in predicting outbreaks of social unrest. As noted in Blair et al. (2017) and Bazzi et al. (2019) a primary challenge in predicting public violence is the lack of quality supporting data. In both Ukraine and Tunisia, unrest was incited by sudden

[1]It should be noted that PRS publishes monthly reports on its surveyed countries but those are only available to its subscribers.

unpredictable events but the dissent amongst the population was there all along but there was no way to detect it. This paper adds to the field by developing a quantitative, efficient, and scalable measure of public dissent using Canada and Kenya as proof of concept case studies.

Related Literature

Traditional methods of evaluating public dissent rely on crowd sourcing, polling, surveying, and event coding; but all have their limitations. Ungar et al. (2012) relies on expert opinion, but instead of just a few experts Ungar et al. employ thousands, using a mixture of crowd sourcing and simplification of complex issues. Ungar et al.'s has their army (over 2000 individuals) of forecasters assign probability estimates to specific events happening within a given time (e.g. Q:'Will there be an assassination attempt on Julius Caesar before March 15^{th}?', A: 'Yes, 42% probability.'), updating their predictions as needed before a deadline. Finally, all predictions are combined to form a single aggregate forecast of an event.

Ungar et al.'s method, and prediction markets in general, are extremely effective in harnessing the wisdom of crowds, but at the same time they are hamstrung by the simplifications needed in order to harness that wisdom. They work best when asking the crowd simple questions, which may not capture all the nuances and complexities necessary to understand an issue. This applies especially when attempting to gauge public dissent in a country. Additionally, in dealing with esoteric issues, there might only be a few experts with area knowledge, which leaves this method vulnerable to the same problems as described by Tetlock (2005). Finally, maintaining and incentivizing a large number of forecasters is likely to be costly and time intensive, as one must first acquire said forecasters and then wait for them to make and adjust their judgments.

Polling and surveying provide flexible formats for assessment and when implemented properly can be effective. Hirose et al. (2013) and Blair et al. (2017) show that assessments of public opinion are predictive of civil violence. Unfortunately, effective polling and surveying are costly and time consuming. Wright et al. (2019) details the intricate steps needed to conduct a rigorous survey to gauge public perceptions of government corruption in Afghanistan. The Afghanistan Nationwide Quarterly Research survey requires significant infrastructure, relationships with local leaders, and surveyors have to be recruited and trained on proper technique. It is not always economically feasible to conduct a high quality survey of public opinion, especially on a continual basis.

There are also unique issue to polling and surveying that are difficult for researchers to overcome. The biggest hurdle is one of honesty since respondents often have little incentive to be truthful, to varying degrees of malevolence. For instance, polls and surveys are susceptible to the 'social desirability bias,' wherein respondents have a tendency to provide what they perceive to be the socially acceptable answer to questions (DeMaio 1984; Stephens-Davidowitz 2017). This could be especially problematic if someone is being asked about popular government entities or policies.

The issue is further intensified by the fact that many places where accurate measurement of public dissent is most needed, might also be places where honest public speech is not safe. According to Freedom House (2017), of the 195 countries evaluated, only 44% were regarded as 'free' in regards to political rights. This means that in most countries a person might be unwilling to provide their honest thoughts to a stranger asking about their government. On the other extreme, respondents may provide strategic answers with the intent of influencing potential policy that may be based on poll results, biasing results (Morgan and Stocken 2008). Finally, results could be biased because respondents provide false information purely for their own trollish amusement (Stephens-Davidowitz 2017).

The rise of the internet and movement of traditional print media onto digital formats has created new data opportunities for the study of public dissent. Event coding works by scraping information from international news sources, specifically targeting reports related to political violence and civil unrest (protests, riots, civil wars, etc). Previously, these databases were meticulously hand coded (Leetaru and Schrodt 2013) but newer automated systems have emerged. The most prominent automatic event coding systems are the Global Data on Events, Location and Tone (GDELT) (Leetaru and Schrodt 2013) and the International Crisis Early Warning System (ICEWS) (O'Brien 2010). These systems are useful in that they are able to quickly process and aggregate the torrent of information the 24 news cycle generates.

GDELT and ICEWS are not without their flaws. Both systems have a tendency to duplicate events[2] and their screening methodology remains opaque (Wang et al. 2016). These automated event coding systems are also flawed in that they rely predominately on online English language news reports. This makes cross country comparisons unfeasible because the underlying data is skewed towards countries with more online English news coverage or are just covered more by English news sources, leaving many developing countries underrepresented. The other issue is that event coding is entirely ex-post. The focus on news media means that very little of the signals and public motivation potentially preceding an

[2]The issue seems less pronounced with ICEWS.

event are captured before reaching headlines.

Social Media Literature

The spread of ICT (information and communication technologies) has provided a wealth of opportunity and new challenges for researchers (Dafoe and Lyall 2015; Weidmann 2015). Social media platforms, such as Twitter, are of particular interest and potential. Twitter, and other micro-blogging websites, allow users to post short messages (tweets) that can be view and shared by other users. These posts can include tags that allow users to link posts with a common theme. This creates a vast network of information that can be freely and publicly observed. With a current active monthly user base of over 300 million people (Twitter, 2016) spread across the world, researchers can tap into the zeitgeist of a population.

Twitter data has already shown to be useful in several areas, often performing better than traditional data sources. Asur and Huberman (2010) were able to use Twitter chatter to predict film box office returns better than the industry standard. Bollen et al. (2011) show that Twitter data can be used to forecast stock market fluctuations. Smith and Wooten (2013) shows that people use Twitter as a source of information and were able to estimate demand for this information. In terms of politics, O'Conner et al. (2010) and Lampos et al. (2013) use Twitter as a more accurate source for political forecasting than traditional polling.

There is also interesting research concerning issues of public dissent using Twitter data. Carley et al. (2013) find that Twitter chatter increases as large scale political events unfold. Carley et al.'s findings demonstrate that there is a tangible connection between real world and online behavior; people are tweeting in response to things that are happening in life. Even in extremely repressive regimes such as Saudia Arabia, people are undeterred in posting online dissent (Pan and Siegel 2020). This point is further reinforced by research suggesting that Twitter can be used to assess reform efforts (Seabold et al. 2015), protest recruitment (González-Bailón et al. 2011), protest participation (Kallus 2014), and forecasting civil unrest events (Ramakrishnan et al. 2014). This line of research has been deemed so promising that the US Department of Defense has funded several ongoing projects in this area (Minerva Initiative 2014).

Some research suggests that the spread of social media platforms increase protests activity by reducing coordination costs (Enikolopov et al. 2017). However in this capacity social media is no different than other ICTs in terms of reducing coordination costs of collective action (Pierskall and Hollenbach 2013; Shapiro and Siegel 2015). The important distinction for researchers is that the content on Twitter is observable whereas other ICT platforms are not. Because the content of a tweet is observable, the intensity of the dissent in the tweet can be ascertained rather than merely the volume of tweets.

The central premise of this paper is that people use online platforms, specifically Twitter, to complain about politics and express dissent against their government. A well functioning government should be like air; if it's working well, no one should be talking about it. The more dysfunctional a government and its institutions are, the more demand there will be for dissent against the government. Social media is merely the means by which they express some of that dissent, not the cause of the dissent. By examining Twitter data directly, we mitigate many of the issues of other measures of public dissent that rely heavily on expert opinion, polls, surveys, or online news. Instead, we are able to develop nearly a real time estimate of public dissent.

A key contribution of this paper is its data collection and processing methodology. Previous research collected Twitter data by focusing on a certain accounts (Zeitzoff 2011), a limited set of specific keywords (Zeitzoff et al. 2015), users with known locations (Steinert-Threlkeld et al. 2015; Korolov et al. 2016), or a specific hashtag (Elson et al. 2012; Larson et al. 2019). There is nothing wrong with these approaches and they are effective in evaluating their events of interest ex-post. However, in order to gauge the boarder notion of public dissent a more expansive approach is needed.

Our process takes an ex-ante approach which involves capturing tweets containing language expressing public dissent based on criteria derived from theory. These tweets are processed through regular expressions to develop a quantitative measure of public dissent at the country level that is compatible with empirical testing and allows for cross-country analysis using the same methodology. This methodology gives researchers more control over the analysis than opaque natural language processing methods. A significant benefit of this approach is its low resource cost; this entire project utilized basic computer hardware and open-source software and had no external funding. Additionally, while the methodology of this paper is designed to collect tweets related to public dissent, the framework is flexible enough that it can be adapted to any area and language where data on public opinion in required.

Theory

For this paper the definition for public dissent and the theoretical model is from Spangler and Smith (2015). There, dissent starts at the individual level and is any action against the government by a member of the public. These actions include contacting representatives, attending protests, voting for opposition parties, and other similar behaviors but also involves publicly pointing out government failures and shortcomings. These actions are aggregated to form public dissent, D_t. This definition of dissent is similar to Ramakrishnan el al.'s definition of 'civil unrest' (2014) but is broader to capture more subtle acts of defiance.

Individual dissent is a reflection of perceived governmental quality. Every day, individuals face societal issues they themselves cannot overcome, but negatively affects their lives. These are issues such as crime, corruption, and other societal problems (usually with public goods characteristics) that are difficult for individuals to secure themselves and are usually provided by governments. However, for whatever reason, the government is unable to address the issues sufficiently for the individual. The simultaneous frustration with governmental expectations and the inability to do anything leads the individual to do the only thing they can do, dissent. Dissenting provides a cathartic release for the individual, making them feel slightly better.

The individual responds to government policy and chooses their level of dissent, $d_{i,t}$, that maximizes their utility. Equation (1) shows the general form of the individual's utility from dissent, with B the benefit from dissent, C the cost, and P_t is the probability of being caught dissenting.

$$\max_{d_{i,t}} U_{i,t} = B(d_{i,t}, x_i, g_t, E) - P_t C(d_{i,t}, A) \tag{1}$$

where:

- $d_{i,t}$ is dissent
- x_i is activism preference
- g_t is benefits from the government
- E is quality of life
- A is the punishment for dissenting

On the benefit side, an important component is an individual's activism preference, x_i. x_i represent how prone individuals are to activism and follows some distribution across a population. For some values of x_i an individual would receive disutility from dissent, so they would elect not to, $d_{i,t} = 0$. Whereas higher values of x_i means that the individual receives utility from dissenting, $d_{i,t} > 0$. This heterogeneity among the population means that each period for given government policy choices there is a spread of people that do not dissent and variation in the level of dissent among those that do.

Other factors influencing the utility from dissent stem from the one's quality of life and the policy choices set by the government. If an individual has a high quality of life, E, and/or receives benefits from the government, g_t; there is less for the individual to be discontent about. E and g_t reduced the utility from dissent an individual receives. The better one's life is, the less there is for one to dissent about.

On the cost side is the likelihood of being caught dissenting, P_t, and the severity of punishment, A, each reducing the utility from dissent. The riskier it is to dissent and the harsher the punishment, the less utility one will receive from it. $P_t = f(D_{t-1}, S_t)$, a function of total dissent from the previous period, D_{t-1} and current security allocations of the government, S_t. D_{t-1} reduces the probability of being caught dissenting (lost in the crowd) while S_t increases P_t.

Maximization of individual utility with respect to $d_{i,t}$ yields $d^*_{i,t}$, individual dissent.

$$d^*_{i,t} = f(x_i, g_t, E_{i,t}, P_t, A) \tag{2}$$

From the individual, finding total dissent is just a matter of consolidating dissent across the population, N_t.

$$D_t = \sum_{i=1}^{N_t} d_{i,t} \tag{3}$$

where D_t is total public dissent from the population.

A core assumption of this paper is that digital behavior is representative of real-world behavior. By this we mean that online dissent and real-world dissent are related to another. The online grumblings captured in our process are the same sentiments people in the real-world have but with the benefit of being easily captured digitally allowing us to form an estimate of total public dissent.

Methods

The spread of social media, specifically Twitter, makes it possible to transition from the theoretical model to the empirical. When people post public messages critical of their government they are voicing dissent; and the more posts they make the more they are dissenting.

Aggregate this across an entire nation and one can observe how dissatisfied the public is and what the major points of contention are. In basic terms, we gather tweets critiquing either the government or some issue under the government's purview, and bundle them together to create our measure of public dissent, $\hat{D}_t$.

Our notion of public dissent is similar to Gurr's (2015) concept of 'relative deprivation,' wherein there is a difference between what their government does and what they expect their government to do. Here we are emphasizing the public's perceived and vocalized failings of the government across a range of issue areas (political, economic, social) to create a single estimate of public dissent. This is different than explicitly capturing public opinion in a poll, because people are freely expressing it without being asked and it's more open ended nature with respect to breadth of potential topics. This measure can also be easily dissected than other measures by pulling apart the aggregation and seeing what the big issues are.

This paper follows a similar method as Smith and Wooten (2013) and Carley et al. (2013), but is expanded to capture the more open ended nature of public dissent. Figure 1 shows the Twitter data collection and processing used in this paper, from Twitter to final dataset used in estimation.

[Figure 1 here]

The first step is to establish a connection to the Twitter API (application program interface), which allows one to tap into Twitter's flow of data. From there it is a matter of pulling the relevant tweets. We created a codex of words expected to be consistent with the language one would use to express dissent against a government. The codex acts as a sorting mechanism, pulling all matching tweets from the stream of data. Any tweet posted, in the entire world, containing at least one of our words is captured.

The sources of public dissent have been well researched. Grossman (1991), Acemoglu and Robinson (2001), MacCulloch (2001 and 2005), Apolte (2012a) explore the role of income inequality. Other approaches concentrate on the role of regime types (Guttman and Reuveny, 2014) or government institutions (Goldstone et al. 2010; Acemoglu et al 2010; Acemoglu et al 2012; Fukuyama 2014). Based on this research, we developed a list of words and phrases we consider relevant to public dissent. These are words such as: 'police', 'rule of law', 'corruption', 'unemployment', as well as the names of important political figures and institutions in our sample countries, Canada and Kenya. The goal is to capture tweets representing fundamental concerns about the state of a country, not the obscure musing of elites. For this project the codex contained 397[3] words and phrases.

The codex has to be limited due to Twitter's internal controls. At maximum only 1% of total Twitter traffic flow can be captured at a time, collecting more severs the connection to the API. While this may seem like a low volume of data, it is still potentially millions of tweets and the majority of posts on Twitter are irrelevant to our purposes. A more expansive codex would have risked running into this roadblock and ceasing the collection process.

While by no means exhaustive, the goal was initially to collect as many tweets as possible that might be expressing dissent against the government. The authors of this paper do not claim to be experts on either Canada or Kenya politics, but extensive research and care was taken to ensure that the codex reflected the contemporary political landscape of each country by examining government websites, news agencies[4], and social media posts.

Initial data collection of tweets containing words from our list began June 13, 2016 and

[3]The codex is available in the supplementary materials.

[4]For Canada: Canadian Broadcast Company, National Post and National Newswatch. For Kenya: Daily Nation and The Star.

ran to September 11, 2016. Aside from technical limitations barring collection for a longer period, there were no a priori expectations of political events happening during this period. Choosing to start our data collection period at a relatively random point lets us to better test our process as a general method of gauging public dissent.

Regular Expressions

The next step is to run our collected tweets, the raw data, through regular expressions (regex). Regular expressions allow us to go beyond a simple word count and instead account for the context of the tweet. Regular expressions work by creating a search criteria. When one of our collected tweets matches the regular expression the tweet is then pulled and the relevant account information pulled (time of posting, number of retweets, age of account, etc). An additional step in this process involved removing potential bots[5] by eliminating any tweet posted within 30 days of account creation. After 30 days Twitter's internal controls are decent at removing bots.

There is not an exact science behind coding regular expressions, it requires balancing specificity and generality. Code too specifically and nothing will match the criteria. Code too generally and the risk of false positives increases. As a rule, our regular expressions err towards specificity. Since this is a comparative study between Canada and Kenya the number of regular expressions for each country needed to be roughly approximate. Table 1 shows the count of regular expressions by category for each country and a brief list of what each category entails. For this project, 474 unique regular expression were used.[6]

[5]Fake accounts controlled by computers.

[6]The full regular expression list is available in the supplementary materials.

[Table 1 here]

Each matched regular expression, $\chi_{i,t}$, in a tweet is counted and used to estimate individual dissent,$\hat{d_{i,t}}$.

$$\hat{d_{i,t}} = \sum \chi_{i,t} \tag{4}$$

An advantage of this discrete system is that we can easily employ nonparametric statistical tests such as the Mann–Whitney and Kolmogoro-Smirnov tests to compare sample dissent distributions. Additionally, a single tweet can contain multiple regular expressions depending on its content. The more expressive a tweet, the more an individual is dissenting. Take this example tweet in Figure 2 admonishing an unspecified political leader:

[Figure 2 here]

A regular expression meant to capture the dissent in this tweet would look like:

```
([^.?!]*)(\b(Leader)\b)([^.?!]*)(corrupt)
```

This regular expression would capture any tweet with 'Leader' and 'corrupt' in the same sentence. To capture situations in which 'corrupt' comes before 'Leader' in a sentence, the regular expression would need to be reversed.

```
([^.?!]*)(\b(corrupt)\b)([^.?!]*)(Leader)
```

To increase computational efficiency, regular expressions can be expanded to include words that convey similar sentiments and ideas.

```
([^.?!]*)(\b(Leader)\b)([^.?!]*)((corrupt(ion)?)|(crook(ed)?)|(criminal))
```

This regular expression now captures any tweet with 'Leader' along with the words corrupt, corruption, crook, crooked, and criminal. This tweet also contains other inflammatory statements we would want to capture with regular expressions as shown in Table 2.

[Table 2 here]

Counting the number of matched regular expression in the above tweet, then $\hat{d_{i,t}} = 4$. This tweet is used as an example of a single tweet containing multiple points of interest, but this is the minority of cases. During the time of study Twitter had a 140 character limit, so most tweets only match a single regular expression. Another advantage of estimating dissent in this fashion instead of counting tweets or users, is that this accounts for situations in which there is a small but intense minority voicing dissent.

Estimating dissent with regular expressions this way is inline with event coding systems discussed previously. ICEWS uses the Conflict and Mediation Event Observations (CAMEO) scale which classifies event types and places them on a -10 to 10 intensity scale; with -10 being open warfare and 10 being a peace treaty (Gerner et al. 2002). Zeitzoff (2011) employed a modified CAMEO intensity scale for coding tweets in their study of the 2008-2009 Gaza Conflict. Our method takes a step back from the CAMEO system and instead of scoring large scale global events, we are simply scoring word groupings.

One unique aspect of Twitter, is the ability for other users to repost tweets by others users. Here, these 'retweets' are essentially other Twitter users expressing the same level of dissent as the original tweeter. Therefore, for the purposes of dissent, the number of retweets acts as a multiplier. Including retweets also allows us to partially account for more influential users, who would likely have their messages retweeted more than others.

Individual dissent is aggregated across the Twitter population, $N_{Twitter}$, to form estimates of total public dissent in a country, $\hat{D}_t$ over whatever time interval is desired: daily, weekly, monthly, etc.

$$\hat{D}_t = \sum_{i=1}^{N_{Twitter}} \hat{d_{i,t}} \tag{5}$$

Controlling for context using regular expressions is important for several reasons. First, location information is only known if the user voluntarily provides it on their Twitter profile, which relatively few do. While there is a tendency for political discussions to remain with the same national network (Zeitzoff et al. 2015), explicitly coding regular expressions to focus on Canadian or Kenyan issues ensures we analyze the right tweets. Even if these are tweets are coming from people outside the country, they are still discussing issues unique to our sample countries. They could be reporters, scholars, tourists, or expatriates; more than likely they are discussing something relevant.

Second, some words have different cultural meanings that might create bias if only a simple word count was employed. For example, one of our code words is 'anarchy' and in Kenya it is used in the traditional manner of discussing issues involving lack of government and lawlessness. However, in Canada the vast majority of tweets containing 'anarchy' were discussing the TV show *Sons of Anarchy*. The potential for misidentification is why single word regular expressions were used sparingly in this paper. Single word regular expressions were only used when collecting tweets in languages other than English (Swahili and French) or with very specific terms used only in a negative context (e.g. 'nairobbery'[7] or '#cdnpoli'[8]). The issue of translation should be minimal since in both Canada (Poblete et al. 2011) and

[7] A colloquial expression for the high crime rate in Kenya's capital.
[8] A popular hashtag used for discussing Canadian politics.

Kenya (The Economist 2014) the predominant language of choice on Twitter is English.

An alternative method of text analysis is sentiment analysis which works by evaluating small bodies of text and determining the overall emotional content. The important words in a text such as the adjectives, adverbs, and verbs are assigned an emotional value based on a predetermined subjective lexicon (Pang and Lee, 2008). The count of positive and negative words are differenced to give overall sentiment. If sentiment > 0 then positive, if sentiment < 0 then negative, and 0 is neutral. Sentiment analysis is a useful tool and can be used to understand public feelings concerning a subject and is used for evaluating geopolitical events (Overbey et al. 2017) and determining public opinion of politicians (Ringsquandl and Petkovic 2013).

An issue with sentiment analysis is its tendency towards binary response: positive or negative. More nuanced assignments that use more of the emotional spectrum are possible but some texts may be assigned multiple emotions or even contradictory ones depending on their complexity which confounds conversion to a quantitative measure.

Another issue with sentiment analysis is that it is based on a subjective lexicon. It many not be appropriate to apply a lexicon to language made outside of it's original cultural context (Silge and Robinson 2018). Recall the cross-national issue with the word 'anarchy'. The regular expressions count captures some of the emotion embodied in the text, but our approaches gives a more robust and empirically translatable answer. By using regular expressions we are able to tell not just when people are angry but the degree of their anger.

By not fully incorporating sentiment analysis, we do open ourselves to the potential criticism that even with our regular expressions we are not able to fully control and quantify the tone of a tweet. For example, how do we deal with sarcasm. To answer this criticism, we feel that any tweet containing a political message, even one clearly satirical in nature, does not happen in a vacuum. The tweet authors have encountered something in their life that causes them to respond. The fact that they've responded sarcastically is just a choice of phrasing and is inconsequential for our purposes. It only matters that they posted the tweet. The same logic can be applied to people tweeting in defense of their government and institutions. Again, this is not happening in a vacuum; these people are reacting to something and we are capturing that in their tweet.[9] However, as a robustness check we include some sentiment analysis as part of our empirical findings which show that most of our collected tweets were negative in nature.

Case Studies: Canada and Kenya

Given that this research paper is meant as a proof of concept, the scope of the empirical analysis was limited to two sample countries: Canada and Kenya. Each country was selected for their similarities and differences. First, English is the major language of politics, education, government, and most importantly the internet in both countries. Having a shared language significantly reduces the potential for misidentification translation would entail.

[Figure 3 here]

Second, the populations of each country are prolific users of Twitter. In Canada there are over 7 million monthly active users on Twitter (Statista, 2017), 19% of the total population. In Kenya there are an estimated 700 thousand monthly active users (Kemibaro, 2014), 1.4% of the total population. This means that Twitter provides an easy way of surveying the political moods of large sections of the Canadian and Kenyan populations. Figure 3 shows

[9] This logic also applies to critiquing retweets.

the distribution and concentration of the analyzed tweets for users with known locations. The distribution of tweets follows the major population centers of each country.

Finally and most importantly, there are different a priori expectations of each's level of public dissent. Canada is a country that often ranks among the top of nations in terms of development, citizen happiness, and governmental transparency. Conversely, Kenya is a developing country with a history of political instability, corruption, and ethnic tension. Most notably, there was substantial post election violence in 2007 after the election of Mwai Kibaki as president. This unrest resulted in the deaths of 1,200 people and displaced hundreds of thousands (Blair 2016).

Figure 4 shows the count of political violence incidents in each from 2014 to 2017 and Table 3 shows contrasting statistics related to political stability for each country. The data in Figure 4 comes from the Political Instability Task Force (PITF) worldwide atrocities dataset (Goldstone et al. 2010).

[Figure 4 here]

From Figure 4, we observe that Canada had only one instance of political violence from 2014 to 2017.[10] Conversely, Kenya had 73 incidents of political violence over the same time period, one of which was the assassination of a member of parliament.[11] Table 3 contains relevant comparative 2016 statistics between Canada and Kenya. Measurements of the economy, political rights, corruption, and quality of life point towards drastic differences Canada and Kenya. Given the differences between Canada and Kenya shown in Figure 4 and Table 3, a measure of public dissent should reflect these substantial differences.

[10] A shooting at a mosque in Quebec City in January 2017.

[11] Kenyan MP George Muchai.

[Table 3 here]

Data and Estimation

Summary statistics from the Twitter analysis are presented in Table 3. $\hat{d_{i,t}}$ refers to dissent from a single tweet while $\hat{D}_t$ is a weekly aggregation, $\hat{D}_t = \sum \hat{d_{i,t}}$. From June 13, 2016 and to September 11, 2016 we collected 45,246 tweets from Canada and 29,257 from Kenya. As expected most tweets matched a single regular expression.

Dissent

Table 4 contains summary statistics from the Twitter data. An interesting result are the parallels in dissent between the two countries. Median individual dissent is the same for both countries at 1, both express dissent early in the week, and there is a mutual concern regarding immigration and refugees (Table 6). However, beyond these initial similarities there are stark differences.

In nominal terms, there were more Canadian tweets expressing dissent than Kenyan. This result is because Canada has a monthly active user population roughly ten times greater than Kenya. However, a more accurate view is to look at the proportion of the user population dissenting. Only 0.64% of Canadian Twitter users expressed some level of dissent during this period. Meanwhile 4.16% of Kenyan Twitter users expressed dissent, meaning relative to their user populations significantly more Kenyans expressed dissent.

[Table 4 here]

A necessary assumption we make is that dissent scales with user population. If Canada and Kenya had exactly the same user population, the proportion of the population expressing dissent would still be 0.64% and 4.16%. This is a needed in order to obtain relative estimates of public dissent. Under this assumption, we normalize total dissent by user population, which increases Kenyan median $\hat{D}_t$ from 2,010 to 16,924, which is over four times higher than Canada's. Figure 5 shows this difference across the entire sample period, the dotted green line representing the nominal data while the solid green line represents the scaled. Looking at the scaled data, we observe that there is far more dissent in Kenya than Canada over the period observed.

[Figure 5 here]

Sentiment analysis using the Bing lexicon (Silge and Robinson 2018) shows that most tweets from both countries had an overall negative context. Given that the regular expressions were designed to capture tweets related to dissent, a predominately negative response is expected. The minority of positive tweets were either emotionally complex or possibly pro-government but are still relevant to our purposes as explained earlier.

An issue with the sentiment analysis is that a sizable percentage of tweets, 35.7% for Canada and 48.6% for Kenya, were unable to be classified, meaning they used words that were not in the lexicon. The inability to classify so many tweets demonstrates the limitations of sentiment analysis when applying a lexicon outside its native culture.

Table 5 contains results from a Mann-Whitney U test and a Kolmogorov-Smirnov test. Each tests whether two independent samples are pulled from populations with the same distribution. In this context we are testing whether individual dissent in Canada and Kenya are distributed the same. In each test, we reject the null hypothesis of matching distributions of dissent. This indicates the presence of fundamental differences in dissent behavior between the two countries.

[Table 5 here]

A possible reason we see more dissent in Kenya is that the summer of 2016 was the start of the 2017 presidential election. Elections have a tendency to raise the political awareness of a population and highlight shortcomings of the ruling regime, potentially increasing the level of dissent in a country. Given Kenya's past history with election violence, tensions were likely high and we observed this through Twitter. Table 6 shows the top 5 issues for both countries; in Kenya the top two issues are election related.

The issue of tribalism includes tweets that specifically cited the problem of tribalism or made derogatory comments about members of Kenya's various tribes. Kenyan politics is divided predominately along tribal lines, with the Kikuyu being the largest and predominate holders of political power. The Luo are another large tribal group and represent the primary political opposition to the Kikuyu dominated government. The post election violence in 2007 was primarily between these groups.

Prior to the data collection period beginning, there were mass protests throughout Kenya in May 2016 concerning the belief by opposition members that the two election commissions, the Independent Electoral and Boundaries Commission (IEBC) and the Ethics and Anti-Corruption Commission (EACC), are biased in favor of the ruling party (The Economist, 2016). At the same time there were a series of suspected extrajudicial killings by police (The BBC, 2016). We see public anger from these issues carry over into our observation period, with many people decrying such injustices and instances of police killings. Another issue of major concern for Kenyans was the Nairobi slum of Kibera, one of the largest slums in the world. This is an area with poor access to civil services, extreme poverty, and high crime rates.

[Table 6 here]

In Canada the major issues of concern are generally first world issues: cost of living, taxes, and generic political kvetching. Generic politics is at the top due to the popular hashtag '#cdnpoli'. Canadian users often attach this hashtag when discussing any issue related to Canadian politics. The topics tagged are quite diverse and often contain other regular expressions but the popularity of '#cdnpoli' warranted being singled out. Concern for terrorism stems from the murder of a Canadian held hostage by Islamic extremists in the Philippines and the Trudeau administration's refusal to pay ransom (Vice News 2016). The issue of immigration and refugees rose to prominence due to Canada resettling thousands of Syrian refugees (Beauchamp 2016).

The generality of the main topics captured is important because specific demographic information for either country is limited. Only US Twitters users have been extensively studied, showing a tendency to be younger and better educated (Wojuik and Hughes 2019). It is possible these trends hold elsewhere, but we don't necessarily need them to. Table 6 shows the general nature of the issues of the captured tweets in both countries. These are not esoteric topics of elites, but represent core national issues. This suggests we are picking up on what the broader population in each country thinks are important issues. So while we can't ensure the measure is completely representative, we have taken steps to show that it is still substantive.

One interesting trend of our analysis can be seen in Figures 5. In mid August, there is a bump in dissent for both countries. Examination of news articles in each country during this period reveals no significant events that would explain the spike. However, broadening the scope to the world level does uncover an explanatory event: The 2016 Summer Olympics in Brazil.

[Figure 6 here]

Search history for 'Olympics' from Google Trends (Figure 6) for each country, confirms that citizens from each country were interested in the Olympics. Given that the Olympics are a major global event, it's logical that Canadian and Kenyan would be interested. However, that relationship alone does not explain why we see a spike in dissent during the Olympics.

A possible explanation is that there may be a spillover effect from interest in the Olympics. Since people use Twitter as a news source (Smith and Wooten 2013; Twitter 2016), it is entirely plausible that they logged on to Twitter to see/discuss their country's performance in the Olympics and then stayed to discuss political events relevant to their own country. Given that social media users tend to use the platform for all aspects of their life (Valenzuela 2013), this kind of topic drift should be expected. Since the regular expressions used were coded to capture exclusively Canadian and Kenyan issues, this is the most plausible explanation. This reveals an unexpected sensitivity to outside large scale events this type of analysis might have.

To further contrast the differing levels of public dissent, Figure 7 shows Kenya's public dissent indexed against Canada's. As before, the solid green line represents the scaled data and the dotted line the nominal. Over the period of analysis, Kenya has on average 500% more public dissent than Canada.

To put this estimate into a real world context, in May and early June 2016 there were large scale protests across Kenya. These protests were lead by the opposition CORD party, including presidential candidate Raila Odinga, and were against perceived corruption by election commissions. The protests turned violent and police cracked down and fired teargas into the crowd (The Economist 2016). In contrast, the Canadian equivalent would be Justin Trudeau being teargassed by RCMP mounties before the 2015 Canadian election; a scenario that is difficult to imagine.

[Figure 7 here]

Canada's developed economy, efficient government, and strong civil society likely factor heavily into this result. While showing that Canada has less public dissent than Kenya may seem like an obvious result, it is nonetheless important. The necessary disparity between the estimates signifies that this is a valid method of cross country analysis.

Empirical Validation

To add statistical credibility to this method of assessing public dissent, we test the relationship between our estimate of public dissent and recorded civil unrest events at the daily interval. Civil unrest events are things such as riots, protests, demonstrations, and other similar events that one would expect to be a product of public dissent. Using the ICEWS dataset, we record the number of civil unrest incidents that occurred for each country. ICEWS is the only available dataset that records civil unrest events for both Canada and Kenya. Because both countries predominately use English, the issues noted earlier regarding event coding systems are less a concern. Summary statistics for the data are in Table 7.

[Table 7 here]

We then preform a Granger Causality test to see if online public dissent portents civil unrest events in the real world. A Granger Causality test works by comparing two linear models: a nested auto-regressive model with k lags and a full auto-regressive model that also includes the independent variable of interest with k lags of each. For a given lag order, an independent variable x Granger-causes y if it's inclusion in a model preforms better than the nested auto-regressive model (Granger, 1969). While we are employing a *causality* test, we are not using it in an attempt to say online posts create real world actions. Instead we are using these tests to show that our estimate is emblematic of the same underlying issues causing people to take to the streets.

$$\begin{aligned} \text{Nested: } y_t =& \alpha_0 + \alpha_1 y_{t-1} ... + \alpha_k y_{t-k} \\ \text{Full: } y_t =& \alpha_0 + \alpha_1 y_{t-1} ... + \alpha_k y_{t-k} + \beta_1 x_{t-1} ... + \beta_k x_{t-k} \end{aligned} \tag{6}$$

For testing, daily Kenyan scaled dissent is linear and daily Canadian dissent is changed using the inverse hyperbolic sine (IHS) transformation, $log(\hat{D}_t + (\hat{D}_t^{\ 2} + 1)^{\frac{1}{2}})$. In both countries the daily series of civil unrest incidents and estimated public dissent are stationary, thus requiring no other transformations. The IHS transformation has the same interpretation as the standard logarithmic transformation but has the benefit of being defined at zero (Burbidge et al., 1988) which is necessary because there are some days without observed dissent. The differences in specification is likely due to the structural and distributional differences in dissent noted previously in Tables 3 and 5.

Table 8 contains the results from the Granger tests for both Canada and Kenya. Results from an Wald-test show that there is a statistically significant relationship between our measure of public dissent and civil unrest incidents. In Canada our measure of dissent is able to signal civil unrest events one day in advance and in Kenya it is two. This result is also unidirectional for both, testing for civil unrest incidents Granger causing public dissent tweets did not yield significant results under any specification.

[Table 8 here]

It would seems that there is evidence that our measure of public dissent is tapping into the real feelings of the populace of both countries. Recall the contrasting realities between Canada and Kenya in Figure 4 and Table 3. Measurements of the economy, political rights, corruption, and quality of life point towards drastic differences between Canada and Kenya. We believe our measure of public dissent is capturing these same issues, but in a more holistic and efficient way than other methods. Instead of relying on quarterly surveys or year end estimates, we are able to see daily fluctuations in public dissent. The potential for rapid response makes this a far more useful measure in identifying which countries are on the verge of political crisis before unraveling and should be included as a factor in studies trying to predict violent outbreaks. Furthermore, this measure of public dissent is flexible enough that it can be scaled to whatever time interval is needed for other quantitative studies. We believe that this paper has made a strong case for the use of Twitter data to assess public dissent in a country. At the very least we have presented a line of research that warrants more exploration, especially in research focusing on the prediction of civil unrest.

Conclusion

This paper has presented a new method of estimating public dissent. By collecting tweets expressing dissent against we are able to obtain estimates of public dissent for Canada and Kenya. After normalizing for population difference, our estimates show that Kenya likely has substantially higher public dissent than Canada and that our estimates are predictive of civil unrest events. Our methodology allows for easy and comparative measurement of public dissent in a way that is scalable across countries. Furthermore, this methodology can be employed to study the intensity of public opinion towards other subjects.

Not only do we observe the expected difference between Canada and Kenya, but they are different by orders of magnitude. This is an obvious and necessary result. Had their been only a marginal difference or none at all, our methodology would be frivolous. Our results are in line with other methods and data, highlighting the validity of this process. While this paper used data from Twitter as its basis, the methodology could easily be employed to other similarly structured social media platforms.

Knowing what countries are potentially in crisis is important for several reasons. First, being able to intervene before a country falls apart is far easier than trying to reassemble the broken pieces. Second, these failed states often become havens for criminals, terrorists, and other misanthropes. Third, instability in one country has an unfortunate tendency to leak beyond its borders (Braha 2012). As the whole world becomes more integrated knowing where crisis may emerge becomes more critical.

As important as this research is, it is not impervious to criticism. First, the principle short coming of this paper is the paper is its limited scope and time span. Resource and technical limitations reduced the scope to only two countries and prevented a longer period of data collection. More countries would have increased validity and with a longer collection period it is likely we could have been able to able to capture potential seasonal effects, such as election cycles. A longer time series would also allow for robust forecasting methods to be employed. However, given our limitations, this paper still displays a successful proof of concept.

Second, this method of estimating public dissent is highly dependent upon the social media population of a country. The smaller the population using Twitter in a country, the less useful this estimate will be. However, growing the user base, especially in developing countries, is a strategic goal of Twitter (Twitter 2016). The steady rise in smart phone usage across the world (Poushter 2016) should aid Twitter's market penetration.

Third, some governments are manipulating social media for their own benefit (King 2016; Helmus et al. 2018). This is a concern[12] but one that is likely overstated. Only a few countries in the world have both the desire and the means to negatively affect social media platforms outside their own territory. A social media company's entire business is predicated upon its ability to maintain and grow their user population. If consumers do not feel that they are getting an honest experience on a platform, they will go elsewhere. Thus it becomes the goal of multibillion dollar firms to counter the machinations of rogue states. Twitter has made overtures to ensure a more authentic experience for users by purging millions of fake accounts from the site (Confessore and Dance 2018). Internally it might actual behoove authoritarian regimes to mitigate their manipulation and censorship of social media. Restricting social media does not eliminate social unrest, it merely moves it to more opaque places. Social media provides an easy sources of domestic surveillance allowing regimes to identify potential problem areas (Lorentzen 2014; Qin et al 2017). It's better to know what the people are thinking than to be caught off guard such as in the Tunisia and Ukraine cases.

The final critique of using social media data is that we have selection bias because these platforms tend to have a younger user base. This is certainly the case with Twitter users in

[12]We do not suspect that our sample countries were affected by such targeting.

Canada (Insights West 2016) and Kenya (Simon et al. 2014), where the average Twitter user is in their twenties. However, Rothchild (2015) shows that a properly statistically weighted non-representative sample can still be effective for estimation. Additionally, this potential bias may even be an asset that enhances our measure. The demographic group most likely to use social media, the young, are also the most likely to advocate for large scale political change. Thus by forming the basis of our measure on dissent on social media, we're actually able to account for the group most likely to take to the streets.

The goal of this paper is not to replace traditional means of measuring public dissent. Instead, it is meant as an enhancement incorporating societal and technological changes to increase the accuracy of estimates. Substantive expertise is still necessary to identify key issues, know what language people use for expressing dissent, and how to properly judge the validity of those statements. However, using this framework would help reduce the subjective nature of interpreting public dissent. This paper has served its purpose of providing proof of concept for a new efficient means of estimating public dissent. Given more attention and resources, we believe this could be a valuable avenue of further research.

References

Acemoglu, D. and J. A. Robinson (2001). A Theory of Political Transitions. *American Economic Review 91*(4), 938 – 963.

Acemoglu, D., J. A. Robinson, and D. Woren (2012). *Why nations fail: the origins of power, prosperity and poverty*, Volume 4. SciELO Chile.

Asur, Sitaram and Huberman, Bernardo A (2010). Predicting the future with social media. In *Web Intelligence and Intelligent Agent Technology (WI-IAT), 2010 IEEE/WIC/ACM International Conference on*, Volume 1, pp. 492–499. IEEE.

Bazzi, S., R. A. Blair, C. Blattman, O. Dube, M. Gudgeon, and R. M. Peck (2019). The promise and pitfalls of conflict prediction: Evidence from colombia and indonesia. Technical report, National Bureau of Economic Research.

Blair, Robert A and Blattman, Christopher and Hartman, Alexandra (2017). Predicting local violence: Evidence from a panel survey in liberia. *Journal of Peace Research 54*(2), 298–312.

Bollen, Johan and Mao, Huina and Zeng, Xiaojun (2011). Twitter mood predicts the stock market. *Journal of computational science 2*(1), 1–8.

Braha, Dan (2012). Global civil unrest: contagion, self-organization, and prediction. *PloS one 7*(10), e48596.

Burbidge, John B and Magee, Lonnie and Robb, A Leslie (1988). Alternative transformations to handle extreme values of the dependent variable. *Journal of the American Statistical Association 83*(401), 123–127.

Carley, Kathleen M and Pfeffer, Jurgen and Liu, Huan and Morstatter, Fred and Goolsby, Rebecca (2013). Near real time assessment of social media using geo-temporal network analytics. In *Advances in Social Networks Analysis and Mining (ASONAM), 2013 IEEE/ACM International Conference on*, pp. 517–524. IEEE.

Confessore, Nicholas and Dance, Gabriel J.X. (2018, Jul). Battling fake accounts, twitter to slash millions of followers. *New York Times*.

Dafoe, Allan and Lyall, Jason (2015). From cell phones to conflict? reflections on the emerging ict–political conflict research agenda. *Journal of Peace Research 52*(3), 401–413.

DeMaio, Theresa J (1984). Social desirability and survey. *Surveying subjective phenomena 2*, 257.

Edmund Blair (2016). Kenyan political unrest raises fears of new flare-up in 2017 vote. *Reuters*.

Elson, Sara Beth and Yeung, Douglas and Roshan, Parisa and Bohandy, Susan R and Nader, Alireza (2012). *Using social media to gauge Iranian public opinion and mood after the 2009 election*. Rand Corporation.

Enikolopov, Ruben and Makarin, Alexey and Petrova, Maria (2017). Social media and protest participation: Evidence from russia. *Available at SSRN 2696236*.

Freedom House (2017). *Freedom in the World 2017*. Rowman & Littlefield.

Fukuyama, F. (2014). *Political order and political decay: From the industrial revolution to the globalization of democracy*. Macmillan.

Gerner, D. J., P. A. Schrodt, O. Yilmaz, and R. Abu-Jabr (2002). Conflict and mediation event observations (cameo): A new event data framework for the analysis of foreign policy interactions. *International Studies Association, New Orleans*.

Goldstone, J. A., R. H. Bates, D. L. Epstein, T. R. Gurr, M. B. Lustik, M. G. Marshall, J. Ulfelder, and M. Woodward (2010). A Global Model for Forecasting Political Instability. *American Journal of Political Science 54*(1), 190–208.

Goldstone, Jack A and Bates, Robert H and Epstein, David L and Gurr, Ted Robert and Lustik, Michael B and Marshall, Monty G and Ulfelder, Jay and Woodward, Mark (2010). A global model for forecasting political instability. *American Journal of Political Science 54*(1), 190–208.

González-Bailón, Sandra and Borge-Holthoefer, Javier and Rivero, Alejandro and Moreno, Yamir (2011). The dynamics of protest recruitment through an online network. *Scientific reports 1*.

Granger, C. W. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, 424–438.

Grossman, H. I. (1991). A General Equilibrium Model of Insurrections. *American Economic Review 81*(4), 912 – 921.

Gurr, Ted Robert (2015). *Why Men Rebel*. Routledge.

Guttman, J. M. and R. Reuveny (2014). On Revolt and Endogenous Economic Policy in Autocratic Regimes. *Public Choice 159*(1-2), 27 – 52.

Helmus, Todd C and Bodine-Baron, Elizabeth and Radin, Andrew and Magnuson, Madeline and Mendelsohn, Joshua and Marcellino, William and Bega, Andriy and Winkelman, Zev (2018). *Russian Social Media Influence: Understanding Russian Propaganda in Eastern Europe*. Rand Corporation.

Hirose, Kentaro and Imai, Kosuke and Lyall, Jason (2013). Can civilian attitudes predict civil war violence? *Social Science Research Network,[Online]. Available: at SSRN: http://ssrn.com/abstract 2446168*.

Howell, Llewellyn D (1998). *The Handbook of Country and Political Risk Analysis*. the PRS Group.

Insights West (2016). 2016 canadian social media monitor. Technical report, Insights West.

Kallus, Nathan (2014). Predicting crowd behavior with big public data. In *Proceedings of the 23rd International Conference on World Wide Web*, pp. 625–630. ACM.

King, Gary and Pan, Jennifer and Roberts, Margaret E (2017). How the chinese government fabricates social media posts for strategic distraction, not engaged argument. *American Political Science Review 2*.

Korolov, Rostyslav and Lu, Di and Wang, Jingjing and Zhou, Guangyu and Bonial, Claire and Voss, Clare and Kaplan, Lance and Wallace, William and Han, Jiawei and Ji, Heng (2016). On predicting social unrest using social media. In *2016 IEEE/ACM international conference on advances in social networks analysis and mining (ASONAM)*, pp. 89–95. IEEE.

Kuran, Timur (1989). Sparks and prairie fires: A theory of unanticipated political revolution. *Public choice 61*(1), 41–74.

Lampos, Vasileios and Preotiuc-Pietro, Daniel and Cohn, Trevor (2013). A user-centric model of voting intention from social media. In *ACL (1)*, pp. 993–1003.

Larson, J. M., J. Nagler, J. Ronen, and J. A. Tucker (2019). Social networks and protest participation: Evidence from 130 million twitter users. *American Journal of Political Science 63*(3), 690–705.

Leetaru, Kalev and Schrodt, Philip A (2013). Gdelt: Global data on events, location, and tone, 1979–2012. In *ISA annual convention*, Volume 2, pp. 1–49. Citeseer.

Lorentzen, Peter (2014). China's strategic censorship. *American Journal of Political Science 58*(2), 402–414.

MacCulloch, R. (2001). What Makes a Revolution? Technical report, ZEI working paper.

MacCulloch, R. (2005). Income Inequality and the Taste for Revolution. *Journal of Law and Economics 48*(1), 93 – 123.

Minerva Initiative (2014). Minerva Research Summaries & Resources Fall 2014. Technical report, US Department of Defense.

Morgan, John and Stocken, Phillip C (2008). Information aggregation in polls. *The American Economic Review 98*(3), 864–896.

O'Brien, Sean P (2010). Crisis early warning and decision support: Contemporary approaches and thoughts on future research. *International studies review 12*(1), 87–104.

O'Connor, Brendan and Balasubramanyan, Ramnath and Routledge, Bryan R and Smith, Noah A (2010). From tweets to polls: Linking text sentiment to public opinion time series. *ICWSM 11*(122-129), 1–2.

Overbey, Lucas A and Batson, Scott C and Lyle, Jamie and Williams, Christopher and Regal, Robert and Williams, Lakeisha (2017). Linking twitter sentiment and event data to monitor public opinion of geopolitical developments and trends. In *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction and Behavior Representation in Modeling and Simulation*, pp. 223–229. Springer.

Pan, Jennifer and Siegel, Alexandra A (2020). How saudi crackdowns fail to silence online dissent. *American Political Science Review 114*(1), 109—-125.

Pang, Bo and Lee, Lillian (2008). Opinion mining and sentiment analysis. *Foundations and Trends® in Information Retrieval 2*(1–2), 1–135.

Pierskalla, Jan H and Hollenbach, Florian M (2013). Technology and collective action: The effect of cell phone coverage on political violence in africa. *American Political Science Review 107*(2), 207–224.

Poblete, Barbara and Garcia, Ruth and Mendoza, Marcelo and Jaimes, Alejandro (2011). Do all birds tweet the same?: characterizing twitter around the world. In *Proceedings of the 20th ACM international conference on Information and knowledge management*, pp. 1025–1030. ACM.

Poushter, Jacob (2016). Smartphone ownership and internet usage continues to climb in emerging economies. *Pew Research Center 22*.

PRS Group (2011). *The 2011 Political Risk Yearbook: Middle East North Africa*, Volume 2. The PRS Group.

PRS Group (2013). *The 2013 Political Risk Yearbook: East Europe*, Volume 7. The PRS Group.

Qin, Bei and Strömberg, David and Wu, Yanhui (2017). Why does china allow freer social media? protests versus surveillance and propaganda. *Journal of Economic Perspectives 31*(1), 117–40.

Ramakrishnan, Naren and Butler, Patrick and Muthiah, Sathappan and Self, Nathan and Khandpur, Rupinder and Saraf, Parang and Wang, Wei and Cadena, Jose and Vullikanti, Anil and Korkmaz, Gizem and others (2014). 'beating the news' with embers: forecasting civil unrest using open source indicators. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pp. 1799–1808. ACM.

Ringsquandl, M. and D. Petkovic (2013). Analyzing political sentiment on twitter. In *AAAI Spring Symposium: Analyzing Microtext*, pp. 40–47.

Seabold, Skipper and Rutherford, Alex and De Backer, Olivia and Coppola, Andrea (2015). The pulse of public opinion: using twitter data to analyze public perception of reform in el salvador.

Shapiro, Jacob N and Siegel, David A (2015). Coordination and security: How mobile communications affect insurgency. *Journal of Peace Research 52*(3), 312–322.

Silge, Julia and Robinson, David (2018). *Text Mining with R: A Tidy Approach*. O'Reilly.

Simon, Tomer and Goldberg, Avishay and Aharonson-Daniel, Limor and Leykin, Dmitry and Adini, Bruria (2014). Twitter in the cross fire—the use of social media in the westgate mall terror attack in kenya. *PloS one 9*(8), e104136.

Smith, Ben and Wooten, Jadrian (2013). Pundits: The confidence trick: Better confident than right? *Significance 10*(4), 15–18.

Spangler, Ethan and Smith, Ben (2015). A dynamic model of political stability and dissent. In *2015 Southern Economic Association Annual Conference*.

Statista (2017). Number of twitter users in canada from 2012 to 2020 (in millions). Technical report, Statista.

Steinert-Threlkeld, Z. C., D. Mocanu, A. Vespignani, and J. Fowler (2015). Online social networks and offline protest. *EPJ Data Science 4*(1), 19.

Stephens-Davidowitz, Seth (2017). *Everybody Lies*. Harper Collins.

Tetlock, Philip (2005). *Expert Political Judgment: How good is it? How can we know?* Princeton University Press.

The BBC (2016). Kenyans in fear of police 'death squads'. `www.bbc.com/news/world-africa-36733808`. Accessed: 2018-12-23.

The Economist (2014). A million conversations now: Twitter in kenya. *The Economist Group*.

The Economist (2016). Of kicking and kikuyu. *The Economist Group*.

Tullock, Gordon (1971). The paradox of revolution. *Public Choice 11*(1), 89–99.

Tullock, Gordon (1974). *The social dilemma: The economics of war and revolution*. University publications.

Twitter (2016). Twitter annual report 2016. `http://bit.ly/2yPxFDd`.

Ungar, Lyle and Mellors, Barb and Satopää, Ville and Baron, Jon and Tetlock, Phil and Ramos, Jaime and Swift, Sam (2012). The good judgment project: A large scale test of different methods of combining expert predictions. Technical report, AAAI Technical Report.

Valenzuela, Sebastián (2013). Unpacking the use of social media for protest behavior: The roles of information, opinion expression, and activism. *American Behavioral Scientist 57*(7), 920–942.

Vice News (2016, July). Kenyans in fear of police 'death squads'. `https://bit.ly/2Sheuvv`. Accessed: 2018-12-23.

Wang, Wei and Kennedy, Ryan and Lazer, David and Ramakrishnan, Naren (2016). Growing pains for global monitoring of societal events. *Science 353*(6307), 1502–1503.

Weidmann, Nils B (2015). Communication, technology, and political conflict: Introduction to the special issue. *Journal of Peace Research 52*(3), 263–268.

Wojcik, S. and A. Hughes (2019). Sizing up twitter users. *Washington, DC: Pew Internet & American Life Project. Retrieved May 1*, 2019.

Wright, Austin L. and Condra, Luke N. and Claudy, Sarah (2019). scandals and electoral accountability evidence from a natural experiment. *Working Paper*.

Zeitzoff, Thomas (2011). Using social media to measure conflict dynamics: An application to the 2008–2009 gaza conflict. *Journal of Conflict Resolution 55*(6), 938–969.

Zeitzoff, Thomas and Kelly, John and Lotan, Gilad (2015). Using social media to measure foreign policy dynamics: An empirical analysis of the iranian–israeli confrontation (2012–13). *Journal of Peace Research 52*(3), 368–383.

	Kenya	Canada	
Political	170	142	Political figures & institutions, elections, corruption
Economic	16	20	Unemployment, poverty, recession
Quality of Life	15	14	Health, environment, pollution
Society	32	34	Cultural issues, inequality
Security	12	12	Terrorism, crime, national security
Total	245	229	

Table 1: Regular Expressions by Category

([^.?!]*)(\b(Leader)\b)([^.?!]*)((idiot(ic)?)|(stupid(ity)?)|(incompetent))
([^.?!]*)(\b(Leader)\b)([^.?!]*)((corrupt(ion)?)|(crook(ed)?)|(criminal))
([^.?!]*)(\b(Leader)\b)([^.?!]*)(tyrant)
([^.?!]*)(\b(Leader)\b)([^.?!]*)(impeach(ed)?)

Table 2: Example Regular Expressions

	Canada	Kenya	Source
GDP per captia (2010 $US)	50,407.3	1,143.4	World Bank
Human Development Index	0.922	0.585	UN
Corruption Index	82	26	Transparency International
Political Rights	1	4	Freedom House
Life Expectancy	82.3	67	World Bank

Table 3: Comparative Statistics 2016

	Canada	Kenya
$N_{Twitter}$	45,136	29,175
% User Population	0.64%	4.16%
Tweets	45,246	29,257
Median $\hat{d_{i,t}}$	1	1
Sentiment Analysis	75.1% Negative	68.9% Negative
Most Common Day	Tuesday	Monday
Median Weekly $\hat{D}_t$	4,596	2,010/16,924*

*=scaled by user populations

Table 4: Twitter Data Summary Statistics

Mann-Whitney U Test
P-value 0.000
Kolmogoro-Smirnov Test
P-value 0.0482

Table 5: Statistical Distribution Tests

	Canada	**Kenya**
1.	#cdnpoli	Tribalism
2.	Terrorism	Election Commissions
3.	Immigration/Refugees	Kibera/Slums
4.	Inflation/Prices	Immigration/Refugees
5.	Taxes	Injustice

Table 6: Top 5 Topics

Canada	Min	Max	Mean
Civil Unrest Events	0	72	10.78
Daily $\hat{D}_t$	0	6,894	540.7
Kenya	Min	Max	Mean
Civil Unrest Events	0	34	7.47
Daily $\hat{D}_t^*$	0	45,072	2,917
*=scaled			

Table 7: Empirical Testing Summary Statistics

Civil Unrest Incidents	**Specification**	**Lag Order**	**F Statistic**
Canada	$log(\hat{D}_t + (\hat{D}_t^2 + 1)^{\frac{1}{2}})$	1	3.793*
Kenya	$\hat{D}_t$	2	5.749***

***=1% significance, *=10% significance

Table 8: Granger Causality Test

Figure 1: Twitter Data Collection and Processing

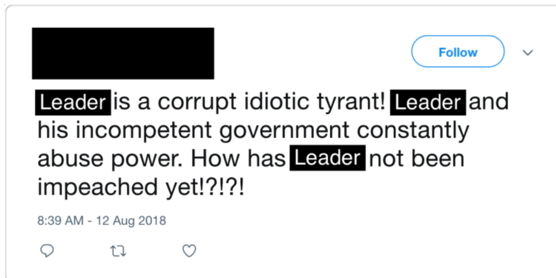

Figure 2: Example Tweet

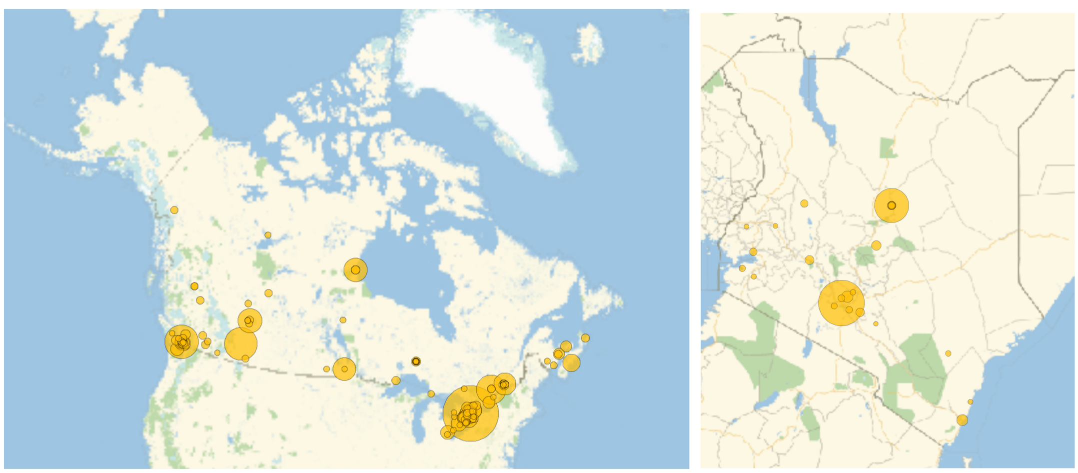

Figure 3: Tweet Distribution and Densities in Canada and Kenya

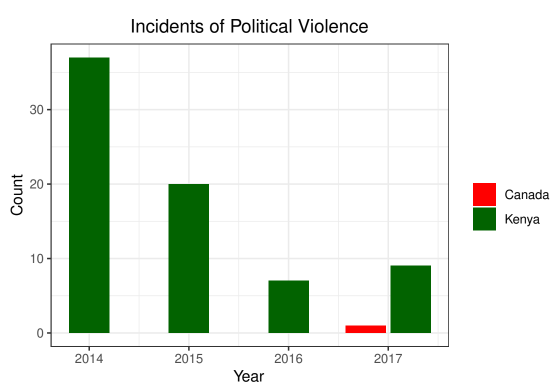

Figure 4: Incidents of Political Violence 2014-2017

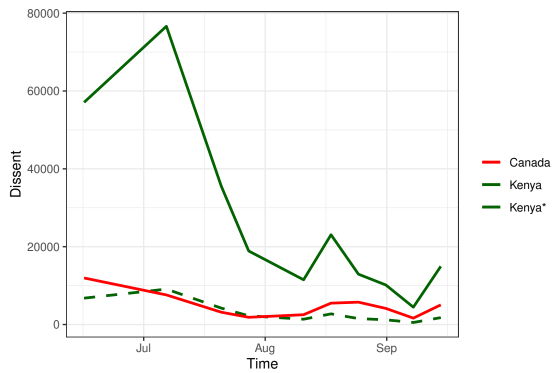

Figure 5: Weekly Total Dissent

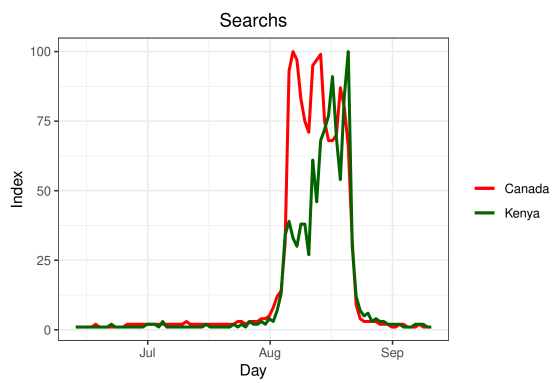

Figure 6: Daily Internet Searches for 'Olympics'

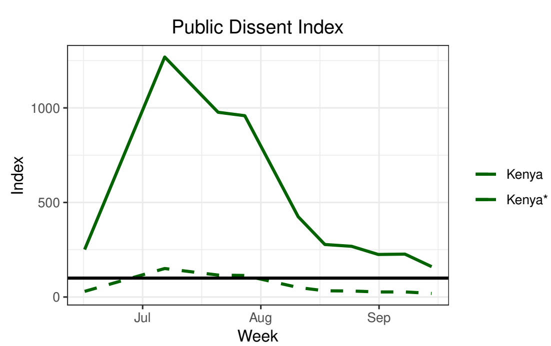

Figure 7: Weekly Public Dissent Index