Common Knowledge and Protest

James Fowler, Delia Mocanu, Zachary Steinert-Threlkeld, Alex Vespignani

June 27, 2014

Abstract

Does the spread of information affect protests? If so, does it matter how that information spreads? We argue that two key features help individuals translate that information into protest. First, individuals are more likely to protest if they know that others know that they share these sentiments; in other words, common knowledge amongst individuals can lead to protest. Second, individuals are more likely to believe that common knowledge exists when those from whom they learn their information are themselves influential. To measure protests, we use machine-coded events data from the Global Dataset on Events, Location, and Tone (Leetaru & Schrodt 2013). GDELT identifies 120,000 unique protests in Algeria, Bahrain, Egypt, Iraq, Jordan, Kuwait, Libya, Lebanon, Morocco, Oman, Qatar, Saudi Arabia, Syria, Tunisia, the United Arab Emirates, and Yemen from November $1st$, 2010 and December 31^{st}, 2011. To measure common knowledge and individuals' influence, we take advantage of geolocated data from Twitter. Across this study's 16 countries, we have collected 17,000,000 individual messages and analyzed each one for key features of common knowledge creation and individual influence. We find strong support for the importance of common knowledge. More common knowledge created on Day 1 correlates strongly with more protest on Day 2, even after controlling for temporal and country variation. Moreover, it matters who creates that common knowledge. When individuals who are influential - who have many followers reading their messages - contribute to common knowledge, more protest is seen the following day. As far as we are aware, this is the first paper which analyzes protest dynamics on a daily level across such a large time frame and number of countries. It also appears to be the first time data from Twitter is incorporated into a regression model.

WORKING PAPER - DO NOT CITE
Prepared for the 2014 Political Networks Conference

1 Introduction

Does the spread of information affect protests? If so, does it matter how that information spreads? An affirmative answer to these questions is not immediately obvious. Individuals can know that others harbor resentment against a regime's repression or disappointment over economic growth but not decide to protest, both of which could cause an individual to protest. We argue that two key features help individuals translate that information into protest. First, individuals are more likely to protest if they know that others know that they share these sentiments; in other words, common knowledge amongst individuals can lead to protest. Second, individuals are more likely to believe that common knowledge exists when those from whom they learn their information are themselves influential. That is, individuals weight information by the identity of the person spreading, and a key component of the spreader's identity is that person's popularity (how many other people learn the same information from the spreader). Taken together, more common knowledge, and more common knowledge generated by popular individuals, should lead to more protest.

The decision to protest does not occur in a vacuum. Instead, individuals look around themselves, incorporate some new information, and update their willingness to protest; if that willingness surpasses an internal threshold, they protest (Granovetter 1978). Scholars have established a long list of explanatory variables that correlate with protest - disappointing economic growth (Gurr 1971, Brancati 2013), ethnic tension (Posen 1993, Cederman, Girardin & Gleditsch 2009), or regional diffusion (Weyland 2012), for example. Implicit to these theories is the presence of information that economic growth should have been higher, there exists a reason to dislike non-coethnics, or that similar countries have recently experienced protest. The existence of individuals' having acquired the information that leads them to develop a desire to protest is therefore foundational to these theories, but how they acquire that knowledge remains under-explored.

A surge in protests in the last few years has led to a renewed interest in understanding the causes of protest, the variables that influence their ebb and flow, and factors which determine their ultimate success or failure. Perhaps the most prominent large scale, widespread protests of the previous two decades are those collectively known as the Arab Spring. The phrase "Arab Spring" refers to the series of protests which started in Tunisia on December 14^{th}, 2010 (leading to the resignation of that country's president), slowly spread to neighboring countries over the following

6 weeks, inspired massive turnout in Egypt that caused Hosni Mubarak to resign on February 11^{th}, 2011. Soon saw anti-regime protests occur throughout North Africa and the Middle East (Weyland 2012, Hale 2013).

The Arab Spring is an ideal case to test the effect of common knowledge and individual influence for four reasons. First, studies whose data come from one country risk attributing outcomes to variable(s) that may in fact be explained by country characteristics and not the variable(s) itself. Second, the events of the Arab Spring occurred at different times and lasted for many months, providing substantial variation on our dependent and independent variables. Third, protests are more difficult in authoritarian regimes than in democracies (Kricheli, Magaloni & Livne 2011), biasing against finding any results. The results we do find can therefore be expected to be more pronounced in more democratic regimes. Fourth, the Arab Spring is perhaps the only transnational protest event which has occurred after the arrival of tools which allow us to accurately measure our variables. While the theory we develop should apply to protests in 1848 or 1989, the data do not, and probably never will, exist to test common knowledge and individual influence in those periods.

To measure protests, we use machine-coded events data from the Global Dataset on Events, Location, and Tone (Leetaru & Schrodt 2013). GDELT reads news stories, extracts the main actors and action in that story, and generates a data point for each event. GDELT splits events into one of 20 categories, one of which is protest. We extract all protests from Algeria, Bahrain, Egypt, Iraq, Jordan, Kuwait, Libya, Lebanon, Morocco, Oman, Qatar, Saudi Arabia, Syria, Tunisia, the United Arab Emirates, and Yemen between November 1^{st}, 2010 and December 31^{st}, 3011. GDELT has identified over 120,000 unique protests during the period of this study, and only 25% of our country-days contain no protests. All data are at a daily level.

To measure common knowledge and individuals' influence, we take advantage of geolocated data from Twitter. Twitter is a popular microblogging service through which individuals compose and consume short messages. Conventions on Twitter, specifically hashtags and retweets, facilitate the creation of common knowledge. Twitter also has a social network component, as individuals can follow others' accounts and have their own followers. Like all social networks, Twitter exhibits an extreme inequality in degree distribution, which is to say that there are a small number of users with very many followers and very many users with few followers. We only analyze tweets from users who allow GPS coordinates on their messages as well as those who report their location, ensuring that our results reflect domestic processes and are not biased by individuals external to the events (Lotan, Ananny, Gaffney, Boyd, Pearce & Graeff 2011, Aday, Freelon, Farrell, Lynch & Sides 2012). Across this study's 16 countries, we have collected 17,000,000 individual messages and analyzed each one for key features of common knowledge creation and individual influence. While we know the second the tweet was authored, we have aggregated all tweets to the country-day level.

Measuring the amount of common knowledge created per day, the influence of individuals who create that knowledge, and the number of protests on subsequent days, we find strong support for the importance of common knowledge. The more common knowledge that is created on Day 1 correlates strongly with more protest on Day 2, even after controlling for temporal and country features. Moreover, it matters who creates that common knowledge. When individuals who are influential - who have many followers reading their tweets - contribute to common knowledge, more protest is seen the following day. As far as we are aware, this the first paper which analyzes protest dynamics on a daily level across such a large time frame and number of countries and is the first time data from Twitter is incorporated into a regression model.

2 Common Knowledge, Influential Individuals, and Protest

Common Knowledge

Individuals face a variety of challenges when deciding whether or not to protest. They have to estimate the risk of arrest or violent state repression, and they have to weigh that against the potential benefits of any change in policy that may result (Davenport 2007). They would like to know how many other individuals are going to protest, or at least whether the crowd will be large or small. They will also want to know the composition of those who join - is it mostly hardcore activists, an average citizen, or some combination thereof (Lohmann 1994)? Because the answers to these questions depend to a large extent on how many people joint the protests, protests are commonly thought of as coordination games. Protesting *en masse* decreases the cost of repression the individual will face and increases the probability of achieving policy change. The individual wanting to protest is therefore strongly incentivized to make sure to coordinate his protest action with other individuals who want to protest or already are.

Common knowledge facilitates coordination, including protest coordination. A key for protests to emerge and sustain themselves is therefore the creation of common knowledge within a population about the regime's popularity and the details of subsequent protests (Richards 2001). Common knowledge is the set of public information which everyone knows, everyone knows that everyone else knows this knowledge, everyone knows that everyone knows that everyone knows this knowledge, and so on. Knowing that everyone else knows what I know facilitates behaviors that have positive externalities, whether that behavior is the adoption of a new technology or joining a protest (Chwe 2003). For example, there is no point to buying a fax machine if you have the only machine, so a widespread, expensive advertising campaign tells the individual that many people know about the machine and are likely to have it, making you more likely to buy it. Similarly, one does not want to protest if he will be the only one protesting; the costs are too high and the benefits are certainly lower. But if it is common knowledge that others want to protest, an individual is much more likely to protest, *even without knowing specifically how many others will protest or what the state's response will be.*

It should be noted before proceeding further that this paper uses "common knowledge" in this paper to refer to what is commonly called "almost common knowledge". True common knowledge is of the sort where Player 1 knows Fact A, Player 2 knows Fact A, Player 1 knows Player 2 knows Fact A, Player 2 knows Player 1 knows Fact A, Player 1 knows Player 2 knows Player 1 knows Fact A, Player 2 knows Player 1 knows Player 2 Fact A, so on *ad inifnitum.* While common knowledge can attain in formal models, it is hard to see how true common knowledge ever attains in a population. Instead, what actually exists is certainly closer to "almost common knowledge". 30 individuals may each know that the others know to protest at 6 p.m. every Monday, and each may even know that the others know that everyone knows. But does each person know that the other 29 know that he knows that they know? Or that the other 29 knows that he knows that they know that he knows? Strictly speaking, probably not, because the group would spend all its time confirming the infinite recursiveness of its common knowledge and so never actually protest (Rubinstein 1989).

Instead, this paper assumes that individuals will act as though there is common knowledge if a personal threshold of perceived common knowledge is surpassed. In the fax machine example, the early adopter does not need to know that everyone saw the advertisement and that everyone knows that everyone saw the advertisement and that everyone knows that everyone knows that everyone knows that everyone saw the advertisement, so on forever. The early adopter will assume that a large number of people saw the advertisement, some of whom will also purchase the fax machine. Likewise for protest: an individual does not need to know that everyone else in the country has heard the call to protest, only that enough people have heard it to make it likely that enough people will protest. For the rest of the paper, the use of "common knowledge" is therefore synonymous with this form of "almost common knowledge".

The observation that common knowledge is necessary for protest is not new, and many theories of protest implicitly incorporate it. Timur Kuran argues that authoritarian regimes maintain power by convincing everyone to display false preferences; not knowing others' true preference, each individual thinks his or her actual preference is far from average and so does not protest. An exogenous shock causes individuals to reveal their true preferences, and a bandwagon effect soon leads to widespread protests that can overthrow a regime (Kuran 1989). In other words, authoritarian regimes work to create false common knowledge about their popularity, but once individuals realize that this common knowledge is wrong, they quickly shift to a new set of common knowledge which can lead to the overthrow of the regime. Susan Lohmann theorizes that the participation of activist moderates in protests provides information to potential protestors about a regime's true popularity, and a large enough turnout of activist moderates will convince moderates to protest, leading to a new regime (Lohmann 1994). In this model, everyone can observe the ebb and flow of protests, and this cadence creates common knowledge amongst a population; if enough activist moderates join the protests, the set of common knowledge shifts away from one of regime popularity and strength.

The centrality of common knowledge for protestors leads to the first hypothesis:

H1 *As common knowledge about protests increases, protest is more likely to occur.*

Individuals need to coordinate where and when to protest, and focal points provide a common means to coordinate those details (Schelling 1960). Focal points can inform individuals both about where and when to convene. Public areas (central squares, parks), large gatherings (sporting events), or prominent buildings (Grand Central Station) thus provide convenient, commonly known points where individuals are likely to gather (since they gather there already). Knowing when to protest follows a similar logic: if there is a period in time when large amounts of people will be somewhere, like at a church service or soccer match, then it makes sense to protest around that time. A famous example of a focal point in place and time is that of the Nikolai Church in Leipzig, East Germany; it hosted peace prayers every Monday from 5:00-6:00 p.m., and the protests that would lead to the downfall of the German Democratic Republic started every Monday following those prayers (Lohmann 1994). Elections also provide a convenient focal point in time: if a ruler cancels planned elections, that cancellation provides a public signal about the regime's performance, and the date of the election becomes a focal point to coordinate protests (Tucker 2007, Fearon 2011, Little 2012, Hyde & Marinov 2014).

Common knowledge of focal points is easier to obtain than common knowledge that others want to protest. That common knowledge must be created through communication, the conveyance of information between two participants. It can be one way (television or radio) or interactive, as in a conversation. These two broad types of communication, interactive and not, roughly correspond to mass media communication and interpersonal communication. Media disseminates a common set of information to a large audience, contributing to the common knowledge of the members of that audience (Anderson 1991). The economies of scale of mass media - the ability to reach a large audience through the control of a few pieces of infrastructure - makes it amenable to centralized control. Being amenable to centralized control, many states find it useful to censor or own media in their territory (Egorov, Guriev & Sonin 2009). Doing so gives them the ability to directly influence individuals' common knowledge and inhibit the creation of alternative sets of facts, hindering the spread of anti-regime messages (Edmond 2013). States find this ability particularly useful when they want to create compliance with their rule without resorting to violence (Warren 2014).

Because the economies of scale to controlling interpersonal communication are much lower than for mass media, it is more difficult for authoritarian regimes to control interpersonal communication. Economies of scale in interpersonal communication are much lower than mass media for two reasons: by definition, it requires person to person contact, and each point of contact risks degradation or misinterpretation of the original message. The decentralized and unreliable nature of interpersonal communication makes it difficult to impossible for an authoritarian regime to control, but it is just as difficult for individuals to create common knowledge amongst themselves.[1] For example, when a potential protestor learns that his colleague is going to protest tomorrow, he does not know how many other people the colleague has told or how the colleague knows about the protest (perhaps he is a government informant).

If individuals have complete isolation (no connections to each other) or complete integration (everyone interacts with everyone), common knowledge cannot be formed. On one extreme, if no one interacts with anyone else, each person is an island of knowledge, and common knowledge can never arise. On the other hand, everyone interacting with everyone else means that each individual is exposed to the same set of knowledge. What that knowledge is therefore "common", its content will never change - there will never be shifts to new common knowledge equilibria - because each individual receives the same countervailing information (Watts 2002, Centola & Macy 2007).

In actuality, social networks contain a few individuals who communicate with, and therefore transmit knowledge to, many more people than the median person. These popular individuals therefore have a much greater influence on common knowledge than the average person. The presence of highly connected individuals means that a social network does not have complete integration, which is to say that the network contains different sets of knowledge. While there may be pockets of common knowledge (cliques), there does not exist *ex ante* the same set of knowledge one would find in a fully-connected network.

Moreover, these highly connected individuals have a greater influence on common knowledge creation than the average individual precisely because they are connected to so many people. The logic is very similar to what makes mass media appealing to rulers: if one individual could reach millions of individuals only through their social network, they would be just as influential as a regime's nightly news program. Very few individuals do have that level of influence. Those that do tend to be charismatic heads of large organizations which are actively suppressed or co-opted; Mohamed Badie, the head of Egypt's Muslim Brotherhood, and Fethullah Gülen, a Turkish iman exiled in Pennsylvania, epitomize this kind of person. Though most well-connected individuals do not have this level of influence, they nonetheless spread information to more people than most people do. Communication coming from these highly connected individuals should therefore have a greater

[1] This is why mass media is so powerful: convincing the citizen to support the regime obviates the need to monitor its citizens' every interaction for anti-regime information.

influence on common knowledge than those who are less connected (Christakis & Fowler 2010, Borge-Holthoefer, Rivero & Moreno 2012, Jalili 2012, Garcia-Herranz, Moro, Cebrian, Christakis & Fowler 2014).

The importance of social networks for spreading information, and the influence in those networks of highly-connected individuals, leads to the second main hypothesis:

H2 *As more influential members create common knowledge, protest becomes more likely.*

Scope Limitations

This theory takes no account of the content of messages, which is to say that the type of common knowledge is not relevant at this stage. Citizens deciding whether or not to protest need to generate three main kinds of common knowledge: (1) knowledge that others share their anti-regime beliefs (Kricheli, Magaloni & Livne 2011), (2) knowledge about the logistics of protest (where, when, what to bring, etcetera), and (3) knowledge about the protests as they occur (the level of repression, the size of turnout, and the composition of that turnout, primarily). Each of those facets of common knowledge have their own impact on subsequent protests; the first two should increase it, the latter may decrease it depending on various factors. The theory as circumscribed here - that more hashtags and retweets create more common knowledge and so incite more protest - therefore assumes that the first two's effects outweigh the third's (possibly) negative effect.

Theorizing about what kinds of protest communication exist and to what extent they affect subsequent protests is important but outside the scope of this paper. In other words, this theory makes no claim as to the content of protest communication. "Protest communication" is a broad category that encompasses very specific types of communication. "The government is looting the people", "Down with the secularists", "Meet me at the roundabout after prayers", or "The police aren't arresting people" have very different meanings. The first two are general expressions of discontent that lets others know about general levels of dissatisfaction in a country; its effect on protest should be weak and delayed. The third specifically focuses on coordinating an upcoming protest and should therefore strongly correlate with subsequent events. The fourth concerns a protest ongoing and should convince more people to join the protest, though it may not spur more protests than are already occurring. These sample messages are all clearly about protest, so an increase in their prevalence should correlate, to a first approximation, to more subsequent protest.

3 Previous Work

When protests start is a complex function of factors as wide ranging as the resources available to individuals (McCarthy & Zald 1977, Jenkins 1983), their political opportunities (McCammon, Campbell, Granberg & Mowery 2001, Meyer 2004), economic grievances (Gurr 1971, Fearon & Laitin 2003), and ethnicity (Posen 1993, Cederman, Girardin & Gleditsch 2009). The decision to protest is further influenced by the repression strategies a state undertakes (Davenport 2007). Moreover, determining exactly when and why protests first occur is an active area of research (Weyland 2012), and most theories of protest do not endogenize the initial protest (Kuran 1989, Lohmann 1994, Rasler 1996, Moore & Newman 2000). For a more thorough review than what can be covered here, see Benford (2000).

Scholars of protest and revolution has also become to incorporate the Arab Spring into these theories. Kurt Weyland argues that protestors observed the flight of President Ben Ali from Tunisia and used cognitive shortcuts to update their beliefs about the feasibility of domestic regime overthrow. These shortcuts, specifically the availability and representative heuristics, caused individuals in 2011, as well as in 1848, to overestimate their odds of similar success: “they strike at inopportune moments and challenge autocrats who sit in the saddle more firmly than these challengers believed” (Weyland 2012, pg. 922). Comparing the diffusion of regime change in the Arab Spring to similar diffusion in 1848, 1989, and 1998-2005, Henry Hale concludes that structural conditions and a lack of focal points for internal elite defection explains the spread of protests across much of the Arab world in 2011 (Hale 2013). Global exogenous shocks, such as food price volatility, may also be a contributing factor to the widespread emergence of protests (Lagi, Bertrand & Bar-Yam 2011).

This is also the first paper, as far as we are aware, which uses Twitter data in a regression framework. Previous papers that use similar data have either kept their analysis at a descriptive level or case studies. Owing to the novelty of the data source, the earliest work on Twitter focused on understanding characteristics of the network and patterns of communication. These papers have analyzed how individuals engage in conversation (Honeycutt & Herring 2009, Boyd, Golder & Lotan 2010), the spread of hashtags and when users adopt new ones (Romero, Meeder & Kleinberg

2011), the characteristics of hashtag trends (Becker, Gravano & Naaman 2011), and what can be learned from retweets (Kwak, Lee, Park & Moon 2010, Nagarajan, Purohit & Sheth 2010, Suh, Hong, Pirolli & Chi 2010). A large literature has arisen to study how Twitter is used during crises such as earthquakes (Mendoza, Poblete & Castillo 2010), wildfires (Vieweg, Hughes, Starbird & Palen 2010), and contagious disease outbreaks (Chew & Eysenbach 2010). Many papers also provide descriptive statistics about Twitter use during protests. The earliest protests in which Twitter appears to have been used are Moldova (Mungiu-Pippidi & Munteanu 2009) and Iran (Burns & Eltham 2009, Rahimi 2011). Twitter was also a key tool during the 2009 G20 meeting in Pittsburgh (Earl, McKee Hurwitz, Mejia Mesinas, Tolan & Arlotti 2013), the United Nation's climate summit in Copenhagen (Segerberg & Bennett 2011), Thailand's 2010 street protests (Bajpai 2011), and Occupy Wall Street and its offshoots (Lowe, Theocharis & W. van Deth 2013, Tremayne 2014).

There is a growing body of work that specifically analyzes the use of social media during the Arab Spring, which we extend in 3 ways. First, we present the first work which incorporates Twitter data into regressions, which means we present the first work which concretely connects online activity to offline outcomes. While it is understood that the online and offline spheres are connected, the work that has established that connection has relied on fieldwork methods and surveys (Breuer 2012, Tufekci & Wilson 2012). Second, these studies have focused on communication within one country, primarily Tunisia (Breuer 2012, Lim 2013), Egypt (Solomon 2012, Starbird & Palen 2012, Hanna 2013), or both (Lotan et al. 2011, Aday et al. 2012). Our paper analyzes 16 countries from the Middle East and North Africa, including those, such as Lebanon and the United Arab Emirates, which experience very little to no protest. Third, most of these studies analyze Twitter data without reference to the location of users. Their conclusions are therefore about Twitter globally, which makes it difficult to connect any behaviors to outcomes in specific countries. Restricting our sample to only messages from users in our 16 countries expands the phenomena we can study, allowing us to study not Twitter *qua* Twitter but Twitter as a measurement of other variables of interest (common knowledge and individual influence).

Within political science, this paper joins a growing body of quantitative work at the intersection of information and communication technology and state repression. Jan Pierskalla and Florian Hollenbach find that, in Africa, cell phone coverage increases the probability of violent conflict (Pierskalla & Hollenbach 2013). Jacob Shapiro and Nils Weidmann find the opposite effect in

Iraq; using time-variant data on new cell phone coverage, they find that the provision of cellular coverage decreases insurgent violence (Shapiro & Weidmann 2011). Gary King, Jennifer Pan, and Molly Roberts measure censorship on Chinese blogs; they find that Chinese censors target posts which could generate collective action but are more permissive of writings critical of the Communist Party (King, Pan & Roberts 2013).

The analysis with the most data about new technology and the Arab Spring comes from the United States Institute of Peace's *Blogs and Bullets II: New Media and Conflict After the Arab Spring* (Aday et al. 2012). Using data from Twitter and bit.ly, they identify roughly 35,000 links which received almost 9.5 million clicks from individuals in Bahrain, Egypt, Libya, and Tunisia from January 1^{st} to Aprilst, 2011.[2] Aday and his coauthors find that links receive the most attention from individuals outside of the Middle East and North Africa, and only Bahrain received a moderate amount of attention from individuals within the region. Attention given is spiky, and they suggest that social media is used a "megaphone" to attract international attention.

This paper builds on these works theoretically and empirically. Theoretically, the temporal resolution and ability to read individuals' communication allows the paper to measure the hypothesized mechanism, common knowledge creation, in a way that most previous work has not because of temporal or content limits. Empirically, because of the difficulty of obtaining daily temporal resolution, previous work has has to focus on only those data, the telecommunications data, for which the detailed resolution was available. In their report, Aday and his co-authors describe the limits of previous approaches: "We can offer remarkable new insights into the consumption of information, but we cannot demonstrate that consuming this information led to changes in political attitudes and behavior. Unsurprisingly, we find that media data tell us most about ... media" (Aday et al. 2012, pg. 15). This paper records protest at the daily level and aggregates Twitter messages to the same timescale, allowing it to connect two previous unconnected domains of behavior.

[2] bit.ly is a company which shortens links users sends. For example, www.nytimes.com/world/obama-putin-ukraine-tension could become bit.ly/Qdk4 (fictional example). Shortening is useful for space limited services such as Twitter.

4 Data

We calculate the dependent variable with data from the Global Dataset on Events, Location, and Tone (Leetaru & Schrodt 2013). GDELT is a machine-coded events dataset that codes dyadic events, including protests, from publicly available news reports. (GDELT codes 20 main categories of events; for a full description of each kind, see the codebook for the Conflict and Mediation Event Observations dataset (Gerner, Schrodt, Abu-Jabr & Yilmaz 2002). Each GDELT row records a primary actor, the primary actor's action (the event), and the actor receiving the action, in addition to metadata such as date, GPS coordinates of the actors and events, the tone of the article, and how many articles write about the event. For example, a row might show that Egyptian police (primary actor) in Alexandria (location) increased their alert status (event) in response to protestors (actor receiving action). Multiple reports are aggregated into one event so that each entry records a unique event; the number of news stories and the number of news organizations writing about the event are separately recorded.

One GDELT event category is "Protest". we therefore selected all protest events where each actor was in one of the 16 countries of this study. and we focused on November 1st, 2010 - December 31st, 2011. Protests did not start in Tunisia until the middle of December 2010, so the month of November provides a baseline against which to measure subsequent events. The ending date was chosen because the data for the main independent variable ends on that day.

Since multiple, sometimes hundreds, of protests can occur in any country and are recorded as separate events, we aggregate each country's protests to the daily level. Whereas GDELT's row is location-day, the row for the dependent variable in my dataset represents country-day because of this aggregation. Specific locations were thrown out for two reasons. First, news reports are often so vague as to not permit coding of events at a level lower than country. For example, GDELT records 29,035 protests in Egypt for this study, but only 21,055 can identify the city in which they occurred; for Yemen, the numbers are 12,977 and 10,951 respectively. Second, the independent variable has similar levels of geographic resolution. Running the analysis at the city-day level is ideal, but there are not enough data, especially for the independent variables, to permit that.

The independent variables come from Twitter, a global social network with over 500 million users generating almost 500 million tweets (messages) daily. Anyone with an internet connection or phone can access it, and most users create and consume content using their mobile devices. Only China and North Korea have completely blocked access to it. Twitter the company does not edit or censor its users' tweets, so the content of the network reflects what individuals are discussing at any moment.[3] To collect the data, 10% of all tweets were collected through Twitter's streaming API from October 18^{th}, 2010 through May 17^{th}, 2012 (Mocanu, Baronchelli, Perra, Vespignani, Goncalves & Zhang 2013). GPS coordinates and user-reported location were used to find individuals from this paper's 16 countries, and only tweets from November 1^{st}, 2010 through December 31^{st}, 2011 were kept. The result is 17,364,873 geocoded tweets; see Table **??** for more detail by country.

Twitter was chosen as a data source for four reasons. First, it has become a tool for citizens to gather and disseminate information in information-scarce environments such as authoritarian regimes. It therefore is a critical component of many protest movements, starting with the 2009 Iran election protests and continuing through the Ukraine civil conflict (Radsch 2012, Tufekci & Wilson 2012, Faris 2013). Second, it provides some of the best temporal resolution of any data source. Though one could theoretically trace communication by the second, studies more commonly look at patterns over periods of hours or days; this study aggregates Twitter messages to the daily level. It is therefore one of the few sources available to researchers interested in dynamic processes - protests, financial markets, advertising, etcetera - that can provide micro-level information on these processes. Third, state actors belatedly realized the power of social media, leaving it unregulated; lack of regulation made social media an attractive tool for anyone seeking independent information, and the information contained in social media therefore more closely reflected the offline world than did official news sources (Hamdy & Gomaa 2012).

Finally, there are other digital data sources one could use instead of Twitter. Facebook is also a common organizational and information tool for individuals in authoritarian regimes, and it was used extensively in the Arab Spring (Hanna 2013). Its data are much less accessible than Twitter's. While both Facebook and Twitter provide an application programming interface (API) for third parties, Facebook's requires a researcher to write an application that each user then must permit to access its social network data. Twitter's two APIs, on the other hand, are much more open. The obstacles to retrieving data from Facebook are thus much higher: a researcher would have

[3] Twitter will censor tweets to comply with countries' laws. For example, it has censored a neo-Nazi group's tweets in Germany.

to create an application that a large proportion of users then install or have direct access from Facebook to its data. One could also find individuals' blogs and scrape them for their content. Scraped blog content was not chosen because it would provide less data than Twitter. Bloggers and Twitterers both produce content, but Twitter makes it easier to identify individuals' accounts, observe interaction among individuals, and geolocate individuals.

The basic function on Twitter is to tweet (write a short message), and a user's followers will see this tweet. While a tweet can be just a message, it often has other characteristics that give it different functionality. Individuals will affix a hashtag - the # sign - to the front of a word to associate it with a certain conversation, e.g. "Protestors in #tahrir #egypt show the weakness of dictators". If a different user then searches for messages containing "#tahrir" or "#egypt", this tweet will be returned; employing a hashtag therefore makes the information in one's message more likely to spread beyond just one's social network (Romero, Meeder & Kleinberg 2011). A message that directly refers to another user by name is said to contain a user mention. If a user writes, "@ZacharyST, you'll like this story", @ZacharyST will receive a personal notification about the message; a tweet with a user mention is still viewable by the followers of the original author. A retweet is when a user shares the message of someone he follows with those who follow him; it is akin to forwarding an e-mail to everyone in your address book. Finally, tweets often contain links to photos and articles. We take advantage of these different kinds of tweets, and what they imply for common knowledge creation, to test our hypotheses.

Dependent Variable

The dependent variable - $Protests_{i,t}$ - measures the number of protests in country i on day t.

Since the period of study encompasses 426 days, there are 426 * 16 = 6,816 country-days to analyze. Figure 1 shows the evolution of protests in Egypt, with the y-axis representing what percent of Egypt's events are protests. The first protests on January 25^{th} are evident, as are subsequent major events such as a protest against the Supreme Council of Armed Forces (Friday of Cleaning) and the first elections at the end of the year.

One concern with the dependent variable is that it measures news coverage of protests, not actual protests. This concern is misplaced, for two reasons. First, GDELT has records for protests on days when Western media were not focusing on the Arab Spring. While these records of course

Figure 1: GDELT Recording of Protests

(a) Protest Dynamics in Egypt

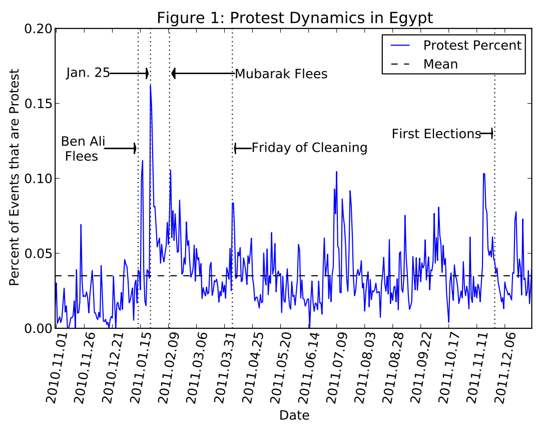

(b) Protest Dynamics in the UAE

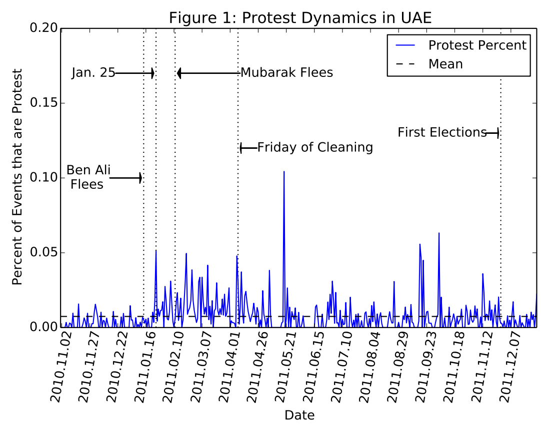

come from news articles, they probably come from sources more focused on, and therefore providing more accurate tallies of protest in, the Middle East and North Africa than the *New York Times*, *Associated Press*, or *BBC*. The extensive temporal recording of protests suggests that subsequent increases in recorded protests are true protests and not the result of having more reporters in those countries. Second, GDELT eliminates all records it creates that have the same values for date, source, target, and event, removing about 20% of initial observations (Leetaru & Schrodt 2013, pg. 21). When these duplicates are found, the event's values for number of articles, number of mentions, and number of sources are updated. The remaining rows should therefore each represent unique protests. This paper's dependent variable is the count of rows per day per country, not the count of articles per day per country.[4]

Figure 1 reinforces this first point. The figure shows that GDELT records protests every day in Egypt, a high activity country, and almost every day in the United Arab Emirates. It also shows spikes in protests around dates that are understood to have experienced a lot of protest, and these spikes occur only around relevant events for Egypt and the United Arab Emirates. In Egypt, there is no uptick in protests when Tunisian President Ben Ali flees, but subsequent protests clearly track major events. The chart for the UAE shows some evidence of heightened protest activity, but none,

[4]As a further validation check, the authors are building their own web scraper to download articles from Reuters, BBC, Voice of America, and United Press International. These articles will be used to generate protest data using Phil Schrodt's TABARI software. Those results will be presented at the 2014 APSA meeting in Washington D.C.

with a possible exception of January 25^{th}, track the events in Egypt.

Independent Variables

Each of the 17,000,000 tweets is examined for receives a dummy for whether it contains a hashtag, a user mention, a link, or is a retweet. If it contains none of these features, it is labeled a basic tweet. Aggregated to the day level, these measures allow us to create our three measures of common knowledge creation.

The first measure, *Hashtags* $\%_t$, is the percent of a day's tweets which contain a hashtag. The second, *Retweets* $\%_t$, measures the percent of a day's tweets with are retweets. Hashtags and retweets are used concurrently to spread knowledge of an event. When a user writes a message, she can insert a hashtag - the # sign - before a word, signifying that the message is about the word with the hashtag in front of it. For example, the message "You guys! We're about to head to a meeting point in front of Marie Louis store in batal ahmed street #jan25" contains the "#jan25" hashtag, a common hashtag used to talk about protests in Egypt, even after January 25^{th}. This hashtag signifies that the author's message is about the events of that day (January 25^{th} was the first day of major protests in Egypt). Users quickly converge on a few hashtags to use for an event, whether that event is a protest, sporting event, or pop culture meme (Bruns & Burgess 2011, Lehmann, Gonçalves, Ramasco & Cattuto 2012).

Moreover, Twitter makes it very easy to find all tweets containing a hashtag. A user interested in upcoming protests could therefore search, from her phone or a computer, for "#jan25", "#egypt", or other hashtags and retrieve every tweet containing those hashtags. That person is therefore quickly exposed to vastly more information than she could gain from traditional interpersonal communication, and she knows that everyone else searching those hashtags will see the same tweets. She is therefore confident that when she reads about the meeting in Batal Ahmed street, many others have read about it as well, and others who search for "#jan25" know that others have seen that tweet as well. The prevalent use of hashtags, convergence to very few during major events, and ease of finding information related to the hashtag make tweets with hashtags a key component of common knowledge creation.

Retweets also promote common knowledge creation. An example of a retweet is, "RT @Ekramibrahim: Police, specially in civil clothes are holding electricity sticks. #jan25". Ekramibrahim is the author of the message after the colon, but the person who sent the tweet read Ekramibrahim's message and retweeted it to her followers. This secondary message, akin to forwarding an e-mail, is the retweet, and the reader knows it was seen by at least the followers of Ekramibrahim and the person who retweeted it. A message can be retweeted an infinite amount of times, though a user who sees a retweet only knows that at least one person retweeted it. (With one extra click, a user can see how many times the original tweet was retweeted, but there is no way for the researcher to observe if a user knows how many times a tweet was retweeted.) More retweets of one tweet therefore means that more people have seen the same set of information, and more retweets in general mean that a greater proportion of information on Twitter is shared information. By spreading the same set of information across many individuals, retweets facilitate the creation of common knowledge.

Retweets are often used to spread breaking news and recruit individuals to political causes, both of which should affect how many protests occur (**?**). They perform this action not just by containing pertinent text but by featuring hashtags that make their content easier to find by other interested individuals and containing links that point to news and photo sharing websites. Links which point to personal media sites and trivia sites have very low retweet rates, further suggesting that retweeting serves to spread information of interest to a large audience (Suh et al. 2010). Retweeted tweets have a short shelf life, further suggesting the immediate effects they have: 50-60% of a tweet's eventual retweets occur in the initial hour, 75-100% in the first day (Kwak et al. 2010, Liere 2010). When examining the density of retweet networks, retweets that contain links to outside information form denser networks that those focused on social action or identity making (Nagarajan, Purohit & Sheth 2010). The specific behavior of retweets leads Kwak et. al to conclude, "The strength of Twitter as a medium for information diffusion stands out by the speed of retweets" (Kwak et al. 2010, pg. 599).

$Hashtags\ \%_t$ and $Retweets\ \%_t$ therefore capture what percent of a day's Twitter messages are designed to contribute to common knowledge. Because each is associated with generating common knowledge, we expect that protests becomes more likely as each measure increases. Figure 2 shows $Hashtags\ \%_t$ and $Retweet\ \%_t$ for Egypt; that the measures appear to vary with important events there suggests they capture some component of common knowledge creation. The percentages can add to more than one.

Figure 2: Common Knowledge Creation in Egypt, Tweet Type

(a) Hashtags (b) Retweets

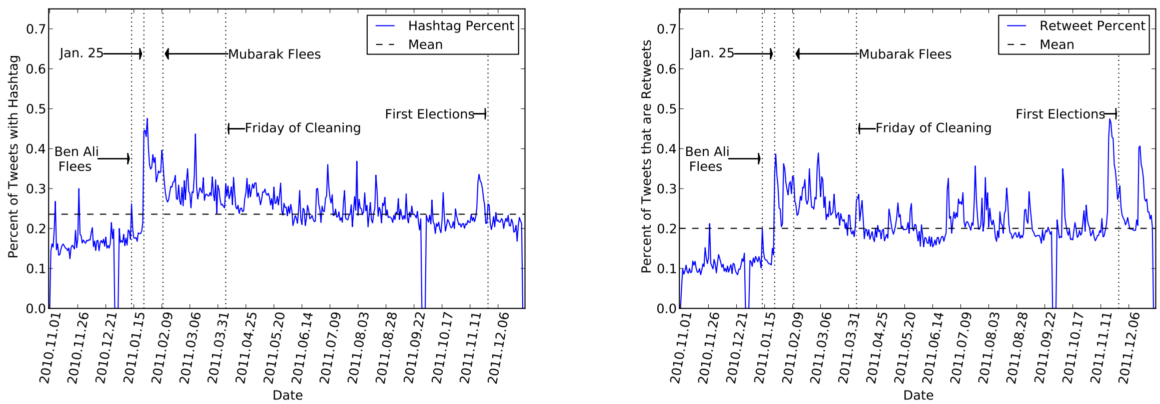

Those two measures represent our first measurements of common knowledge, but they are imprecise. For example, many hashtags are not related to protest activity at all; most, in fact, are not even political. An increase in tweets with hashtags could therefore capture more people tweeting about a sporting event or a viral meme, not an upcoming protest event. At the extreme, 1,000 tweets could have 1,000 unique hashtags, and *Hashtags* $\%_t$ does not distinguish between that or 1,000 tweets using the same hashtag. To distinguish between 1,000 tweets with 1,000 different hashtags and 1,000 tweets with an identical hashtag, we have created *Hashtag* $Gini_t$. A Gini measure captures the amount of inequality in a sample; a Gini of 1 means that one observation in that sample has all the objects being counted, while a score of 0 means every observation has the same count of objects. While the most famous application of the Gini score is income inequality, we use it to measure "hashtag inequality." As *Hashtag* $Gini_t$ approaches 1, tweets from that day use the same hashtag more frequently; as it approaches 0, each tweet with a hashtag is more likely to be the only tweet with that hashtag. High levels of hashtag inequality therefore mean more of one day's tweets use the same hashtag, which we take as evidence of individuals coordinating on one or a few hashtags. As *Hashtag* $Gini_t$ increases, we therefore expect common knowledge to have increased, leading to a subsequent increase in protest.

Whereas the percentage measures tell us simply about increases, these measures tell us about concentration. Yet, as with the percentage measures, the Gini measures could be capturing individuals' common knowledge creation not related to protests. As with Figure 2, Figure 3 provides

Figure 3: Common Knowledge Creation in Egypt, Gini Coefficients

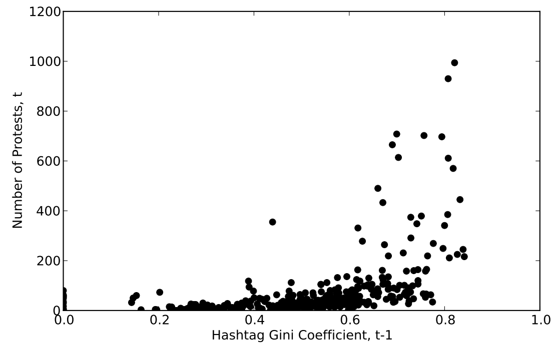

suggestive evidence that this is not the case. Figure 3 shows the correlation between a day's Gini coefficient and the next day's number of protests. There is a clear positive relationship between $Hashtag\ Gini_{t-1}$ and $Protest_t$.

These three variables - $Hashtags\ \%_t$, $Retweets\ \%_t$, and $Hashtag\ Gini_t$ - test Hypothesis 1, that more common knowledge leads to more protest.

To test Hypothesis 2, we construct measures of the influence of users who use hashtags or retweet and compare them to all users. We measure this by recording the number of followers each author of a tweet has, which provides an estimate of how many people saw that tweet. While the production of a tweet does not mean that all of a person's followers have seen the message (Aday et al. 2012) and cannot reveal an individual's centrality in a network, this measure provides a rough approximation for the level of influence each individual has. From this out-degree measure, we measure the average out-degree of all tweets with a hashtag, the average out-degree of all retweets, and the average out-degree of all users, per day. Since hashtags and retweets create common knowledge and previous research has shown that popular users use hashtags more than average ones (Garcia-Herranz et al. 2014), we are interested in the ratio of the hashtag out-degree and retweet out-degree to the average user's out-degree. We call these measures $Hashtag\ Influence_t$ and $Retweet\ Influence_t$, respectively. Figure 4 shows their behavior in Egypt.

These measures capture actions that other scholars have demonstrated occur, especially in times of crisis, and operate through generating common knowledge and subsequent action. For example, in the case of the 2010 Yushu Earthquake in China, users immediately took to Sina-Weibo (China's

Figure 4: Individuals' Influence on Common Knowledge Creation in Egypt

(a) More Influential People use Hashtags (b) Less Influential People Retweet

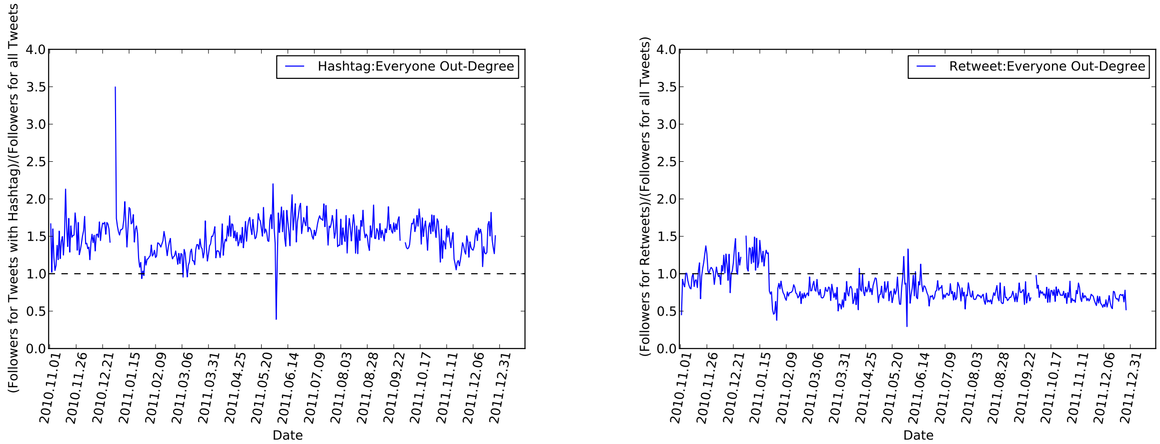

Twitter) to express opinions, a wide range of emotions, discuss responses, and spread information. The messages that were re-posted (equivalent to retweeting) were much more likely to have discussed actions individuals could take and provide updates on the crisis (Qu, Huang, Zhang & Zhang 2011). During two natural disasters in the United States, the Oklahoma grassfires of April 2009 and Red River Floods of March-April 2009, a similar surge of information occurred, and retweets differed from normal tweets two ways. Retweets were 45% more likely to provide "situational awareness" updates than normal tweets, and they were 74% more likely to reference a specific place (Vieweg et al. 2010). More subtly, the most popular retweets were ones with broad appeal (photo sharing, links to real time updates), but the most popular retweets among locals contained more "locally relevant information" (Starbird & Palen 2010). Moreover, users of Twitter quickly converge to a set of hashtags for an event, and political hashtags are explicitly used by individuals with an interest in the events to which that hashtag refers. Engagement with the hashtag inserts the user into an online public community similar to what occurs through offline public conversations in semi-public spaces such as coffee shops or taxi cabs (Bruns & Burgess 2011, Lehmann et al. 2012, Lim 2012). In the case of protests, convergence to one or a few hashtags occurs quickly, and that hashtag is used to coordinate protests, spread information about police action, and elicit outside support (Bajpai 2011, Earl et al. 2013, Lowe, Theocharis & W. van Deth 2013). Communication inspired by contemporary events also has more hashtags (and links) than normal tweets (Becker, Gravano & Naaman 2011).

Some work has started to look at retweets and hashtags during the Arab Spring, with most analyses focusing on Egypt. An analysis of messages with the keywords "#sidibouzid", "tunisia", "#jan25", and "egypt" during two weeks in January 2011 finds that the most popular retweet networks are started by activists or bloggers (who tend to be in support of protests) retweeting activists or journalists, suggesting that these three kinds of users play a key role in generating common knowledge (Lotan et al. 2011). Focusing on users who were demonstrably in Egypt, other research finds that the most retweeted tweets during the protests in Tahrir Square focus on first-hand accounts of violence, detained peers, requests for soldiery, political humor, and tactical information. Tweets with very specific information about breaking events - "The army has allowed a group of 100 Mubaraks thugs into Qasr el-Nil bridge now. #Jan25" - were also retweeted, though at a lower rate than ones with a more general purpose (Starbird & Palen 2012). A survey of individuals in Tahrir Square from late January through February 2011 finds that 13% of attendees used Twitter to communicate about protests, and those that did use it to communicate about protest were more likely to protest than those who did not. In addition, 48.2% of individuals at Tahrir Square actively disseminated videos and pictures of the events on social media (Tufekci & Wilson 2012).

5 Research Design

The estimated model is:

$$Protests_{i,t} = \beta_0 + \beta_1 * \Omega_{i,t-1} + \boldsymbol{\beta} * \mathbf{X} + \epsilon_{i,t} \tag{1}$$

where Ω represents the independent variable of interest in each model, $\mathbf{X}$ represents a series of controls, and ϵ is a stochastic error term. Because the dependent variable is a count of protests, it is an integer always greater than or equal to 0. The appropriate model is therefore a variant of a count model, in this case a negative binomial model.

The negative binomial was chosen for two reasons. First, the dependent variable is overdispersed. Figure 5 demonstrates the overdispersion. Since a Poisson model assumes that a distribution's mean and variance are the same, it is not the right model for these data. Second, a zero-inflated negative binomial was not chosen because the observed zeroes are not inflated by false negatives. Zero-inflated models assume that the zeros come from two sources, true zeroes and zeroes that are recorded as such because the true outcome is unknown. For GDELT, it is much more likely that observed zeroes are true zeroes. As Figure 1 showed, GDELT records protests on most country-days - 75.96% of them - in this sample, and it records fewer protests in countries where we expect there to have been fewer protests. While it is possible that there are some false negatives in the GDELT data, it is doubtful those are a significant proportion of the 24.04% of country-days with zero protest.

Figure 5: Distribution of Protests per Day

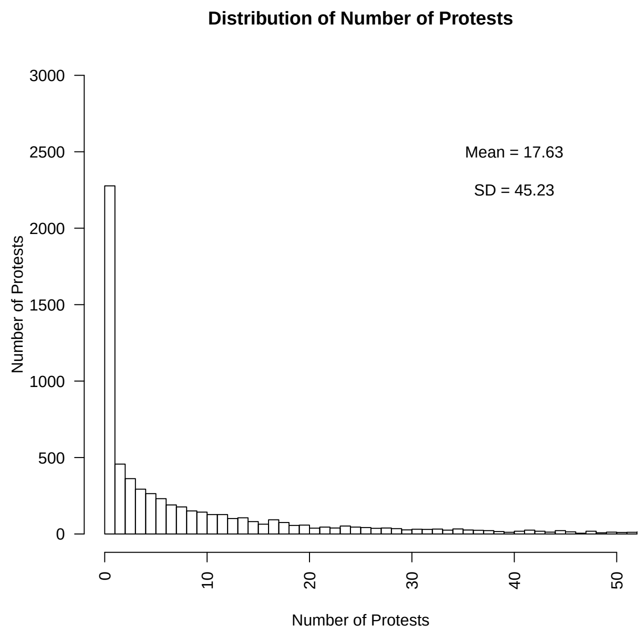

We control for the percent of a day's tweets which contain links, user mentions, or are basic; these measures are *Link %*, *Retweet %*, and *Basic %*, respectively. All models include a lagged dependent variable to capture serial correlation between days, as each days' level of protest does not appear to be independent of the previous day's. Most models also control for a country's population, measured in 2010, to control for the fact that large countries will have more protests on average. The only models which do not control for population are those run with country fixed effects, as those are perfectly collinear with population. Models with country fixed effects also include day fixed effects. The fixed effects capture country and day specific influences on the dependent variable that we cannot measure directly. For example, Egypt, Syria, Yemen, Libya, and Tunisia experienced most of the protests observed in this study; without controlling for the unique reasons they experienced more protest, we risk capturing a latent effect that drives both common knowledge creation and protests. Similarly, the first months of 2011 experienced higher levels of protest, across almost all the countries, than later in the year. If those days coincidentally record higher levels of *Hashtag %* or *Hashtag Gini*, a latent effect not related to common knowledge creation or individual influence could lead to a spurious correlation.

The models contain few control variables, for two reasons. First, we are wary of curve fitting, a possibility especially acute given the high dimensionality of our data. We strive to follow Achen's Rule of Three, and the models which contain more than 3 variables show similar effects as more parsimonious ones (Achen 2002). Second, there are very few measures which vary at the daily level and are comparable across all countries in this study. For example, Navid Hassanpour has measured the dispersion of protests in Cairo for 18 days after President Mubarak disconnected internet and cell phone service (Hassanpour 2014). While we would prefer to include daily-level measures of protest location or size of protests, such measures are unavailable at the scale of our study and infeasible to code by hand. Moreover, many factors that may correlate with protests - demographics (Nordå s & Davenport 2013), sectarian divisions (Matthiesen 2013), or civil-military relations (Quinlivan 1999) - do not vary on a daily level. Controlling for those factors via fixed effects is therefore the best we can do.[5]

Table 1 provides more detail on the primary variables, while Table 2 shows the variables by country.

All models are run with country-clustered standard errors. All independent variables are lagged by one day to capture the sequential nature of the hypotheses. If both sides of the model were measured during the same time period, we could most likely capture correlation caused by a latent variable and not any connection between common knowledge, individuals' influence, and protest.

[5]In models not shown here, we substitute more precise measures for fixed effects. To more precisely tease apart country-level effects, we control for whether a country is a monarchy, its population, its unemployment rate, its Freedom of the Press, and other measures. The substantive results for common knowledge do not change.

Table 1: Summary Statistics of Key Variables

Statistic	N	Mean	St. Dev.	Min	Max
Protest	6,816	17.626	45.226	0	994
Hashtag %	6,816	0.212	0.120	0.000	1.000
Hashtag Gini	6,816	0.278	0.201	−0.000	0.841
Hashtag Influence	6,434	1.367	0.691	0.010	8.442
Retweet %	6,816	0.134	0.092	0.000	1.000
Retweet Influence	6,211	0.848	0.556	0.003	7.834
Link %	6,816	0.315	0.181	0.000	1.000
Mention %	6,816	0.335	0.132	0.000	1.000
Population	6,816	18,344.450	19,472.190	1,251.513	78,075.710

Table 2: Main Variables by Country

Country	Protests	Hashtag %	Hashtag Gini	Hashtag Influence	Retweet %	Retweet Influence	Link %	Mention %	Population
Egypt	29035	0.24	0.51	1.50	0.80	0.20	0.27	0.33	78,075,710
Syria	24684	0.37	0.50	0.99	1.00	0.14	0.58	0.18	21,532,650
Yemen	12977	0.26	0.31	1.32	0.86	0.11	0.58	0.17	22,763,010
Libya	11146	0.27	0.28	1.65	1.34	0.16	0.33	0.27	6,040,610
Tunisia	9974	0.27	0.27	1.11	0.96	0.12	0.42	0.35	10,549,100
Bahrain	7313	0.28	0.51	1.30	0.78	0.30	0.18	0.29	1,251,510
Iraq	5525	0.19	0.25	1.26	0.94	0.08	0.28	0.32	30,962,380
SaudiArabia	5321	0.14	0.41	1.08	0.78	0.15	0.16	0.43	27,258,390
Jordan	3830	0.24	0.29	1.16	0.72	0.12	0.43	0.32	6,046,000
Lebanon	3212	0.22	0.26	1.17	0.62	0.14	0.27	0.35	4,341,090
Algeria	1723	0.23	0.05	2.23	0.77	0.13	0.45	0.29	37,062,820
Morocco	1722	0.20	0.18	1.24	0.85	0.12	0.34	0.38	31,642,360
Oman	1243	0.10	0.02	1.06	0.60	0.05	0.31	0.46	2,802,770
Kuwait	1053	0.09	0.09	1.12	1.18	0.06	0.13	0.49	2,991,580
UAE	833	0.16	0.25	2.02	0.82	0.12	0.20	0.38	8,441,540
Qatar	546	0.14	0.27	1.60	0.52	0.15	0.10	0.36	1,749,710

6 Results

We run three sets of models to test our hypotheses.

The models in Table 3 test Hypothesis 1, that more common knowledge leads to more protest. Models 1-3 present each of our three measures of common knowledge creation - emphHashtag $\%_{t-1}$, *Retweet* $\%_{t-1}$, and *Hashtag* $Gini_{t-1}$ - separately. Each correlates positively with protest the next day with a p value of less than .01. Model 4 puts the percentage measures of common knowledge creation in one model and includes control for links, mentions, and basic tweets. emphHashtag $\%_{t-1}$ and *Retweet* $\%_{t-1}$ are still positive at the same level of significance, though their effects are reduced. Model 5 adds *Hashtag* $Gini_{t-1}$; it retains its level of significance but causes the other two measures to become indistinguishable from zero. Model 6 is the same as Model 5 except with time and country fixed effects. The substantive results from Model 5 stay the same.

Table 3: Coordination on Hashtags Correlates with Protest

	DV: $Protests_{i,t}$					
	Protest					
	(1)	(2)	(3)	(4)	(5)	(6)
Hashtag Gini$_{i,t-1}$	2.404***				1.938***	2.418***
	(0.268)				(0.489)	(0.517)
Hashtag %$_{i,t-1}$		4.371***		2.918***	1.121	0.160
		(0.614)		(0.626)	(0.769)	(0.585)
Retweet %$_{i,t-1}$			2.942***	1.358***	0.133	0.760
			(0.599)	(0.469)	(0.564)	(0.504)
Link %$_{i,t-1}$				0.075	0.532	0.559
				(0.638)	(0.459)	(0.408)
Mentions %$_{i,t-1}$				−0.386	−0.283	0.084
				(0.310)	(0.271)	(0.135)
Basic Tweet %$_{i,t-1}$				−1.626	−1.453	0.078
				(1.300)	(1.172)	(0.464)
Population	0.009***	0.013***	0.012***	0.013***	0.009**	
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	
Protest$_{i,t-1}$	0.024***	0.023***	0.029***	0.022***	0.020***	0.011***
	(0.001)	(0.002)	(0.001)	(0.002)	(0.002)	(0.001)
Constant	0.871***	0.541***	1.066***	1.153*	1.005**	−0.234
	(0.168)	(0.204)	(0.191)	(0.633)	(0.461)	(0.496)
Day FE	No	No	No	No	No	Yes
Country FE	No	No	No	No	No	Yes
Observations	6,800	6,626	6,626	6,626	6,626	6,626
Log Likelihood	−21,679.180	−21,234.610	−21,507.970	−21,170.120	−21,057.670	−19,602.460
AIC	43,366.360	42,477.220	43,023.940	42,356.240	42,133.350	40,078.930

Note: *p<0.1; **p<0.05; ***p<0.01
Population is measured in millions.

Population is positively correlated with subsequent protest for all models in Table 3 in which it is included; in all models except Model 5, the chance this result was found accidentally is less than 1%. In all models, the lagged dependent variable correlates positively and is significant at the 1% level.

The models in Table 4 test Hypothesis 2, that more influential members creating common knowledge correlates with more subsequent protest. Population and the previous day's number of protests have the same effects as in the tests for the first hypothesis, but the effects for individuals' influence are more contingent. In Models 1, 3, and 4, *Hashtag Influence*$_{t-1}$ is never statistically significant but is consistently less than 0. In Models 2, 3, and 4, *Retweet Influence*$_{t-1}$, is positive with a p-value less than .01. To investigate whether a country's population is driving the results, we drop it from Model 4, but the results hold. Once we control for unobserved country and day effects, however, *Hashtag Influence*$_{t-1}$ and *Retweet Influence*$_{t-1}$ switch places.

Finally, Table 5 combines the models from Tables 3 and 4. The main findings hold. *Hashtag*

Table 4: Influence of Information Spreaders Correlates with Protest

	DV: $Protests_{i,t}$				
	Protest				
	(1)	(2)	(3)	(4)	(5)
Hashtag Influence$_{i,t-1}$	−0.054		−0.070	−0.019	0.129***
	(0.093)		(0.090)	(0.092)	(0.034)
Retweet Influence$_{i,t-1}$		0.231***	0.239***	0.227***	−0.012
		(0.078)	(0.077)	(0.076)	(0.035)
Population	0.011**	0.011**	0.011**		
	(0.005)	(0.005)	(0.005)		
Protest$_{i,t-1}$	0.031***	0.031***	0.031***	0.033***	0.014***
	(0.002)	(0.002)	(0.002)	(0.001)	(0.001)
Constant	1.555***	1.304***	1.388***	1.510***	−0.354
	(0.200)	(0.207)	(0.227)	(0.209)	(0.301)
Day FE	No	No	No	No	Yes
Country FE	No	No	No	No	Yes
Observations	6,626	6,626	6,626	6,626	6,626
Log Likelihood	−21,619.810	−21,598.540	−21,594.880	−21,682.900	−19,790.260
AIC	43,247.610	43,205.080	43,199.770	43,373.790	40,446.530
Note:	*p<0.1; **p<0.05; ***p<0.01				
	Population is measured in millions.				

$Gini_{t-1}$ positively correlates with next day's protest and does so at a high level of statistical significance. *Retweet Influence*$_{t-1}$ and *Hashtag Influence*$_t$ maintain the same signs and switch significance at the same time, and *Popluation* and the lagged dependent variable continue to be strong predictors of protest.

Because the model is a negative binomial, the coefficients are not directly interpretable. To gain a more intuitive understanding of the magnitude of the common knowledge effects, we plot the marginal effects on protest of *Hashtag Gini*$_{t-1}$ and *Hashtag Influence*$_{t-1}$. We use Model 3 from Table 5 and hold all variables at their mean. Figure 6 shows the expected number of protests as *Hashtag Gini*$_{t-1}$ and *Hashtag Influence*$_{t-1}$. Moving from a Gini coefficient corresponding to all tweets using different hashtags to one where all tweets use the same hashtag correlates with 31 more subsequent protests.[6] A day where users with no followers are the only ones to use a hashtag(s) to one where those use them are 8 times as influential as those who do not corresponds to 10 more protests.

To put these numbers in context, the median number of observed protests is 5 (the average, 17.625, is much higher because of highly contentious days). The median of *Hashtag Gini*$_{t-1}$ is

[6] There were days with a Gini coefficient of 0, implying no hashtags were used multiple times, but low values are an artifact of days with little data. The highest observed Gini is .84 on January 20^{th}, 2011 in Egypt for #Egypt.

Table 5: Full Model

	DV: $Protests_{i,t}$		
	Protest		
	(1)	(2)	(3)
$Hashtag\ Gini_{i,t-1}$	1.925***	2.354***	2.455***
	(0.489)	(0.432)	(0.519)
$Hashtag\ \%_{i,t-1}$	1.178	0.554	0.165
	(0.760)	(0.633)	(0.580)
$Retweet\ \%_{i,t-1}$	0.164	−0.158	0.834*
	(0.562)	(0.589)	(0.473)
$Hashtag\ Influence_{i,t-1}$	−0.048	−0.007	0.143***
	(0.081)	(0.079)	(0.040)
$Retweet\ Influence_{i,t-1}$	0.142***	0.139***	0.028
	(0.054)	(0.054)	(0.039)
$Link\ \%_{i,t-1}$	0.601	0.641	0.491
	(0.484)	(0.548)	(0.380)
$Mentions\ \%_{i,t-1}$	−0.227	−0.319	0.063
	(0.282)	(0.307)	(0.136)
$Basic\ Tweet\ \%_{i,t-1}$	−1.273	−1.509	0.092
	(1.146)	(1.376)	(0.460)
Population	0.009**		
	(0.004)		
$Protest_{i,t-1}$	0.020***	0.021***	0.010***
	(0.002)	(0.002)	(0.001)
Constant	0.858*	1.101**	−0.519
	(0.483)	(0.550)	(0.496)
Day FE	No	No	Yes
Country FE	No	No	Yes
Observations	6,626	6,626	6,626
Log Likelihood	−21,042.980	−21,106.380	−19,581.940
AIC	42,107.970	42,232.770	40,041.890

Note: *p<0.1; **p<0.05; ***p<0.01
Population is measured in millions.

Figure 6: Marginal Effects of Common Knowledge and Influential Individuals

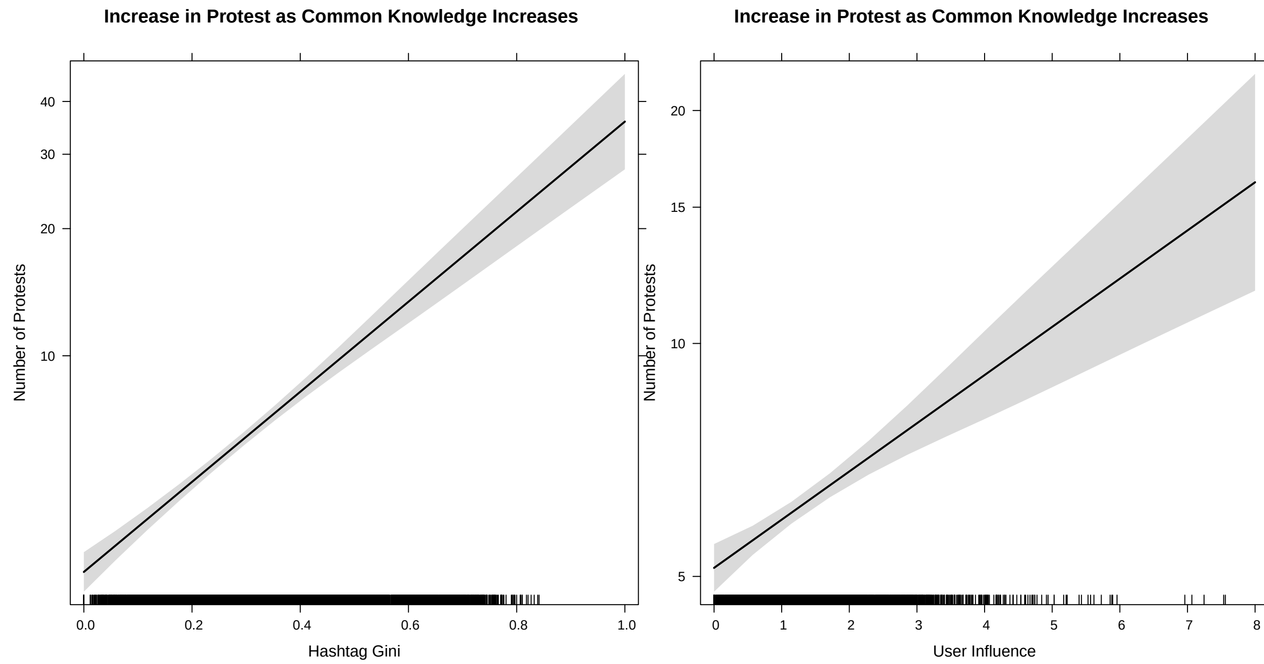

.255 and the standard deviation is .201 therefore correspond to an increase from 5 to 8 observed protests, or 60% of the median number of protests per day.

Discussion

Based on the role of hashtags, retweets, and popular users, the following dynamic drives protest communication. Though some users are certainly discussing events before they become widely talked about, especially in the case of planned events like protests, most individuals are not involved in the conversation until the event becomes widely known (Lehmann et al. 2012). Individuals who have the most professional interest in the news - journalists, bloggers, and activists - then become main propagators of information to a larger community (Lotan et al. 2011, Aday et al. 2012). Through retweets and hashtags, these main propagators generate common knowledge about events as they occur and encourage individuals to participate or lend support. The use of hashtags insures that individuals most likely to participate in the event are exposed to the shared information. Consumers of this information then fall into two main groups, non-locals and locals, with locals expressing much more interest in specifics and non-locals consuming general information.

Technology which facilitates interpersonal communication should therefore also facilitate common knowledge and, in turn, protests. Indeed, it is the ability of the internet to facilitate information sharing between individuals that causes many to believe the internet has had a transformative effect on politics (Garrett 2006, Shirky 2008). This paper focuses on one particular platform, Twitter, to argue that social media generates common knowledge, facilitating protests. Twitter is a social network focused on messaging between the user and the user's followers; it is primarily accessed through phones, even phones without internet connections, making it omnipresent to its users.

There is a longstanding debate on the effect that the internet and cell phones have on collective action (Garrett 2006, Shapiro & Weidmann 2011, Pierskalla & Hollenbach 2013). While these technologies are certainly not a panacea and are susceptible to adaptive regimes with enough resources to control their effects (Morozov 2012), this paper suggests that these technologies, or at least Twitter, do impact protests. They are not the causes of protests or the reason some movements succeed where others fail. They do, however, create common knowledge that otherwise would not have existed, leading to more protests than would have occurred without their existence.

This paper does not seek to uncover the causes of the Arab Spring or provide a unifying framework for understanding countries' outcomes. Rather, it shows how individuals can mobilize for protest and how certain types of individuals facilitate this mobilization. The arguments presented here can also apply in other countries and should have been a factor in waves of protest prior to the advent of the internet and cellular communication.

This paper has found that Twitter promotes the creation of common knowledge about authoritarian regimes. This common knowledge makes the coordination of protests easier, leading to more protests than would otherwise have occurred without it. Common knowledge is created primarily through hashtags and retweets, as they facilitate the spread of information to a wide audience who otherwise may have only had private information about their preferences and regime policies. While the users of Twitter may not be representative of the population from which they are drawn, especially in developing countries, that results are still found suggests room for further theoretical refinement.

It is possible that what matters is not common knowledge spread across the population at large but common knowledge amongst a small pool of initial protestors. This observation dovetails with

Lohmann's focus on activist moderates (Lohmann 1994). In her model, the key piece of information future protestors seek is the number of activist moderates who protested in the previous period. Most of the countries in this study had small cadres of activists who would protest, and get arrested, often. Their protests provided no new information to the mass of citizens who wanted a regime change but were worried their sentiment was not widely shared. Twitter then may have served as a way to connect the always activists to those with slightly higher thresholds of action. Once the first wave of activist moderates created common knowledge, they protested in large enough numbers to convince the rest of society to reveal their true preferences. Most people then bandwagoned once the first activist moderates had common knowledge, and regimes soon found themselves facing widespread protests (Kuran 1989).

One area that remains to be explored is whether or not information and communication technologies have economies of scale which make them more difficult it for a repressive state to control. This study suggests they do - the number of signals available to individuals has increased dramatically, making regime control of information much more difficult (Edmond 2013). On the other hand, the events of this study, and protests in countries such as Venezuela or Ukraine, may mean that controlling these technologies require large investments in human and financial capital that few regimes can afford. Those that make the investment to control digital technologies, such as Iran and China, have been able to monitor and censor the online arena much the same way they have done with mass media (Howard 2010, King, Pan & Roberts 2013). Future research should investigate the relationship between the state and control of ICT in greater detail.

References

Achen, Christopher H. 2002. "Toward a New Political Metholodgy: Microfoundations and ART." *Annual Review of Political Science* 5(1):423–450.

Aday, Sean, Deen Freelon, Henry Farrell, Marc Lynch & John Sides. 2012. New Media and Conflict After the Arab Spring. Technical report United States Institute of Peace Washington D.C.: .

Anderson, Benedict. 1991. *Imagined Communities: Reflections on the Origin and Spread of Nationalism*. 2nd ed. New York City: Verson.

Bajpai, Kartikeya. 2011. A Framework for Analyzing Collective Action Events on Twitter. In *Information Systems for Crisis Response and Management*. Number May Portugal: .

Becker, Hila, Luis Gravano & Mor Naaman. 2011. "Hip and Trendy: Characterizing Emerging Trends on Twitter." *Journal of the American Society for Information Science and Technology* 62(5):902–918.

Benford, Robert D. & David A. Snow. 2000. "Framing Processes and Social Movements: An Overview and Assessment." *Annual Review of Sociology* 26:611–639.

Borge-Holthoefer, Javier, Alejandro Rivero & Yamir Moreno. 2012. "Locating privileged spreaders on an online social network." *Physical Review E* 85(6).

Boyd, D., S. Golder & G. Lotan. 2010. "Tweet, Tweet, Retweet: Conversational Aspects of Retweeting on Twitter." *2010 43rd Hawaii International Conference on System Sciences* pp. 1–10.

Brancati, D. 2013. "Pocketbook Protests: Explaining the Emergence of Pro-Democracy Protests Worldwide." *Comparative Political Studies* 20(10):1–28.

Breuer, Anita. 2012. The Role of social media in mobilizing political protest: Evidence from the Tunisian Revolution. Technical report German Development Institute Bonn: .

Bruns, Axel & Jean E Burgess. 2011. The Use of Twitter Hashtags in the Formation of Ad Hoc Publics. In *6th European Consortium for Political Research Conferenc*. Number August Reykjavik: pp. 1–10.

Burns, Alex & Ben Eltham. 2009. "Twitter Free Iran: An Evaluation of Twitter's Role in Public Diplomacy and Information Operations in Iran's 2009 Election Crisis.".

Cederman, Lars-Erik, Luc Girardin & Kristian Skrede Gleditsch. 2009. "Ethnonationalist Triads: Assessing the Influence of Kin Groups on Civil Wars." *World Politics* 61(3):403–437.

Centola, Damon & Michael Macy. 2007. "Complex Contagion and the Weakness of Long Ties." *American Journal of Sociology* 113(3):702–734.

Chew, Cynthia & Gunther Eysenbach. 2010. "Pandemics in the age of Twitter: Content Analysis of Tweets during the 2009 H1N1 Outbreak." *PloS one* 5(11):e14118.

Christakis, Nicholas A & James H Fowler. 2010. "Social network sensors for early detection of contagious outbreaks." *PloS ONE* 5(9):e12948.

Chwe, Michael Suk-young. 2003. *Rational Ritual: Culture, Coordination, and Common Knowledge.* Princeton: Princeton University Press.

Davenport, Christian. 2007. "State Repression and Political Order." *Annual Review of Political Science* 10(1):1–23.

Earl, Jennifer, Heather McKee Hurwitz, Analicia Mejia Mesinas, Margaret Tolan & Ashley Arlotti. 2013. "This Protest Will Be Tweeted: Twitter and protest policing during the Pittsburgh G20." *Information, Communication & Society* 16(4):459–478.

Edmond, C. 2013. "Information Manipulation, Coordination, and Regime Change." *Review of Economic Studies* 80(4):1422–1458.

Egorov, Georgy, Sergei Guriev & Konstantin Sonin. 2009. "Why Resource-poor Dictators Allow Freer Media: A Theory and Evidence from Panel Data." *American Political Science Review* 103(04):645.

Faris, David. 2013. *Dissent and Revolution in a Digital Age: Social Media, Blogging and Activism in Egypt.* I.B. Tauris.

Fearon, J. D. 2011. "Self-Enforcing Democracy." *The Quarterly Journal of Economics* 126(4):1661–1708.

Fearon, James D. & David D. Laitin. 2003. "Ethnicity, Insurgency, and Civil War." *The American Political Science Review* 97(1):75–90.

Garcia-Herranz, Manuel, Esteban Moro, Manuel Cebrian, Nicholas A. Christakis & James H. Fowler. 2014. "Using friends as sensors to detect global-scale contagious outbreaks." *PloS ONE* 9(4):e92413.

Garrett, R. Kelley. 2006. "Protest in an Information Society: a review of literature on social movements and new ICTs." *Information, Communication & Society* 9(2):202–224.

Gerner, Deborah J., Philip A. Schrodt, Rajaa Abu-Jabr & Omur Yilmaz. 2002. Conflict and Mediation Event Observations (CAMEO): A New Event Data Framework for the Analysis of Foreign Policy Interactions. In *Annual Meeting of the International Studies Association.*

Granovetter, Mark. 1978. "Threshold Models of Collective Behavior." *American Journal of Sociology* 83(6):1420–1443.

Gurr, Ted. 1971. *Why Men Rebel.* Princeton: Princeton University Press.

Hale, Henry E. 2013. "Regime Change Cascades: What We Have Learned from the 1848 Revolutions to the 2011 Arab Uprisings." *Annual Review of Political Science* 16(1):331–353.

Hamdy, Naila & Ehab H. Gomaa. 2012. "Framing the Egyptian Uprising in Arabic Language Newspapers and Social Media." *Journal of Communication* 62(2):195–211.

Hanna, Alexander. 2013. "Computer-aided Content Analysis of Digitally Enabled Movements." *Mobilization: An International Quarterly* 18(4):367–388.

Hassanpour, Navid. 2014. "Media Disruption and Revolutionary Unrest: Evidence From Mubarak's Quasi-Experiment." *Political Communication* 31(1):1–24.

Honeycutt, Courtenay & Susan C. Herring. 2009. Beyond Microblogging : Conversation and Collaboration via Twitter. In *Proceedings of the 42nd Hawaii International Conference on System Sciences*. pp. 1–10.

Howard, Philip. 2010. *The Digital Origins of Dictatorship and Democracy*. Oxford: Oxford University Press.

Hyde, Susan D. & Nikolay Marinov. 2014. *Information and Self-Enforcing Democracy: The Role of International Election Observation Information and Self-Enforcing Democracy*. Vol. 68.

Jalili, Mahdi. 2012. "Social power and opinion formation in complex networks." *Physica A* pp. 1–8.

Jenkins, J Craig. 1983. "Resource Mobilization Theory and the Study of Social Movements." *Annual Review of Sociology* 9(1983):527–553.

King, Gary, Jennifer Pan & Margaret E. Roberts. 2013. "How Censorship in China Allows Government Criticism but Silences Collective Expression." *American Political Science Review* 107(02):326–343.

Kricheli, Ruth, Beatriz Magaloni & Yair Livne. 2011. "Taking to the Streets: Theory and Evidence on Protests under Authoritarianism.".

Kuran, Timur. 1989. "Sparks and Prairie Fires: A Theory of Unanticipated Political Revolution." *Public Choice* 61(1):41–74.

Kwak, Haewoon, Changhyun Lee, Hosung Park & Sue Moon. 2010. What is Twitter, a Social Network or a News Media ? In *International World Wide Conference*. pp. 591–600.

Lagi, Marco, Karla Z. Bertrand & Yaneer Bar-Yam. 2011. "The Food Crises and Political Instability in North Africa and the Middle East.".
URL: *http://arxiv.org/abs/1108.2455*

Leetaru, Kalev & Philip Schrodt. 2013. GDELT: Global Data on Events, Language, and Tone, 1979-2012. In *International Studies Association*. San Diego, CA: .

Lehmann, Janette, Bruno Gonçalves, José J. Ramasco & Ciro Cattuto. 2012. "Dynamical Classes of Collective Attention in Twitter." *Proceedings of the 21st International Conference on World Wide Web - WWW '12* pp. 251–260.

Liere, Diederik Van. 2010. How Far Does a Tweet Travel? Information Brokers in the Twitterverse. In *Modeling Social Media*. Toronto: pp. 1–4.

Lim, Merlyna. 2012. "Clicks, Cabs, and Coffee Houses: Social Media and Oppositional Movements in Egypt, 2004-2011." *Journal of Communication* 62(2):231–248.

Lim, Merlyna. 2013. "Framing Bouazizi: 'White lies', hybrid network, and collective/connective action in the 2010-11 Tunisian uprising." *Journalism* 14(7):921–941.

Little, Andrew T. 2012. "Elections, Fraud, and Election Monitoring in the Shadow of Revolution." *Quarterly Journal of Political Science* 7(3):249–283.

Lohmann, Susanne. 1994. "The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989-91." *World Politics* 47(1):42–101.

Lotan, Gilad, Mike Ananny, Devin Gaffney, Danah Boyd, Ian Pearce & Erhardt Graeff. 2011. "The Revolutions Were Tweeted: Information Flows During the 2011 Tunisian and Egyptian Revolutions Web." *International Journal of Communications* 5:1375–1406.

Lowe, Will, Yannis Theocharis & Jan W. van Deth. 2013. "Using Twitter to Mobilise Protest Action: Transnational Online Mobilisation Patterns and Action Repertoires in the Occupy Wall Street, Indignados and Aganaktismenoi.".

Matthiesen, Toby. 2013. *Sectarian Gulf: Bahrain, Saudi Arabia, and the Arab Spring that Wasn't.* Stanford: Stanford University Press.

McCammon, Holly J., Karen E. Campbell, Ellen M Granberg & Christine Mowery. 2001. "How Movements Win: Gendered Opportunity Structures and U.S. Women's Suffrage." *American Sociological Review* 66(1):49–70.

McCarthy, John D & Mayer N Zald. 1977. "Resource Mobilization and Social Movements: A Partial Theory." *American Journal of Sociology* 82(6):1212–1241.

Mendoza, Marcelo, Barbara Poblete & Carlos Castillo. 2010. Twitter Under Crisis: Can we trust what we RT? In *Social Media Analytics*. Washington D.C.: pp. 71–79.

Meyer, David S. 2004. "Protest and Political Opportunities." *Annual Review of Sociology* 30(1):125–145.

Mocanu, Delia, Andrea Baronchelli, Nicola Perra, Alessandro Vespignani, Bruno Goncalves & Qian Zhang. 2013. "The Twitter of Babel: Mapping World Languages through Microblogging Platforms." *PLOS One* 8(4):e61981.

Moore, Cristopher & M.E.J. Newman. 2000. "Epidemics and percolation in small-world networks." *Physical Review E* 61(5):5678–82.

Morozov, Evgeny. 2012. *The Net Delusion: The Dark Side of Internet Freedom.* New York City: PublicAffairs.

Mungiu-Pippidi, Alina & Igor Munteanu. 2009. "Moldova's "Twitter Revolution"." *Journal of Democracy* 20(3):136–142.

Nagarajan, Meenakshi, Hemant Purohit & Amit Sheth. 2010. A Qualitative Examination of Topical Tweet and Retweet Practices. In *International AAAI Conference on Weblogs and Social Media*. pp. 1–4.

Nordå s, Ragnhild & Christian Davenport. 2013. "Fight the Youth: Youth Bulges and State Repression." *American Journal of Political Science* 57(4):n/a–n/a.

Pierskalla, Jan H. & Florian M. Hollenbach. 2013. "Technology and Collective Action: The Effect of Cell Phone Coverage on Political Violence in Africa." *American Political Science Review* 107(02):207–224.

Posen, Barry R. 1993. "The Security Dilemma and Ethnic Conflict." *Survival* 35(1):27–47.

Qu, Yan, Chen Huang, Pengyi Zhang & Jun Zhang. 2011. Microblogging after a Major Disaster in China: A Case Study of the 2010 Yushu Earthquake. In *Computer Supported Cooperative Work*. Hangzhou, China: pp. 25–34.

Quinlivan, James T. 1999. "Coup-proofing: Its Practice and Consequences in the Middle East." *International Security* 24(2):131–165.

Radsch, Courtney C. 2012. "Unveiling the Revolutionaries: Cyberactivism and the Role of Women in the Arab Uprisings.".

Rahimi, Babak. 2011. "The Agonistic Social Media: Cyberspace in the Formation of Dissent and Consolidation of State Power in Postelection Iran." *The Communication Review* 14(3):158–178.

Rasler, Karen. 1996. "Concessions, Repression, and Political Protest in the Iranian Revolution." *American Sociological Review* 61(1):132–152.

Richards, John E. 2001. "Institutions for Flying: How States Built a Market in International Aviation Services." *International Organization* 55(4):993–1017.

Romero, Daniel M, Brendan Meeder & Jon Kleinberg. 2011. Differences in the Mechanics of Information Diffusion Across Topics: Idioms, Political Hashtags, and Complex Contagion on Twitter. In *International World Wide Web Conference*. pp. 695–704.

Rubinstein, Ariel. 1989. "The Electronic Mail Game: Strategic Behavior Under "Almost Common Knowledge"." *The American Economic Review* 79(3):385–391.

Schelling, Thomas C. 1960. *The Strategy of Conflict*. Cambridge: Harvard University Press.

Segerberg, Alexandra & W. Lance Bennett. 2011. "Social Media and the Organization of Collective Action: Using Twitter to Explore the Ecologies of Two Climate Change Protests." *The Communication Review* 14(3):197–215.

Shapiro, Jacob N. & Nils B. Weidmann. 2011. "Talking About Killing: Cell Phones, Collective Action, and Insurgent Violence in Iraq.".

Shirky, Clay. 2008. *Here Comes Everybody: The Power of Organizing Without Organizations*. New York City: Penguin Books.

Solomon, Samuel. 2012. Cascades in Cairo: The Role of Facebook and Twitter in the Egyptian Revolution of 2011 PhD thesis Georgetown University.

Starbird, Kate & Ley Palen. 2010. Pass It On ?: Retweeting in Mass Emergency. In *Information Systems for Crisis Response and Management*. Number December 2004 Seattle: pp. 1–10.

Starbird, Kate & Leysia Palen. 2012. (How) Will the Revolution be Retweeted? Information Diffusion and the 2011 Egyptian Uprising. In *CSCW*. Seattle: .

Suh, Bongwon, Lichan Hong, Peter Pirolli & Ed H. Chi. 2010. Want to be Retweeted? Large Scale Analytics on Factors Impacting Retweet in Twitter Network. In *IEEE Second International Conference on Social Computing*. pp. 177–184.

Tremayne, Mark. 2014. "Anatomy of Protest in the Digital Era: A Network Analysis of Twitter and Occupy Wall Street.".

Tucker, Joshua a. 2007. "Enough! Electoral Fraud, Collective Action Problems, and Post-Communist Colored Revolutions." *Perspectives on Politics* 5(03):535.

Tufekci, Zeynep & Christopher Wilson. 2012. "Social Media and the Decision to Participate in Political Protest: Observations From Tahrir Square." *Journal of Communication* 62(2):363–379.

Vieweg, Sarah, Amanda L. Hughes, Kate Starbird & Leysia Palen. 2010. Microblogging During Two Natural Hazards Events : What Twitter May Contribute to Situational Awareness. In *Human Factors in Computing Systems*. Atlanta: pp. 1079–1088.

Warren, T. Camber. 2014. "Not by the Sword Alone: Soft Power, Mass Media, and the Production of State Sovereignty." *International Organization* 68(01):111–141.

Watts, Duncan J. 2002. "A simple model of global cascades on random networks." *Proceedings of the National Academy of Sciences* 99(9):5766–5771.

Weyland, Kurt. 2012. "The Arab Spring: Why the Surprising Similarities with the Revolutionary Wave of 1848?" *Perspectives on Politics* 10(04):917–934.