American Political Science Review (2017) 111, 2, 379–403
doi:10.1017/S0003055416000769

Spontaneous Collective Action: Peripheral Mobilization During the Arab Spring

ZACHARY C. STEINERT-THRELKELD *University of California, Los Angeles*

Who is responsible for protest mobilization? Models of disease and information diffusion suggest that those central to a social network (the core) should have a greater ability to mobilize others than those who are less well-connected. To the contrary, this article argues that those not central to a network (the periphery) can generate collective action, especially in the context of large-scale protests in authoritarian regimes. To show that those in the core of a social network have no effect on levels of protest, this article develops a dataset of daily protests across 16 countries in the Middle East and North Africa over 14 months from 2010 through 2011. It combines that dataset with geocoded, individual-level communication from the same period and measures the number of connections of each person. Those on the periphery are shown to be responsible for changing levels of protest, with some evidence suggesting that the core's mobilization efforts lead to fewer protests. These results have implications for a wide range of social choices that rely on interdependent decision making.

INTRODUCTION

Large groups of people acting without centralized leadership can organize protests. Protests occur as a result of decentralized coordination of individuals, and this coordination helps explain fluctuating levels of protest. Individuals in the core of a social network—those such as activists, members of the media, or civil society organizations—do not mobilize protests. Instead, those on the periphery of the network communicate with each other about the near future (where and when to protest) as well as events as they unfold (the presence of police, what the police are doing, supplies needed, and so on). While those at the center of the network do engage in the same behavior as others, their effect is washed out in comparison to that of the masses they try to lead. I call the ability of the periphery to mobilize *spontaneous collective action.*

There exist two competing explanations for how individuals decide to undertake action. Whether deciding to vote (Downs 1957; Quattrone and Tversky 1988; Riker and Ordeshook 1968), join a political organization (González-Bailón, Borge-Holthoefer, Rivero and Moreno 2011; Klandermans and Oegema 1987; McAdam 1986), or protest (Goldstone 2001; Lichbach 1998; Moore 1995), individuals may decide to do so as a result of effort from centralized, well-connected individuals (the core) or those on the periphery. Those at the center of a social network can provide focal points for action, alternative policies for voters, new information about policies, or demonstrate a regime is weaker than previously thought, all contributing to individuals taking collective action (Dalton, Greene, Beck and Huckfeldt 2002; Gerber, Karlan and Bergan 2006; Shachar and Nalebuff 1999; Taylor 1988). On the other hand, individuals can decide to vote (or protest or join a movement) based on the influence of those they know (Gerber, Green and Larimer 2008; Schussman and Soule 2005), beliefs in their own ability to affect an outcome (Finkel, Muller and Opp 1989; Goldstone 1994; Opp 2012), or from observing the behavior of others (Granovetter 1978; Lohmann 1994). These others are the peripheral members.

This argument is tested using data from the Arab Spring, the protests which started in Tunisia in December 2010 and soon spread through North Africa and the Middle East. The events of the Arab Spring, the most prominent large-scale, widespread protests since the collapse of the Soviet Union, provide an ideal situation in which to test this theory. "Arab Spring" refers to the series of protests which started in Tunisia on December 14, 2010 (leading to the resignation of that country's president), slowly spread to neighboring countries over the following six weeks, and inspired massive turnout in Egypt that caused President Hosni Mubarak to resign on February 11, 2011. This article will show that these protests were not driven by the people who had tried for years to organize them. Instead, they were organized by large groups of individuals discussing amongst themselves where to go, how to get there, when to go, and what was going on once there. This article does not seek to explain the Arab Spring, but it does, in the course of developing the spontaneous collective action theory, present the first large-scale, systematic evidence on how individuals behaved in each country.

To test the core versus peripheral hypotheses, this article connects two large-scale datasets. First, a machine-coded events dataset, the Integrated Conflict Early Warning System (ICEWS), is combed to measure the

Zachary C. Steinert-Threlkeld is Assistant Professor, Department of Public Policy, University of California, Los Angeles, Los Angles, CA 90095, www.zacharyst.com (zst@luskin.ucla.edu).

I would like to thank the anonymous reviewers and editorial staff who provided excellent feedback. This article has also benefitted from the generous attention and brainstorming of Lawrence Broz, James Fowler, Emilie Hafner-Burton, David Lake, and Barbara Walter. Derrick Cogburn, Jesse Driscoll, Erik Gartzke, Philip N. Howard, and Nils Weidmann provided needed encouragement along the way. Nathan Combes, Scott Guenther, Maya Oren, and Aditya Ranganath have known this article since it was just an idea, and our group sessions were integral to its success. Many thanks as well go to Jason Jones, who provided the original idea for measuring coordination. The article benefitted from active participants at many conferences over too many years, as well as the patience of the Human Nature Group and the UCSD International Relations Workshop. This article would not be possible without the assistance of Alessandro Vespignani and Delia Mocanu. All remaining errors are mine.

Received: June 17, 2015; revised: July 2, 2016; accepted: July 19, 2016.

number of protests per day across 16 countries from November 1, 2010 through December 31, 2011. Second, a dataset of geolocated tweets in the same countries from the same period is built. These 13,754,998 tweets show what was being said, when it was being said, where, and how many connections each tweet author had. Combining these datasets and using a wide range of models and operationalizations, mass mobilization is shown to occur through peripheral individuals.

This article proceeds in eight sections. The following section compares and contrasts existing theories of protest mobilization with the one developed here. The article then presents the empirical strategy, followed with the main findings and a battery of robustness checks to reinforce them. A closer look at Egypt presents another strategy for identifying core actors, and the article concludes with final thoughts and suggestions for future research.

THEORY

Coordination drives protest mobilization, and peripheral members of a network drive coordination. Coordination consists of two components.

First, individuals deciding to protest must receive a credible signal that large numbers of people are protesting, suggesting that the cost of protesting is low. Second, once protests have commenced, information about upcoming protests—when a protest will occur, where protesters will convene, routes to take, supplies needed, and so on—needs to be provided. Members on the periphery of a social network better provide both components of coordination than those in the core. Individuals at the core of a network—those connected to many more people than the median person—are socially distant from most of those connections and few in number. This distance attenuates the weight of the signal the core sends (Centola and Macy 2007, 725–6), while their rarity limits the influence of their action and relevance of the protest information they provide.

Networks provide a framework for understanding how a phenomenon spreads between items; when these items are people, the network is a social network, and connections represent two people between whom a phenomenon can spread. These phenomena fall into two categories, simple and complex. A simple contagion is a phenomenon which can spread between individuals after one exposure, such as illness or information about job opportunities (Granovetter 1973). Disease transmission or news are canonical examples: John only needs to meet one person with the flu to catch it, and Jane only needs to talk with one person to learn tomorrow's weather. John does not become more sick from meeting a second infected person, and Jane does not become more knowledgeable receiving the same weather report from a second person. Except for rare cases, simple contagions always spread in a network (Newman 2003). Simple contagions spread quickest when core nodes are affected since those nodes can spread the phenomenon in question to many nodes at once, regardless of the structure of the underlying network (Watts 2004, 257–60). In simple contagion models, diffusion of a phenomenon is less likely when the diffusion starts on the periphery.

Complex phenomena are those whose transmission requires an individual to observe that phenomenon in two or more people. Contact with two or more sources is required when the phenomenon possesses positive externalities, gains credibility or legitimacy when multiple people partake, or have an emotional component (Centola and Macy 2007, 707–8). Models of complex contagion are often called threshold models since they require an individual to be exposed to a defined amount of other people in the network before switching states (Granovetter 1978; Schelling 1978).[1] The existence of thresholds makes the spread of complex contagions less certain, as network structure—the distribution of thresholds—can cause a contagion to stay trapped in one part of the network (Watts 2002). Because contact with more than one source is required for complex contagion's spread, core members do not automatically lead to the spread of the phenomenon in question. The existence of that phenomenon in peripheral parts of the network becomes essential for its spread throughout the network.

Protests are a complex contagion phenomenon because increasing participation makes others more likely to join. Individuals are more likely to protest as others protest, since the cost of protesting decreases as a function of group size. Individuals are especially more likely to protest when they know others who are protesting (Opp and Gern 1993), and those on the periphery of a network are more likely to know others on the periphery than in the core (Kwak, Lee, Park and Moon 2010; McPherson, Smith-Lovin and Cook 2001). Since there exist many more individuals on the periphery of a network than in the core, protest is therefore more likely to occur when those on the periphery of a network mobilize.

The first mechanism through which the periphery of a social network mobilizes protest is through providing a credible signal about participation in the protest. Peripheral members mobilize other participants better than those in the core because they provide a more credible signal that the protest enjoys widespread participation. If a protest is dominated by core members, the signal suggests that the policy disagreement does not affect many people who do not usually protest. This insight is similar to that made by Susanne Lohmann: she argues that unexpected participation of "moderate activists" drives protest mobilization because "extreme activists" always protest, so their participation is not a credible signal about the severity of a grievance

[1] The threshold is sometimes defined as a constant and sometimes as a fraction of network size. This distinction matters for small networks but not large ones. For example, in a network of eight individuals, a threshold of 1/8 does not represent complex contagion because an individual will switch states when only 1 person it knows has has; in a network of 800,000, a threshold of 1/8 would correspond to a late mover. Because mass protest involves large groups of people, the difference between numeric and proportional thresholds is moot. See Centola and Macy 2007 for an extended discussion on the difference between fractional and numeric thresholds.

(Lohmann 1994). A larger than expected turnout of "moderate activists" signals to others that grievances are widely shared, leading to the expectation that one's action will decisively lead to a policy change. In discussing the effect of network structure on collective action, David Siegel explains that:

> [...] the people at the bottom of the network—the proletariat, if you will—can [mobilize] if they have enough connections among themselves. The key here is to obtain a sufficiently large and well-connected group of people at the bottom of the hierarchy who [...] are highly internally motivated to participate. If these requirements are achieved, the bottom of the hierarchy can spur the network on to very high levels of participation. (Siegel 2009, 134–5)

The second mechanism through which the periphery of a social network mobilizes protest is by providing more information about a protest as it unfolds, and this information has the effect of coordinating protestor movement and tactics. One type of information is situational awareness, knowledge about unfolding events, and peripheral members, because of their number, provide this awareness in ways the core cannot. Situational awareness entails knowing the size of the police presence, which routes police block, whether or not police engage with protesters, paths around police, and where other protest groups find themselves. Protesters are also more likely to reach and hold onto their desired site if they can approach it from multiple directions and coordinate their action, as doing so makes it harder for police to contain the protesters (Gunning and Baron 2013, 168–74). But, since there exists a finite supply of core individuals, splitting a protest into subcomponents means that the ability of core members, who are few in number, to control them is lessened.

Moreover, once engagement with state forces commences, order often dissolves; a protest is a quickly shifting series of actions occurring in an area too big to be observed by a few individuals. During a street engagement with government forces, protesters may require reinforcements on some streets and not others, while supplies such as gas masks necessary in one place but not elsewhere. If a group is able to cause police to retreat, communicating that advantage to nearby protesters can provide reinforcements to exploit this development. But relying on core members, who are few in number, to coordinate these reactions decreases the efficacy with which protesters can react to new developments. Situational awareness therefore increases the likelihood of protest success, and situational awareness is increased when information flows from and between as many individuals as possible.

Situational awareness also entails providing logistical support for a protest. During the initial march to a protest site, one key piece of information is what kind of equipment individuals need. Gas masks, onions, and soda mitigate the effect of tear gas, while hammers, slingshots, and shields are necessary if projectiles are to be employed. Individuals also need to know to where these supplies need to be delivered, as some groups of protesters may be marching peacefully while others in different neighborhoods confront the police. The same logic holds once a protest site, such as a city's main square, is established. At this point, the protest site becomes a miniature city; the provision of food, medical supplies, sanitation, communications equipment, and security all require coordination.

A comparison with military tactics clarifies the importance of situational awareness. The chief advantage of German armored divisions at the start of World War II was coordination enabled by new communication technology. Equipped with radios, unit commanders could communicate with their tanks in real time, maintaining tactical awareness throughout a battle and so allowing them to exploit enemy weaknesses or cover their own (Citino 2004). Iraqi forces in the first Gulf War had not learned this lesson: battlefield commands flowed through centralized headquarters in Baghdad, and Coalition forces were able to bomb these facilities, hindering the ability of Iraqi frontline forces to respond to battlefield developments (Press 2001). In protests, it is the police who have traditionally had the coordination advantage because of their distributed communication, while protesters have often lacked a similar ability.

Peripheral individuals are better positioned to coordinate than the core. Even in an authoritarian setting, the existence of widespread discontent is often not a surprise. In Tunisia and Egypt, for example, it was widely known that the regimes were unpopular. In Tunisia, oligarchic elites and weak rule of law alienated large segments of society, from students to the working class, especially outside of Tunis, and desperation suicides were not uncommon events (Al-Zubaidi and Cassel 2013; Breuer 2012). In Egypt, police indiscretion, religious persecution, and economic instability similarly dispirited a majority of the population (Gunning and Baron 2013, 97–127). It was well understood in these countries that dissent was widespread and a minority of a society benefited from current policies at the expense of most others. Widespread, commonly understood dissatisfaction means that latent desires for policy change are known to exist, rendering the task one of coordinating protest. The periphery then drives mobilization because it signals that disparate, numerous groups of individuals are acting on this discontent.

Signalling and situational awareness allow peripheral members to coordinate their action. For example, a message such as "#jan25 protests will take place all throughout cairo, including shubra, mohendessin, in front of cairo university and on arab league street" issued on the morning January 25th, the first major day of protests in Egypt, provides information about where individuals who want to protest can join others (Idle and Nunns 2011, 33). Information less explicitly about coordination can also have a coordinating effect. A large amount of the communication leading up to a protest focuses on supplies needed, how to dress, how to behave towards the police, and the identity of protesters. This communication does not tell people when or where to go, but it helps them estimate levels of support in the population and danger they may face

(Gerbaudo 2012; Lohmann 1994). The more people that provide this information, the easier protest coordination becomes.

The importance of signalling and protest information leads to a primary hypothesis that can be broken down into constituent parts.

H1 As coordination from the periphery of a social network increases, more protests should occur.
H2 As coordination from the core of a social network increases, there should be no change in the number of protests.

The importance of peripheral participation as a signal of broad support is found in the experience of Egyptian mobilization on January 25. As groups of protesters marched through outlying neighborhoods, they urged onlookers to leave their shops, apartments, and workplaces. Many did, and the protest size snowballed (Cambanis 2015, 51). Protesters also emphasized the different parts of society they represented, with particular care taken to recruit outside of the middle class as well as emphasize independence from the Muslim Brotherhood (Gunning and Baron 2013, 180). The initial mobilization therefore included youth, members of football fan clubs, the poor and working class, in addition to individuals who were habitual protesters. Moreover, habitual protesters situated in the core of the Egyptian social network had tried to initially protest on January 18; only five activists protested, reflecting the importance of mobilization from the periphery (Gunning and Baron 2013, 91).

That peripheral members of a social network provide more information than the core finds support in other settings as well. In a study of diffusion on Facebook, Bakshy et al. (2012) find that weak ties are responsible for most information diffusion because they are more numerous than strong ties (individuals who interact frequently), just as those on the periphery are more numerous than those in the core. Recruitment to Spain's indignados movement, which started less than four months after Egypt's first protests, was characterized by individuals' exposure to the same information from different sources (González-Bailón, Borge-Holthoefer, Rivero and Moreno 2011). Adoption of political attitudes is also increased after exposures from different sources (Romero, Meeder and Kleinberg 2011), and controlled experiments have confirmed the importance of multiple sources of exposure for changing health attitudes (Centola 2010). Complex contagion also drove mobilization processes during the collapse of the Soviet Union (Opp and Gern 1993) and the American Civil Rights movement (McAdam 1986), though scholars at the time did not use that language. Finally, this spontaneous collective action might be the process which drives mass urban mobilizations that bring together disparate groups of individuals who previously did not engage in antiregime behaviors (Beissinger 2013; Tufekci 2014).

That protest is a complex contagion explains why many states have large domestic intelligence apparatuses and fear mass public gatherings. If an individual desiring to protest is concerned that sharing that information will lead to punishment, individuals are less likely to form connections with other individuals. In network terms, there will be fewer bridges between communities, inhibiting the spread of protest mobilization information. (If protest were a simple contagion phenomenon, a small number of protesters could have a large effect, and governments would have to make the costs of protest very high to prevent any display of antiregime sentiment.) Large public gatherings therefore provide one of the few occasions individuals have of bridging their immediate social communities; these bridges may cause individuals' protest thresholds to be surpassed, and a chain reaction of protests may ensue. For example, protests in Egypt against the Iraq War and marking the third anniversary of the Second Intifada led to the first large-scale public chants against Hosni Mubarak and started the process by which previously disconnected groups of individuals began to coordinate their antiregime actions (Gunning and Baron 2013, 39–47). In Russia in 1917, an industrial lockout, International Women's Day, and military leave brought together tens of thousands of workers, women, and disgruntled soldiers into the streets of St. Petersburg; the Romanovs fled four days later (Kuran 1989, 63). China even allows criticism of government officials and policy so long as it does not lead to appeals for collective action (King, Pan and Roberts 2013).

That protest is a complex contagion phenomenon also does not mean core members are unimportant in terms of protest mobilization. There are at least three mechanisms by which core members can facilitate protests: convincing individuals to blame their dissatisfaction on government policies, revealing the state is weaker than commonly believed, and fostering group identity. First, a core member can help those on the periphery ascribe their policy dissatisfaction to specific policies of those in power because the information to assign blame is a simple contagion phenomenon. As Javeline summarizes: "individuals faced with any grievance should be more likely to protest if they can make specific attributions of blame for the grievance and that one mechanism by which entrepreneurs [core individuals] might solve collection action problems is by first solving blame attribution problems" (Javeline 2003, 119). Second, core members can engage in violence which, if not terminated, reveals that antiregime preferences are widespread and the regime may be weak (Bueno de Mesquita 2010). Third, core members can create norms of solidarity, causing individuals to calculate their participation based on group gains (Goldstone 1994). Once individuals see themselves as part of a larger group, the benefits of protest increase while the costs decrease, making them more likely to mobilize when the opportunity arises.

These core-based mechanisms are not related to protest mobilization, however, as they occur before mass protests. They predispose individuals to be ready to mobilize, but they do not directly mobilize. In the language of Timur Kuran, they cause preferences to change, but they do not provide the initial spark

(Kuran 1989, 63–6). The theory of spontaneous collective action also treats the spark as exogenous.

Scope Conditions

There are two primary scope conditions to the application of the theory to the Arab Spring. First, a country's regime type may determine whether or not protest is a complex contagion phenomenon. Arrests, perfunctory trials, and long jail terms were standard state practices from Morocco to Bahrain (Bellin 2012; Gunning and Baron 2013; Khatib and Lust 2014), making it difficult for some core social network members to organize collective action. Second, mobilization is bounded by the costs a state imposes on protesting. Libya, Saudi Arabia, Bahrain, Syria, and Egypt in 2013 engaged in sustained violent repression of collective action, with heterogenous outcomes.

Authoritarian regimes are likely to repress individuals who impugn them, as targeted repression is a more effective tactic than indiscriminate killing (Siegel 2011). Arbitrary jailing, torture, forced exile, and threats to family are all common tactics used to silence antiregime individuals. In countries where those who desire policy change and are central to a network are routinely intimidated or silenced, they may not have the ability or desire to engage in coordination activities, and coordination would necessarily occur through those on the periphery. Moreover, in countries tolerant of mass gatherings, individuals may have lower thresholds of participation since they do not fear repression. If an individual does not expect protest to be large to be safe, he or she may join a protest alone or after hearing about it from a core social network member. In these cases, protest is more likely to be a simple contagion event and so be more affected by core members of a social network.

Second, any state can stop protests if it is willing to impose high enough costs (Blaydes and Lo 2011). In March of 1988 in Burma, protests started over an event just as random as a Tunisian fruit vendor lighting himself on fire: a youth arrested for fighting other youth was released from jail through political connections. Tension boiled over the summer, a general strike started on August 8, and the state engaged in ambiguous amounts of repression. On September 18, repression became less ambiguous as a result of an army coup; the ensuing repression resulted in at least 1,000 deaths in Rangoon, 3,000 nationwide (Ferrara 2003). Protests, which had been stronger throughout August but were tapering by mid-September, ceased. In 1989, a protest movement in China grew over the course of several months; by the end of May, Beijing hosted 250,000 soldiers, and multiday violent repression began on June 3. That repression, in conjunction with the arrest of party leaders and Communist Party reformers, squelched the movement. In 2011 in Egypt, individuals soon realized that the armed forces were not going to repress protests, yet in August 2013, the Egyptian army massacred hundreds, perhaps thousands, of pro-Muslim Brotherhood supporters who were staging long-term sit-ins at two Cairo squares after the July 3 coup against President Mohamed Morsi (Meirowitz and Tucker 2013). Repression against secular activists also increased, with those continuing to protest facing lengthy jail sentences or death (Mackey 2015). Bahrain's security forces killed protesters at the Pearl Roundabout, after welcoming a coalition of forces from Gulf states; leaders of al-Wefaq, the main Shia opposition party that participated in government before the start of protests, are now in jail, and the party's leader faces a four year sentence for inciting violence against the monarchy (Kerr 2015). While a state faces internal and external costs from repression, the ultimate success of any protest mobilization depends on the state's willingness to repress.

DATA

The Integrated Conflict Early Warning System, a machine-coded events dataset that reads newspaper articles, provides the dependent variable, number of protests, across 16 countries in the Middle East and North Africa from November 1, 2010 through December 31, 2011(Boschee, Lautenschlager, O'Brien, Shellman, Starz and Ward 2015). ICEWS codes 20 categories of events of increasing severity, from public statements through unconventional mass violence. All events coded as protests in one of these 16 countries country are kept. These countries are Morocco, Algeria, Tunisia, Libya, Egypt, Lebanon, Syria, Jordan, Saudi Arabia, Oman, Yemen, Bahrain, the United Arab Emirates, Kuwait, Iraq, and Qatar; Israel is excluded. Figure 1 shows the ICEWS recording of protest in two high-protest (Egypt, Bahrain) and low-protest (Morocco, Qatar) countries.

Social media data are ideal for understanding protest, for three reasons. First, in states that control information disseminated through newspapers, radio, and television, social media are one of the few independent sources of information (Edmond 2013). Social media have therefore become a tool for citizens to gather and disseminate information in information-scarce environments such as authoritarian regimes. In this way, social media may have a similar effect as independent media (Egorov, Guriev and Sonin 2009) or the disclosure of economic data by an autocrat (Hollyer, Rosendorff and Vreeland 2015). Second, 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 and Gomaa 2012). Social media has therefore become a critical component of many protest movements, starting with the 2009 Iran election protests and continuing through the Ukraine civil war (Burns and Eltham 2009; Rahimi 2011). Third, it provides the best temporal resolution of any data source. It is therefore one of the few sources available to researchers interested in dynamic processes that can provide microlevel information on these processes.

FIGURE 1. Protests per Million Inhabitants

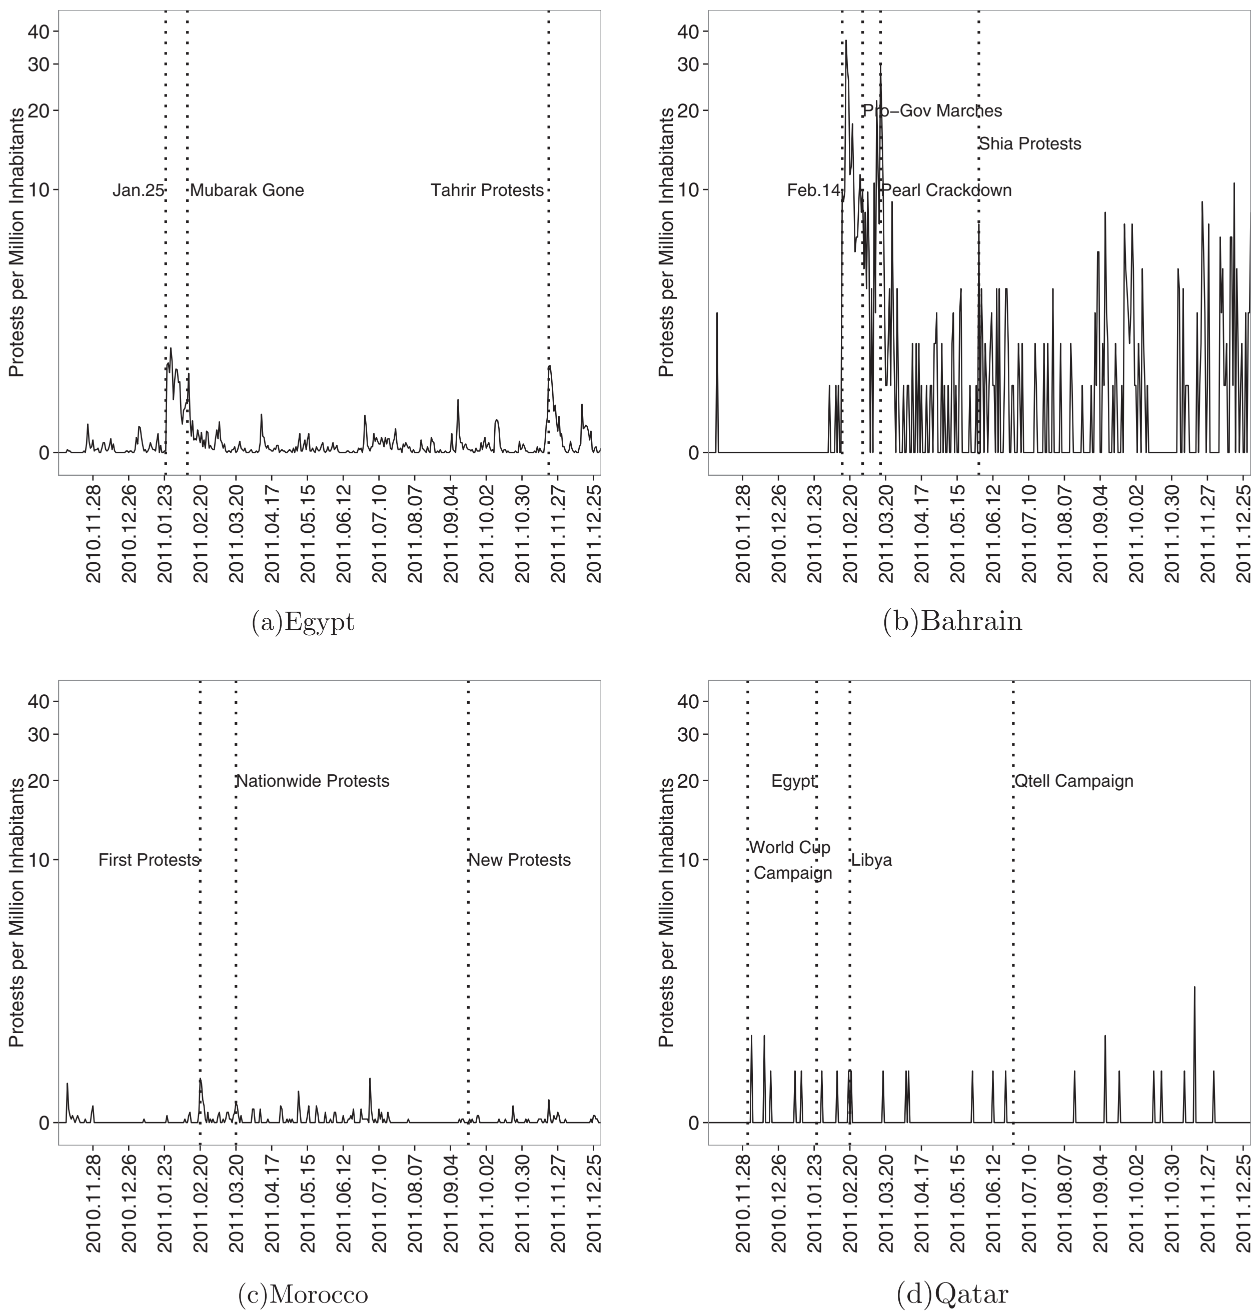

Notes: This figure shows that ICEWS captures different levels of intensity of each country's protests, both temporally and in cross section. Egypt, which experienced sustained, widespread protest, has the most recorded protests of any country in the dataset, but it has fewer per person than Bahrain. Morocco had a sustained protest campaign that did not mobilize as many people as Egypt or Bahrain, and Qatar experienced no protests. ICEWS' count of protests also varies during days they are expected to.

Twitter, a global social media platform, provides data on daily, individual-level communication. It is a global social network with over 500 million users generating almost 500 million daily messages (tweets). Anyone with an internet connection or phone can access it, and most users create and consume content using their mobile devices; contrary to popular belief, one can compose and consume tweets from any kind of phone, though smartphones greatly facilitate the process. The company does not edit or censor its users' tweets, so the content of the network reflects what individuals are discussing at any moment.[2] Only China and North Korea have completely blocked access to it, though

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

countries have temporarily blocked it at different times.

There are four reasons to prefer Twitter as a data source to other social media platforms. First, it is one of the most used social media platforms, usually second only to Facebook (Duggan and Smith 2013). Second, it is often used during crisis events to disseminate information, including during protests (Earl, McKee Hurwitz, Mejia Mesinas, Tolan and Arlotti 2013; Tonkin, Pfeiffer and Tourte 2011). Third, though it is used to discuss political events such as protests, it is also used to engage in quotidian topics like celebrity gossip, the weather, and sports (Boyd, Golder and Lotan 2010; Sinha, Dyer, Gimpel and Smith 2013). In the sample of tweets used later to train a naïve bayes classifier, *almost 75% were not about political events*. Fourth, Twitter provides a large amount of its data through two programming interfaces, making Twitter data easier to obtain than Facebook's. While other sites with social networking components, such as YouTube or reddit, are also relatively easy to gather data from, none are also used as comprehensively as Twitter.

Moreover, the norms of communication on Twitter makes it the most reliable way to measure coordination across so many countries and days. There are four ways a user can modify a plaintext tweet. The most common is the # symbol, known as the hashtag. Individuals will affix a hashtag to the front of a word to associate it with a certain conversation, e.g., "Eyewitnesses: NDP thugs throwing molotov cocktails inside the Egyptian Museum. I repeat NDP thugs, NOT anti-Mubarak protesters. #Jan25 #fb". If a different user then searches for messages containing "#Jan25" or "#fb", 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 and Kleinberg 2011). 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 and Burgess 2011; Lehmann, Gonçalves, Ramasco and Cattuto 2012).

Twitter makes it easy to find all tweets containing a hashtag. A user interested in upcoming protests could therefore search, from her smartphone or a computer, for "#jan25", "#egypt", or other hashtags and retrieve every tweet containing those hashtags.[3] 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 the key coordination mechanism.

Twitter data come courtesy of researchers at Northeastern University's Laboratory for the Modeling of Biological and Socio-Technical Systems (Mocanu, Baronchelli, Perra, Vespignani, Goncalves and Zhang 2013; Steinert-Threlkeld, Mocanu, Vespignani and Fowler 2015). The tweets involved in this analysis were extracted from Twitter's 10% API, an unbiased sample of 10% of all public activity on the platform. There are two ways in which country of origin was identified. First, if a Twitter user has enabled location sharing, the tweet will have GPS coordinates, and Twitter will assign a two letter country code to those tweets. Each tweet is then read for a two letter code corresponding to one of the 16 countries and saved if there is a match. Second, users can report their location as part of their profile, and that location is reported as metadata with each tweet. The user-reported location is then compared to a dictionary of cities and country names to assign each tweet to a city or country; if that location is part of this study, the tweet is saved.[4] Unlike previous studies that analyze contentious events, tweets in this dataset were not selected based on hashtags, providing a representative sample of what Twitter users were actually talking about, e.g., protests or the weather, during this period. Only 19.74% of all tweets in this sample contain a hashtag, and most are apolitical.

1.95% of tweets in this sample have GPS coordinates, with the location information of the rest coming from user-reported location. These numbers correspond with other work that finds more than an order of magnitude more tweets when using self-reported location (Leetaru, Wang, Cao, Padmanabhan and Shook 2013). It is worth nothing that tweets in the United States with GPS coordinates exhibit bias towards urban areas, nonwhites, and high-income groups (Malik, Lamba, Nakos and Pfeffer 2015), and there is some evidence that users of Twitter in Egypt tend to be well-off individuals in cities (Tufekci and Wilson 2012). Malik et al. (2015) do not include tweets with user-reported location and Tufekci and Wilson (2012) do not ask whether users geotag their tweets, so it is unclear if using user-reported location removes these biases.

For a comprehensive of review of Twitter as a data source and a tutorial on how to use it for social science research, see Steinert-Threlkeld (2017).

MEASURES

Coordination

A Gini coefficient for hashtags operationalizes coordination. The Gini coefficient, which ranges from zero to one, usually measures income inequality, but it can

started to delete accounts from the Islamic State of Iraq and Syria that are deemed to incite violence.

[3] This search is not case sensitive: a user searching for "#jan25" will see the same results as one searching for "#Jan25". The searched term will not return tweets that use the character string without a hashtag, e.g., a tweet that says "police thugs r everywhere in egypt jan25" will not show up in search results.

[4] For more detail, see the Materials and Methods section of Mocanu et al. (2013). That article uses only tweets from that stream with GPS coordinates for its analysis, whereas I use tweets with GPS coordinates or user-reported location because there were not enough tweets with GPS coordinates in the countries in this study.

be used on any orderable discrete quantity. Instead of measuring wealth per person, it here measures occurrences per hashtag per day per country; a one means that one hashtag accounts for all hashtags used in that country on that day, a zero that every observed hashtag occurs the same number of times. This measure is labeled $Coordination_{i,t}$ for the rest of the article, and the Supplementary Material provides a graphical explanation of the operationalization.

Equation (1) shows this calculation. For each day t in each country i, there exist n unique hashtags. $Coordination_{i,t}$ counts the number of times each hashtag j occurs (c_j) and uses those counts to calculate the Lorenz Curve of (hashtag) inequality, for $i = 1, ..., 16$ and $j = 1, ..., 426$.

$$Coordination_{i,t} = \frac{2\sum_{j=1}^{n} j\,c_j}{n\sum_{j=1}^{n} c_j} - \frac{n+1}{n}. \tag{1}$$

Three other Twitter behaviors that may impact coordination are measured. Retweets, equivalent to forwarding an email to one's entire contact list, can also promote coordination. 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 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 number of times, though a user who sees a retweet only knows that at least one person retweeted it; in practice, most tweets are not retweeted, those that are are not retweeted often, and the retweet rate decays to almost 0% after 24 hours (Kwak, Lee, Park and Moon 2010; Liere 2010; Starbird and Palen 2012).[5]

A message can also contain a user mention or a link. A message that directly refers to another user by name is said to contain a user mention. If a user writes, "@ramezm i noticed a debate: #25jan or #jan25", @ramezm will receive a personal notification about the message; a tweet with a user mention is still viewable by the followers of the original author. Tweets also often contain links to photos and articles, though those messages are rarely retweeted (Suh, Hong, Pirolli and Chi 2010).

Retweets, links, and mentions are not as effective at promoting coordination as hashtags. While more retweets of one tweet means that more people have seen the same set of information, the prevalence of hashtags means the information in a retweet is also available to those searching for hashtags that the retweet happen to contain. The same logic is true of links: if a link is meant to provide coordinating information, it will almost certainly contain a hashtag that is also relevant to coordination. While it is possible that user mentions have a strong coordinating effect outside of their employment of hashtags, they are dyadic and tend to be part of conversations—they are not used to mobilize protesters.

Equation (2) shows the calculation of these other measures of potential coordination. For each day t and country i, the measure counts the number of tweets, K, and tweets with measure M_k. m is a count of tweets with a link, mention of another user, or that are a retweet, depending on the measure in question. Dividing the measure by the number of tweets that day in that country quantifies the amount of other possible coordination that could have occurred in addition to $Coordination_{i,t}$,

$$MPercent_{i,t} = \sum_{i=1}^{16}\sum_{t=1}^{426}\frac{1}{K} * \sum_{k=1}^{K} m_k. \tag{2}$$

Note that $Coordination_{i,t}$ is one variable that encompasses the two mechanisms, signalling and protest information provision. This measurement choice was made for four reasons. First, $Coordination_{i,t}$ should measure information protest information provision because it measures hashtag concentration, and individuals on Twitter use hashtags to quickly identify their tweets as being about a specific topic. During periods of heightened political awareness, the most common hashtags are most likely to be about politics; a tweet with "#jan25" is not likely to be about sports or the weather, for example. While this measure could create false positives—days that appear to have high coordination but are really people talking about something else like a meme or a cultural event—the Supplementary Material's case study shows that this is not the case.

Second, $Coordination_{i,t}$ is preferred to a supervised learning model of protest information provision because it scales easily and is directly comparable across countries. Tweets often contain slang that varies by country, so making a supervised learning model for each country is a large project in its own right. Aside from requiring much more labor, creating a content model, whether supervised or unsupervised, risks constricting results to words or topics that the researcher has an *a priori* expectation will matter (Grimmer and Stewart 2013). A hashtag Gini, on the other hand, is agnostic to what words people say or how many topics they discuss; caring only about the hashtags, it will measure any hashtag used (not just the ones thought of in advance), revealing after the fact which hashtags are most salient.

Third, $Coordination_{i,t}$ is preferred to selecting specific hashtags because it scales easily and is directly comparable across countries. The hashtags used to coordinate an event are different in each country and change over time. Determining which hashtags to use for a given period of time therefore requires subject matter expertise on the digital arena of many countries over many months, which is not feasible for a large cross-national study. Moreover, because

[5] 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.

individuals pool on certain hashtags during protest periods, measuring hashtag concentration picks up on the most common protest hashtags. This behavior is explored in more detail in the Supplementary Material.

Fourth, $Coordination_{i,t}$ is preferred over a direct measure of signalling because it is not clear how to measure signalling on Twitter. The best approach would be to separate out the amount of signalling and protest information tweets which come from the periphery and the core. But it is not clear how to classify a "signal" tweet, and a more precise method of measuring protest information tweets does not scale well, as discussed in the previous paragraph. In addition, the three other coordinating behaviors discussed above can each signal peripheral participation. A link to a news article about the mass protests is not the same as a user saying, "I will protest tomorrow and I have never protested before," even if both tweets come from the same account. Similarly with a photo showing a diverse crowd at a protest. In Table 4, an attempt is made to measure the signal component of coordination, and the next subsection discusses how to tease apart that $Coordination_{i,t}$ which is from the core.

Core Coordination

To measure coordination from the core, one has to first identify individuals at the center the network. Identifying this core is difficult. There are too many users—20,094 in Bahrain and 79,235 in Egypt alone—in the Twitter data to assign manually an identity to each one, and that attempt would result in a low identification rate because Twitter does not require individuals to publicly disclose any identification information. One can measure, however, the number of followers each account has; this measure, in-degree centrality, is not as precise a measure of centrality as those created with complete network data (Kwak, Lee, Park and Moon 2010; Pei, Muchnik, Andrade Jr., Zheng and Makse 2014), but complete network data are not available. Those in the core are therefore approximated based on the distribution of popularity in each country. For the main model, a tweet belongs to a core member if its author's number of followers are at or above the 95th percentile for all users in country i. More formally,

$$Core = \begin{cases} 1 & \text{if } \mathrm{PR}_i(\mathrm{f}) \geq 0.95 \\ 0 & \text{if } \mathrm{PR}_i(\mathrm{f}) < 0.95 \end{cases}, \tag{3}$$

where PR(f) is the percentile ranking of the tweet based on the number of followers.

Previous work that manually identified a random sample of users from Tunisia and Egypt informed the selection of this threshold (Lotan, Ananny, Gaffney, Boyd, Pearce and Graeff 2011). Table 1 compares the number of followers and tweet production for the categories identified in Lotan et al. (2011) with this article's primary popularity threshold; the threshold used later is bolded. In Tunisia, the core measure appears to roughly be most similar to bloggers; in Egypt, to bloggers and activists, though the manually identified accounts in Egypt are much more popular than any of the popularity measures. Mainstream media accounts and employees of mainstream media are the most central in each country and skew the country-level results upwards. The Results section shows that varying the follower threshold does not change the result. The Supplementary Material also shows how tweet production and the ratio of the core's followers to the periphery's followers varies by country and threshold; users at the 95% threshold account for 10% of all tweets in Kuwait, up to 50% in Syria.

Having identified tweets produced from those in the core, one can then identify when the core engages in coordination. Because hashtags are the primary method of coordination and high levels of coordination lead to protest, the percentage of hashtags per country per day produced from the core is interacted that with the coordination measure. The percent of tweets with hashtags that are created in the core is defined as

$$Core\ Coordination_{i,t} = Coordination_{i,t} * \frac{1}{K} * \sum_{k=1}^{K} Core_k * Hashtag_k. \tag{4}$$

For each country i on each day t, each of the K tweets is read for whether it contains a hashtag and is from a core account. The number of those tweets is divided by the number of tweets in that country-day and interacted that with that country-day's level of coordination, resulting in a core coordination measure for that country-day. The regression results leave the constituent parts of the variable as the variable name to ease interpretation; the summary statistics use the shortened name to save space.

In-degree—the number of followers an account has—is chosen to measure core position for three reasons. First, it is the only network centrality measure which does not require complete network data. Twitter imposes limits on how often one can download data, making it impossible to create the adjacency matrix necessary to create measures like Eigenvector or k-core centrality. Second, one could approximate a complete network by inferring ties when a retweet or mention occurs, but that requires pooling data, losing the time component necessary for this theory (Barberá, Wang, Bonneau, Jost, Nagler, Tucker and González-Bailón 2015). Finding "hidden influentials," those who follow many more people than follow them (normal Twitter behavior) and who are mentioned much more than they mention others (which happens for well-known accounts like celebrities or politicians) is another strategy, but it again relies on pooling across time (Gonzalez-Bailon, Borge-Holthoefer and Moreno 2013). Third, this article's theory is about how position in a social network correlates with protest mobilization. If in-degree does not identify those who influence protest, then that means those in the network core do not influence protest. Future work should then

TABLE 1. Comparing Core Measure with Hand-coded Accounts

Group	Accounts	Followers Avg.	Tweet Avg.	Mention %	Retweet %	Hashtag %	Link %
Egypt - Lotan	37	15138.71	949.78	0.40	0.06	0.36	0.32
Mainstream Media	1	103927.00	5281.00	0.00	0.00	0.74	0.70
Nonmedia org.	2	23877.40	457.50	0.32	0.07	0.22	0.52
MSM employee	9	22463.50	650.22	0.41	0.01	0.21	0.21
Blogger	15	8394.17	1070.67	0.52	0.08	0.33	0.22
Activist	10	8036.55	703.40	0.42	0.07	0.28	0.33
Core 99.9 percentile	80	37001.28	924.69	0.33	0.02	0.39	0.44
Core 99 percentile	793	7104.31	736.38	0.45	0.05	0.27	0.32
Core 98 percentile	1585	4033.08	591.71	0.46	0.05	0.25	0.31
Core 97 percentile	2378	2875.48	515.21	0.46	0.04	0.25	0.31
Core 96 percentile	3170	2256.03	453.52	0.45	0.04	0.25	0.30
Core 95 percentile	**3962**	**1868.79**	**409.94**	**0.45**	**0.04**	**0.25**	**0.30**
Blackout	**740**	**8046.33**	**650.95**	**0.22**	**0.05**	**0.24**	**0.61**
Tunisia - Lotan	10	7942.94	248.60	0.33	0.12	0.59	0.56
Mainstream media	2	5604.49	741.00	0.16	0.12	0.77	0.78
MSM employee	1	52503.00	1.00	0.00	0.00	0.00	1.00
Blogger	3	1910.77	258.33	0.57	0.13	0.30	0.20
Activist	4	2496.28	57.00	0.59	0.09	0.36	0.29
Core 99.9 percentile	7	17749.31	206.71	0.25	0.06	0.22	0.68
Core 99 percentile	62	4880.14	410.92	0.37	0.06	0.27	0.55
Core 98 percentile	123	3095.44	444.93	0.44	0.14	0.31	0.43
Core 97 percentile	184	2392.96	374.43	0.45	0.13	0.31	0.40
Core 96 percentile	245	1968.22	337.18	0.47	0.12	0.30	0.37
Core 95 percentile	**307**	**1681.87**	**308.22**	**0.47**	**0.11**	**0.31**	**0.38**

Notes: Categories are borrowed from Lotan et al. (2011). They coded for accounts associated with mainstream media organizations, mainstream new media organizations (news sites that exist only online), mainstream media employees, any organization that is not a media organization (Vodafone and Wikileaks are their examples), bloggers, activists, digerati, political actors, celebrities, researchers, bots, and a residual category. Any of those categories not identified here means that no account from that category was found in the data. The bold rows represent the category used to identify core members. Other categories are used in robustness checks, with no changes to the results. For a discussion of the **Blackout** row, please see the Egypt case study.

find ways to identify individual accounts that are influential; one approach, handcoding accounts by profession, is explored in Table 5.

MODEL

The base model is

$$Protests_{i,t} = \beta_0 + \beta_1 * \Omega_{i,t-1} + \boldsymbol{\beta} * \mathbf{X}_{i,t-1} + Protests_{i,t-1} + \epsilon_{i,t}, \quad \textbf{(5)}$$

where Ω represents the independent variables 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 zero. Since $Protests_{i,t}$ is overdispersed and the zeroes are true zeroes, a negative binomial model instead of a Poisson or zero-inflated negative binomial is used.

Because high levels of coordination are colinear with high levels of hashtag usage, the model controls for the percent of a day's tweets that have hashtags, ensuring that it measures actual coordination and not a coincidental increase in hashtag usage. The models of peripheral coordination control for the percent of a day's tweets which are retweets, contain links, or mention another user because those features may have some coordination effect. The models of core coordination similarly control for the percent of all tweets with at least one hashtag that are from accounts of the core; the percent of all tweets that are retweets which are from core members; and so on for links and mentions.

There are three nonindividual controls: country fixed-effects, a lagged dependent variable, and a lagged measure of the number of repression events as measured by ICEWS. Repression is any event with a CAMEO code of exhibiting military posture (event root code 15), coercion (17), using unspecified unconventional violence (18), a physical assault (182), torture (1822), or death by physical assault (1823).

Every variable on the right-hand side is lagged by one day to mitigate any simultaneity effects. All models include country fixed effects but no day fixed effects, as the latter bias the errors and lead to underestimates of protests. Finally, all models are run with country-clustered standard errors.

Table 2 shows the correlation between the main independent variables, and Table 3 shows the average value of each variable per country (along with each's total tweets and protests).

https://doi.org/10.1017/S0003055416000769 Published online by Cambridge University Press

TABLE 2. Variable Correlation

	Protest$_{i,t}$	Coord.$_{i,t-1}$	Hashtag %$_{i,t-1}$	Retweet %$_{i,t-1}$	Link %$_{i,t-1}$	Mention %$_{i,t-1}$	Protest$_{i,t-1}$	Core Hashtag%$_{i,t-1}$	Core Retweet%$_{i,t-1}$	Core Link%$_{i,t-1}$	Core Mention%$_{i,t-1}$	Repression$_{i,t}$	Core Coord.$_{i,t-1}$
Protest$_{i,t}$	1	0.281	0.218	−0.012	0.111	−0.155	0.785	0.178	0.099	0.149	0.1	0.571	0.309
Coord.$_{i,t-1}$	-	1	0.594	−0.038	0.101	−0.237	0.295	0.535	0.372	0.519	0.349	0.269	0.88
Hashtag %$_{i,t-1}$	-	-	1	0.203	0.427	−0.448	0.231	0.521	0.351	0.405	0.186	0.185	0.649
Retweet %$_{i,t-1}$	-	-	-	1	0.11	−0.233	−0.013	0.122	0.279	0.156	0.098	0.029	−0.051
Link %$_{i,t-1}$	-	-	-	-	1	−0.659	0.106	0.325	0.232	0.23	0.024	0.063	0.291
Mention %$_{i,t-1}$	-	-	-	-	-	1	−0.155	−0.344	−0.265	−0.3	0	−0.126	−0.351
Protest$_{i,t-1}$	-	-	-	-	-	-	1	0.175	0.108	0.152	0.1	0.673	0.317
Core hashtag %$_{i,t-1}$	-	-	-	-	-	-	-	1	0.656	0.702	0.545	0.158	0.729
Core retweet %$_{i,t-1}$	-	-	-	-	-	-	-	-	1	0.583	0.501	0.152	0.467
Core link %$_{i,t-1}$	-	-	-	-	-	-	-	-	-	1	0.49	0.142	0.615
Core mention %$_{i,t-1}$	-	-	-	-	-	-	-	-	-	-	1	0.155	0.417
Repression$_{i,t-1}$	-	-	-	-	-	-	-	-	-	-	-	1	0.271
Core coord.$_{i,t-1}$	-	-	-	-	-	-	-	-	-	-	-		1

TABLE 3. Variables by Country

Country$_i$	Protests$_{i,t}$	Tweets$_{i,t}$	Coord.$_{i,t}$	Hashtag %$_{i,t}$	Retweet %$_{i,t}$	Link %$_{i,t}$	Mention %$_{i,t}$	Core Hashtag %$_{i,t}$	Core Retweet %$_{i,t}$	Core Link %$_{i,t}$	Core Mention %$_{i,t}$	Core Coord.$_{i,t-1}$
Egypt	3,379	3,742,648	0.59	0.21	0.04	0.29	0.42	0.52	0.55	0.47	0.48	0.31
Syria	2,057	229,476	0.60	0.34	0.03	0.61	0.22	0.54	0.49	0.48	0.35	0.37
Yemen	1,885	61,517	0.39	0.25	0.05	0.62	0.19	0.52	0.23	0.41	0.24	0.25
Tunisia	882	228,554	0.37	0.25	0.07	0.43	0.40	0.50	0.61	0.37	0.49	0.19
Bahrain	798	1,056,990	0.53	0.22	0.06	0.17	0.42	0.39	0.28	0.46	0.28	0.22
Libya	663	84,991	0.37	0.24	0.11	0.34	0.32	0.42	0.54	0.37	0.38	0.17
Iraq	585	146,113	0.35	0.18	0.09	0.28	0.35	0.44	0.51	0.31	0.43	0.16
Jordan	511	273,227	0.38	0.21	0.06	0.45	0.36	0.45	0.48	0.42	0.38	0.17
Morocco	298	300,454	0.25	0.18	0.07	0.34	0.43	0.40	0.40	0.33	0.34	0.11
Lebanon	261	522,891	0.36	0.21	0.06	0.28	0.41	0.50	0.61	0.50	0.41	0.18
Algeria	248	7,474	0.08	0.21	0.11	0.44	0.34	0.40	0.43	0.36	0.18	0.04
Kuwait	161	29,838	0.09	0.09	0.02	0.13	0.52	0.09	0.02	0.06	0.21	0.01
Saudi Arabia	156	4,425,797	0.48	0.13	0.05	0.16	0.51	0.46	0.37	0.45	0.44	0.23
Oman	150	8,509	0.02	0.09	0.05	0.31	0.48	0.06	0.02	0.11	0.14	0.00
UAE	58	1,531,524	0.35	0.15	0.07	0.20	0.43	0.36	0.37	0.42	0.31	0.13
Qatar	29	1,104,995	0.34	0.12	0.06	0.10	0.43	0.39	0.35	0.37	0.32	0.14

TESTS

Results

The main results are presented in Table 4. Columns 1 and 2 show coordination only from the periphery, and columns 3 and 4 build the models for core coordination in the same way. The main model, used throughout the rest of the article, is shown in column 5. Across most models, $Coordination_{i,t-1}$, the measure for peripheral coordination, is significant with a p value much less than 0.01. The only other significant variables are a lagged dependent variable (positive), intercept (negative), lagged repression (weakly positive), and noninfluential hashtag percent (positive). Note that Model 5, the full model, suggests that coordination from the core is inversely associated with protests.

Because the model is not linear, coefficients do not directly translate into changes in the outcome variable. The marginal effects of $Coordination_{i,t-1}$ & $Coordination_{i,t-1}$*$Core\ Hashtag\ \%_{i,t-1}$ are shown in Figure 2. Going from no coordination to the maximum observed values leads to about two additional protests, a 400% increase, while there exists no effect for core coordination.

A series of time series diagnostic tests confirm the model specification.[6] A Durbin-Watson test for serial correlation returns a test statistic of 1.9741 and p value of 0.1303, suggesting no serial correlation. The Dickey-Fuller coefficient is −11.25 and has a p value less than 0.01, so the dependent variable is stationary (visual inspection also confirms the stationarity). The Breusch-Pagan test statistic is 539.09 with a p value almost at 0; to control for the heteroskedasticity, I use country-clustered standard errors. Finally, a Lagrange-Multiplier test with the King & Wu test for two-way fixed effects returns a chi-square value of −0.81, so it is safe to avoid using time fixed effects.

Verification

There are three possibilities that may undercut these findings. First, the models may use the wrong measure of coordination, both for the core and the periphery. Second, the operationalization of core members may be wrong. Third, reliance on machine-coded data may bias in favor of finding results.

Figure 3 allays the first concern. To confirm that $Coordination_{i,t}$ measures coordination, a supervised learning model for Egypt and Bahrain was created. Those two countries were chosen because they experienced widespread protest and have too many tweets to code individually. For each country, 3,000 tweets were randomly selected and coded into overlapping categories, one of which was protest information. A naïve bayes classifier was trained on a random 95% of each country's coded tweets; the other 5% was used to validate out of sample performance. This process was repeated 30 times, and the results were averaged into a final model; this process is known as bagging and is akin to bootstrapping in regression. The resulting model is applied to each country's tweets, creating a classification for every single tweet in the sample from Egypt and Bahrain. Once protest information tweets are identified, they are aggregated by country-day and compared to the $Coordination_{i,t}$ measure. That result is shown in Figure 3: there is a strong positive relationship between the measure of coordination and the actual number of protest information tweets. See the Supplementary Material for the codebook and more details on the supervised learning model.

Two placebo tests also address concerns about $Coordination_{i,t}$. The first test shows that the correlation of $Coordination_{i,t-1}$ and $Protest_{i,t}$ decreases substantially as an increasing number of the most popular hashtags on a country-day are removed, suggesting that coordination on nonprotest hashtags does not drive subsequent protest. The second test shows that the correlation of $Coordination_{i,t}$ and $Protest_{i,t}$ peaks with a one-day lag and decreases monotonically before and after. The Placebo Tests subsection addresses these tests in more detail.

Table 5 addresses the second concern. In column 1, accounts from Lotan et al. (2011) in this sample are identified and controlled for. Accounts from activists or bloggers are called "Online Actors," while there are not variables for politician, researcher, digerati, or celebrity accounts because those were not found in the sample from either country. It appears that mainstream media accounts (official accounts of news organizations) do positively correlate with subsequent protest. The main coordination measure is still strongly significant, as is retweet percentage; links appear to decrease in rate leading up to protests. The second and third columns of Table 5 show alternative measures of coordination in the core. $Core\ Reachout\ \%_{i,t-1}$ measures the percent of all retweets and mentions that come from those in the core. Core coordination may occur through those in the core engaging with specific individuals (mentions) or acting as information brokers (finding important tweets and retweeting them), not through hashtags. Column 2 controls for this possibility. On the other hand, the core may have a coordination effect simply by being active leading up to protests; their activity may signal a breakdown of support for the regime, a willingness to incur high personal costs that inspires the periphery, or they just may not use hashtags. There appears to be some effect for tweets from the core—$Core\ Tweet\ \%_{i,t-1}$ is significant at $p \leq 1$—but their tweet activity on high coordination days does not correlate with subsequent protest. In all three models, coordination from the periphery is still significant and coordination from the core is not.

To confirm the 95% threshold used to identity the core, the threshold was varied from the 80th percentile to the 99.9th, and Model 5 from Table 4 is rerun for each threshold. Figure 4 shows how the significance level of $Coordination_{i,t-1}$*$Core\ Hashtag\ \%_{i,t-1}$ varies as the percentile threshold changes; the vertical lines are

[6] A panel OLS model is used for the diagnostics to ease calculation of the test statistics. Robustness checks show that the panel OLS results match those of the negative binomial.

TABLE 4. Peripheral Coordination and Protest

	DV: $Protest_{i,t}$				
	Coordination		Core Coordination		Full Model
	(1)	(2)	(3)	(4)	(5)
$Coordination_{i,t-1}$	1.932*** (0.472)	1.936*** (0.469)	1.809** (0.766)	1.830** (0.774)	2.575*** (0.639)
Hashtag $\%_{i,t-1}$		0.707 (0.539)			0.578 (0.616)
Retweet $\%_{i,t-1}$		0.405 (0.876)			−0.362 (1.019)
Link $\%_{i,t-1}$		−0.536 (0.386)			−0.763* (0.378)
Mention $\%_{i,t-1}$		− 0.858* (0.512)			− 0.921** (0.410)
$Repression_{i,t-1}$	0.020* (0.011)	0.020* (0.012)	0.021* (0.011)	0.022** (0.011)	0.021* (0.012)
$Protest_{i,t-1}$	0.127*** (0.010)	0.121*** (0.011)	0.126*** (0.010)	0.125*** (0.010)	0.116*** (0.010)
Core hashtag $\%_{i,t-1}$			0.600 (0.393)	0.631 (0.441)	0.900** (0.372)
Core retweet $\%_{i,t-1}$				0.063 (0.398)	0.158 (0.413)
Core link $\%_{i,t-1}$				0.450 (0.553)	0.711 (0.565)
Core mention $\%_{i,t-1}$				−0.485 (0.316)	−0.159 (0.258)
$Coordination_{i,t-1}$* core hashtag $\%_{i,t-1}$			−0.313 (1.302)	−0.503 (1.333)	−1.868** (0.942)
Intercept	−0.934*** (0.044)	−0.569** (0.229)	−1.182*** (0.160)	−1.295*** (0.215)	−1.019*** (0.207)
Country FE	yes	yes	yes	yes	yes
N	6,800	6,620	6,800	6,800	6,620
Log likelihood	−8,469.486	−8,296.844	−8,463.422	−8,459.476	−8,280.817

Notes: *$p < 0.1$; **$p < 0.05$; ***$p < 0.01$. Model 3 includes $Coordination_{i,t}$ and $Core\ Hashtag_{i,t-1}$ because those are the constituent parts of the measure of influential coordination.

FIGURE 2. Marginal Effects of Peripheral and Core Coordination

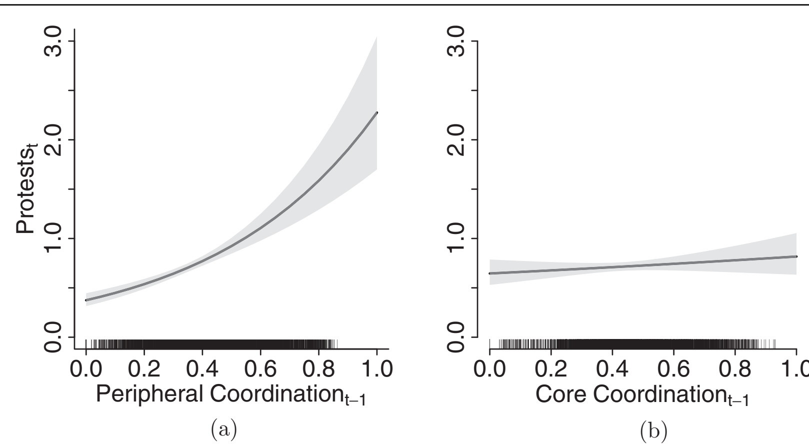

TABLE 5. Robust to Operationalization of Core, Periphery

	$Protest_{i,t}$		
	Core Manual ID (1)	Core Reachout (2)	Core Tweet Share (3)
$Coordination_{i,t-1}$	3.041***	1.678**	2.308***
	(0.419)	(0.729)	(0.772)
Hashtag $\%_{i,t-1}$	−1.728	0.412	0.448
	(7.574)	(0.641)	(0.620)
Retweet $\%_{i,t-1}$	12.705***	−0.130	−0.246
	(1.407)	(0.926)	(0.934)
Link $\%_{i,t-1}$	−1.525	−0.816**	−0.852**
	(1.341)	(0.359)	(0.374)
Mention $\%_{i,t-1}$	−1.596	−0.701	−0.941**
	(6.119)	(0.460)	(0.433)
Mainstream media $\%_{i,t-1}$	18.313***		
	(5.636)		
MSM empl. $\%_{i,t-1}$	62.784		
	(53.143)		
Online actor $\%_{i,t-1}$	−7.310**		
	(3.191)		
Spam $\%_{i,t-1}$	−7.989		
	(17.187)		
$Repression_{i,t-1}$	0.016***	0.020*	0.021*
	(0.005)	(0.011)	(0.011)
$Protest_{i,t-1}$	0.044***	0.118***	0.116***
	(0.009)	(0.010)	(0.010)
Core reachout $\%_{i,t-1}$		1.358	
		(1.247)	
Core tweet $\%_{i,t-1}$			1.078*
Core hashtag $\%_{i,t-1}$		0.416	0.366
		(0.514)	(0.518)
Core retweet $\%_{i,t-1}$		−0.233	0.101
		(0.580)	(0.419)
Core link $\%_{i,t-1}$		0.512	0.365
		(0.528)	(0.650)
Core mention $\%_{i,t-1}$		−1.394	−0.529
		(1.092)	(0.370)
$Coordination_{i,t-1}$*core reachout $\%_{i,t-1}$		0.595	
		(1.074)	
$Coordination_{i,t-1}$*core tweet $\%_{i,t-1}$			−1.276
			(1.246)
Intercept	0.335	−0.848***	−0.832***
	(4.195)	(0.212)	(0.162)
Country FE	yes	yes	yes
N	830	6,620	6,620
Log likelihood	−1,712.510	−8,282.427	−8,282.744

$^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$

at ±1.96 to show significance at the 5% level. Figure 4a shows the result from 10 models, where model 1 uses a cutoff at the 99.9th percentile and model 10 is at the 99th; Figure 4b shows the result from 20 models, where model 1 uses a 99th percentile threshold and the 20th uses the 80th percentile. In all iterations, $Coordination_{i,t-1}$, $Repression_{i,t-1}$, and $Protest_{i,t-1}$ remain significant.

The results in Figure 4 are intriguing. The effect of core coordination hovers around zero for most of the threshold's range and is distinguishable from zero at the 99th and 99.1st percentiles, as well as at the 95th percentile. On the other hand, at the upper extreme of the follower distribution, the 99.7th percentile and above, the sign on *Core* $Coordination_{i,t-1}$ is positive and significant using a 95% confidence interval.

The apparent positive effect from core coordination above the 99.7th percentile should not, however, be assigned much weight, for three reasons. First, these models also find that *Core Link* $\%_{i,t-1}$ is negative and significant, with an effect from half as strong to equal with that of core coordination. Second, the pooled re-

TABLE 6. Core Threshold Descriptive Statistics Across Countries

Country	Group	Accounts	Followers Avg.	Tweet Avg.	Mention %	Retweet %	Hashtag %	Link %
Algeria	99.7th percentile	2	6786.50	1.50	0.33	0.00	0.33	1.00
Bahrain	99.7th percentile	61	15294.42	530.87	0.17	0.03	0.57	0.49
Egypt	99.7th percentile	238	17975.93	992.12	0.37	0.06	0.32	0.42
Iraq	99.7th percentile	14	23258.58	455.00	0.74	0.07	0.29	0.05
Jordan	99.7th percentile	26	9197.23	436.38	0.19	0.03	0.35	0.73
Kuwait	99.7th percentile	10	25240.53	7.00	0.40	0.07	0.07	0.30
Lebanon	99.7th percentile	52	7997.13	550.71	0.19	0.15	0.26	0.63
Libya	99.7th percentile	11	17287.99	276.36	0.39	0.05	0.59	0.25
Morocco	99.7th percentile	30	19132.81	165.03	0.45	0.03	0.27	0.37
Oman	99.7th percentile	2	122889.50	1.50	0.67	0.00	0.00	0.33
Qatar	99.7th percentile	65	23668.00	428.31	0.42	0.06	0.29	0.33
Saudi Arabia	99.7th percentile	266	13375.55	588.77	0.61	0.03	0.19	0.24
Syria	99.7th percentile	11	5484.73	630.27	0.14	0.02	0.39	0.74
Tunisia	99.7th percentile	19	10068.36	377.95	0.34	0.04	0.31	0.63
UAE	99.7th percentile	142	22961.05	261.39	0.33	0.05	0.32	0.36
Yemen	99.7th percentile	5	5132.13	371.60	0.05	0.00	0.33	0.75

FIGURE 3. Verifying Operationalization of Coordination

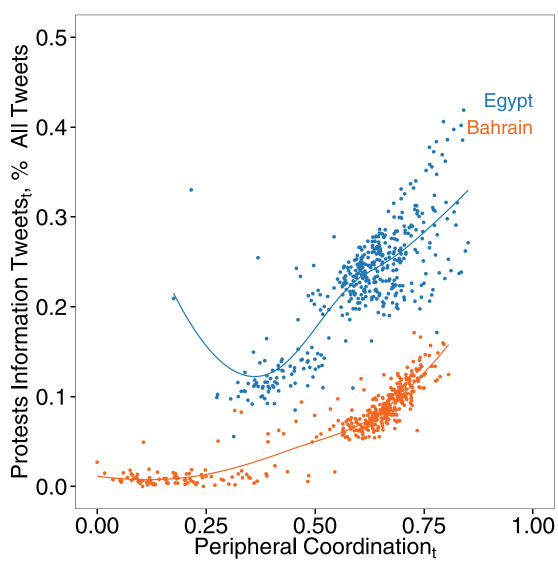

sults are driven by outlier countries with few tweets and users at or above the 99.7th percentile threshold. Whether the threshold is 99.7, 99.8, or 99.9, the effect disappears when the model is rerun using only Bahrain, Egypt, Jordan, Lebanon, Qatar, Saudi Arabia, and the United Arab Emirates, the only countries with more than 10,000 tweets from users at or above the 99.7th percentile. All these countries are also the only ones with more than 25 users at this point in the distribution (except for Morocco, with 30). Because the final dataset is at the country-day level, it does not distinguish between a day in Egypt that may have 500 tweets from 20 accounts in the 99.7th percentile core from one in Algeria that has 1 tweets from 1 of the 2 accounts above the same cutoff. The resulting models therefore overweight Algerian core users. Rerunning the main model with only the seven countries just described therefore provides a more accurate understanding of the dynamics this far into the followers' distribution, and models on these seven countries show no effect for members of the core. Third, it is more likely than not that these accounts represent institutions such as news outlets or nonprofit organizations than people.[7] As Table 6 shows, these accounts are frequent tweeters, and those tweets are more likely to contain hashtags or links. Such behavior is most similar to how news organizations use Twitter (Lotan, Ananny, Gaffney, Boyd, Pearce and Graeff 2011); Table 1 in the Measures section compares the 99.9th percentile core users to confirmed news accounts in Tunisia and Egypt, showing similarity between the two. The probability that these accounts are news organizations is further increased by rerunning the main model using only Arabic tweets. Using only Arabic tweets when the core is defined at the 99.7th, 99.8th, or 99.9th thresholds, $Coordination_{i,t-1}{*}Core\ Hashtag\ \%_{i,t-1}$ is not statistically significant. Overall, the positive effect suggested in Figure 4a is probably driven by a few media accounts in countries with less Twitter data than others in the sample.

Table 7 verifies the ICEWS dependent variable. ICEWS relies on news reports, and these reports have well-known biases in coverage (Davenport and Ball 2002; Eck 2012; Herkenrath and Knoll 2011). Machine-coded events data can suffer from event duplication (Caren 2014; Hammond and Weidmann 2014), though one of ICEWS' strengths is its focus on event deduplication. The results could therefore be driven by news media's bias towards major, unexpected events and duplicated events. In the first column of Table 7, ICEWS' count of public statements is the dependent variable. If ICEWS simply picks up news activity, it should record

[7] I cannot know for sure because the data are anonymized.

FIGURE 4. Change in Effect Size as Function of Core Threshold

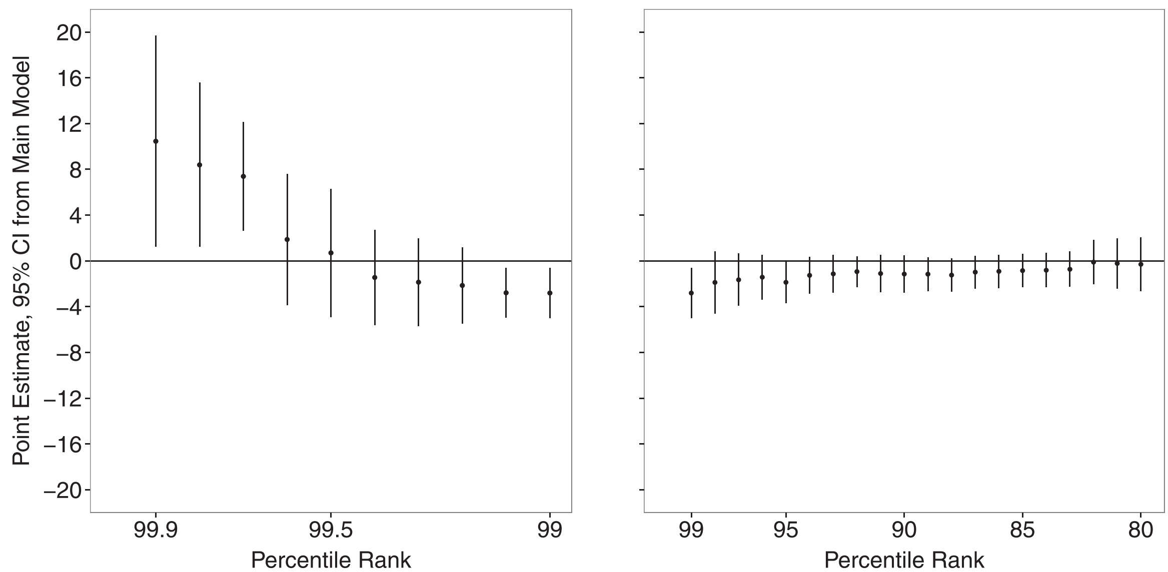

(a) Top 1% (b) 99th - 80th Percentile

a surge in public statements along with protest, and coordination will then correlate with public statements. Column 1 shows that coordination does not correlate with public statements, suggesting that $Protest_{i,t}$ actually captures protest.[8]

The second column shows that the results do not appear to be driven by duplication. The dependent variable is $Protest\ Rate_{i,t}$, the number of protests on a country-day divided by the number of ICEWS events at that time. If ICEWS duplicates, then the protest rate should not change across the sample and there will be no correlation between $Coordination_{i,t-1}$ and $Protest\ Rate_{i,t}$. $Coordination_{i,t-1}$ is still significant and $Coordination_{i,t-1}{}^{*}Core\ Hashtag\ \%_{i,t-1}$, in line with the main results.[9]

Columns 3 and 4 in Table 7 present the final verification of ICEWS' protest count. Since newspapers overreport major events and ICEWS overreports newspapers, it is possible that the results are driven by the upper end of the protest distribution. Column 3 drops all protest-days that have protests three standard deviations above the country's average, and column 4 drops all protest-days in the upper quartile of each country's protest distribution. In both cases, $Coordination_{i,t-1}$ is significant while elite coordination is not. The main finding of this article, that coordination occurs along the periphery of a network, is not an artifact of using machine-coded data.[10]

Further verification of the dependent variable is shown in Figure 5, which shows that ICEWS' count of protests strongly correlates with hand-coded data. The Armed Conflict Location and Event Dataset (ACLED) is hand-coded and contains data on the number of riots and protests in Algeria, Egypt, Libya, Morocco, and Tunisia for 2010 and 2011 (Raleigh, Linke, Hegre and Karlsen 2010). ACLED provides greater event granularity than the Social Conflict in Africa Dataset, another hand-coded events dataset that contains protests (Hendrix, Hamner, Case, Linebarger, Stull and Williams 2012). The two measures have a Pearson's correlation coefficient of 0.468.

The Supplementary Material contains further analysis. To make sure the model is not sensitive to specific countries, the main model is run while throwing out countries that may have overly influenced results. Removing the five countries with the highest levels of protests per capita, the results hold. Removing the five countries with the lowest levels of protests per capita, the results hold. Removing the five countries with the most tweets per capita, the results hold. Removing the five countries with the fewest tweets per capita, the results hold. I also include a fixed effect for Fridays, as protests commonly followed Friday prayers; there is a Friday effect, but

[8] Columns 1 and 2 use an ordinary-least squares estimator because the dependent variable is no longer a count.

[9] The Supplementary Material visualizes the protest rate and shows that it varies in tandem with real-world events.

[10] Two new machine-coded projects, the University of Illinois' Social, Political, Economic Event Database (SPEED) and the Open Event Data Alliance's Phoenix project, are exciting events data projects. SPEED combines machine-coded data with human verification to achieve human-level accuracy with machine-coded breadth (Nardulli, Althaus and Hayes 2015). The Phoenix project, an open source system associated with Pennsylvania State University and Parus Analytical Systems, is a major evolution of the TABARI system. Phoenix's main advantage over ICEWS, which also uses a heavily modified version of TABARI, is that it releases new data daily, while ICEWS releases monthly on a one year delay. As of this writing, SPEED's public data only go through 2005, and Phoenix's data starts on June 20, 2014. Phoenix's Github page is https://github.com/openeventdata.

TABLE 7. Robust to Dependent Variable

	Public Statements Rate$_{i,t}$	*Protest Rate$_{i,t}$*	*Protest$_{i,t}$*	
	All (1)	All (2)	Drop 3 SD (3)	Drop Top Quarter (4)
Coordination$_{i,t-1}$	0.047	0.100***	2.052***	1.498*
	(0.039)	(0.036)	(0.567)	(0.794)
Hashtag %$_{i,t-1}$	−0.051***	0.058	1.127*	0.408
	(0.013)	(0.026)	(0.582)	(0.579)
Retweet %$_{i,t-1}$	0.002	−0.035	−0.628	−0.755
	(0.032)	(0.035)	(0.749)	(0.587)
Link %$_{i,t-1}$	0.029**	−0.035	−0.069	0.671
	(0.013)	(0.025)	(0.385)	(0.464)
Mention %$_{i,t-1}$	0.003	0.003	0.242	0.726
	(0.021)	(0.021)	(0.600)	(0.561)
Public statements rate$_{i,t-1}$	0.099***			
	(0.013)			
Repression rate$_{i,t-1}$		0.047**		
		(0.020)		
Repression$_{i,t-1}$			0.036***	0.056***
			(0.016)	(0.020)
Protest rate$_{i,t-1}$		0.356***		
		(0.035)		
Protest$_{i,t-1}$			0.122***	0.063***
			(0.011)	(0.010)
Core hashtag %$_{i,t-1}$	−0.00002	0.029***	0.301	−0.625
	(0.031)	(0.018)	(0.271)	(0.697)
Core retweet %$_{i,t-1}$	0.011	−0.001	0.112	−0.122
	(0.007)	(0.014)	(0.295)	(0.117)
Core link %$_{i,t-1}$	0.036	−0.001	0.972**	1.252**
	(0.023)	(0.014)	(0.405)	(0.622)
Core mention %$_{i,t-1}$	−0.004	−0.006	−0.227	−0.136
	(0.025)	(0.017)	(0.282)	(0.333)
Coordination$_{i,t-1}$*core hashtag %$_{i,t-1}$	−0.021	−0.068	−1.291	1.119***
	(0.059)	(0.064)	(0.973)	(0.343)
Intercept	0.069***	0.041**	−2.062***	−3.783***
	(0.021)	(0.016)	(0.368)	(0.497)
Model	OLS	OLS	Neg. Binom.	Neg. Binom.
Country FE	yes	yes	yes	yes
N	6,620	6,620	6,471	2,916
R^2	0.091	0.221		
Adjusted R^2	0.088	0.218		
Log likelihood			−7,244.759	−1,765.218

*$p < 0.1$; **$p < 0.05$; ***$p < 0.01$

it does not change the effects of coordination or core coordination. These results suggest that the coordination patterns are widespread throughout the dataset and not dependent on a few countries or specific days.

To address concerns about machine-coded data, it uses the ACLED measures as a dependent variable and shows that coordination may still occur without the core's coordination, though ACLED's little variation on the dependent variable means most results do not attain traditional levels of statistical significance. In results not presented, ICEWS is shown to correlate with GDELT (Leetaru and Schrodt 2013), another machine-coded events dataset; their Pearson correlation coefficient is 0.785. All models presented here and in the Supplementary Material were rerun using GDELT, and all results hold. Finally, the Supplementary Material shows that the number of protests recorded on a country-day in ICEWS correlates with the estimated number of people who protested in a country on that day, as recorded by the Social Conflict in Africa Dataset (Saleyhyan and Hendrix 2012).

To make sure that domestic actors seeking to draw international attention do not drive coordination, the Twitter-based variables are generated using only Arabic tweets. The results hold.

The Supplementary Material also shows the internal validity of the dependent variable and treatment. ICEWS records different numbers of protest across countries and days in ways that accord with those countries' protest periods. For example, ICEWS

FIGURE 5. ICEWS Correlates with Handcoded Data

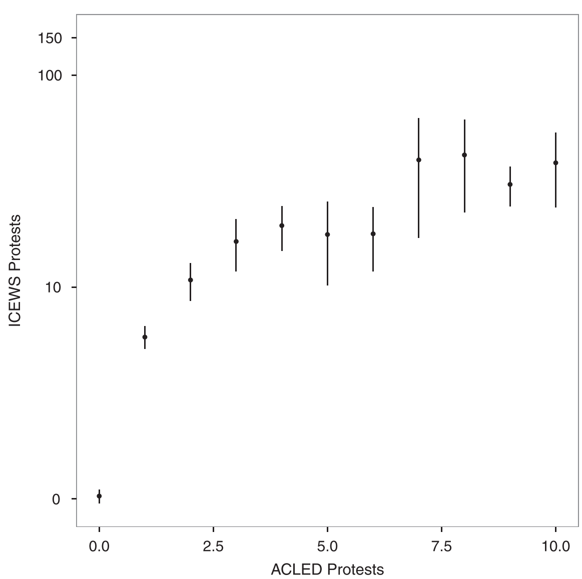

FIGURE 6. *Coordination$_{i,t}$* Peaks with One-day Lag

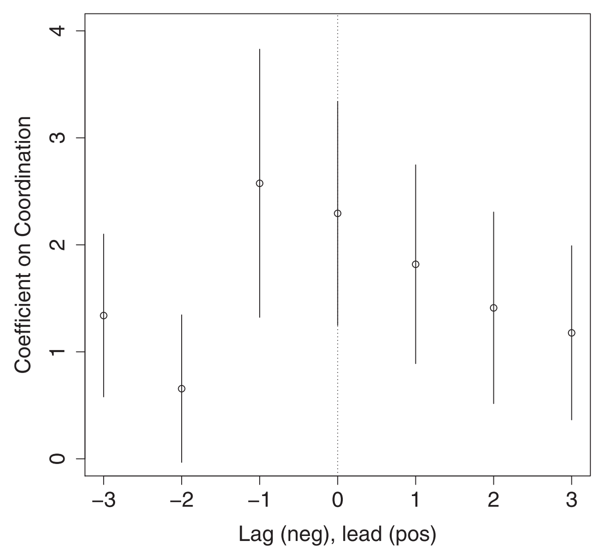

records Morocco's initial protests on February 20, 2011 as well as subsequent flares throughout the year. It reveals, however, that Morocco experienced fewer protests than Egypt and Bahrain, reflecting perhaps the regime's greater legitimacy or more adroit handling of the protests (Benchemsi 2014). ICEWS shows that Qatar experienced very little protest, with the recorded incidences probably reflecting noise in the coding process. Moreover, tracking the most common hashtags per day shows that they were protest related hashtags. In other words, *Coordination$_{i,t}$* measures protest hashtags, not popular hashtags about nonprotest activity.

Placebo Tests

Two placebo tests confirm that the results on coordination are not a result of measurement error.

First, the coordination measure is calculated for each country-day while excluding the top 5, 10, and 20 most common hashtags for each day. The resulting coordination is therefore the coordination that occurs on less common hashtags, which are more likely to be nonprotest hashtags. (Hashtags were not manually identified so that the measure could scale easily across countries and days.) If coordination on less common hashtags correlates with protest as much as coordination on all hashtags, then the operationalization has mistaken chatter focused on nonprotest hashtags with protest coordination.

The results in Table 8 confirm that the original *Coordination$_{i,t}$* measures protest coordination. The table shows the results of models where an increasing number of the most common hashtags per day—5, 10, then 20—are excluded from Coordination$_{i,t}$. The coefficient on Coordination$_{i,t}$ ranges from 39.3% to 43.6% smaller (it is 2.575 when not excluding hashtags) and decreases as an increasing number of hashtags are removed. At the same time, the coefficient on *Hashtag %$_{i,t-1}$* is now significant in each model (it was not significant in the original model), increases as more hashtags are removed, and almost triples is size compared to the full model. Not measuring the coordinating effect of the most popular hashtags pushes the correlation to *Hashtag %$_{i,t-1}$*. Note as well that the signs, coefficient size, and results on the other coefficients are very similar to their values in the original model.

Second, the main model, Model 5 from Table 4, is rerun with different lags and leads on *Coordination$_{i,t-1}$*; if *Coordination$_{i,t-1}$* does not measure coordination, there should be no change in coefficient size as the lags and leads change. In fact, as Figure 6 shows, the effect size is much larger for a one-day lag than it is for two- or three-day lags, and thereafter it decreases monotonically. The models suggest a positive correlation between a day's protest and future correlation, but this correlation is never as strong as it is for a one-day lag, and it decreases as the future moves further away. These results are consistent with the theory of spontaneous collective action, as the theory does not say protest cannot affect future coordination, only that past coordination affects future protest. As Table 3 of the Supplementary Material shows, controlling for past protests does not change the results on *Coordination$_{i,t-1}$*.

EXOGENOUS IDENTIFICATION OF THE CORE IN EGYPT

This section takes advantage of a sudden increase in the difficulty of accessing the internet in Egypt to identify core members of the Egyptian social network. Those who could access the internet during this period are more likely to be in the core than those who could not, so tracking their communicative behavior throughout

TABLE 8. Coordination in the Hashtag Long Tail

	$Protest_{i,t}$		
	Remove Top 5 (1)	Remove Top 10 (2)	Remove Top 20 (3)
$Coordination_{i,t-1}$	1.563**	1.537**	1.450**
	(0.687)	(0.686)	(0.721)
Hashtag $\%_{i,t-1}$	1.305*	1.423*	1.534**
	(0.687)	(0.727)	(0.677)
Retweet $\%_{i,t-1}$	−0.388	−0.473	−0.505
	(0.967)	(0.980)	(0.952)
Link $\%_{i,t-1}$	−0.772**	−0.771**	−0.776**
	(0.386)	(0.378)	(0.372)
Mention $\%_{i,t-1}$	−0.884**	−0.889**	−0.877**
	(0.382)	(0.389)	(0.387)
$Protest_{i,t-1}$	0.121***	0.122***	0.122***
	(0.011)	(0.011)	(0.011)
$Repression_{i,t-1}$	0.023*	0.023*	0.022*
	(0.013)	(0.013)	(0.013)
Core hashtag $\%_{i,t-1}$	0.666	0.692	0.687
	(0.512)	(0.512)	(0.510)
Core retweet $\%_{i,t-1}$	0.244	0.249	0.234
	(0.451)	(0.461)	(0.456)
Core link $\%_{i,t-1}$	0.632	0.633	0.639
	(0.549)	(0.539)	(0.535)
Core mention $\%_{i,t-1}$	−0.259	−0.286	−0.303
	(0.247)	(0.242)	(0.238)
$Coordination_{i,t-1}$*core hashtag $\%_{i,t-1}$	−1.085	−1.185	−1.152
	(0.986)	(0.946)	(0.907)
Intercept	−0.953***	−0.976***	−0.988***
	(0.214)	(0.197)	(0.183)
Country FE	yes	yes	yes
N	6,620	6,620	6,620
Log likelihood	−8,303.888	−8,305.026	−8,306.342

*$p < 0.1$; **$p < 0.05$; ***$p < 0.01$

the study should more precisely identify any role for the core. Those in the core according to this identification strategy also do not lead to protest mobilization.

Egypt's January 25 protests surprised everyone—activists, bystanders, and state authorities—with its large mobilization and brief occupation of Tahrir Square. The Mubarak regime had spent the previous days denying that the events in Tunisia would spread to their country, despite a spate of imitation immolations (Khalil 2011, 127). Many Muslim Brotherhood leaders, despite not having sanctioned the protests, were summarily jailed, as the government assumed only it could mobilize such a crowd. With the next major protest called for January 28 after Friday prayers, the government suspended cell phone service and internet access just after midnight on January 28 (the morning of the 28th). The government appears to have figured that people would not protest if they could not communicate with one another; the plan backfired, as Egyptians had no way to communicate except by going outdoors. Instead, the blackout led to more protests (Hassanpour 2014).

Contrary to conventional wisdom, digital communication was not completely severed. One internet service provider, Noor, functioned through the end of January 31; it provided connectivity for critical government functions, the Cairo stock exchange, and several international hotels (Glanz and Markoff 2011). Individuals who knew to go to hotels or who had friends with access to Noor could therefore use Twitter; even on the one day without any internet access, February 1, one could use landlines, dial-up modems, and Google's "Speak to Tweet" service to get online (Gunning and Baron 2013, 286). The blackout therefore increased the cost of accessing the internet, limiting it to those with expertise or social connections with those who still had access.

Anyone observed tweeting from Egypt between January 28 through February 1 can therefore reasonably be classified as belonging to the core of Egypt's network, regardless of their Lotan et al. coding or number of followers. This dataset observed 740 accounts that used Twitter from Egypt during the blackout, with a maximum of 338 tweeting on February 1. In terms of Twitter behavior, they are most similar to Egyptian bloggers and activists or those in the 99th percentile of the follower distribution. These users, who I call the *blackout core*, have an average of 8046 followers,

FIGURE 7. Blackout Influentials and Protest (a) Blackout Influentials and Protest (b) Blackout, Nonblackout Influentials

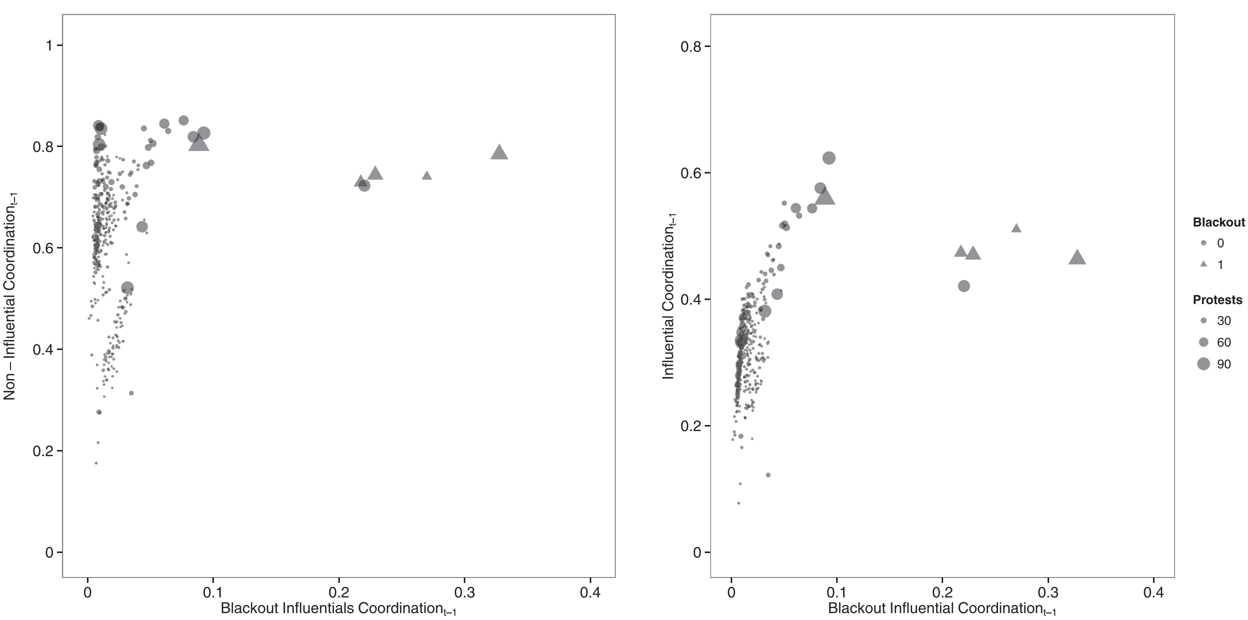

(a) (b)

and are responsible for an average of 650 tweets and 12.05% of all tweets. How they use Twitter differs, however: they retweet less often (4.74% of their tweets are retweets) than bloggers and activists (8% and 7%) but about as often as those in the 99th percentile, and they mention other users very infrequently—at 21.66%, less frequently than any other group in the sample. They use hashtags less frequently, in 21.66% of tweets, than bloggers, activists, or those in the 99th percentile. Yet 59.36% of their tweets contain a link, more than any other Egyptian group except Mainstream Media. That the blackout identification accords with the follower-based measure of influence used throughout the article provides reassurance about the validity of those measures.

Figure 7 shows how the blackout accounts' coordination correlates with $Coordination_{i,t}$ and $Core\ Coordination_{i,t}$ and protest (size of each point). $Blackout\ Core\ Coordination_t$ was calculated the same way as $Core\ Coordination_{i,t}$ except that having tweeted during the blackout, not number of followers, is the grouping variable. For both measures, shown respectively in Figures 7a and 7b, there is strong correlation in the early part of new measure and the other two coordination variables. The day with the lowest level of lagged blackout coordination is January 28, the first day of the blackout; this result makes sense since January 27 was a more representative sample of Egyptian Twitter users than later days would be. The key is to pay attention to the distribution of the size of the points: days with many protests occur across a range of $Blackout\ Coordination_t$ values.

Replicating the main models from the Main Results section confirms that peripheral coordination drives protest mobilization. Table 9 shows these results: $Blackout\ Coordination_{i,t-1}$ is not significant in any model. The only stable result from the models is $Blackout\ Mention_{i,t-1}$, which is positive and significant. Though they infrequently mention other accounts, in comparison to the other categories used to delineate the core, they are more likely to do so during protest events.

DISCUSSION

This article claims that coordination occurs through those with few social connections, and this coordination leads to protest mobilization. These peripheral network individuals outweigh those in the core because protests diffuse through a complex contagion process, a process which, in the context of protests, requires distributed coordination to spread. These results join 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 and 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 and 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 and Roberts 2013).

TABLE 9. Blackout Accounts do not Provide Coordination

	Protest$_t$				
	(1)	(2)	(3)	(4)	(5)
Coordination$_{t-1}$	0.592	4.420***	4.677***	4.283***	4.904***
	(0.447)	(0.952)	(1.037)	(1.036)	(1.176)
Hashtag %$_{t-1}$		−11.399***			−10.446***
		(2.550)			(2.944)
Retweet %$_{t-1}$		14.286*			13.877
		(7.767)			(10.516)
Link %$_{t-1}$		−2.863			2.556
		(2.606)			(3.033)
Mention %$_{t-1}$		−7.604**			−8.928***
		(3.086)			(3.363)
Protest$_{t-1}$	0.049***	0.034***	0.043***	0.046***	0.043***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.008)
Repression$_{t-1}$	0.012	0.014	0.011	0.017	0.008
	(0.018)	(0.018)	(0.018)	(0.018)	(0.017)
Blackout coordination$_{t-1}$			−63.086***	−59.331***	−44.432**
			(17.901)	(22.485)	(22.651)
Blackout hashtag %$_{t-1}$			45.164***	22.291	19.221
			(13.042)	(18.716)	(19.351)
Blackout retweet %$_{t-1}$				50.241**	28.044
				(22.693)	(26.984)
Blackout link %$_{t-1}$				−3.888	−19.470***
				(4.163)	(5.792)
Blackout mention %$_{t-1}$				18.620**	31.304***
				(7.738)	(8.516)
Constant	0.776***	4.493*	−1.942***	−1.721**	3.390
	(0.263)	(2.299)	(0.692)	(0.698)	(2.488)
Country FE	no	no	no	no	no
N	425	415	415	415	415
Log likelihood	−1,130.493	−1,080.934	−1,088.982	−1,080.382	−1,070.603

*$p < 0.1$; **$p < 0.05$; ***$p < 0.01$

Focusing on modern authoritarian countries may restrict the generalizability of this finding. It may also be that modern telecommunication technologies—phones, fax machines, and mass media, not just social media—facilitates mobilization by making it easier for peripheral members to learn of other protests or coordinate without core social network actors (Hale 2013; Warren 2014; Weyland 2012); if that is the case, core social network members may have driven mobilization in earlier centuries. In countries where the state engages in less repression, especially before the start of protests, it may be that core social network members mobilize the periphery. Core network members may also behave in ways not recorded in events or social media data. Each possibility represents an intriguing avenue for subsequent research.

This article uses Twitter data to make claims about off-Twitter behavior, but the usage of Twitter data introduces two sampling concerns. First, services exist which allow individuals to buy followers, so those accounts here identified as belonging to the core may therefore not be true core members. In fact, the accounts most likely to buy followers are those desiring to be in the core and so behaving like core members: politicians, celebrities, or small businesses (Stringhini, Wang, Egele, Kruegel, Vigna, Zheng and Zhao 2013). Moreover, these types of accounts are likely to engage in more tweet and hashtag production than others, increasing the amount of coordination coming from the core (Wu, Hofman, Mason and Watts 2011). While the buying of followers does not appear to have been a systematic practice in any of the countries during this study's period, that practice would bias in favor of finding core coordination.

Second, it is possible that individuals in the offline core select out of the Twitter core to avoid state repression, so the peripheral coordination may derive from peripheral Twitter accounts that belong to core members outside of Twitter. Though possible, work that has manually identified accounts in Tunisia and Egypt has been able to identify accounts belonging to political activists, suggesting online evasion may not be extensive (Lotan, Ananny, Gaffney, Boyd, Pearce and Graeff 2011). Identification of users from Lotan et al. (2011) in these data show that their tweets correlate negatively with subsequent protest (see Table 5, Model 1). Ethnographic work from Egypt shows that certain individuals believed their social network accounts to be monitored, but they did not respond by creating more social media accounts (Gunning and Baron 2013, 284–7).

Despite the reliance on social media data, this article does not address whether they, or telecommunications more broadly, affect protest. On one hand, social media may increase subsequent protest if it causes more individuals to learn about the state's actions and those individuals protest when they would not have without the knowledge-providing role of social media. Yet the knowledge-providing role could have counterbalancing effects: as more people learn the resolve of the state against protesters, fewer individuals may protest than otherwise would have. Appropriately answering this question requires data with very precise location information, preferably with temporal variation of social media presence. These data exist and have been used to test violence in Africa (Pierskalla and Hollenbach 2013) and Iraq (Shapiro and Siegel 2015; Shapiro and Weidmann 2011), but the results are contradictory.

The article assumes that behavior on online networks parallels that of offline interpersonal ones, allowing researchers to make inferences on heretofore hidden behaviors; research comparing behavior on Twitter to behavior off it lends credence to this assumption. Large analyses of Twitter show that user behaviors are the same as those observed offline: the distribution of followers follows a power-law distribution (Kwak, Lee, Park and Moon 2010), users connect to people who look like them (homophily) (Zamal, Liu, and Ruths 2012), friends offline follow each other on Twitter (Xie, Li, Zhu, Lim and Gong 2012), Dunbar's number applies to users on Twitter (Dunbar, Arnaboldi, Conti and Passarella 2015), and interaction decreases with geographic distance (Kulshrestha, Kooti, Nikravesh and Gummadi 2012). Substantively, survey work on activists in the U.S. Civil Rights movement and East Germany's 1989 protests shows that knowing others who had protested is the strongest correlate with a respondent's decision to participate (McAdam 1986; Opp and Gern 1993), and surveys of protesters in Tahrir Square show that individuals learned about the protests primarily through interpersonal connections or satellite television (Tufekci and Wilson 2012). Using Twitter data allows the researcher to observe these patterns across more spaces and time than before, but using the data does not reveal whether Twitter, and online social networks more generally, affects offline interpersonal networks.

No work has been able to show if social media *cause* protest, as it is very difficult to know which countries or regions of countries do or do not have a social media platform and then compare those areas to similar places without social media. Because of the difficulty of isolating social media's affect, this article has chosen not to ask that question. The point of using social media data is to better understand our world. Social media data, especially that which is publicly available, resolves the temporal resolution problem facing previous work, but connecting those data with detailed spatial data is still a challenge. Because of limits in the data for protests and the paucity of tweets from these 16 countries with GPS coordinates, for example, analysis here was restricted to the country level.

This article demonstrates the contributions big data can make to understanding processes of social influence in social networks. Researchers have begun to understand how these data can provide new insights into political phenomena such as voting (Bond, Fariss, Jones, Kramer, Marlow, Settle and Fowler 2012) or ideological sorting (Barbera 2015). These data primarily come from online social networks such as Facebook or Twitter, though anonymized call records, YouTube, or discussion boards (Nielsen 2012) are often used. While an ideal research design would randomly assign treatments in order to measure influence (Aral and Walker 2012; Muchnik, Aral and Taylor 2013), observational studies on these new datasets may allow researchers to observe and measure otherwise hidden mechanisms (Grimmer 2015; Monroe, Pan, Roberts, Sen and Sinclair 2015; Shah, Cappella and Neuman 2015). For example, studies of protest mobilization have relied on *post hoc* surveys, but social media allow one to observe how individuals behaved, not how individuals reported they behaved.

Properly used, social media data should become another tool for researchers, and it is most likely to generate knowledge when used as a window into already existing processes (Bennett and Segerberg 2013). It is not clear that social media create new behaviors or fundamentally change social relations. Its main effect is to lower the cost of communication, and lowering the cost of communication also lowers the cost of data gathering. But lower costs do not clearly favor one group of actors over another: the printing press created Martin Luther's *95 Theses* as well as Russia's *Pravda*, and states have learned how to use the internet and social media to repress (Hoffman 2015; Rod and Weidmann 2015). Using social media data to understand social behavior is therefore the main benefit of "big data." If social scientists have been stuck looking for keys under a streetlight, they now have access to stadium lights. Even stadium lights leave much of the world in the shadows.

REFERENCES

Al-Zubaidi, Layla, and Matthew Cassel, eds. 2013. *Diaries of an Unfinished Revolution: Voices from Tunis to Damascus*. New York: Penguin Books.

Aral, Sinan, and Dylan Walker. 2012. "Identifying Influential and Susceptible Members of Social Networks." *Science* 337 (6092): 337–41.

Bakshy, Eytan, Cameron, Marlow, Itamar, Rosenn, and Lada, Adamic. 2012. "The Role of Social Networks in Information Diffusion." In *Proceedings of the 21st International Conference on World Wide Web*. Lyon: 519–28.

Barbera, Pablo. 2015. "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data." *Political Analysis* 23 (1): 76–91.

Barberá, Pablo, Ning Wang, Richard Bonneau, John T. Jost, Jonathan Nagler, Joshua Tucker, and Sandra González-Bailón. 2015. "The Critical Periphery in the Growth of Social Protests." *PloS ONE* 10 (11): 1–15.

Beissinger, Mark R. 2013. "The Semblance of Democratic Revolution: Coalitions in Ukraine's Orange Revolution." *American Political Science Review* 107 (3): 574–92.

Bellin, Eva. 2012. “Reconsidering the Robustness of Authoritarianism in the Middle East: Lessons from the Arab Spring.” *Comparative Political Studies* 44 (2): 127–49.

Benchemsi, Ahmed. 2014. “Morocco’s Makhzen and the Haphazard Activists.” In *Taking to the Streets: The Transformationof Arab Activism*, eds. Lina Khatib and Ellen Lust. Baltimore: Johns Hopkins University Press, 199–236.

Bennett, W. Lance, and Alexandra Segerberg. 2013. *The Logic of Connective Action*. Cambridge, UK: Cambridge University Press.

Blaydes, L., and J. Lo. 2011. “One Man, One Vote, One Time? A Model of Democratization in the Middle East.” *Journal of Theoretical Politics* 24 (1): 110–46.

Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D.I. Kramer, Cameron Marlow, Jaime E. Settle, and James H Fowler. 2012. “A 61-Million-Person Experiment in Social Influence and Political Mobilization.” *Nature (London)* 489 (7415): 295–8.

Boschee, Elizabeth, Jennifer Lautenschlager, Sean O’Brien, Steve Shellman, James Starz, and Michael Ward. 2015. “ICEWS Coded Event Data.” http://dx.doi.org/10.7910/DVN/28075

Boyd, Dannah, Scott Golder, and Gilad Lotan. 2010. “Tweet, Tweet, Retweet: Conversational Aspects of Retweeting on Twitter.” In *43rd Hawaii International Conference on System Sciences*. IEEE, 1–10.

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, and Jean E Burgess. 2011. “The Use of Twitter Hashtags in the Formation of Ad Hoc Publics.” In *6th European Consortium for Political Research Conference*. Reykjavik: 1–10.

Bueno de Mesquita, Ethan. 2010. “Regime Change and Revolutionary Entrepreneurs.” *American Political Science Review* 104 (3): 446–66.

Burns, Alex, and Ben Eltham. 2009. “Twitter Free Iran: An Evaluation of Twitter’s Role in Public Diplomacy and Information Operations in Iran’s 2009 Election Crisis.” In *Communications Policy & Research Forum 2009*. Sydney: 298–310.

Cambanis, Thanassis. 2015. *Once Upon a Revolution: An Egyptian Story*. New York: Simon and Schuster.

Caren, Neal. 2014. “It is Time to get Rid of the E in GDELT.” http://badhessian.org/2014/05/it-is-time-to-get-rid-of-the-e-in-gdelt/

Centola, Damon. 2010. “The Spread of Behavior in an Online Social Network Experiment.” *Science* 329 (September): 1194–8.

Centola, Damon, and Michael Macy. 2007. “Complex Contagion and the Weakness of Long Ties.” *American Journal of Sociology* 113 (3): 702–34.

Citino, Robert. 2004. “Beyond Fire and Movement: Command, Control and Information in the German Blitzkrieg.” *Journal of Strategic Studies* 27 (2): 324–44.

Dalton, Russell J., Steven Greene, Paul Allen Beck, and Robert Huckfeldt. 2002. “The Social Calculus of Voting: Interpersonal, Media, and Organizational Influences on Presidential Choices.” *The American Political Science Review* 96 (1): 57–73.

Davenport, Christian, and Patrick Ball. 2002. “Views to a Kill: Exploring the Implications of Source Selection in the Case of Guatemalan State Terror, 1977-1995.” *Journal of Conflict Resolution* 46 (3): 427–50.

Downs, Anthony. 1957. *An Economic Theory of Democracy*. New York: Harper and Row.

Duggan, Maeve, and Aaron Smith. 2013. “Social Media Update 2013.” Technical Report, Pew Research Center.

Dunbar, R. I. M., Valerio Arnaboldi, Marco Conti, and Andrea Passarella. 2015. “The Structure of Online Social Networks Mirrors those in the Offline World.” *Social Networks* 43: 39–47.

Earl, Jennifer, Heather McKee Hurwitz, Analicia Mejia Mesinas, Margaret Tolan, and Ashley Arlotti. 2013. “This Protest Will Be Tweeted: Twitter and Protest Policing during the Pittsburgh G20.” *Information, Communication & Society* 16 (4): 459–78.

Eck, Kristine. 2012. “In Data we Trust? A Comparison of UCDP GED and ACLED Conflict Events Datasets.” *Cooperation and Conflict* 47 (1): 124–41.

Edmond, Chris. 2013. “Information Manipulation, Coordination, and Regime Change.” *Review of Economic Studies* 80 (4): 1422–58.

Egorov, Georgy, Sergei Guriev, and Konstantin Sonin. 2009. “Why Resource-poor Dictators Allow Freer Media: A Theory and Evidence from Panel Data.” *American Political Science Review* 103 (4): 645–68.

Ferrara, F. 2003. “Why Regimes Create Disorder: Hobbes’s Dilemma During a Rangoon Summer.” *Journal of Conflict Resolution* 47 (3): 302–25.

Finkel, Steven E., Edward N. Muller, and Karl-Dieter Opp. 1989. “Personal Influence, Collective Rationality, and Mass Political Action.” *American Political Science Review* 83 (3): 885–903.

Gerbaudo, Paolo. 2012. *Tweets and the Streets: Social Media and Contemporary Activism*. London: Pluto Press.

Gerber, Alan, Dean Karlan, and Daniel Bergan. 2006. “Does The Media Matter? A Field Experiment Measuring the Effect of Newspapers on Voting Behavior and Political Opinions.” Working Paper.

Gerber, Alan S., Donald P. Green, and Christopher W. Larimer. 2008. “Social Pressure and Voter Turnout: Evidence from a Large-Scale Field Experiment.” *American Political Science Review* 102 (1): 33–48.

Glanz, James, and John Markoff. 2011. “Egypt Leaders Found ’Off’ Switch for Internet.” http://www.nytimes.com/2011/02/16/technology/16internet.html?_r=1&hp=&pagewanted=all

Goldstone, Jack A. 1994. “Is Revolution Individually Rational?: Groups and Individuals in Revolutionary Collective Action.” *Rationality and Society* 6 (1): 139–66.

Goldstone, Jack A. 2001. “Toward a Fourth Generation of Revolutionary Theory.” *Annual Review of Political Science* 4: 139–87.

González-Bailón, Sandra, Javier Borge-Holthoefer, Alejandro Rivero, and Yamir Moreno. 2011. “The Dynamics of Protest Recruitment through an Online Network.” *Scientific Reports* 1 (197): 1–7.

Gonzalez-Bailon, Sandra, Javier Borge-Holthoefer, and Yamir Moreno. 2013. “Broadcasters and Hidden Influentials in Online Protest Diffusion.” *American Behavioral Scientist* 57 (7): 943–65.

Granovetter, Mark. 1973. “The Strength of Weak Ties.” *American Journal of Sociology* 78 (6): 1360–80.

Granovetter, Mark. 1978. “Threshold Models of Collective Behavior.” *American Journal of Sociology* 83 (6): 1420–43.

Grimmer, Justin. 2015. “We Are All Social Scientists Now: How Big Data, Machine Learning, and Causal Inference Work Together.” *PS: Political Science & Politics* 48 (1): 80–3.

Grimmer, Justin, and Brandom M. Stewart. 2013. “Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts.” *Political Analysis* 21 (3): 267–97.

Gunning, Jeroen, and Ilan Zvi Baron. 2013. *Why Occupy a Square: People, Protests and Movements in the Egyptian Revolution*. London: Hurst Publishers.

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): 331–53.

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

Hammond, Jesse, and Nils B. Weidmann. 2014. “Using Machine-coded Event Data for the Micro-level Study of Political Violence.” *Research & Politics* 1 (2): 1–8.

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

Hendrix, Cullen, Jesse Hamner, Christina Case, Chris Linebarger, Emily Stull, and Jennifer Williams. 2012. “Social Conflict in Africa: A New Database.” *International Interactions* 38 (4): 503–11.

Herkenrath, M., and A. Knoll. 2011. “Protest Events in International Press Coverage: An Empirical Critique of Cross-national Conflict Databases.” *International Journal of Comparative Sociology* 52 (3): 163–80.

Hoffman, Mike. 2015. “US Air Force Targets and Destroys ISIS HQ Building Using Social Media.”

Hollyer, James R., B. Peter Rosendorff, and James Raymond Vreeland. 2015. “Transparency, Protest, and Autocratic Instability.” *American Political Science Review* 109 (4): 764–84.

Idle, Nadia, and Alex Nunns, eds. 2011. *Tweets from Tahrir*. New York: OR Books.

Javeline, Debra. 2003. “The Role of Blame in Collective Action: Evidence from Russia.” *American Political Science Review* 97 (1): 107.

Kerr, Simeon. 2015. “Bahrain Main Shia Opposition Leader Jailed for 4 Years.”

Khalil, Ashraf. 2011. *Liberation Square: Inside the Egyptian Revolution and the Rebirth of a Nation*. New York: St. Martin’s Press.

Khatib, Lina, and Ellen Lust, eds. 2014. *Taking to the Streets: The Transformation of Arab Activism*. Baltimore: Johns Hopkins University Press.

King, Gary, Jennifer Pan, and Margaret E. Roberts. 2013. “How Censorship in China Allows Government Criticism but Silences Collective Expression.” *American Political Science Review* 107 (2): 326–43.

Klandermans, Bert, and Dirk Oegema. 1987. “Potentials, Networks, Motivations, and Barriers: Steps Towards Participation in Social Movements.” *American Sociological Review* 52 (4): 519–31.

Kulshrestha, Juhi, Farshad Kooti, Ashkan Nikravesh, and Krishna P Gummadi. 2012. “Geographic Dissection of the Twitter Network.” In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*. 202–9.

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, and Sue Moon. 2010. “What is Twitter, a Social Network or a News Media?” In *International World Wide Conference*. 591–600.

Leetaru, Kalev H., Shaowen Wang, Guofeng Cao, Anand Padmanabhan, and Eric Shook. 2013. “Mapping the Global Twitter Heartbeat: The Geography of Twitter.” *First Monday* 18 (5-6): 1–33.

Leetaru, Kalev, and Philip Schrodt. 2013. “GDELT: Global Data on Events, Language, and Tone, 1979–2012.” International Studies Association.

Lehmann, Janette, Bruno Gonçalves, José J. Ramasco, and Ciro Cattuto. 2012. “Dynamical Classes of Collective Attention in Twitter.” *Proceedings of the 21st International Conference on World Wide Web - WWW ’12*, 251–60.

Lichbach, Mark I. 1998. “Contending Theories of Contentious Politics and the Structure-Action Problem of Social Order.” *Annual Review of Political Science* 1: 401–24.

Liere, Diederik Van. 2010. “How Far Does a Tweet Travel? Information Brokers in the Twitterverse.” In *Proceedings of the International Workshop on Modeling Social Media*. Toronto: 1–4.

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, and Erhardt Graeff. 2011. “The Revolutions Were Tweeted: Information Flows During the 2011 Tunisian and Egyptian Revolutions Web.” *International Journal of Communications* 5: 1375–406.

Mackey, Robert. 2015. “Egyptian Police Attack Marchers Carrying Flowers to Tahrir Square, Killing Protester.” http://www.nytimes.com/2015/01/25/world/middleeast/egypt-woman-is-killed-in-peaceful-cairo-protest.html

Malik, Momin M., Hemank Lamba, Constantine Nakos, and Jurgen Pfeffer. 2015. “Population Bias in Geotagged Tweets.” In *9th International AAAI Conference on Weblogs and Social Media*. 18–27.

McAdam, Doug. 1986. “Recruitment to High-Risk Activism: The Case of Freedom Summer.” *American Journal of Sociology* 92 (1): 64–90.

McPherson, Miller, Lynn Smith-Lovin, and James M. Cook. 2001. “Birds of a Feather: Homophily in Social Networks.” *Annual Review of Sociology* 27 (August): 415–44.

Meirowitz, Adam, and Joshua A. Tucker. 2013. “People Power or a One-Shot Deal? A Dynamic Model of Protest.” *American Journal of Political Science* 57 (2): 478–90.

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

Monroe, Burt L., Jennifer Pan, Margaret E. Roberts, Maya Sen, and Betsy Sinclair. 2015. “No! Formal Theory, Causal Inference, and Big Data Are Not Contradictory Trends in Political Science.” *PS: Political Science & Politics* 48 (01): 71–4.

Moore, Will H. 1995. “Rational Rebels: Overcoming the Free-Rider Problem.” *Political Research Quarterly* 48 (2): 417–54.

Muchnik, L., S. Aral, and S. J. Taylor. 2013. “Social Influence Bias: A Randomized Experiment.” *Science* 341 (6146): 647–51.

Nardulli, Peter F., Scott L. Althaus, and Matthew Hayes. 2015. “A Progressive Supervised-learning Approach to Generating Rich Civil Strife Data.” *Social Methodology* 45 (1):148–83.

Newman, M. E. J. 2003. “The Structure and Function of Complex Networks.” *SIAM Review* 45 (2): 167–256.

Nielsen, Richard. 2012. “Jihadi Radicalization of Muslim Clerics.” Working Paper

Opp, Karl-Dieter. 2012. “Collective Identity, Rationality and Collective Political Action.” *Rationality and Society* 24 (1): 73–105.

Opp, Karl-Dieter, and Christiane Gern. 1993. “Dissident Groups, Personal Networks, and Spontaneous Cooperation: The East German Revolution of 1989.” *American Sociological Review* 58 (5): 659–80.

Pei, Sen, Lev Muchnik, José S. Andrade Jr., Zhiming Zheng, and Hernán A. Makse. 2014. “Searching for Superspreaders of Information in Real-world Social Media.” *Scientific Reports* 4 (c): 5547.

Pierskalla, Jan H., and Florian M. Hollenbach. 2013. “Technology and Collective Action: The Effect of Cell Phone Coverage on Political Violence in Africa.” *American Political Science Review* 107 (2): 207–24.

Press, Daryl G. 2001. “The Myth of Air Power in the Persian Gulf War and the Future of Warfare.” *International Security* 26 (2): 5–44.

Quattrone, George A., and Amos Tversky. 1988. “Contrasting Rational and Psychological Analyses of Political Choice.” *The American Political Science Review* 82 (3): 719.

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): 158–78.

Raleigh, Clionadh, Andrew Linke, Havard Hegre, and Joakim Karlsen. 2010. “Introducing ACLED: An Armed Conflict Location and Event Dataset: Special Data Feature.” *Journal of Peace Research* 47 (5): 651–60.

Riker, William H., and Peter C. Ordeshook. 1968. “A Theory of the Calculus of Voting.” *The American Political Science Review* 62 (1): 25–42.

Rod, Espen Geelmuyden, and Nils B. Weidmann. 2015. “Empowering Activists or Autocrats? The Internet in Authoritarian Regimes.” *Journal of Peace Research* 52 (3): 338–51.

Romero, Daniel M., Brendan Meeder, and Jon Kleinberg. 2011. “Differences in the Mechanics of Information Diffusion Across Topics: Idioms, Political Hashtags, and Complex Contagion on Twitter.” In *Proceedings of the 20th International Conference on the World Wide Web*. 695–704.

Salehyan, Idean Cullen S., Hendrix Jesse, Hamner, Christina, Case, Christopher, Linebarger, Emily, Stull, and Jennifer, Williams. 2012. “Social conflict in Africa: A new database.” *International Interactions* 38 (4): 503–11.

Schelling, Thomas C. 1978. *Micromotives and Macrobehavior*. New York: W.W. Norton & Company.

Schussman, Alan, and Sarah Anne Soule. 2005. “Process and Protest: Accounting for Individual Protest Participation.” *Social Forces* 84 (2): 1083–108.

Shachar, Ron, and Barry Nalebuff. 1999. “Follow the Leader: Theory and Evidence on Political Participation.” *The American Economic Review* 89 (3): 525–47.

Shah, Dhavan V., Joseph N. Cappella, and W. Russell Neuman. 2015. “Big Data, Digital Media, and Computational Social Science: Possibilities and Perils.” *The ANNALS of the American Academy of Political and Social Science* 659 (1): 6–13.

Shapiro, Jacob N., and David A Siegel. 2015. “Coordination and Security: How Mobile Communications affect Insurgency.” *Journal of Peace Research* 52 (3): 312–22.

Shapiro, Jacob N., and Nils B. Weidmann. 2011. “Talking About Killing: Cell Phones, Collective Action, and Insurgent Violence in Iraq.” Working Paper

Siegel, David A. 2009. “Social Networks and Collective Action.” *American Journal of Political Science* 53 (1): 122–38.

Siegel, David A. 2011. "When Does Repression Work? Collective Action in Social Networks." *The Journal of Politics* 73 (04): 993–1010.

Sinha, Shiladitya, Chris Dyer, Kevin Gimpel, and Noah A Smith. 2013. "Predicting the NFL Using Twitter." Working Paper.

Starbird, Kate and Leysia Palen. 2012. "(How) Will the Revolution be Retweeted? Information Diffusion and the 2011 Egyptian Uprising." In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work*. Seattle:.

Steinert-Threlkeld, Zachary C. 2017. *Twitter as Data*. Under Review.

Steinert-Threlkeld, Zachary C., Delia Mocanu, Alessandro Vespignani, and James Fowler. 2015. "Online Social Networks and Offline Protest." *EPJ Data Science* 4 (1): 19.

Stringhini, Gianluca, Gang Wang, Manuel Egele, Christopher Kruegel, Giovanni Vigna, Haitao Zheng, and Ben Y. Zhao. 2013. "Follow the Green: Growth and Dynamics in Twitter Follower Markets." In *Proceedings of the 2013 Conference on Internet Measurement Conference*. 163–76.

Suh, Bongwon, Lichan Hong, Peter Pirolli, and 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*. 177–184.

Taylor, Michael. 1988. "Rationality and Revolutionary Collective Action." In *Rationality and Revolution*, ed. Michael Taylor. Cambridge, UK: Cambridge University Press, 63–97.

Tonkin, Emma, Heather D. Pfeiffer, and Greg Tourte. 2011. "Twitter, Information Sharing and the London Riots?" *Bulletin of the American Society for Information Science and Technology* 38 (2): 49–57.

Tufekci, Zeynep. 2014. "Social Movements and Governments in the Digital Age: Evaluating a Complex Landscape." *Journal of International Affairs* 68 (1): 1–18.

Tufekci, Zeynep, and Christopher Wilson. 2012. "Social Media and the Decision to Participate in Political Protest: Observations From Tahrir Square." *Journal of Communication* 62 (2): 363–79.

Warren, T. Camber. 2014. "Not by the Sword Alone: Soft Power, Mass Media, and the Production of State Sovereignty." *International Organization* 68 (01): 111–41.

Watts, Duncan J. 2002. "A Simple Model of Global Cascades on Random Networks." *Proceedings of the National Academy of Sciences* 99 (9): 5766–771.

Watts, Duncan J. 2004. "The New Science of Networks." *Annual Review of Sociology* 30 (1): 243–70.

Weyland, Kurt. 2012. "The Arab Spring: Why the Surprising Similarities with the Revolutionary Wave of 1848?" *Perspectives on Politics* 10 (04): 917–34.

Wu, Shaomei, Jake M. Hofman, Winter A. Mason, and Duncan J. Watts. 2011. "Who Says What to Whom on Twitter." *Proceedings of the 20th International Conference on World Wide Web - WWW '11*, 705–14.

Xie, Wei, Cheng Li, Feida Zhu, Ee-Peng Lim, and Xueqing Gong. 2012. "When a Friend in Twitter is a Friend in Life." In *Proceedings of the 4th Annual ACM Web Science Conference*. New York: ACM Press, 344–347.

Zamal, Faiyaz Al, Wendy Liu, and Derek Ruths. 2012. "Homophily and Latent Attribute Inference: Inferring Latent Attributes of Twitter Users from Neighbors." In *Proceedings of the Sixth International AAAI Conference on Weblogs and Social Media*. 387–90.