

Tents, Tweets, and Events: The Interplay Between Ongoing Protests and Social Media

Marco T. Bastos, Dan Mercea, Arthur Charpentier

Accepted for publication in the Journal of Communication

(pre-publication version: some changes still possible)

Abstract

Recent protests have fuelled deliberations about the extent to which social media ignites popular uprisings. In this paper we use time-series data of Twitter, Facebook, and onsite protests to assess the Granger-causality between social media streams and onsite developments at the Indignados, Occupy, and Brazilian Vinegar protests. After applying a Gaussianization procedure to the data, we found that contentious communication on Twitter and Facebook forecasted onsite protest during the Indignados and Occupy protests, with bidirectional Granger-causality between online and onsite protest in the Occupy series. Conversely, the Vinegar demonstrations presented Granger-causality between Facebook and Twitter communication, and separately between protestors and injuries/arrests onsite. We conclude that the effective forecasting of protest activity likely varies across different instances of political unrest.

Keywords: Social Media, Contentious Politics, Granger causality test, Occupy, Indignados, Vinegar Protests

In this article we scrutinize the often debated role of network communication in the 2011 Spanish May 15 Indignados, the 2011 Occupy movement, and the 2013 Vinegar protests in Brazil. Following results from recent investigations (Bakker & de Vreese, 2011; Borge-Holthoefer et al., 2011; Valenzuela, Arriagada, & Scherman, 2012), which positively correlate Internet use with different forms of political participation, we tracked Facebook and Twitter streams related to the aforementioned protests. Although these movements have sought to congregate and take action in physical spaces (Castells, 2012), a linear relationship from digital communication to onsite activity has been reported by scholarship discussing their spread (Tremayne, 2014; Vasi & Suh, 2013) and ranks (Anduiza, Cristancho, & Sabucedo, 2013, p. 10). In our turn, we seek to move this scholarship forward by probing for evidence of any elapsing connection between network communication and on onsite events occurring in the course of collective action.

The first of the three case studies was the Spanish demonstrations that started in Madrid on May 15, 2011 as a protest against welfare cuts, the political establishment and the runaway financial system (Castells, 2012). At the end of the initial demonstration, protestors blocked a major avenue in Madrid and subsequently clashed with the police. After that violent incident, a group of one hundred protestors headed to Puerta del Sol, the city's main square, where overnight camping was organized. During the next few days, the protests and night-time camp-outs spread to more than 30 cities across Spain. The Occupy protests, our second case study, started on September 17, 2011 when Adbusters launched the proposal for a peaceful demonstration to "occupy" the global financial center at Wall Street (Moynihan, 2011). An estimated one thousand people attended the first day of the protest, reportedly inspired by the Spanish uprisings and the events of the "Arab Spring." On September 23, demonstrators began camping in Zuccotti Park. The following day, demonstrations intensified when protestors marched uptown and instances of police brutality were broadcast

on YouTube and television news programs. Subsequently, the Occupy movement spread to cities across the United States. On October 15, one month before New York protestors were forced out of Zuccotti Park, similar demonstrations had happened in 951 cities in 82 countries.

The third and last case study is that of the Vinegar protests in Brazil. The social unrest was initially sparked by opposition to bus and underground fare rises in June 2013. However, the target of contention rapidly shifted onto the running costs of infrastructure projects associated with international sport events, such as the Confederations Cup, the World Cup, and the Summer Olympics (Singer, 2014). Protestors raised conflicting demands encapsulated in concomitant calls for improvements in public services, lower taxation, and expanded welfare benefits. The first large protest was held at the beginning of June, and on June 17, an estimated 250,000 protestors took to the streets of major cities across the country. Protest marches turned violent and urban riots ensued in a number of Brazilian cities. The demonstrations were subsequently dubbed Vinegar in reference to the sixty protestors arrested for carrying vinegar allegedly used as an antidote to the tear gas and pepper spray deployed by the police.

In what follows, we set the theoretical groundwork for our investigation by mapping the immediate field of research. To that end, we reflect on key empirical findings and attendant claims that motivated this study. The subsequent two sections are dedicated to describing the procedures for data collection and aggregation as well as the method employed for the time series analysis. In the final two sections we present the results and discuss the implications of our findings to theories of political protest.

Previous Work

Despite the growing amount of empirical research on network communication linked to physical protests (Borge-Holthoefer, et al., 2011; Gaffney, 2010; González-Bailón, Borge-Holthoefer, Rivero, & Moreno, 2011; Kavanaugh, Yang, Li, & Ed Fox, 2011), and a large body of literature proposing competing accounts about contentious communication (Bennett & Segerberg, 2013; Morozov, 2011), the impact of networked communication on onsite protesting activity remains to be conclusively evidenced. This debate has ranged from the argument that innovative modes of mobilization, organization, and collective action are generated with social media (Howard & Hussain, 2013) to the counterclaim that digital activism has no bite as aggrieved populations confine their outrage to the computer screen (Morozov, 2011).

Yet, the tactical deployment of digital communications by activists has been portrayed as an up-scaling of interest and participation in contentious politics (Earl & Kimport, 2011, pp. 10-14; Lim, 2013; Valenzuela, 2013). What such research on contentious politics has so far done less rigorously (Earl, McKee Hurwitz, Mejia Mesinas, Tolan, & Arlotti, 2013) has been to employ tried instruments such as statistical tests and time series analyses to verify the direction and effect between online communication and onsite protest. Indeed, while there is growing convergence in the knowledge that networked communication on social media is germane to protest participation (Howard & Hussain, 2013, p. 65; Valenzuela, 2013, p. 935), interactions between social media usage and onsite activities have mainly been the object of sample-based quantitative content analyses (Earl, et al., 2013) or ethnographic observations (Gerbaudo, 2012). Moreover, the prevalent interest in those accounts has rested with ascertaining types of communication behavior conducive to political protest (e.g. using social media for news consumption) or embodied in political protests — e.g. tweeting hashtagged action updates or expressing emotional support (Bastos, Raimundo, & Travitzki, 2013).

There are notable recent attempts to measure the influence of online activism on the spread of offline protests (Bastos, Recuero, & Zago, 2014; Jungherr & Jürgens, 2013; Vasi & Suh, 2013). Vasi and Suh (2013) applied event history analysis to Internet search and social media data related to the Occupy Wall Street movement. The authors reported that Internet searches had a direct influence on the emergence of online activism and an indirect influence on the spread of onsite protests. Facebook and Twitter information streams were also found to positively affect the spread of offline protests over time, while cities that experienced online activism were also likely to experience actual occupations. This article complements and extends this treatment by proposing to untangle the relationship between certain forms of onsite activity (e.g. camping out) and networked communication pertaining to it. Specifically, the question we seek to address with this paper is whether the protests of the Indignados, Occupy, and Vinegar movements were followed by commensurate Facebook and Twitter activity; whether they evolved coextensively by exhibiting bidirectional determination (feedback) between onsite and online protest activity, or finally, whether the networked communication on Twitter and Facebook had any bearing on developments at the street protests.

Protest Diffusion: A Tale of Linking Paths

A key reference point for this undertaking has been the literature on protest diffusion. The outbreak of protest has been documented extensively as a contagious occurrence over space and time (Montagna, 2010; Tarrow, 2005, 2011). In Lichbach's (1985) seminal time-series study of political protests, the hypothesis of a random occurrence of protest was rejected as in postwar United Kingdom protest erupted following a concerted build-up over time. Protests occurred in a process of contagion and diffusion of ideas and behaviors that amplified and sustained collective action (Givan, Soule, & Roberts, 2010, pp. 4-7).

Although the literature on the diffusion of contention is sizeable, it has by-and-large been restricted to a focus on traditional channels of communication between prior and potential sites of contention (Vasi & Suh, 2013). To tackle our aim of disentangling the intricate relationship between onsite activity and networked communication, we consequently turned to studies of message diffusion about real-world phenomena on social media (Kallus, 2014; Ko, Kwon, Kim, Lee, & Choi, 2014; Russell Neuman, Guggenheim, Mo Jang, & Bae, 2014). Groshek (2011) performed Granger causality tests on cross-national time-series data to investigate the effects of broadcast media on democratic growth and reported that mass media diffusion Granger-caused democracy only in countries where sociopolitical instability levels were higher and mass media were more prevalent. Yamaguchi et al. (2013) investigated the impact of Twitter messages and internet forums on a signature-collecting campaign supporting traditional Japanese medicine. The authors reported that 78% of the signatures were affected by online activity and that the Twitter effect was smaller than the internet forum (26% and 52%, respectively), although Twitter probably triggered the initial bursts of signatures (Yamaguchi, et al., 2013).

Following a similar interest, Jungherr and Jürgens (2013) examined variations in online data to detect traces of offline phenomena. The authors measured recurrent dynamics of online data and argued that information deliberately published by users on social networking and microblogging services enables the documentation of user activity online as well as in their physical surroundings, as long as the latter are referenced or traceable to a physical location with metadata (Jungherr & Jürgens, 2013, p. 596). In particular, these may be purposeful reports of embodied action representing, *inter alia*, the real-time coverage of law enforcement activity at the site of a protest (Earl, et al., 2013; Gerbaudo, 2012), a well-documented activity undertaken for its potential to fan the flames of contention (Tarrow, 1998), or communication relating to protest tactics (Theocharis, 2013).

Research Hypotheses

Cognate research has shown that both Twitter and Facebook activity were strongly associated with police-protestor interactions (Caren & Gaby, 2011), such as the arrest of more than a hundred protestors at the Occupy Boston encampment on the evening of October 10, 2011 and the detention of twenty-three people at the Occupy Denver encampment on October 14. In the run up to this analysis we noted that Twitter hashtag activity linked to the Occupy movement peaked on October 1, when over 700 individuals were arrested on the Brooklyn Bridge; and on October 15, when hundreds of simultaneous protests were held around the globe. This is in line with Caren and Gaby (2011) and Earl et al. (2013) assessment that both Twitter and Facebook streams were strongly connected to onsite events. However, those authors did not produce a systematic analysis of this relationship; nor did they evaluate multiple instances of political unrest.

In order to test the relevance of the above insights for protest diffusion, the present analysis foregrounds the question of the directionality of a causal link between digital communication on social media and physical protest. Therefore, our primary objective was to scrutinize the interactions between online and onsite activity within the timeframe of the three instances of political upheaval. To this end, we tested the hypothesis that the outbreak of online protest activity at one point in time can be used for prediction of future outbreak of onsite protest activity (H1a), so that the temporal diffusion of protest-related networked communication may contribute to the onset of onsite protests. Conversely, networked communication may only bear on onsite activity indirectly, as a non-relational mechanism that at best ripples through communication networks generating a self-referential digital echo inconsequential to physical participation (H1b).

If the latter postulate stands in contrast to earlier studies (Fisher & Boekkooi, 2010; Tufekci & Wilson, 2012), this may be attributable to differences in i) the nature of empirical

data and ii) the levels of analysis. To briefly unpack these points, the two referenced articles relied on self-reported user behavior captured either with quantitative or qualitative interviews whereas the present analysis draws on aggregated “big data,” which are large datasets of digital data employed to identify behavioral patterns. Secondly, and as an effect of the divergent data sources, this investigation departs from those previous treatments in that it does not generalize collective behavior from individual conduct. Instead, we cross-check media accounts of collective behavior in onsite protest with large datasets of contentious networked communication (Colbaugh & Glass, 2012; Lerman, Galstyan, Ver Steeg, & Hogg, 2011). In this, our study builds on the methodological tradition of protest event analysis (Koopmans & Rucht, 2002) and our methods, as discussed more extensively below, draw on press reports of protests to map, analyze, and interpret their occurrence over time (Koopmans & Rucht, 2002, p. 231).

Our second objective has been informed by Manuel Castells’ (2012) chronicle of the Indignados demonstrations. Castells remarked that Twitter was instrumental to the establishment of encampments in key locations such as at Puerta del Sol in Madrid or in Catalunya Square in Barcelona. In his turn, Castells (2012) highlighted the fundamental part that onsite interaction between law enforcement and protestors played in the fate of that contentious action as well as in the networked communication around it. Taking this cue, we examined social media streams not only against accounts of the number of protestors attending demonstrations, but also against the number of protestors setting up or taking down protest encampments, and the number of protestors injured or arrested by the police during the three instances of political unrest. We designated those involved in such strenuous action as high-functioning political activists (Bobel, 2007). We hypothesized a Granger-causality from the intensity of social media communication to high-functioning political activism

measured by police and encampment activity (H2a), and from high-functioning political activism to the intensity of social media communication (H2b).

Lastly, Earl et al. (2013) contended that Facebook activity tends to be high in the run-up to a physical protest (when used to rally participants) and in the aftermath (when used as an arena for post-hoc reflection on the protest). By contrast, Twitter usage was argued to coincide with onsite action and on-the-ground coordination of protests (Earl, et al., 2013, p. 3). The third objective was therefore to elucidate the relationship between the social networking sites Twitter and Facebook during the three instances of political unrest by employing the Granger causality test to forecast the intensity and directionality between tweets and Facebook posts. Specifically, we hypothesized that in the course of the protests Facebook pages impacted the activity of Twitter streams (H3a). However, the opposite relationship might also hold true due to the use of Twitter for streaming the events as they unfold (H3b). Finally, we tested the Granger-causality hypotheses in both directions for the pairs of variables and reported feedback when the results confirmed bidirectional Granger-causality between the variables.

Research Data

We tracked around one-hundred Twitter hashtags associated with the Indignados, Occupy, and Vinegar protests (roughly 35 hashtags per event) and another one-hundred Facebook pages and groups dedicated to the events (see the Appendix for the list of Twitter hashtags and Facebook pages). While the Indignados dataset encompasses political demonstrations that took place in Spain in May 2011, the Occupy dataset includes a number of locations in the United States and major cities across the world (i.e. Amsterdam, Berlin, Dublin, Frankfurt, London, Paris, Tokyo, and Toronto). Given the purpose of this investigation, we focused on information streams associated with cities rather than conceptual tags such as

#nolesvotes (*don't vote for them*) and #notenemosmiedo (*we are not afraid*) in the Indignados dataset; #occupytheworld and #occupytogether in the Occupy dataset; and #vemprarua (*take to the streets*) and #todarevolucaocomeca (*every revolution begins with*) in the Vinegar dataset. The rationale for this procedure was rooted in the anticipation that activity on city-related hashtags and Facebook groups, unlike conceptual tags, would more closely relate to onsite demonstrations (Thorson et al., 2013, p. 3).

The reported Twitter data was collected from the publicly available Twitter stream. It was retrieved with an authenticated user account running yourTwapperKeeper (O'Brien III, 2010) that connected to the Twitter Streaming API. Private user information was excluded from the analysis which was run solely based on the unique identification of each tweet and Facebook post rather than usernames. We expect the selection of Twitter hashtags to have rendered a representative, if biased (Morstatter, Pfeffer, Liu, & Carley, 2013), sample of the full stream because the data requested for this analysis is well below the 1% threshold of the entire public stream allowed by the Twitter Streaming API (Driscoll & Walker, 2014). The Facebook Graph API imposes no such limits, so the data was collected via a series of requests on the API. Therefore, and unlike the Twitter data gathered for this study, Facebook posts and the number of protestors attending demonstration were collected at the end of the research period.

Facebook and Twitter data for the Vinegar protests were gathered from their very onset. Data collection for the Indignados dataset started on May 17-19, and although the movement generally peaked only on June 19, the first mass mobilizations in Madrid had already begun on May 15. Additionally, while half the Occupy dataset includes the entire period of the demonstrations, the other half presents an average delay of 12 days.¹ Even though the Occupy protests peaked as early as October 1, the movement as a whole climaxed on October 15 when simultaneous events were held in many countries. We believe these

small delays in the archiving processes do not affect the results reported in this study, mostly because the movements grew in popularity and scope after the first demonstrations (Bastos, et al., 2014; Caren & Gaby, 2011; González-Bailón, et al., 2011). Furthermore, our methodology focuses on the distribution of protest activity at different points in time, so shortfalls at the beginning or the end of the period should be of minor effect to the time series analyses (Bressler & Seth, 2011).

Data relating to the number of protestors attending demonstrations or camp-outs at protest locations, as well as data on the number of protestors injured or arrested by the police, were collected from press reports about the episodes (see the Appendix for a table with the number of reports collected from each outlet). This approach was grounded in the protest event analysis tradition, which draws on newspaper articles on a contentious gathering (e.g. demonstrations, marches, or strikes) as the main unit of analysis. Without making claims to the representativity of the data for a universe of protest events, the method enables the collection and cross-validation of media accounts to determine the number of participants at protest events (Koopmans & Rucht, 2002, p. 238). Other information of interest to the research may be gathered at the same time (Koopmans & Rucht, 2002, p. 240) and was duly recorded here, i.e. the number of protestors and individuals camped-out, injured, or arrested while attending the demonstrations.

Reports from media outlets accounting for the number of protestors attending or camping-out at protest locations were often found to be conflicting. When the figures reported in multiple press accounts differed substantially, we calculated the mean of all numbers provided by the press. Whenever media reports failed to produce unambiguous information regarding the number of arrested protests attending Occupy demonstrations, we reverted to relevant information issued by the Occupy Arrests organization (Ernesto, 2011). The resulting dataset comprise the number of tweets, Facebook posts, and protestors engaged

in political activities associated with demonstrations in the cities where protests took place. Table 1 shows the summary statistics of Indignados, Occupy, and Vinegar and indicates the minimum, maximum, median, mean, and standard deviation of the dataset.

INSERT TABLE 1 HERE

There are some asymmetries relating to the data available for each instance of political unrest. Firstly, the Occupy series does not include the number of injuries. We also have not found consistent and reliable sources of information regarding the number of camped-out and injured protestors in the cities that experienced Occupy movements. Secondly, the Indignados is the only series that includes data about camped-out protestors. Thirdly, Vinegar is the only series that contains data about protestors injured during demonstrations. These differences did not preclude our analyses as all series included a minimum of two variables of online protest activity (tweets and Facebook posts) and two variables of onsite protest activity, the first being the number of protestors onsite and the second being the number of protestors arrested during demonstrations. This latter variable relates to high-functioning political activism, and whenever available we also tested a fifth variable associated with high-functioning political activism (protestors camping-out or injured during demonstrations).

At first call, we found that onsite activity was sparse, intermittent, and seldom spanned the entire period of the analysis. We addressed this shortcoming by aggregating the variables for protest activity online and onsite by city (i.e. Madrid, New York, or São Paulo) and subsequently creating single datasets for the Indignados, Occupy, and Vinegar protests. This approach allowed for the creation of a complete series associated with each instance of political arrest because the observations were first aggregated by geographic location and subsequently by the overarching political movement. In order to control for the highly skewed pattern of protest activity online and onsite, and the different location where events

took place, we only aggregated data for cities that experienced protesting activity onsite and online (on Twitter and Facebook), so venues with high activity online but no activity onsite (or vice-versa) were not included in the aggregated data. This empirical observation is corroborated by the distinction between hyper-local and the network communication embedded in ground-level activity within the Occupy movement (Bennett, Segerberg, & Walker, 2014). We took this point as a cautionary note regarding the gamut of local operations which a design such as ours can capture, as much ground-level activity does not register in the network communication.

Our processed data thus include the daily number of tweets, Facebook posts, protestors, and high-functioning political activists attending street demonstrations (see the Appendix for the breakdown of the variables considered). Figure 1 shows a histogram of Twitter, Facebook, and protesting activity in logarithmic scale for the three instances of political unrest. The bar charts on the left show the daily aggregated data, and the plot on the right summarizes the hourly aggregated data for each event.

INSERT FIGURE 1 HERE

Research Methods

In order to measure predictive causal connectivity across the instances of political unrest, we modeled the data as stochastic time series and performed a Granger causality test. Originally developed for economic time series data by Wiener (1956) and Granger (1969), and since then applied to time series data of many different domains, the Granger causality test offers a data-driven, theoretically sound, and easy to apply statistical time series approach to causal inference based on prediction (Bressler & Seth, 2011; Schelter, Winterhalder, & Timmer, 2006). The null hypothesis of no Granger-causality is rejected only if no lagged values of the explanatory variable have been retained in the regression. Yet, Granger-causation is only

equivalent to causation under the assumption that there are no other potential causes.² This is the first attempt to apply the Granger causality test to this type of series, and in this study we tested the two-way, paired relationship between six numeric variables, namely: tweets, posts, protestors, camped-out, arrested, and injured protestors.

The assumption underlying the Granger causality test is that the explanatory variable Granger-causes the outcome variable whenever there is a non-expected output that leads to an increase in the outcome variable. This framework states that a process X is considered a cause of another process Y if the knowledge about the past of X significantly improves the prediction of the future of Y , as opposed to the prediction based only on the knowledge about the past of Y . Notably, our variables of interest relate to complex social events that might impose violations to the abovementioned assumption of non-competing causes. However, we are convinced of the appositeness of the Granger causality test for our data. This is because the method enabled testing the extent to which the past of protest-related online activity contains information that helps predict the future of onsite protest activity (and the other way around) more accurately than using only the past of one of the variables. Translated to the events studied in this paper, we propose political messages on social media Granger-cause onsite activity if social media spikes are followed by a corresponding increase in the volume of protestors attending onsite demonstrations, so that online protest activity Granger-causes onsite protest activity.

In order to perform the pairwise Granger causality test for the time series, we relied on the R platform for statistical computing (R Development Core Team, 2014) and performed a VAR lag order selection criteria to choose the best lag length for the VAR time series model (Schwarz, 1978; Tschernig & Yang, 2000). We decided to adopt VARs of order 1³ to avoid misinterpreting the results. When dealing with causality for higher order VARs, it is difficult to analyze the causality implied by the lag. For instance, a causality with a VAR(3) means

that there could be a causality effect tomorrow, the day after tomorrow, or even the day after that without allowing for conclusions on the lag. Moreover, for a vast majority of pairs, order 1 is the optimal order for an AIC criteria procedure *VARselect* (Lütkepohl, 2007).⁴ Therefore, the results reported in this paper are based on a VAR(1) model and the non-causality reported in the analyses should be understood as non-causality from one day to the next day. Twitter and Facebook data are highly skewed with heavy tails considerably affecting the estimation of correlation (i.e. autocorrelation). In such cases extreme variations might appear correlated due to tail, but not to regular dependence. In order to reject causal effects due to the skewed distribution of the series, and seeing that the Granger causality test assumes a bivariate Gaussian distribution, we relied on a semi-parametric transformation to correct from the non-normality of the individual time series (Sanggyun & Brown, 2010).

A pure nonparametric transformation was considered in Hiemstra & Jones (1994) with a formal testing procedure. In fact, it is possible to consider a semi-parametric transformation to obtain individual Gaussian time series (Eichler, 2010). This approach was discussed by Liu et al. (2009) and Liu et al. (2012) as a method to assess causality between two time series i and j . The method is based on the following procedure: for time series i , find an empirical marginal distribution function based on the ranks \hat{F}_i and \hat{F}_j , $\hat{F}_i(x) = \frac{1}{T+1} \sum_{t=1}^T \mathbf{1}(X_{i,t} \leq x)$, and similarly for j , and subsequently map the observation into the $[0,1]$ copula space (Taamouti, Bouezmarni, & El Ghouch, 2014), $\hat{U}_{i,t} = \hat{F}_i(X_{i,t})$ and $\hat{U}_{j,t} = \hat{F}_j(X_{j,t})$. Finally, we define $\tilde{X}_{i,t} = \Phi^{-1}(\hat{U}_{i,t})$ and perform standard Granger causality tests on permutations of pairs of online and onsite protest activity $(\tilde{X}_{i,t}, \tilde{X}_{j,t})$. We considered pairwise causality to avoid misinterpretation — i.e. the extension of Granger’s bivariate causality on VAR to higher dimensions as discussed extensively in the statistical literature (Granger & Lin, 1995; Jea, Lin, & Su, 2005; Lin, 2007).

Results

The first of our Granger causality tests was performed in order to determine if protest-related social media activity on Twitter and Facebook provided significant information for forecasting protest-related onsite activity during the Indignados in Spain, the Occupy, and the Vinegar protests in Brazil.⁵ Figure 2 shows the results of the Granger causality test across the three instances of political unrest with the F-statistic and p values between the pairwise variables. The test statistic in Figure 2 is colored in a gradient heat map, with significant test statistics displayed in white boxes. The results indicate that online communication on Twitter and Facebook predicted onsite protest activity in the Indignados ($p < .00$ and $p < .01$, respectively) and the Occupy datasets ($p < .00$ and $p < .00$, respectively), with bidirectional Granger-causality between online and onsite protest activity in the Occupy series ($p < .00$ for all pairwise variables). In the case of the Vinegar protests, the direction of the prediction was only from online to online and onsite to onsite variables — from Facebook to Twitter ($p < .04$) and from protestors to injuries and arrests ($p < .04$ and $p < .01$, respectively). Finally, the results of the Granger test for lag 1 indicated that tweets Granger-caused Facebook posts in the Indignados and the Occupy series ($p < .00$ for both variables).

INSERT FIGURE 2 HERE

Next we relied on arrests and camped-out protestors as measures of high-functioning political activism to test H2. We found Granger-causality between online and onsite protest activity in the Indignados series, as tweets Granger-caused camped-out protestors, and in the Occupy series, as arrests Granger-caused tweets ($p < .04$ for both variables). The results indicate that the Granger-causality between online and onsite protest activity varies considerably across the individual instances of political unrest. While the online and onsite series in the Vinegar dataset evolved mostly self-referentially, the online and onsite series in

the Occupy dataset presented a significant level of feedback, with Twitter and Facebook both predicting and being predicted by protestors attending demonstrations onsite.

The Granger causality tests for the Indignados dataset, on the other hand, pointed to a one way relationship from online to onsite protest activity, with both Twitter and Facebook predicting protests and Twitter also predicting camp-outs. These patterned Granger-causalities observed in Figure 2 also shed light on different adoptions of social media platforms and enables us to reflect on H3. The pivotal role played by Twitter for the coordination of local logistics, particularly in the organization of encampments in the Indignados movement, or acting as a live-feed among US occupiers (Castells, 2012, p. 172), is consistent with the results shown in Figure 2. In fact, the Indignados and the Occupy series suggest that one could forecast onsite protests by monitoring the use of Twitter, and to a lesser extent, of Facebook posts. Vinegar is the only series where the Granger causality test yielded no significant results for the relationship between online and onsite protest, with statistically significant predictions only within online and onsite activity, but not across the two modalities of protest activity. Table 2 provides a breakdown of the Granger causality tests and displays only significant results with p value and the F-statistic.

INSERT TABLE 2 HERE

Firstly, the results were consistent with H1a for the Indignados and the Occupy series, as the outbreak of online protest activity contains information that helps predicting the future of onsite protest activity. We therefore rejected H1b and concluded that contentious communication is Granger-causal of physical participation in demonstrations. In the Occupy series, we also found bidirectional Granger-causality between online and onsite protest activity, thus further confirming H1a and thereby rejecting the hypothesis that online communication is inconsequential to onsite protest activity (H1b). Secondly, hypotheses H2a and H2b were both partially confirmed, as tweets were found to be Granger-causal of

camped-out activity in the Indignados series and arrests Granger-caused tweets in the Occupy series. Likewise, protest activity was found to Granger-cause high-functioning political activity in the Vinegar series, as the increase of protesting activity predicts clashes with the police measured by arrested and injured Vinegar protestors.

Thirdly, hypothesis H3b was confirmed in the Indignados and the Occupy series. In that instance, tweets Granger-caused Facebook posts with highly significant p values. Nonetheless, hypothesis H3a could not be rejected, as Facebook posts were found to Granger-cause tweets in the Vinegar series. Embedding these findings back into their context, we interpret the partial confirmation of hypotheses H2 and H3 as further evidence that political activists relied on social media platforms, and arguably successfully, to organize rallies and encampments (camp-outs), particularly in the Indignados and the Occupy series. Moreover, we note that the streaming of incidents of police harassment and/or police brutality (arrests and injuries) during the Occupy protests had immediate impact on social media and looped back on demonstrations onsite (Castells, 2012; Earl, et al., 2013).

Notwithstanding the abovementioned caveats, the results of the Granger causality test lend some verification to the claim that online activity can predict onsite activity. Except for the Vinegar series, protesting activity was found to be predictive of multiple instances of social media activity in the Indignados and Occupy series. Moreover, we also found evidence of onsite protest activity predicting online activity in the Occupy series with highly significant p values both for Twitter and Facebook streams. Figure 3 shows the direction of Granger-causalities across the instances of political protest considered in this study, with social media activity depicted in blue, high-functioning political activism shown in red, and onsite protest activity colored green.

FIGURE 3 HERE

Lastly, we note that the results of the Granger causality test were only used to assess the direction of the predictive relationship between online and onsite protest activity. In this sense, we argue that Twitter and Facebook activity observed during the political events considered in this study are informative of (and consequently forecast) next day's protest activity. Our results thereby indicate that the increase of protest-related communication on social media can be a means to forecast onsite protest activity.

Discussion and Further Research

The Indignados, Occupy, and Vinegar political protests were largely organized by grassroots activists working in central city locations over weeks or months. These political movements operated in a horizontal, consensus-based decision-making mode enacted in assembly meetings in which face-to-face interaction was the primary means of communication and a central platform for advocating participatory democracy (Mercea, Nixon, & Funk, 2013). On the other hand, these movements also relied on social media to recruit participants and enhance mobilization (González-Bailón, et al., 2011), resulting in a great deal of discussion about the extent to which social media aid in igniting popular protests (Gerbaudo, 2012). In this article we tested this hypothesis and found compelling evidence that online protest activity is informative of and forecasts onsite protest activity across multiple instances of political unrest. In the remainder, we consider the wider implications of these findings.

We have shown that in the case of the Occupy movement there was a feedback in the prediction of online and onsite protest activity (both on Twitter and Facebook); that in the case of the Indignados movement networked communication was predictive of onsite protest activity (both on Twitter and Facebook); and that the Vinegar protests presented no Granger-causality from online to onsite protest activity. Instead there were online to online interaction (from Twitter to Facebook) and from onsite to onsite protest activity (protestors to injuries

and arrests). In short, in two of the three cases we found significant Granger-causality from Twitter and Facebook protest activity to demonstrations onsite and from physical protests to social media platforms, thus establishing a feedback loop from social media activity to onsite protests and back to social media.

For the reasons outlined above, we view these findings as an initial and limited confirmation by way of a large-scale cross-national study that online and onsite protest activity can be used for mutual predictions (Earl, et al., 2013; Fisher & Boekkooi, 2010; Tufekci & Wilson, 2012). Moreover, on the basis of the Indignados and the Occupy evidence we would submit as a basis for further testing the claim that the outbreak of onsite protest activity can be forecasted by related streams of information on Twitter and Facebook. Consequently, we would propose that wider comparative inquiries verify whether the increase of political messages on social media associated with a specific protest movement constitutes a fertile basis for forecasting the direction of future activity in the same political movement.

In view of the Granger-causalities observed in this study, and recognizing the contextual variability that may further bear on the relationship between online and onsite protest activity, Twitter and Facebook are likely to have amplified demonstrations through continuous networked communication that feeds into the process of participant recruitment. This is a contentious assertion as it stands at odds with claims of political disengagement and the ineffectiveness of social media communication in promoting onsite protest participation, a presumed state of affairs derided as “slacktivism” (Morozov, 2011).

The results of the Indignados series showed that we can expect more protestors attending demonstrations whenever there is a rise in the number of messages related to the demonstrations on Twitter and Facebook. This directional relationship from social media communication to demonstrations may be difficult to grasp if one expects the staging of

demonstrations to trigger communication on Twitter and Facebook; or, in the context of the rising penetration of mobile internet technologies, if one anticipates social media communication to be predominately driven by the reporting of onsite activity (Earl, et al., 2013). With this analysis we believe we may aid with the systematic substantiation of the notion that social media may be a central component of the panoply of tools for organizing and publicizing protest activities before they take place. Additionally, the test of the Occupy series also revealed significant feedback between online and onsite protest activity, likely sustained by different affordances of Facebook and Twitter. While the former is used as a forum for group-bounded interaction where the general public has limited to no access, Twitter streams are notoriously open, easy to monitor, and may be preferred for real-time or post-hoc announcements of protest actions. Nonetheless, the results remain silent about the negotiation of social media usage at the levels of groups and individual protestors who may calibrate it in light of recognized threats intrinsic to social media communication such as surveillance (Mercea, 2012).

Our findings are supported by the image of a vast media ecology which spawned around the Occupy protests and that integrates broadcast and social media in global hubs of communication. Elsewhere, it was posited that social media were entwined in a dense, multi-layered matrix of *stitching mechanisms* (Bennett, et al., 2014, p. 234) that testifies to the enmeshment of online and onsite activities. The latter contributed to the organization of activist groups in disparate geographic locations by facilitating the syndicated creation and coordination, the deployment of resources, and the strategies or meanings utilized in the protests. Nonetheless, as already argued, such ostensible symbiosis was in no way complete and immutable, with many onsite activities not being chronicled online and networked communication simultaneously acting as a bridge for emotional support or political deliberation across protest sites (Mercea, et al., 2013). The results from the Indignados series

appeared consistent with this description. Findings for the Vinegar protest, while requiring further in-depth qualification, allude to a dissimilar balance between local and hyper-local communication and highlight the existence of alternative media ecologies lacking effective integration between broadcast and social media (Saad-Filho, 2013). Lastly, the bidirectional Granger-causality in the Occupy series provides evidence to the claim that political constituents are becoming increasingly connected as individuals rather than as members of a community or group (Bennett & Segerberg, 2013; Flanagin, Stohl, & Bimber, 2006).

In making these claims, we nevertheless highlight that Facebook data collected for this study is restricted to updates posted on the stream of pages, groups, and communities associated with the political movements (see the Appendix for the full list of Twitter hashtags and Facebook groups and pages), consequently encompassing only content made available on Facebook's public stream. Second, Twitter and Facebook user-base (Pew Research Center, 2012, 2013) are most likely not representative of the demographics of citizens that engaged in political demonstrations during the Indignados, Occupy, and Vinegar protests. Therefore, the present results pertain to a publically active contingent of social media users that communicated about the protests rather than the population at large.

It must also be noted that even though we compiled the data by number of tweets, posts, and protestors posting messages and attending demonstrations in the cities where demonstrations took place, the analysis reported in this paper is based upon daily aggregate data per event (Indignados, Occupy, and Vinegar). The Granger causality test requires a long time series, and the data aggregated per city where events took place did not provide a sufficient number of time points as required for the test. In other words, each instance of political unrest was aggregated over all spatial locations instead of their regional disaggregations. This kind of aggregation is customary in the literature on time series analysis when the aim is to show temporal variations in events at the macro level that transcends

individual instances of multi-site protest events. Lichbach (1985, p. 589) commented that although protest results from the actions of specific actors acting at specific times, in specific places, and in distinctive ways, the empirical study of protest over time cannot be quite so microscopic. Instead, it inevitably relies on information gleaned from reports about the episodes, usually taken from journalistic sources. With this in mind, we claim that although perfectible to better account for contextual variability, our approach is appropriate to these series. We also highlight that if the interplay between onsite and online protest activity could be investigated at more disaggregated levels — such as city or neighborhood level — more significant relationships would probably be unveiled. Further research drawing on more detailed data is necessary to determine the relative strategies of social media users in reference to instances of political unrest.

In the last instance, this study should aid with further untangling the interplay between online and onsite protest activity. Principally, it has enabled us to look beyond the vexing debate of whether social media activity has caused political unrest in the past five years, or whether social upheaval gave the tone to social media communication. In short, the reported findings provide substantive grounds to move beyond an “either or” view towards an integrated approach spanning both dimensions of time and space to account for the networked communication of political unrest in the 21st century. By showing that the Granger-causality between online and onsite contentious actions varies considerably across different instances of political unrest, we have provided the necessary empirical baseline to move cognate scientific debates beyond predominant questions of directionality and onto questions of magnitude of mutual elapsed effects of onsite activity and related networked communication.

Acknowledgements

This research was supported by the São Paulo Research Foundation (Grants 10/06243-4 and 11/22495-6) and conducted while the first author was a postdoctoral fellow at the University of São Paulo. The authors are thankful to the anonymous reviewers and the editor for the valuable comments, as well as to Marcus Ohlström for the insightful discussions and conversations.

References

- Anduiza, E., Cristancho, C., & Sabucedo, J. M. (2013). Mobilization through Online Social Networks: the political protest of the indignados in Spain. *Information, Communication & Society*, 1-15.
- Bakker, T. P., & de Vreese, C. H. (2011). Good News for the Future? Young People, Internet Use, and Political Participation. *Communication Research*. doi: 10.1177/0093650210381738
- Bastos, M. T., Raimundo, R., & Travitzki, R. (2013). Gatekeeping Twitter: message diffusion in political hashtags. *Media, Culture & Society*, 35(2), 260-270. doi: 10.1177/0163443712467594
- Bastos, M. T., Recuero, R., & Zago, G. (2014). Taking tweets to the streets: A spatial analysis of the Vinegar Protests in Brazil. *First Monday*. doi: 10.5210/fm.v19i3.5227
- Bennett, W. L., & Segerberg, A. (2013). *The Logic of Connective Action: Digital Media and the Personalization of Contentious Politics*. Cambridge: Cambridge University Press.
- Bennett, W. L., Segerberg, A., & Walker, S. (2014). Organization in the crowd: peer production in large-scale networked protests. *Information, Communication & Society*, 17(2), 232-260. doi: 10.1080/1369118x.2013.870379

- Bobel, C. (2007). 'I'm not an activist, though I've done a lot of it': Doing Activism, Being Activist and the 'Perfect Standard' in a Contemporary Movement. *Social Movement Studies*, 6(2), 147-159.
- Borge-Holthoefer, J., Rivero, A., García, I., Cauhé, E., Ferrer, A., Ferrer, D., et al. (2011). Structural and Dynamical Patterns on Online Social Networks: The Spanish May 15th Movement as a Case Study. *PLoS ONE*, 6(8), e23883. doi: 10.1371/journal.pone.0023883
- Bressler, S. L., & Seth, A. K. (2011). Wiener–Granger causality: a well established methodology. *Neuroimage*, 58(2), 323-329.
- Caren, N., & Gaby, S. (2011). Occupy Online: Facebook and the Spread of Occupy Wall Street. *Social Science Research Network*. doi: 10.2139/ssrn.1943168
- Castells, M. (2012). *Networks of outrage and hope*. Cambridge: Polity Press.
- Colbaugh, R., & Glass, K. (2012). Early warning analysis for social diffusion events. *Security Informatics*, 1(1), 1-26.
- Driscoll, K., & Walker, S. (2014). Working Within a Black Box: Transparency in the Collection and Production of Big Twitter Data. *International Journal of Communication*, 8, 20.
- Earl, J., & Kimport, K. (2011). *Digitally Enabled Social Change: Activism in the Internet Age*. Cambridge: MIT Press.
- Earl, J., McKee Hurwitz, H., Mejia Mesinas, A., Tolan, M., & Arlotti, A. (2013). This Protest will be Tweeted. *Information, Communication & Society*, 16(4), 459-478. doi: 10.1080/1369118x.2013.777756
- Eichler, M. (2010). *Graphical Gaussian modelling of multivariate time series with latent variables*. Paper presented at the International Conference on Artificial Intelligence and Statistics.

Ernesto, C. (2011). OccupyArrests.com Retrieved 15/08/2012, 2012, from

<http://stpeteforpeace.org/occupyarrests.sources.html>

Fisher, D. R., & Boekkooi, M. (2010). Mobilizing Friends and Strangers: Understanding the role of the Internet in the Step It Up day of action. *Information, Communication & Society*, 13(2), 193-208.

Flanagin, A. J., Stohl, C., & Bimber, B. (2006). Modeling the Structure of Collective Action. *Communication Monographs*, 73(1), 29-54.

Gaffney, D. (2010). *#IranElection: quantifying online activism*. Paper presented at the Web Science Conference 2010, Raleigh, NC, USA.

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

Givan, R. K., Soule, S. A., & Roberts, K. M. (2010). Introduction: The dimensions of diffusion. In R. K. Givan, S. A. Soule & K. M. Roberts (Eds.), *The diffusion of social movements: actors, frames, and political effects* Cambridge Cambridge University Press.

González-Bailón, S., Borge-Holthoefer, J., Rivero, A., & Moreno, Y. (2011). The Dynamics of Protest Recruitment through an Online Network. *Sci. Rep.*, 1. doi: 10.1038/srep00197

Granger, C. W. (1969). Investigating causal relations by econometric models and cross-spectral methods. *Econometrica: Journal of the Econometric Society*, 424-438.

Granger, C. W., & Lin, J.-L. (1995). Causality in the long run. *Econometric Theory*, 11(03), 530-536.

Groshek, J. (2011). Media, Instability, and Democracy: Examining the Granger-Causal Relationships of 122 Countries From 1946 to 2003. *Journal of Communication*, 61(6), 1161-1182. doi: 10.1111/j.1460-2466.2011.01594.x

- Hiemstra, C., & Jones, J. D. (1994). Testing for Linear and Nonlinear Granger Causality in the Stock Price-Volume Relation. *The Journal of Finance*, 49(5), 1639-1664.
- Howard, P. N., & Hussain, M. M. (2013). *Democracy's Fourth Wave?: Digital Media and the Arab Spring*. Oxford: Oxford University Press.
- Jea, R., Lin, J.-L., & Su, C.-T. (2005). Correlation and the time interval in multiple regression models. *European Journal of Operational Research*, 162(2), 433-441. doi: <http://dx.doi.org/10.1016/j.ejor.2003.07.020>
- Jungherr, A., & Jürgens, P. (2013). *Forecasting the pulse: how deviations from regular patterns in online data can identify offline phenomena*. Paper presented at the Internet Research 14.0.
- Kallus, N. (2014, April 7–11). *Predicting Crowd Behavior with Big Public Data*. Paper presented at the 23rd International World Wide Web Conference, Seoul, Korea.
- Kavanaugh, A., Yang, S., Li, L. T., & Ed Fox, S. S. (2011). Microblogging in crisis situations: Mass protests in Iran, Tunisia, Egypt. *Proceedings of the Conference on Human Factors in Computing Systems CHI*(6), 1-7.
- Ko, J., Kwon, H. W., Kim, H. S., Lee, K., & Choi, M. Y. (2014). Model for Twitter dynamics: Public attention and time series of tweeting. *Physica A: Statistical Mechanics and its Applications*, 404, 142-149. doi: <http://dx.doi.org/10.1016/j.physa.2014.02.034>
- Koopmans, R., & Rucht, D. (2002). Protest event analysis. In B. Klandermans & S. Staggenborg (Eds.), *Methods of Social Movement Research* (pp. 231-259). Minneapolis: Minnesota University Press.
- Lerman, K., Galstyan, A., Ver Steeg, G., & Hogg, T. (2011, 17-21 July). *Stochastic Models of Social Media Dynamics*. Paper presented at the 5th International AAAI Conference on Weblogs and Social Media (ICWSM), Barcelona, Spain.

- Lichbach, M. (1985). Protest: Random or Contagious?: The Postwar United Kingdom. *Armed Forces & Society*, 11(4), 581-608. doi: 10.1177/0095327x8501100407
- Lim, M. (2013). Framing Bouazizi: 'White lies', hybrid network, and collective/connective action in the 2010–11 Tunisian uprising. *Journalism*. doi: 10.1177/1464884913478359
- Lin, J.-L. (2007). *Notes on testing causality*. National Chengchi University. Institute of Economics, Academia Sinica. Retrieved from http://faculty.ndhu.edu.tw/~jlin/files/causality_slide.pdf
- Liu, H., Lafferty, J., & Wasserman, L. (2009). The nonparanormal: Semiparametric estimation of high dimensional undirected graphs. *the Journal of machine Learning research*, 10, 2295-2328.
- Liu, Y., Bahadori, T., & Li, H. (2012, July). *Sparse-GEV: Sparse Latent Space Model for Multivariate Extreme Value Time Serie Modeling*. Paper presented at the 29th International Conference on Machine Learning (ICML-12), Edinburgh, Scotland, GB.
- Lütkepohl, H. (2007). *New introduction to multiple time series analysis*. Heidelberg: Springer.
- Mercea, D. (2012). Digital prefigurative participation: The entwinement of online communication and offline participation in protest events. *New Media & Society*, 14(1), 153-169. doi: 10.1177/1461444811429103
- Mercea, D., Nixon, P. G., & Funk, A. (2013). Reflections from Occupy the Netherlands *Politics and the Internet in Comparative Context: Views from the Cloud* (pp. 232). London: Routledge.
- Montagna, N. (2010). The making of a global movement: cycles of protest and scales of action. *The Sociological Review*, 58(4), 638-655.
- Morozov, E. (2011). *The Net Delusion: How Not to Liberate the World*. London: Allen Lane.

- Morstatter, F., Pfeffer, J., Liu, H., & Carley, K. M. (2013). Is the sample good enough? comparing data from twitter's streaming api with twitter's firehose. *Proceedings of ICWSM13*.
- Moynihan, C. (2011, 17/09/2011). Wall Street protest begins, with demonstrators blocked, *The New York Times*. Retrieved from <http://cityroom.blogs.nytimes.com/2011/09/17/wall-street-protest-begins-with-demonstrators-blocked/>
- O'Brien III, J. (2010). yourTwapperKeeper. Retrieved from <https://github.com/540co/yourTwapperKeeper>
- Pew Research Center. (2012). Social media and political engagement. In L. Rainie, A. Smith, K. L. Schlozman, H. Brady & S. Verba (Eds.). Washington, D.C.: Pew Internet & American Life Project.
- Pew Research Center. (2013). The Demographics of Social Media Users, 2012. In M. Duggan & J. Brenner (Eds.). Washington, D.C.: Pew Research Center's Internet & American Life Project.
- R Development Core Team. (2014). R: A Language and Environment for Statistical Computing (Version 3.0.3). Vienna, Austria: CRAN. Retrieved from <http://www.R-project.org>
- Russell Neuman, W., Guggenheim, L., Mo Jang, S., & Bae, S. Y. (2014). The Dynamics of Public Attention: Agenda-Setting Theory Meets Big Data. *Journal of Communication*, 64(2), 193-214. doi: 10.1111/jcom.12088
- Saad-Filho, A. (2013). Mass Protests under 'Left Neoliberalism': Brazil, June-July 2013. *Critical Sociology*, 39(5), 657-669. doi: 10.1177/0896920513501906
- Sanggyun, K., & Brown, E. N. (2010, 14-19 March 2010). *A general statistical framework for assessing Granger causality*. Paper presented at the 2010 IEEE International

- Conference on Acoustics Speech and Signal Processing (ICASSP), Dallas, Texas, USA.
- Schelter, B., Winterhalder, M., & Timmer, J. (2006). *Handbook of Time Series Analysis*: Wiley. com.
- Schwarz, G. (1978). Estimating the dimension of a model. *The annals of statistics*, 6(2), 461-464.
- Singer, A. (2014). Social and Political Complexion of the June Events. *New Left Review*, 85(1), 19-37.
- Taamouti, A., Bouezmarni, T., & El Ghouch, A. (2014). Nonparametric estimation and inference for conditional density based Granger causality measures. *Journal of Econometrics*.
- Tarrow, S. (1998). *Power in Movement: Social Movements and Contentious Politics*. Cambridge: Cambridge University Press.
- Tarrow, S. (2005). *The new transnational activism*. Cambridge: Cambridge University Press.
- Tarrow, S. (2011). *Power in Movement*. Cambridge: Cambridge University Press.
- Theocharis, Y. (2013). The Contribution of Websites and Blogs to the Students' protest Communication Tactics During the 2010 UK University Occupations. *Information, Communication & Society*, 16(9), 1477-1513.
- Thorson, K., Driscoll, K., Ekdale, B., Edgerly, S., Thompson, L. G., Schrock, A., et al. (2013). YouTube, Twitter and the Occupy Movement: Connecting content and circulation practices. *Information, Communication & Society*, 16(3), 421-451.
- Tremayne, M. (2014). Anatomy of protest in the digital era: A network analysis of Twitter and Occupy Wall Street. *Social Movement Studies*, 13(1), 110-126.
- Tschernig, R., & Yang, L. (2000). Nonparametric lag selection for time series. *Journal of Time Series Analysis*, 21(4), 457-487.

- Tufekci, Z., & Wilson, C. (2012). Social media and the decision to participate in political protest: Observations from Tahrir Square. *Journal of Communication*, 62(2), 363-379.
- Valenzuela, S. (2013). Unpacking the Use of Social Media for Protest Behavior The Roles of Information, Opinion Expression, and Activism. *American Behavioral Scientist*, 57(7), 920-942.
- Valenzuela, S., Arriagada, A., & Scherman, A. (2012). The Social Media Basis of Youth Protest Behavior: The Case of Chile. *Journal of Communication*, 62(2), 299-314. doi: 10.1111/j.1460-2466.2012.01635.x
- Vasi, I. B., & Suh, C. S. (2013). *Protest in the internet age: Public attention, social media, and the spread of 'Occupy' protests in the United States*. Paper presented at the Politics and Protest workshop.
- Wiener, N. (1956). The theory of prediction. *Modern mathematics for engineers*, 165-190.
- Yamaguchi, R., Imoto, S., Kami, M., Watanabe, K., Miyano, S., & Yuji, K. (2013). Does Twitter Trigger Bursts in Signature Collections? *PLoS ONE*, 8(3), e58252. doi: 10.1371/journal.pone.0058252

Table 1 Parameter Estimates for Daily Data of Indignados, Occupy, and Vinegar

		Date	Tweets	Posts	Protestors	Camped-out	Arrests
Indignados	Min	17 May 2011	1944	772	25	50	3
	Median	NA	11024	3356	2000	425	5
	Mean	NA	27012	4177	15877	882	9
	Max	29 Jun 2011	198095	12647	212000	4775	26
	SD	NA	45766	2766	40921	1416	9
	NA's	0	0	0	14	34	39
Occupy	Min	24 Sep 2011	2271	3469	1	NA	1
	Median	NA	74007	13393	1900	NA	31
	Mean	NA	86265	15727	6724	NA	82
	Max	29 Dec 2011	414408	47431	56200	NA	724
	SD	NA	67973	8533	13490	NA	143
	NA's	0	0	0	46	NA	18
Vinegar	Min	13 Jun 2013	8525	57860	4000	1	3
	Median	NA	20696	239827	52500	4	31
	Mean	NA	61746	253166	204046	32	59
	Max	29 Jun 2013	294257	573286	1569150	100	325
	SD	NA	91445	127050	423983	45	97
	NA's	0	0	0	4	11	7

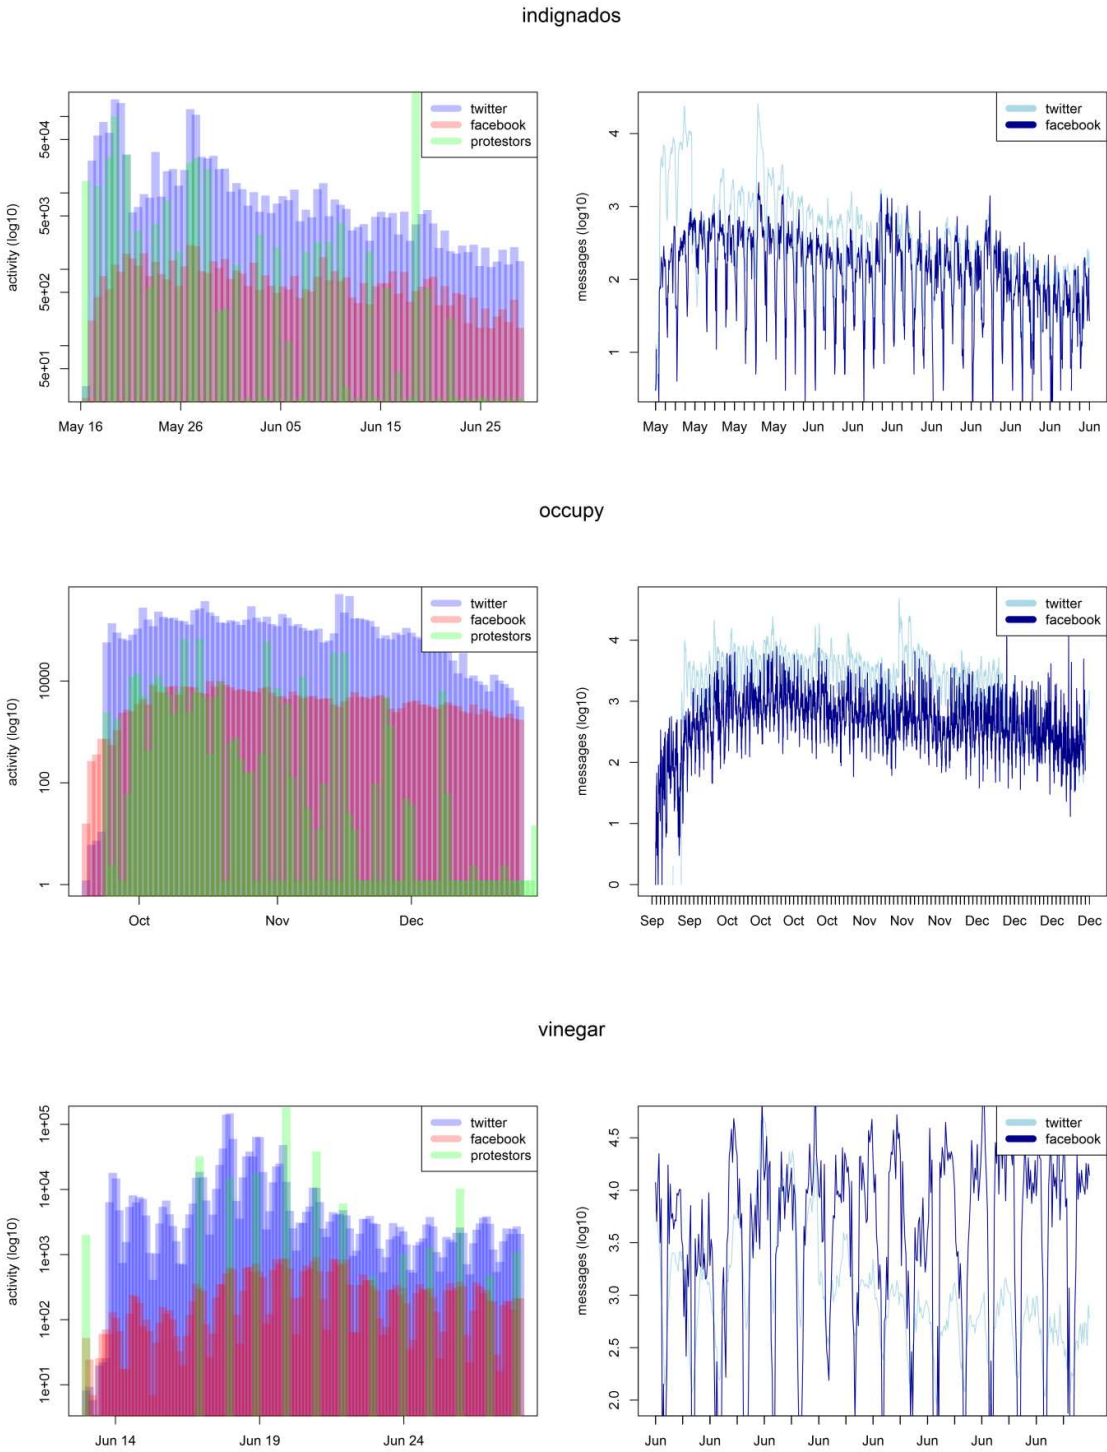


Figure 1 Histogram of Twitter, Facebook, and protest activity (left) and plot of the hourly aggregated data of Twitter and Facebook (right)




Figure 2 F-statistic and p values of Granger causality test between the variables (predictors on vertical and predicted on horizontal axis)

Table 2: Statistically significant Granger-causality relationships between onsite and online protest activity

	From	To	<i>p</i> value	F-statistic
Indignados	Facebook →	Protestors	0.006935	7.681863265
	Twitter →	Camped-out	0.040705	4.327367764
	Twitter →	Facebook	0.002285	9.932696869
	Twitter →	Protestors	0.000028	19.79784760
Occupy	Arrests →	Twitter	0.043938	4.114710565
	Facebook →	Protestors	0.003091	8.986728723
	Protestors →	Facebook	0.004278	8.366341747
	Protestors →	Twitter	0.002806	9.172262406
	Twitter →	Facebook	0.000107	15.67231216
Vinegar	Twitter →	Protestors	0.000597	12.20012279
	Facebook →	Twitter	0.043679	4.495825716
	Protestors →	Arrests	0.008828	8.016421032
	Protestors →	Injuries	0.044385	4.463507130




Figure 3 Directionality of Granger-causality between tweets, Facebook posts, and camped-out, injured, and arrested protestors that participated in demonstrations (all three instances of political unrest considered)

Notes

¹ The average delay is relative to September 17, 2011 when Adbusters launched the proposal to occupy Wall Street. Most Occupy movements are offshoots of the seminal “Occupy Wall Street,” and therefore occurred after this initial event.

² Granger-causality is nonetheless useful for forecasting synchronous events even in the presence of unknown causes — i.e. the tested variables X and Y are influenced by variable Z that is not supported by the data. In this case, we might conclude that Y Granger-causes X , although X might actually be driven by the influence of Z over Y and subsequently over X . Translated to the cases studied in this paper, it is possible that press coverage of onsite protests influenced social media that later influenced onsite protests, but the test measures only the Granger-causality between social media and onsite protest activity. Nonetheless, we expect the results of the test to quantify the extent to which the past of protest-related social media activity contains information that helps predict the future of onsite protest activity (and the other way around) more accurately than using only the past of one of the variables.

³ In this case, VARs of order 1 means we are dealing with 1-day lag between online to onsite events (and vice-versa).

⁴ We also performed a multivariate time series analysis with a VAR(1) model on the entire set of variable as a regression model of values at time t and variables observed at time $t-1$, with explanatory variables indicating possible and instantaneous correlation. However, and contrary to the pairwise approach, coefficients in the autoregressive matrix cannot be interpreted as valid because it is not possible to isolate individual effects or interpret a coefficient as the impact of an impulse shock on one variable at time $t-1$ (with the remaining variables remaining unchanged). In this case, the causality can actually occur through another channel — i.e. a shock on X_2 at time $t-1$ might indicate a shock on X_3 at time $t-1$ or t can

point to a shock on XI , so the causality interpretation might appear as valid with results that are likely to be misleading.

⁵ We selected a $p < .05$ as level of significance for the tests, so even in cases where H_0 holds true, the causality can still be rejected with $\frac{1}{20}$ chance with a 5% criteria. We accepted the null hypothesis of Granger-causality and drew the conclusion that the first order difference of variable Y Granger-causes the first order difference of variable X whenever $p < .05$. When only one relationship was significant, we reported unidirectional Granger-causality. When the relationship was significant both ways, we reported bidirectional Granger-causality (feedback). If neither of them was significant, we rejected the null hypothesis of Granger-causality and concluded that the variables are independent.