

Content and Network Dynamics Behind Egyptian Political Polarization on Twitter

Javier Borge-Holthoefer

Walid Magdy
Kareem Darwish
Qatar Computing Research Institute
Doha, Qatar

Ingmar Weber

{jborge, wmagdy, kdarwish, iweber}@qf.org.qa

ABSTRACT

There is little doubt about whether social networks play a role in modern protests. This agreement has triggered an entire research avenue, in which social structure and content analysis have been central –but are typically exploited separately.

Here, we combine these two approaches to shed light on the opinion evolution dynamics in Egypt during the summer of 2013 along two axes (Islamist/Secularist, pro/anti-military intervention). We intend to find traces of opinion changes in Egypt’s population, paralleling those in the international community –which oscillated from sympathetic to condemnatory as civil clashes grew. We find little evidence of people “switching” sides but observe clear changes in volume with both pro- and anti-military camps becoming more active at different stages.

Our work contributes new insights into the dynamics of large protest movements, specially in the aftermath of the main events –rather unattended previously. It questions the standard narrative concerning a simplistic mapping between Secularist/pro-military and Islamist/anti-military. Finally, our conclusions provide empirical validation to sociological models regarding the behavior of individuals in conflictive contexts.

Author Keywords

Egypt; Twitter; mobilization; polarization; opinion switch

ACM Classification Keywords

J.4 Social and Behavioral Sciences: Sociology

INTRODUCTION

When Time magazine nominated the anonymous protester as the Person of the Year in 2011, it was already clear that online technologies occupied (and occupy) a central spot as the backbone of modern civil protests. Beyond that cycle of turmoil, subsequent events have only reinforced such impression. Public availability and down-to-the-second temporally resolved data have given Twitter an incomparable advantage

in addressing important questions, for instance how social networks actually facilitate the emergence and diffusion of protests [18], or how their contents change in time.

This is of particular importance in the Middle-East and North Africa (MENA) region, where chronic political unrest is combined with a speeded adoption of online technologies. Furthermore, the rapid growth of social networking sites since 2011 shifts away from typical social and entertainment uses of online media, towards those that are more political and civic [21]. Indeed, recent turmoil in MENA countries has provided good case studies for online interaction: Iran (2009), Tunisia, Lybia, Egypt and Bahrain (2011), Palestine and Israel (ongoing Gaza conflict) and Syria (2011-present), have witnessed social revolts, violence and even some government changes. And in all of those, online social networks have been claimed to have a key role.

Since 2011, Egypt’s society is known to be highly polarized with the two dominant poles typically labeled as “Islamist” and “Secularist”. On top of this, the country has experienced much political upheaval during the past year. In such heavily polarized atmosphere, the movement named “*Tamarod*” (meaning *Rebellion* in Arabic) was inceptioned in Egypt aiming to overthrow president Morsi a few months into his presidency. The efforts of the movement culminated in mass demonstrations and counter-demonstrations on June 30, 2013 and resulted in military intervention (coup or revolution¹) to displace Morsi on July 3, 2013. The military takeover and ensuing opposition to it has led to significant tension and much violence. As a result, a new pro- vs. anti-military intervention dimension has emerged. This latter axis has been widely considered in mass media to be aligned with the previous polarization, in the sense that Islamist/Secularist coincide with anti-military/pro-military, respectively. However, during the violent aftermath of the army’s ascent to power, some observers speculated that some of the supporters of the military intervention might have changed their opinion, due to what is mostly viewed –including international observers– as excessive use of force against protesters, with at least several hundred killed.

In this paper, we examine how these events have transpired on Twitter and some of the underlying phenomena. One angle that we are particularly interested in is one of changes in volume or “loudness” of different political camps over time

¹Different political camps use different terminology to refer to the events surrounding the military intervention on July 3, 2013.

and if these were driven by opinion changes. This question is partly motivated by the fact that “the winds turned” several times for the Muslim Brotherhood, and foreign governments were also struggling to decide which side to support [1].

In particular, we attempt to address some inter-related research questions which –we hypothesize– can be tackled using large amounts of Twitter data. First, we examine whether “opinion switching” between both camps can be detected and estimated. If so, we can move on and question further whether perceived opinion changes are due to (i) people actually switching sides, or (ii) the relevant camps becoming increasingly louder or more muted.

In quantifying the answer to these questions, we learn about the underlying psychological and sociological mechanisms which are at play in conflictive contexts. For instance, it might be more beneficial –or less costly– for an individual to withdraw from an ongoing conflict than to actually switch sides: with a similarly minded social neighborhood, an actor is faced with the “volunteer’s dilemma” (in this case, who switches side first); from which the rational outcome is to become mute [10].

Beyond partial accounts, we approach these questions both from the structural (network) side and the semantic (content) side. Concerning content, we used a set of manually labeled hashtags to build a pro- or anti-military-intervention classifier, which is used to classify tweets based on its textual content. On the network side, we used retweets of hand-labeled seed users to derive a Secular vs. Islamist leaning for users. We apply our methodology on a set of nearly 6 million Arabic tweets crawled between June 21 and September 30, 2013. From these, we reconstructed the network of Twitter users who authored tweets in the collection, and we crawled all meta information for 120,000 users along with their latest 3,200 tweets prior to December 2013. All the tweets and Twitter user IDs in our collection are publicly available², so that researchers can replicate the exact dataset we used for possible future studies.

We are well aware that studies conducted solely on Twitter data have certain limitations and drawbacks. From a general point of view, Twitter can not be regarded as *the* voice of civil society. As a matter of fact, Twitter (and online social networks in general) have been adopted by a minority in these societies. Furthermore, the adoption of these new technologies are not uniformly spread, which implies that the data are, most probably, biased. And yet, we are confident that our study stands on a reliable ground. Though limited, Twitter is a valuable monitor of social dynamics and it is doubtless an influencing actor in modern societies, partially shaping the flow of information among individuals and, in many occasions, performing as an agenda-setting actor. Methodologically, our longitudinal approach tracks not absolute values (raw signal), but rather relative changes (first derivative of the signal) which diminish the impact of inherent biases. Whereas biases are expected to strongly impact the absolute

²<http://alt.qcri.org/~wmagdy/EgyMI.htm> Only IDs are provided in order not to violate Twitter’s terms and conditions.

level of, say, Morsi support on Twitter, trends and the direction of *changes* in this level are expected to be more robust.

We find that despite –or because of– the dramatic and violent events there is very little evidence of users changing sides or “switching”. We look at switching between Secularist and Islamist camps and between pro-military and anti-military camps. Our network and content analyses indicate that less than 5% of users switched sides. Instead, the main narrative seems to be one of pro-military intervention and Secular users being dominant in terms of volume leading up to July 3, and anti-military intervention and Islamist users gaining in volume afterwards. Furthermore, in contradiction to the dominating narrative in mass media, the correlation between being a secular and a supporter of military intervention is far from perfect. However, some correlation was noticed between being an Islamist and against the military intervention.

EGYPT’S COUP D’ETAT: TIMELINE

In this section, we summarize the major events that unfolded in Egypt during the summer of 2013. In general, we can distinguish three major periods: the lead up to the June 30 protests; the military intervention on July 3; and the aftermath of the intervention, including the violent crackdown against the Rabia sit-in. The chronology should help in understanding the conflict in the time period that we cover. We also invite the reader to have a look at a Wikipedia article on the topic³ and an Al-Jazeera interactive timeline tool for the Egyptian turmoil⁴ for more information.

The major events were as follows:

- **Late April, 2013:** A movement named *Tamarod* started in an effort to force the Egyptian president Mohamed Morsi out of office. The movement, which called for mass protests on June 30, 2013, claimed to have collected 22 million signatures by June 29. It was aided by strong media support and major power and gas crises⁵. *Tamarod* is often branded as a Secularist or pro-military movement.
- **June 28 and 30:** Large anti-Morsi protests take place in Cairo and across the country. Counter-demonstrations in support of Morsi also begin with major sit-ins in Rabia and Nahda squares. The sit-ins continue non-stop until they are disbanded on August 14. Morsi supporters are often branded as Islamists.
- **July 1:** The military issues a two-day ultimatum to the Egyptian president to accept the demands of *Tamarod* and to call for early presidential elections.

³http://en.wikipedia.org/wiki/2013_Egyptian_coup_d%27%C3%A9tat

⁴<http://www.aljazeera.com/indepth/interactive/2013/08/2013817122637981237.html>

⁵<http://www.spiegel.de/international/world/egyptian-army-gives-morsi-of-muslim-brotherhood-a-48-hour-ul.html>

- **July 2:** The president, in an address to the nation, concedes to some demands from Tamarod. However, he refuses to step down or to call for early elections ⁶.
- **July 3:** General Sisi, the head of the military, orchestrates a coup to overthrow the president, suspend the constitution, disband the elected legislature, arrest Morsi’s most prominent supporters, and shutdown pro-Islamist television stations. A new interim president is appointed. Protests in support of Morsi continue and grow.
- **July 8:** Clashes between security forces and anti-coup protestors erupt in front of a National Guard compound leading to many deaths.
- **July 24:** General Sisi asks his supporters to demonstrate, so as to obtain a “mandate” to fight “possible terrorism”.
- **July 26:** Pro-military protests take place in Tahrir square in response to General Sisi’s demands.
- **July 27:** Clashes between security forces and anti-coup protestors erupt near Rabia square, leading to the death of dozens (mostly demonstrators).
- **August 14:** Security forces crack-down on the sit-ins, leading to the death of hundreds of anti-coup protesters. Demonstrators attempt to establish new sit-ins, but these are as well violently disbanded.
- **August 16:** Major anti-coup protests take place, leading again to the death and arrest of many protesters.

Besides the previous specific dates, protests and clashes are intensified on Fridays (prayer day for Muslims) with continuing sit-ins in Rabia and Nahda squares up until August 14.

RELATED WORK

For a helpful contextualization of our work, we focus on three research avenues. First, a general framework is that of temporally resolved social dynamics, i.e. works in which scholars attempt to track the evolution of social phenomena (as opposed to a static snapshot). Within this area, we highlight the studies that zoom in on political events: the growth of grassroots movements, recruitment processes and (bi)polarization of opinions. A second major area –indeed with much work devoted to it– is that of the Arab world, digital media and/or the political/civic use of it. Third, we briefly discuss sociological accounts of opinion dynamics –“switches” and silencing.

In the first line of research, it is possible to identify some works which provide longitudinal accounts of political (protest) events. Some examples are the work by Borge and collaborators [9] (devoted to the growth of the “Indignados” movement in Spain) or González-Bailón *et al.* [19] (recruitment processes again in the 2011 protests in Spain). But closer to the current approach, we pay attention to Weber *et al.*’s work [36, 37], who explore the bipolar political scenario in Egypt (Secular vs. Islamist). To do so, they track the

retweeting behavior from certain seed users as a signal to deduce a user’s political orientation. Furthermore, they transfer the polarity from users to hashtags based on usage patterns – such that a hashtag is tagged as Islamist, Secular or “neutral”. In doing so, the authors provide a means to assess real-time online polarization. They provide anecdotal evidence that the increases in this polarity score anticipate periods of political violence.

Needless to say, polarization has been a focus of interest before. Starting with the NOMINATE score [32], to quantify the statement that U.S. politics follows a 1-dimensional left-to-right schema, much research has followed. Focusing on digital media and U.S.-centered politics (i.e. left-vs.-right polarization), Adamic *et al.* [4] and Conover *et al.* [12] capitalized on data from the blogosphere and Twitter respectively. The time-resolved retweet graphs constructed from our data set present some similarities to Adamic’s work; and yet, the temporal dimension is missing in these articles.

Second, there is a large body of work on the so-called Arab Spring, some of which zoom in on the Egyptian Revolution in particular. These depart from the present work insofar they typically include some discussion on the role of social media during the protests and the revolution, but do not use data from social networks (nor Twitter in particular) to approach polarization or major events. A notable exception is work by Mostak [29], who tests the political hypotheses that “Islamism in an ideology of the poor” with online data. To approach the question, he looks for geographical correlations between census information for income and house value/size, and estimates for how much Islamist activity on Twitter there is originating from a particular region. Down to conflictive situations, we find a plethora of works devoted to the Arab Spring –as the expression “Twitter Revolution” was coined in mass media: some examples can be found in [24, 34, 6, 31, 23, 7], often from a qualitative point of view or relying on surveys. With the use of data from blogs, Al-Ani *et al.* [5] explore alternative news sources –beyond the government-supplied versions of events. However, none of them quantifies online polarization, or provide a longitudinal point of view. Other countries and conflicts have of course caught the attention of researchers, such as Tunisia [39] and Palestine [38]. Closer to this work, Choudhary *et al.* [11] performed time-resolved sentiment and response analysis to Twitter activity during the events in 2011, that eventually led to the displacement of Mubarak and the onset of a transitional period. Apparently, they used English tweets coming from Egypt and relevant news sources. In our work, we focus on Arabic tweets –clearly underrepresented in scholarly articles– to better model people in the region.

Finally, we find literature that focuses on the monitoring of opinion changes and the existence of online “silent crowds”. Far from the political debate, some researchers have tried to tackle the fundamental differences between the “vocal minority” and the “silent majority” [30]. In this work, a *caveat* is placed to consider the significant differences between those highly active online media users –who dominate the ongoing discussions over social networks– and the vast majority of

⁶<http://english.ahram.org.eg/NewsContent/1/64/75538/Egypt/Politics-/Egypts-Morsi-defies-calls-to-step-down,-offers-opp.aspx>

people who participate much less frequently. The same line of argumentation is found in Venkataraman *et al.* [35]. A step closer to our politically motivated work, Lin *et al.* [25] exploit sentiment analysis to track political opinions in U.S. – although their method is intended for general purposes. Common to our framework, a real-time, evolving opinion-shift account is offered. Finally, on a more theoretical ground, the work by Chenoweth *et al.* [10] suggests that there are indications that, in a conflictive environment, members of opposing groups do not necessarily switch sides to tip the balance in favor of one of the groups, but they may merely withdraw their support.

DATA COLLECTION

Collecting Relevant Egyptian Tweets

We collected Arabic tweets in the period between June 21, 2013 and September 30, 2013. This period covers the major events in the aforementioned chronology. We used Twitter4J APIs⁷ to collect the tweets matching the query “*lang:ar*”, (so as to retrieve Arabic tweets). On average, the number of collected tweets per day was 3 million. The Arabic text of the collected tweets was pre-processed using state-of-the-art normalization technique for social Arabic text [15] to facilitate later filtering. The normalization includes letter normalization, diacritics removal, decorative text replacement and word elongation compression.

To extract the tweets that are relevant to Egypt, we constructed a rich set of Boolean queries that cover different facets of Egyptian politics. The set contained 112 requests that represent key players in the Egyptian political sphere: government representatives, political parties, opposition leaders, media anchors and the Twitter accounts of some relevant players. Each Boolean request is prepared to handle different spellings of names, acronyms, and nick-names used by different political groups. As such, an enriched set of queries underlies every request in order to retrieve relevant tweets from the Arabic news platform TweetMogaz⁸ [26, 16]. TweetMogaz automatically generates news about Egypt, taking Twitter as the source. By default, the platform uses 84 queries and applies an adaptive tracking algorithm to enrich them, relying on ongoing news [27]. For this work, we manually revised the automatically enriched keywords to ensure a high precision and recall. On top of that, we also considered context keywords, i.e. relevant to that particular period: “revolution”, “June 30”, “coup”, “Rabia” and “massacre”. This extra set was picked by two native Egyptian, who regarded these words as most meaningful to the studied turmoil episode.

The number of extracted tweets matching those queries during the mentioned time window was 5.9 million. To ascertain the accuracy of the matching items, we manually judged a set of 500 random tweets, from which only 5 were assessed as not relevant, rendering a 99% estimated accuracy. Figure 1 shows the volume of tweets collected on the filtered set over time. As expected, major events led to peaks in the number of

tweets on the days they occurred. There is clear mapping between the peaks in Figure 1 and the events listed in the timeline section (see above), which hints at the fact that Twitter is –among others– a faithful and sensitive monitor of real-world (offline) life.

Collecting Users’ Network Data

We identified Twitter accounts that authored/retweeted more than 10 tweets in our collection, to find out their network information for network analysis. In doing so, we guarantee that the monitored users did engage, at some time or another, the ongoing crisis in Egypt. This led to the identification of nearly 120,000 Twitter accounts, out of roughly 1 million in our collection of Egyptian tweets. For each user, we extracted their network information, which comprehends the number of followers, friends and, most importantly, the last 3,200 tweets they authored or retweeted. These latest tweets allowed us to measure the network dynamics of each user through their activities including retweets, mentions, or replies to other accounts. The resulting information exceeded 750GB of user data and more than 300 million tweets for these users. Table 1 summarizes the obtained tweets data collection.

CONTENT-BASED ANALYSIS OF TWEETS

In this section, we describe the analysis we conducted on tweets, based on their textual content. To this end, we built a supervised classifier that groups tweets according to their content, i.e. whether they express support (pro-MI: pro-military intervention) or opposition (anti-MI) to the military intervention, plus a “neutral” category as well. We use this classifier to measure the response to the events in the time window of interest. Later, we deepen our analysis to track for possible opinion changes at the individual level.

Classifying Tweets as Pro/Anti-Military Intervention

The design of the classifier comprehends three main steps. First, we constructed lists of seed hashtags that are either exclusively or predominantly used by either of the two groups. Then, we expanded those lists to identify more tweets that could potentially belong to either group. Finally, we employed a supervised method to classify the tweets.

For the first step (seed lists), we extracted the most frequent hashtags in our collection, which were then tagged as either most likely pro-MI or anti-MI group. The rest of hashtags were neglected. Briefly, we assumed pro-MI hashtags to express support for the army, calls for a regime change or a *revolution* (as opposed to *coup*), opposition to Morsi or his regime or attacks against the Muslim brotherhood. Conversely, the group of hashtags categorized as anti-MI expressed maligning the members of the Supreme Council of Armed Forces (SCAF), the usage of *military coup* (as opposed to *revolution*), expression of support for Morsi or his regime, or rejection to the crackdown against the sit-ins. The lists contained the highest frequent 150 hashtags for each polarity, see Table 2 for some examples.

With the seed lists at hand, we heuristically classified tweets in our collection as belonging to the pro-MI or anti-MI groups if they exclusively contained hashtags belonging to the corresponding lists, for either of the groups. If a tweet contained no

⁷<https://code.google.com/p/tweet4j/>

⁸<http://www.tweetmogaz.com>

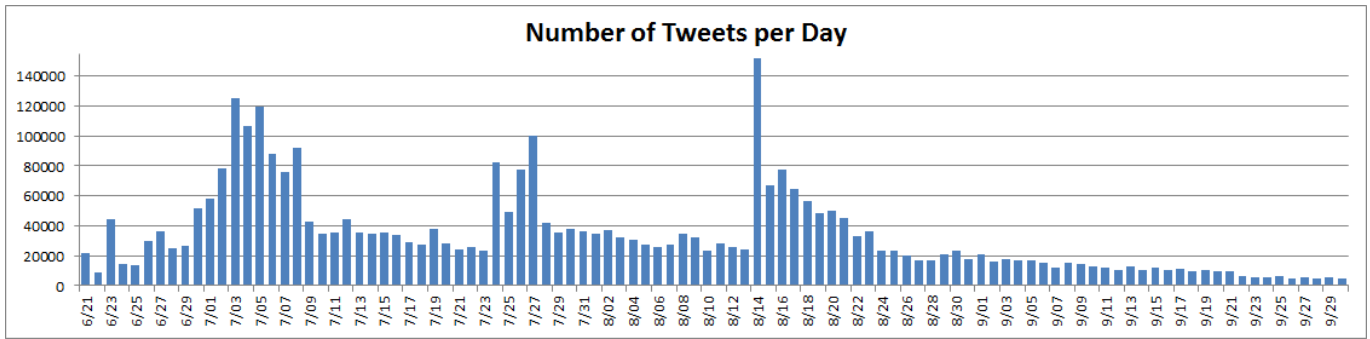


Figure 1. Distribution of the collected tweets about Egypt over time

Data	Description	Included Information	Size
Egyptian tweets	Egyptian tweets between June 21 and October 1, 2013	Tweet ID, text, Twitter user, and Timestamp	5,902,086 tweets
User Data	Profiles of Twitter users with more than 10 tweets in the Egyptian tweets set.	Username, screen name, ID, number of followers and followees. List of last 3,200 posted tweets by the account (prior to the mid of December 2013)	121,003 profiles. Recent 3,200 tweets per profile

Table 1. Tweets data collection used in our study

hashtags indicating polarity or contained hashtags from both lists, it was discarded. Using the filtered tweets, we identified new words and hashtags that exclusively appeared in either group, but not in both. We recursively added these words and hashtags to their respective lists. We repeated this expansion step four times. The resultant lists for the pro-MI or anti-MI group lists contained 1,105 and 1,487 words and hashtags respectively. Exploiting these final filtering lists, we constructed a subset of the tweet collection containing 12% of the tweets in the whole dataset.

Finally, from the filtered tweets subset we randomly selected 500 tweets containing hashtags from each list. These 1,000 tweets were manually verified and judged as exclusively belonging to pro-MI, anti-MI or neutral. Judgments were performed by two Egyptian Arabic speakers, who are keenly aware of the situation in Egypt. The inter-annotator agreement between them was above 93%. In case of disagreement, the annotators discussed over the tweets until consensus was reached. After this process, the agreement exceeded 96%. For the remaining tweets (33), one of the judgments was selected at random. Table 3 shows the distribution of the 1,000 tweets over the three groups. We used a multi-class Support Vector Machine (SVM) classifier to classify tweets as pro-MI, anti-MI, or neutral. We employed the SVMLight Multiclass implementation [13] for classification. We used the following features: word unigrams, word bigrams, and hashtags. We performed 20 fold cross validation over our 1,000 labeled tweets where each fold (containing 5% of the examples) was used for testing and the 19 other folds (containing 95% of the examples) were used for training. The average accuracy for the folds was 87.0% with a standard deviation of 3.4%. Table 4 shows the classifier confusion between the guessed and actual tags.

Pro-MI		Anti-MI	
Hashtag	Translation	Hashtag	Translation
#ارحل_يا_مرسي	Step down Morsi	#noscaf	No SCAF (army)
#السيسي_البطل	Sisi the hero	#السيسي_خائن	Sisi traitor
#ثوره_30_يونيو	30 June revolution	#ضد_الانقلاب	Against coup
#الاخوان_الكاذبين	Muslim brotherhood liars	#منبحة_رابعة	Rabia massacre

Table 2. Examples of seeding hashtags used for Pro/Anti-military classifier

Pro-MI	Neutral	Anti-MI
35%	9%	56%

Table 3. Distribution of tags in training/test set

Though the undetermined tweets may contain pro-MI or anti-MI tweets, our interest focuses on finding the distribution of tweets belonging to both groups over time. Thus we favored precision. Figure 2 plots the proportion of tweets that were classified as opposing (anti-MI: blue) or supporting (pro-MI: orange) the military intervention over time. As illustrated in Figure 2, before the military intervention on July 3 the proportion of tweets belonging to the pro-MI group was dominant (55% to 75%). After July 3, pro-MI tweets witnessed a significant drop, while anti-MI tweets rose substantially. This can be attributed to many different reasons that we attempt to analyze in the next section.

Given the content of our filtered collection, we analyze the top used hashtags by each group over time. We do so in order to understand the topics of interest for each of the groups and how they evolved in response to major events. We identified the 30 hashtags that achieved the highest per-day volumes for any of the days in the full period, for both groups. We eliminated hashtags, such as #Egypt, that were mentioned everyday, achieving the highest aggregate volume over all the

Truth	guess		
	Pro-MI	Neutral	Anti-MI
Pro-MI	0.817	0.064	0.119
Neutral	0.125	0.736	0.139
Anti-MI	0.054	0.022	0.924

Table 4. Confusion between truth and guessed labels for classes

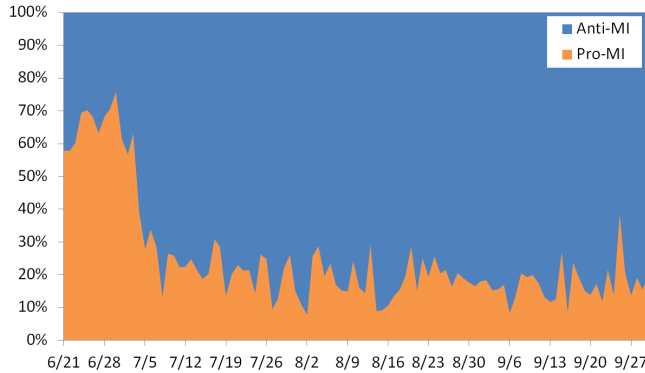


Figure 2. Proportion of tweets per day according to polarity of supporting/opposing military intervention

days –but never exhibited a burst in volume for any given day. Using this scheme, a hashtag that is mentioned 10,000 times in one day is more important than another that is mentioned 1,000 times daily over a span of 100 days. A hashtag showing a spike in volume would likely indicate changes in topics or interest for groups. Figure 3 shows the number of occurrences of the identified 30 hashtags in the period of study. The upper panel shows the most frequent hashtags on the pro-MI group, while the lower panel is devoted to the most frequent ones on the anti-MI group. Surprisingly, the overlap between hashtags representing the topics discussed by each group is almost nonexistent. Only four (out of 30) hashtags are in common in the two groups (*#Morsi*, *#Sisi*, *#Rabia*, and *#Beltagy*). However, the use of these hashtags were different in scale, time periods, and sentiment.

Upon inspection of the tweets and hashtags of the pro-MI group, they are focused more on individuals and organizations. The top hashtag is *#Morsi*. The use of the tag peaked around his ousting on July 3, and the relative volume of the hashtag for the rest of the period continued to be high. Similarly, tweets critically mentioning *#Ikhwan* (Muslim Brotherhood) persisted with relatively high volume for the whole period. *#Tamarod* peaked around Morsi’s deposition, but dwindled quickly afterwards. Other pro-MI tweets belonging to persons or organizations included support for *#Sisi* (the head of the military), attacks directed at *#Tawakkol_Karman* (Yemeni Nobel prize laureate who was critical of the military takeover) or joy at the arrest of *#Beltagy* (former parliamentarian and a prominent anti-MI figure). Other prominent hashtags can be found related to criticism towards the sit-ins in *#Rabia*, with a peak on August 14, and of the *#subway_sit-in*, with a peak on September 15.

Anti-MI tweets and hashtags instead focused mostly on events, particularly those involving violence and loss of life (*#National_Guard_Massacre*, July 8), *#Ra-*

bia_Massacre, *#Rabia*, *#Rabia_Adawia_Massacre*, *#Rabia_Crackdown_Massacre* (July 27 and August 14) and *#Dilga* (September 16). They also include hashtags for announced protests on specific days or weeks, with activities against the military (*#Defeating_the_Coup* on July 8; *#Egypt_against_Coup*, August 2; *#Last_Friday*, August 8; or *#People_Protecting_their_Revolution*, September 6). The most mentioned individuals in tweets where *#Morsi*, mostly around the time of his ouster, general *#Sisi*, distributed over the entire period, and *#Beltagy*, around the time of his arrest.

Detecting Users Polarity Switching

With a reliable classifier at hand, which provides content-based polarity classification, we can observe the changes in polarity at the Twitter user-level over time. A user polarity switch means that a user who was tweeting pro-MI tweets at t_0 moved to authoring anti-MI tweets at $t_1 > t_0$, or vice-versa. We apply this analysis to understand whether this was the reason behind the large decline in the number of pro-MI tweets after the military intervention on July 3 as shown in Figure 2 and the increase in anti-MI tweets.

Two hypotheses are compatible with the observed pattern changes, namely: people switched camps in response to events; or, simpler, one camp became quieter while the other becoming louder. We envisage that the latter is the mechanism behind the phenomenon –in accordance with previous literature [10]. And yet empirical evidence is lacking in a violent context, thus the relevance of the present work.

Detecting switch in user’s political leaning requires a minimum number of tweets of that user over the studied period to monitor his leaning over time. The confidence of detecting a switch in leaning increases the larger the number of pro/anti-MI tweets are examined for a given user: more tweets supply more evidence of the user’s leaning. Unfortunately, most of the users in our collection did not have large number of tweets classified as pro/anti-MI. Therefore, we relied on a small portion of users who had at least n classified tweets, $n \in \{5, 10, 15, 20\}$. Examining leaning switch of users based on only 5 tweets might be inaccurate, however this leads to the examination of more users in our test set; whereas examining users with at least 20 tweets in our collection leads to a smaller test set, though yielding a higher confidence.

To consider a user polarity switch, we examined the first and last third of tweets for the user. If the average polarity changed from the first third compared to the last third, we assume that the user swapped her leaning. Otherwise, we consider that the user held a coherent (stable) opinion over time. Figure 4 reports the percentage of used switched from anti-MI to pro-MI, or vice versa, and presents the number of users examined in the analysis for different values of n . As shown, for all n values, which reflects different confidence levels, the number of user switching sides are minimal, where it did not exceed 5% of the total number of users examined. For example, the largest percentage when $n = 5$, out of 21,312 users examined, only 970 switched from pro-MI to anti-MI, and 280 switched from anti-MI to pro-MI. The percentage of switching is even less for larger n values, indicating that the higher percentages for smaller n might be caused by random

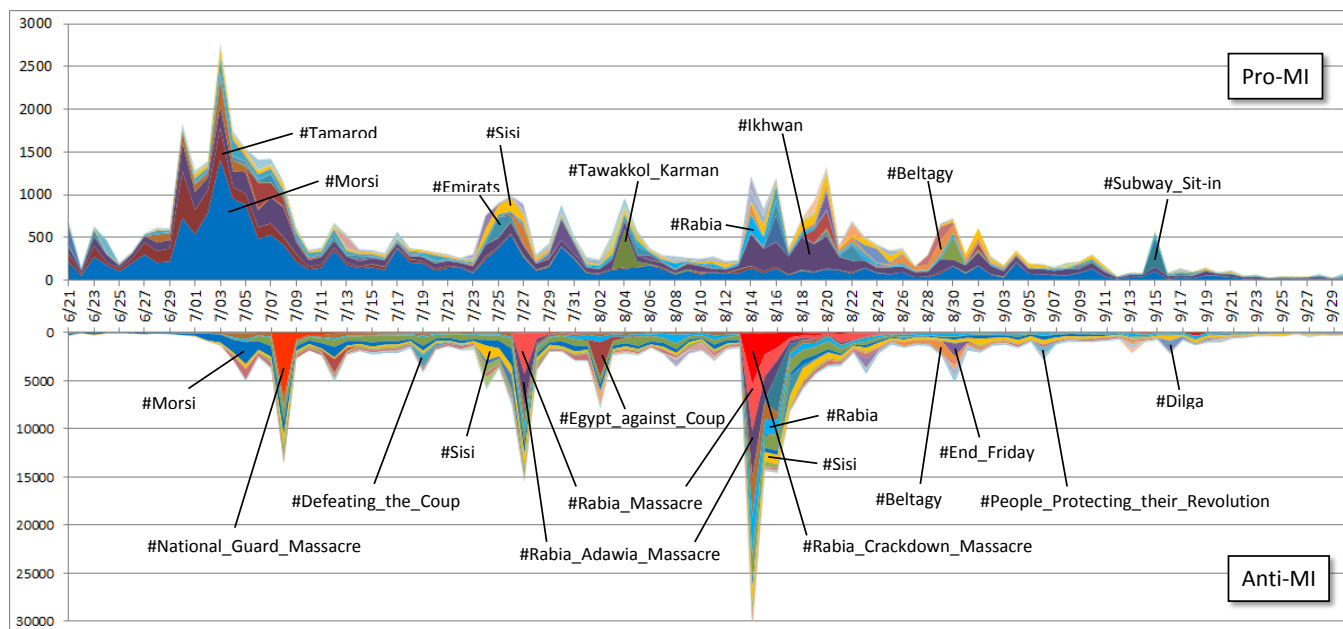


Figure 3. Number of occurrences of the most outstanding 30 hashtags appeared in the pro-MI (upper) and anti-MI (lower) tweets between 21 June and 30 Sep. 2013. Note the scale for the anti-MI plot, which is 10 times the scale of the pro-MI graph. Presented hashtags are translations of the original Arabic ones.

noise in the classifier. In all the switches, the number of users switched against the military intervention are at least triple the number of users switched to support the military intervention.

Our “first third, last third” schema is admittedly an arbitrary one. There are, however, at least two good reasons to proceed in this way. First, an opinion change might not occur in a sudden way, but rather progressively (as events unfold). Thus, it is necessary to focus the attention on the initial opinion state and the final one, with some unobserved time in between. In doing so, we avoid the presence of noise that might “contaminate” our output. For this reason, a $\frac{1}{2} - \frac{1}{2}$ schema would not yield a reliable result (despite its providing larger statistics). In the opposite extreme, a more refined method would be to consider a small initial and final tweeting period (first $\frac{1}{10}$ th, last $\frac{1}{10}$ th, for instance). We would have obtained a more robust result, at the expense of poor statistics.

Table 5 presents some example tweets (translated from Arabic) for users who switched their polarity. Example 1 in the table is for a Twitter user who was calling for protests against Morsi on June 21, and then he was calling for participation in a march against what he called the military coup on July 19. Similarly, the second example is from pro-MI to anti-MI. On July 4, he expressed approval for the overthrow of Morsi, but then on July 24, he attacked General Sisi⁹ for his request for a “mandate”. The final example is for a user who switched from anti-MI to pro-MI, where the user showed some solidarity for protestors against military intervention, then after two and half months he showed support for Sisi and opposition to Morsi.

⁹“Sisi” also means pony in Arabic.

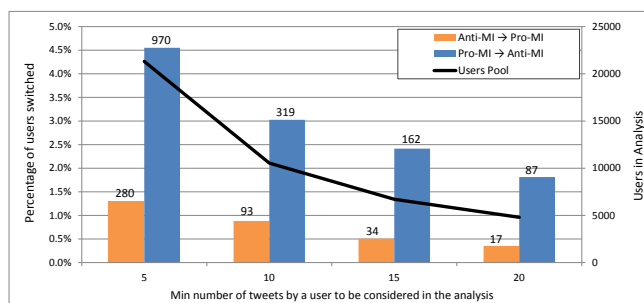


Figure 4. Percentage of users switched from one polarity to the other, and the number of users used for the analysis of switching for different values of n

Although the detected user switches are interesting, they can overall be considered anecdotic: the majority of users tend to express stable opinions. Thus, it can be concluded that the significant changes in polarity proportions in Figure 2 are due to supporters of military intervention becoming quieter and the opponents becoming louder.

NETWORK-BASED ANALYSIS OF TWEETS

From the data collected we built a sequence of temporally-evolving networks (one per day). In these networks, a weighted directed link (i, j, w) is laid between two nodes if user i retweeted a previous tweet by j a number of times w . We used a 3-days overlapping sliding window scheme, where the network on day t is obtained by aggregating 3 days of activity, i.e. retweets from $t - 1$ to $t + 1$; and the network on day $t + 1$ spans from t to $t + 2$, and so on. This has some advantages, namely: a 3-day wide aggregation guarantees that retweeting patterns will reflect a certain stability –away from

Switch	Date	Tweets (English translation)
<i>Pro-MI</i> → <i>Anti-MI</i>	June-21	We will continue to revolt till we reach freedom. Gathering revolution from Alexandria to Cairo to oust Morsi, the sheep.
	July-19	The Mohandseen march is closing the main streets till the police station #No_to_military_coup
<i>Pro-MI</i> → <i>Anti-MI</i>	July-04	Everybody, Morsi is displaced. Damn brotherhood who pretend to be Muslims. Long live Egypt
	July-24	I thought you were a Sisi (meaning: pony), but it turned out that you were a fox #Sisi_is_calling_for_civil_war
<i>Anti-MI</i> → <i>Pro-MI</i>	July-15	#Egypt_coup gas bombs are thrown inside the Fath mosque and people are trapped inside #Egypt_coup
	Sep-28	Please share the photo of Saied Qotb when he was arrested in a box like the one general Sisi put Morsi, the sheep, inside

Table 5. Sample tweets of Twitter users who switched polarity

the noisy immediacy of events; also, the overlapping scheme allows for a finer tracking of ongoing events.

In general, the sequence of retweet-reconstructed networks exhibit some fairly constant global topological properties. As preliminary steps, “silent” nodes were removed and the (weakly) giant connected component was extracted (typically containing over 95% of the original nodes). Density ($\sim 10^{-4}$), average degree ($2 \leq \langle k \rangle_t \leq 5, \forall t$) or clustering ($\sim 10^{-3}$) display small differences across time. The networks are thus very sparse and highly disassortative, in the sense that a few Twitter users get most of the retweeting activity, with hardly any presence of triadic closure.

Community Sizes

In this work, one of the main motivations was to adopt a network approach to track polarity evolution over time. In this direction, we disregard global (like the ones above) or microscopical (node-level) properties, to focus on changes at the meso-scale or group level. To do so, we applied a well-known community detection technique, namely label propagation [33], as implemented in the C-coded *igraph* network analysis package [14]. This algorithm has some highly desirable features. Apart from obtaining significant values in the modularity optimization process (the higher the modularity Q , the better the quality of the resulting partition), label propagation delivers excellent performance in terms of memory and time.

The algorithm runs with some predefined constraints. First, we use a list of seed users for whom the partisan leaning is out of any doubt, both in the Secularist and the Islamist side [37]. Part of the list was taken from Mostak [29] with additional entries compiled in 2013 by an Egyptian expert¹⁰. Second, we limit the number of possible communities to two –since we track for a bipolar political scenario in Egypt. Admittedly, it is possible that the network structures (over time) allow for better, higher- Q partitions with more than two communities, but that is beyond the scope of this work. Also, though Secularists and pro-MI group members do not necessarily align, Islamists were more likely to be in the anti-MI group than Secularists, as we show later. To provide an even more efficient running time, we fed the network at time t with the labels obtained in the previous snapshot $t - 1$, when possible (thousands of users enter and leave the network at each time step, in the latter case nodes are unassigned initially, and the algorithm eventually imposes the corresponding tag on it).

¹⁰Due to largely differing degrees in popularity, the list of seed users boils down to Mohamed ElBaradei (@ElBaradei; Secularist camp) and Muhammad Morsi (@MuhammadMorsi; Islamist camp); the rest of the list has secondary effects on the output of the algorithm.

Under these considerations, the label propagation algorithm yields in general high modularity Q values: the average Q -value per snapshot is $Q = 0.4340$ with very few variation, $\sigma_Q = 0.0150$. This implies that in terms of partition quality, communities remain fairly constant over time. To have a solid standpoint regarding the quality of this result, we have generated (and Q -optimized) 100 surrogate versions of the first snapshot. That particular snapshot has an actual $Q = 0.4431$, whereas its random counterparts achieve an average $Q = 0.0468$ with $\sigma_Q = 0.091$. This means the actual Q corresponds to a z -score of 4.3151, i.e. it indeed represents a significant, far-from-random partition.

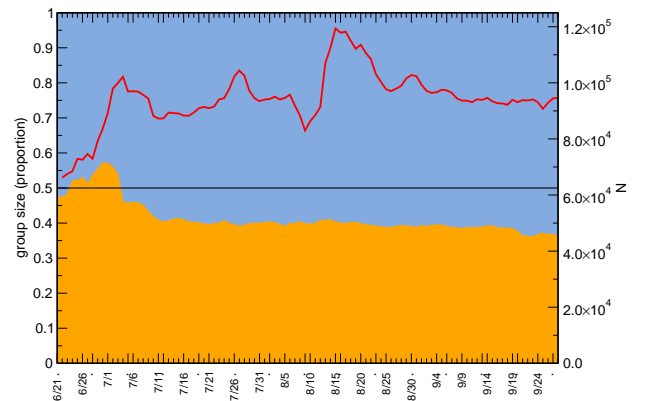


Figure 5. Community sizes corresponding to Islamic (blue) and secular (orange) leanings. The red line tracks the size of the network each day. Dots on the x-axis represent Fridays.

Figure 5 offers a time-resolved view of community size evolution. During the first days of the time window –and up to the military intervention date– Secularists outnumbered Islamists in the retweet network. On July 3, there was a crossover, and the number of Islamists remained larger until the end of the period under study. Even in bursty periods, where the network slices underwent severe changes (note for instance network sizes –red line in Figure 5– around mid-August), Secularists did not (proportionally) recruit newcomers, whereas the Islamist camp did. Unsurprisingly, the network size fluctuated over time, specially as critical events unfold.

Note that it would be overly simplistic to assume “one Twitter user, one vote” and that Twitter user counts perfectly match the proportion in society [28, 22, 17]. This holds both for network-based measures as for content-based measures. Still, the trends in relative sizes of the two camps might be more robust against selection bias and other skews.

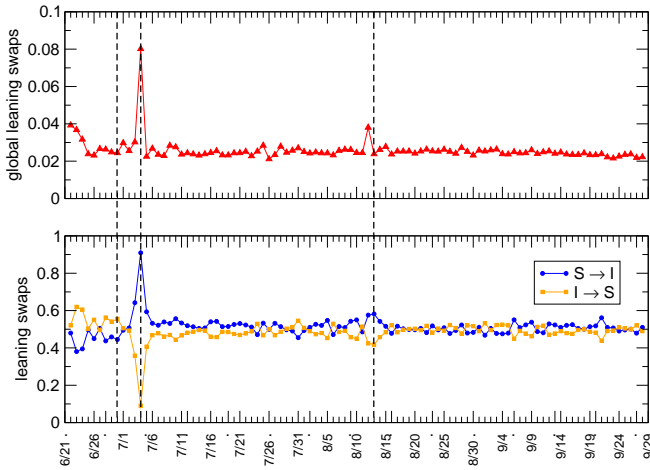


Figure 6. Leaning switches occur at a very low rate (top panel; hardly 3% of switching over time, with a few exceptions on key dates). Also, whenever swapping behavior exists, the direction of these swaps are evenly shared between parties (again, with some exceptions on key dates). Dots on the x -axis represent Fridays.

Global Switching (Islamist vs. Secular)

As stated above, the Islamist dominance of the retweeting scene beyond the 3rd of July is clear. The question remains whether secular sympathizers switched to Islamist positions (as they overwhelmingly belong to the Anti-MI group), thus feeding the growth of the Islamist group as observed in Figure 5. To check this, we measured the global switch ratio (Figure 6, top panel) as the number of observed label switches (from $t - 1$ to t), over the total amount of common nodes in the networks corresponding to $t - 1$ and t . Notably, switching remained at a very low level with a single exception (July 3) when it climbed up to 8%. In the lower panel of the same Figure, we show the proportion of swaps in either direction, across time. Again, most of the time there is an evenly distributed amount of switching in both directions (from Secular to Islamist, and vice versa) with a noteworthy exception on July 3, when 90% of the switches occurred from a secular to an Islamist position. Though this could be due to actual ideological turnover, it might also just point out that the focus of attention –as expressed through retweeting behavior– is temporarily shifted towards the Islamists –whose president is ousted on that day.

Leaning histogram vs. Tweeting Activity

Besides a longitudinal (across time) view of the leaning evolution, we aggregated all this information to provide soft labels per user. This was easily done either for the whole episode (June 21 – September 30) or for some limited period of interest $[t_0, t_f]$. For each user, we calculated his soft label by adding up every time he was assigned an Islamist leaning and normalizing this value by the total number of times he had been present in the network over time. Thus, an $l_i = 1$ label means that user i was tagged as an Islamist for the whole period under consideration. Labels can take values in the interval $0 \leq l \leq 1$. In Figure 7 we placed these soft labels in a histogram (0.05 bin width; note log scale in the histogram counts axis), to get a global view of the leaning

landscape during the unfolding of events. Top panel corresponds to soft labels in the pre-coup period (June 21 – July 3). During this first period, seculars (leaning = 0) and moderate actors ($0 < l < 0.5$) outnumber Islamist ($l = 1$) and pro-Islamic moderates ($0.5 < l < 1$); also, average activity in each of these classes is significantly higher for the former. This scenario drastically changes for the second period (bottom, July 4 – September 30), in which Islamists recruit many new users for their side, as well as the average activity increases. On the secular side, activity drops as the size of the party decreases in relative terms.

In coherence with Figure 5, the number of secular sympathizers is larger (top panel) during the pre-coup period (those with $l < 0.5$), but the opposite is observed for the post-coup period (lower panel). Furthermore, we computed the average strength (number of retweets) per user and per day within each bin (red line in each panel). Clearly, those with a secular leaning are (comparatively) much more active than Islamists during the initial period, but the activity level flattens afterwards –with a slight advantage for the Islamist side.

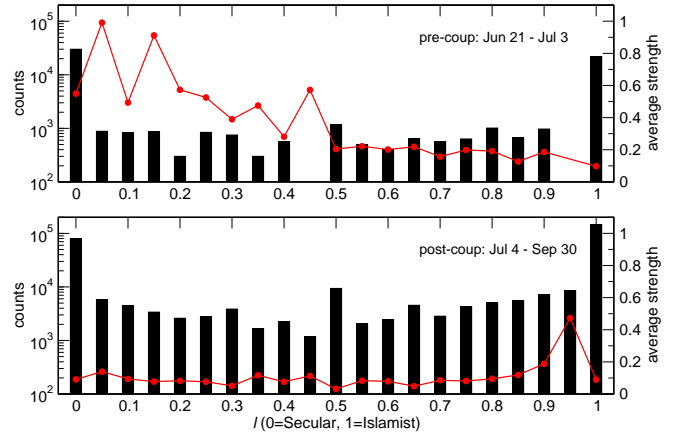


Figure 7. Polarization histograms and average per-day activity (red) for two different periods (pre- and post-intervention, top and bottom panels respectively). Note logarithmic scale for histogram counts (left).

DISCUSSION

At this point, we can attain a global picture of the temporal evolution of the events in Egypt during the time window of interest. On one hand, content-based analysis showed that some group *pro-MI* leaning tweets were flowing in Twitter, and they were especially dominant during the brewing days before the military intervention in early July. Group *anti-MI* activity became dominant soon after July 3, with increasing bursts as violent events unfold. On the other hand, political leaning in the retweet networks underwent a similar transition, i.e. Secularists constitute a larger and more active group up to the coup d’etat, but progressively decrease in number and activity thereafter. These two bipolar schemes (pro/anti-MI, Secular/Islamist) with similar trajectories may lead to the conclusion that Secularists have predominantly a pro-coup standpoint, while Islamist map onto an anti-coup opinion. Such was the general interpretation being projected in classic mass media during the summer of 2013 [2].

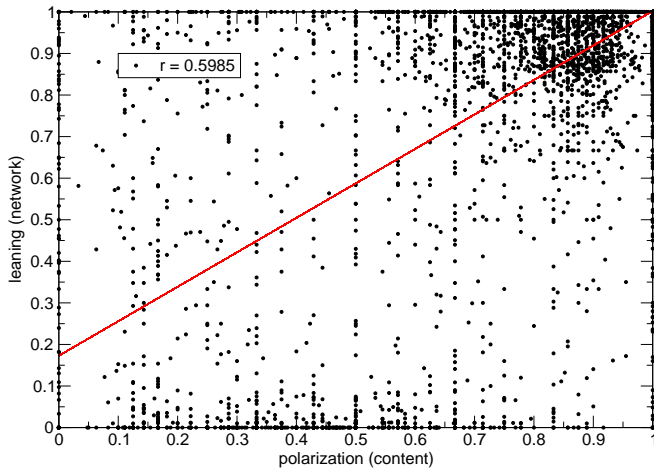


Figure 8. Network-based leaning versus content-based polarity. In the x-axis, lower polarization implies a pro-MI stance, while anti-MI contents are closer to 1. In the y-axis, low leaning values correspond to the Secular camp, whereas Islamists are mapped onto higher ones.

On closer inspection, the scenario turns out to be considerably more complex. If we confront (content-based) group pro/anti-MI polarization scores and (network-based) Secularist/Islamic leaning soft labels for each Twitter users, a relatively strong correlation is found ($r \sim 0.6$, see Figure 8). This means that the leading trend in our analyses (which, noteworthy, have been obtained exploiting different facets of the data) exhibits a significant alignment pro-MI-to-Secularists and anti-MI-to-Islamists; and yet such alignment is far from perfect: ideally, leaning and polarity should meet around the diagonal in the scatter plot. On closer inspection, most of the disagreement occurs on the Secularist side, whereas there seems to be a rather univocal coherence in the Islamist side. Secularists exhibit rather heterogeneous positions regarding the military intervention in the country’s politics.

A possible conclusion is that anti-Morsi protests were in fact quite diverse. In this direction, what we labeled here as Secularists might actually be conformed by a heterogeneous collective, including seculars indeed, as well as people without political affiliations and even small groups of Islamists –not necessarily with a pro-MI leaning. On the whole, Islamists exhibited a solid anti-MI stance as the diverse contrarians progressively withdraw. As facts have evolved during this past year, our conclusions seem to grow stronger: indeed, the provisional government led by General Sisi has prosecuted Islamist figures, but also groups or individuals from other ideological stances [3].

Another observation that can be gleaned from our content- and network-based analyses is that switching between pro- and anti-MI groups, and between Secularists and Islamists, is rather limited. Most of the change in volume in observed polarized tweets seems to be mostly due to each group becoming louder or quieter as time progresses. This is consistent with the work of Chenoweth and Stephan [10] where they observed that members of opposing groups in conflict do not necessarily switch sides to tip the balance in favor of one of the groups, but they may merely withdraw their support.

CONCLUSIONS

In this paper we have analyzed how recent turmoil in Egypt in the period between June 21 and September 30, 2013 has transpired in Twitter. We examine the tweets from the perspective of Secularists vs. Islamists and from the perspective of pro- and anti-military intervention. The main lesson from our observations is that the military takeover on July 3 caused major quantitative (volume of polarized tweets), but not ideological (polarity swaps) shifts among Twitter users. Secularist and pro-military intervention Twitter users were very loud before the take-over, but became increasingly silent afterwards. Conversely, Islamists and anti-military intervention Twitter users become significantly louder after the takeover. Furthermore, a close inspection of our results indicate that the axis Secular/Islamist can not be trivially mapped onto a pro/anti-coup opinion –which was the dominant assumption in traditional media at the time of the military intervention. This is specially true for the Secularist-to-pro-MI projection, which is not as aligned as it was generally presumed. The development of the political situation in Egypt during the past year have confirmed our insights.

To reach such conclusions we have exploited two methods, which rely on different (and non-overlapping) aspects of our dataset: the semantic and the structural dimensions. Noteworthy, the trends observed from these varied methodologies converge, on a qualitative basis.

We are aware of the limitations of the present study. Besides the inherent biases in Twitter data [20], other arguments can be put forth. For instance, it is unlikely that an individual publishes a radical change of opinion through a single medium, i.e. such attitude has an associated social cost. Therefore, the availability of other sources (Facebook, blogs, etc.) would be highly desirable. We stress, however, the precise nature of the events we analyze here. We do not expect, in such bipolar scenario, that a Secularist might turn Islamist (or vice versa). But it is possible –and even *probable*, given the turn of the events– that Seculars who held high expectations for a government change, were afterwards deceived as violence and repression grew. Expressing such discontent is not as costly –if it is at all.

For possible future directions, one could analyze the motivation of users who switched polarity to understand reasons for people to change leaning. We could see some examples of users who withdraw support for the military intervention just after it occurred, others after the violent crackdown of the Rabia sit-in, and others who supported the military-intervention after opposing it in the beginning. This non-trivial switching temporal pattern resembles that of “complex contagion”, for which a rich theoretical literature is available [8], in which early adopters (early switchers, in this case) pave the way for adoption (switching) cascades. Another essential work direction is to add an additional polarization layer to the Islamist/Secular classification. From the results we have obtained, there might be a missing polarity, which could match old regime supporters (Mubarak’s regime). It is known that these were heavily involved in the June 30 protests, and are the strongest supporters of the army. Accounting for them

would be interesting, and can give a more complete image of the situation.

On a general framework, the work contributes to an understanding of bipolarized societies –and these are not limited to Egypt: the underlying questions and methods are more widely applicable. The turmoil in Crimea and Eastern Ukraine, the ongoing conflict in Palestine or the Catholic/Anglican dichotomy in Northern Ireland might be other promising case studies.

REFERENCES

1. These are some news stories that illustrate how, as an example, the US administration has re-evaluated its position towards Mohamed Morsi and the Egyptian army over time:
<http://www.washingtontimes.com/news/2012/jun/24/obama-calls-to-congratulate-morsi-an-islamist-on-w/>
<http://www.theguardian.com/world/2013/jul/03/egypt-obama-us-mohamed-morsi-crisis>
<http://www.theblaze.com/stories/2013/07/02/obama-calls-egypts-morsi-and-says-u-s-does-not-support-either-side>
<http://www.aljazeera.com/indepth/features/2013/07/2013710113522489801.html>
<http://www.forbes.com/sites/dougbandow/2014/01/20/pharaoh-al-sisi-takes-control-in-egypt-obama-administration-sacrifices-security-human-rights-and-democracy>
2. See for instance
<http://www.project-syndicate.org/commentary/how-egypt-can-avoid-the-fate-of-algeria-in-1992-by-lvaro-d-vasconcelos>
3. Prosecution in Egypt beyond the Muslim Brotherhood. Media and secular activists:
<http://www.aljazeera.com/indepth/interactive/2014/07/freeajstaff-days-jailed-journalism-greste-fahmy-mohamed-b-20147166544671591.html>
<http://www.theguardian.com/world/2013/nov/27/egypt-secular-activists-arrested-protest-law>
<http://www.bbc.com/news/world-middle-east-25484064>
4. Adamic, L. A., and Glance, N. The political blogosphere and the 2004 u.s. election: divided they blog. In *LinkKDD@KDD* (2005), 36–43.
5. Al-Ani, B., Mark, G., Chung, J., and Jones, J. The egyptian blogosphere: a counter-narrative of the revolution. In *CSCW* (2012), 17–26.
6. Attia, A. M., Aziz, N., Friedman, B., and Elhusseiny, M. F. Commentary: The impact of social networking tools on political change in egypt’s “revolution 2.0”. *ECRA 10* (2011), 369–374.
7. Azab, N. A. The role of the internet in shaping the political process in egypt. *IJEP 3* (2012), 31–51.
8. Borge-Holthoefer, J., Baños, R. A., González-Bailón, S., and Moreno, Y. Cascading behaviour in complex socio-technical networks. *Journal of Complex Networks 1*, 1 (2013), 3–24.
9. Borge-Holthoefer, J., Rivero, A., García, I., Cauhé, E., Ferrer, A., Ferrer, D., Francos, D., Iñiguez, D., Pérez, M., Ruiz, G., et al. Structural and dynamical patterns on online social networks: the spanish may 15th movement as a case study. *PLoS One 6*, 8 (2011), e23883.
10. Chenoweth, E., and Stephan, M. J. *Why civil resistance works: The strategic logic of nonviolent conflict*. Columbia University Press, 2011.
11. Choudhary, A., Hendrix, W., Lee, K., Palsetia, D., and Liao, W.-K. Social media evolution of the egyptian revolution. *CACM 55*, 5 (2012), 74–80.
12. Conover, M., Ratkiewicz, J., Francisco, M., Goncalves, B., Flammini, A., and Menczer, F. Political polarization on twitter. In *ICWSM* (2011).
13. Crammer, K., and Singer, Y. On the algorithmic implementation of multiclass kernel-based vector machines. *The Journal of Machine Learning Research 2* (2002), 265–292.
14. Csárdi, G., and Nepusz, T. The igraph software package for complex network research. *International Journal Complex Systems* (2006), 1695.
15. Darwish, K., Magdy, W., and Mourad, A. Language processing for arabic microblog retrieval. In *Proceedings of the 21st ACM International Conference on Information and Knowledge management*, ACM (2012), 2427–2430.
16. Elsayy, E., Mokhtar, M., and Magdy, W. Tweetmogaz v2: Identifying news stories in social media. In *CIKM* (2014).
17. Gayo-Avello, D., Metaxas, P. T., and Mustafaraj, E. Limits of electoral predictions using twitter. In *ICWSM* (2011).
18. González-Bailón, S., Borge-Holthoefer, J., and Moreno, Y. Broadcasters and hidden influentials in online protest diffusion. *American Behavioral Scientist 57*, 7 (2013), 943–965.
19. González-Bailón, S., Borge-Holthoefer, J., Rivero, A., and Moreno, Y. The dynamics of protest recruitment through an online network. *Scientific Reports 1* (2011), 197.
20. González-Bailón, S., Wang, N., Rivero, A., Borge-Holthoefer, J., and Moreno, Y. Assessing the bias in communication networks sampled from twitter. *Social Networks 38* (2014), 16–27.
21. Governance, and Innovation Program, D. S. o. G. Arab social media report, vol. 1, issue 2. www.ArabSocialMediaReport.com, 2011.
22. Jungherr, A., Jürgens, P., and Schoen, H. Why the pirate party won the german election of 2009 or the trouble with predictions: A response to tumasjan, a., sprenger, to, sander, pg, & welpe, im “predicting elections with twitter: What 140 characters reveal about political sentiment”. *Social Science Computer Review 30*, 2 (2012), 229–234.

23. Khamis, S. The transformative egyptian media landscape: Changes, challenges and comparative perspectives. *IJOE* 5 (2011), 1159–1177.
24. Lim, M. Clicks, cabs, and coffee houses: Social media and oppositional movements in egypt, 2004–2011. *JC* 62 (2012), 231–248.
25. Lin, Y.-R., Margolin, D., Keegan, B., and Lazer, D. Voices of victory: A computational focus group framework for tracking opinion shift in real time. In *WWW* (2013), 737–748.
26. Magdy, W. Tweetmogaz: a news portal of tweets. In *SIGIR* (2013).
27. Magdy, W., and Elsayed, T. Adaptive method for following dynamic topics on twitter. In *ICWSM* (2014).
28. Metaxas, P. T., Mustafaraj, E., and Gayo-Avello, D. How (not) to predict elections. In *SocialCom-PASSAT*, IEEE (2011).
29. Mostak, T. Social media as passive polling: Using twitter and online forums to map islamism in egypt. Master's thesis, Center for Middle Eastern Studies Harvard University, 2012.
30. Mustafaraj, E., Finn, S., Whitlock, C., and Metaxas, P. T. Vocal minority versus silent majority: Discovering the opinions of the long tail. In *Privacy, security, risk and trust (passat)*, *IEEE 3rd International Conference on Social Computing (SocialCom)*, IEEE (2011), 103–110.
31. Oh, O., Eom, C., and Rao, H. Collective sense-making through the twitter service during the 2011 egypt revolution. In *ICIS* (2012).
32. Poole, K. T., and Rosenthal, H. A spatial model for legislative roll call analysis. *AJPS* (May 1985), 357–384.
33. Raghavan, U. N., Albert, R., and Kumara, S. Near linear time algorithm to detect community structures in large-scale networks. *Physical Review E* 76, 3 (2007), 036106.
34. Starbird, K., and Palen, L. (how) Will the revolution be retweeted?: information diffusion and the 2011 Egyptian uprising. In *CSCW* (2012), 7–16.
35. Venkataraman, M., Subbalakshmi, K., and Chandramouli, R. Measuring and quantifying the silent majority on the internet. In *Sarnoff Symposium (SARNOFF)*, 2012 35th IEEE, IEEE (2012), 1–5.
36. Weber, I., and Garimella, K. R. K. #egypt: visualizing islamist vs. secular tension on twitter. In *ASONAM* (2013), 1100–1101.
37. Weber, I., Garimella, V. R. K., and Batayneh, A. Secular vs. Islamist polarization in Egypt on Twitter. In *ASONAM* (2013), 290–297.
38. Wulf, V., Aal, K., Abu Kteish, I., Atam, M., Schubert, K., Rohde, M., Yerosusis, G. P., and Randall, D. Fighting against the wall: Social media use by political activists in a Palestinian village. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, ACM (2013), 1979–1988.
39. Wulf, V., Misaki, K., Atam, M., Randall, D., and Rohde, M. 'On the ground' in Sidi Bouzid: investigating social media use during the Tunisian revolution. In *CSCW* (2013), 1409–1418.