When Groups Fall Apart: Identifying Transnational Polarization during the Arab Uprisings

Robert Kubinec^{1,*} and John Owen²

¹New York University Abu Dhabi

²University of Virginia

*Corresponding author rmk7@nyu.edu

July 17, 2020

Abstract

It is very difficult to know how international social linkages affect domestic ideological polarization because we can never observe polarization occurring both with and without international connections. To estimate this missing counterfactual, we employ a new statistical method based on Bayesian item-response theory that permits us to disaggregate polarization after the Arab Uprisings into domestic and transnational components. We collected a dataset of Twitter accounts in Egypt and Tunisia during the critical year of 2013, when the Egyptian military overthrew the Islamist President Mohamed Morsi. We find that the coup increased retweets among Egyptian ideological allies by 50 percent each day following the coup and decreased cross-ideological retweets by 25 percent. Tunisian Twitter communities also showed stronger intra-group retweeting although at lower levels than in Egypt. Counter-intuitively, our model shows that the additional polarization in Tunisia after the coup appears to have dampened further polarization among Islamists in Egypt.

Introduction

International diffusion is a perennial and large question in comparative politics, straddling the intellectual boundary with international relations. When we observe similar policies, regimes, social movements, ideologies, or unrest across countries, are they caused by similar conditions in those countries, or have they spread among the countries through some mechanism? The question implicates both our theories of politics and our empirical strategies. Galton's Problem captures the issue: should observations of a phenomenon in several countries at once be treated as independent, or might autocorrelation be at work – for example, through transnational learning and imitation? We can also phrase this issue as a problem of *missing counterfactuals* because it is very difficult to observe diffusion occurring both with and without transnational influence.

In this paper, we present analysis of Twitter data from 2013 in Egypt and Tunisia. We measure the degree of Islamist-secularist and democrat-authoritarian polarization on a given day by the proportion of retweets by ordinary citizens of ideological elites' tweets. A rising proportion of Islamist retweets among Islamists, for example, means that Islamism is becoming more salient for Islamists.

Our dependent variable is change in group polarization. Group polarization is situational and hence may be short-lived (Sunstein 2002); it is different from the stable, long-term social polarization (e.g. Red versus Blue America) that many social scientists study (Moulaert and Sekia 2003). An endogenous process triggered by events whose timing is exogenous, group polarization is a way to conceive of how identities and preferences change in response to stimuli such as a public demonstration or a coup d'état. Of particular interest after the Arab Uprising are (1) polarization along one identity axis across a national population, (2) transnational polarization, i.e., simultaneous polarization along an identity axis in two countries, and (3) direct transnational polarization (within state B responding to events in state A) and indirect transnational polarization (within B

responding to polarization in A).

We choose this time period because Egypt went through a number of political events that seem prime candidates for triggers of group polarization. Most obvious is the July 3, 2013 coup d'état in which secularist military officers overthrew the elected government of Mohamed Morsi of the Muslim Brotherhood. In our models we test for the effect of this coup (whose timing, we assume, is exogenous) on endogenous polarization. We find that it has a statistically discernible long-term effect on the diffusion of group polarization across countries. Overall, retweets within Islamist communities in Egypt and Tunisia increased by approximately 50% and 20% respectively following the coup, while retweets within secularist communities in Egypt and Tunisia increased by approximately 30% and 20%.

In addition to these empirical estimates, we also break new ground by identifying two counterfactual estimates: the *direct* effect of the coup on polarization in both countries and the *indirect* influence of transnational allies on polarization. Most previous studies of polarization focus on the direct effects of domestic political actors polarizing vis-à-vis each other, while any transnational influence is necessarily baked into these estimates. By permitting both of these factors to be estimated separately in a Bayesian multivariate time-series model, we exploit model-based inference to gain insight into unobservable counterfactuals about how polarization occurs both with and without international diffusion.

Counter-intuitively, while we show in this paper that domestic polarization increased, we also show that international connections had a dampening effect among Egyptian Islamists. Through model-based inference, we show that if Egyptian Islamists had lacked knowledge of polarization occurring among their co-religionists in Tunisia, they would have become even more polarized vis-à-vis secularists in their own country. This finding suggests that transnational diffusion does not necessarily increase domestic polarization but rather can serve as an impetus for groups to

strategically de-polarize.

Group Polarization: An Informal Model

Polarization has been studied extensively by social scientists. Much of that work concerns *social* polarization, or segregation into groups that are stable over long periods of time (Mason 2014). By *group* polarization, we mean a process of intense identity and preference segregation that is relatively short-term and subject to reversal or de-polarization, yet politically consequential (Sunstein 2002; Hafer and Landa 2006). Group polarization builds on social polarization, in the sense that it is an activation of a social cleavage that is long-term but latent; it denotes times when social polarization becomes especially salient and political consequential. When two actors polarize in our short-term sense, at time t , both may prefer a 50-50 allocation of goods; at $t + 1$, each may prefer a 60-40 allocation in its favor; at $t + 2$, a 75-25 allocation; and so on.

Stated informally, the basic group polarization model is simple. Assume a population of 100 persons, all belonging to one half or another of x pairs of opposing social groups (democratic or authoritarian, Islamist or secularist, urban or rural, etc.). Fifty are pro-democracy, fifty pro-authoritarian. These groups do not correlate significantly to any other groups; e.g., city dwellers are as likely to be Islamist as secularist. The population thus has cross-cutting cleavages. At time t , the population begins in an unbiased equilibrium, such that, although individuals may identify more strongly with one group than with others, in the population as a whole, no identity axis predominates; social interaction does not skew the distribution of resources, including information, to any of the social groups (Dunning and Harrison 2010). Now suppose that at $t + 1$ three democrats – one Islamist and two secularist, and two urban and one rural – publicly beat an authoritarian. Assuming a low-cost flow of information, that event can trigger the polarization of the population along a democratic-authoritarian ideological axis, such that democrats and authoritarians care

progressively less about where one lives, or mosque-state questions, and progressively more about ideology. The endogeneity of group polarization implies that we can distinguish a *direct* effect from the beating (people observe the beating, feel threatened, and begin to polarize) from an *indirect* effect (people observe their confrères polarizing, and do likewise). If not disrupted or reversed, polarization by definition culminates in inter-group violence. If depolarization occurs, triggered by an (unmodeled) event, the population returns to the status quo ante, in which social polarization along the ideological axis resumes latency.

We expand on the definition of this informal model in Appendix B for the interested reader.

Hypotheses

We propose to test the following hypotheses based on this theory:

- H1 After the coup against Mohamed Morsi, in each country the difference in latent ideological positions between competing groups will diverge in direct reaction to the coup (direct effect).
- H2 After the coup against Morsi, competing groups will show additional influence from the reactions of like-minded allies in other countries (indirect effect).

The event that we can identify a priori as potentially polarizing is the coup that overthrew Egypt's Islamist President Mohamed Morsi on July 13, 2013. As noted earlier, we ascribe an independent effect to the coup because we consider the *timing* of the coup to be exogenous. Suspicions that a coup might occur existed from the time the Egyptian army removed former President Mubarak following widespread protests in 2011. Ongoing tensions between President Morsi and the military establishment suggested that the probability of a coup during 2013 was non-trivial and increased in June (Dempsey 2013). As we have noted, this prior existing polarization undoubtedly

resulted in heightened group salience before the coup. In order for our covariate to be identified, however, we only need to assert that the general public and the Muslim Brotherhood did not know the exact day on which the coup would occur, as our data are measured in days and our model explicitly incorporates over-time trends.

Academic literature on coups notes potential endogeneity, namely that a typical coup is preceded by interactions between the plotters and the government or public, interactions that might lead the latter to anticipate the coup and to prepare for or act to forestall it (Gerling 2017, 6; Meyersson 2016). Plotters have incentives to hint that a coup might be coming, so as to intimidate opponents and fence-sitters, but not to reveal the precise time, place, or personnel involved; governments and their supporters, meanwhile, have incentives to acquire this information (Luttwak 1968, 170-88; Singh 2014, 79-80, 92-96). The Egyptian army officers who overthrew Morsi deliberately created anticipation of a coup, yet left Morsi and his supporters uncertain whether and precisely when it actually would take place. After days of growing demonstrations and counter-demonstrations in several Egyptian cities, General Abdel Fattah al-Sisi issued a 48-hour ultimatum on July 1. But al-Sisi continued to present himself as an honest broker between Morsi and the opposition, and Morsi continued to negotiate, perhaps believing that the ultimatum was a bluff to force a quick resolution. On July 2, elements in the military seized the state's flagship newspaper, Al-Ahram, and announced that they would overthrow Morsi if he did not meet the opposition's demands. Morsi in fact met and exceeded those demands on the same day. Yet, al-Sisi evidently did not convey Morsi's capitulation to the opposition, and executed the coup on July 3 (*Egypt and the Coup: Inside the 11 Days that Toppled Morsi* 2018). No doubt public anticipation of a coup on that day was significant, but al-Sisi maintained enough ambiguity that we can still assume that no one outside the circle of plotters knew for certain when it would take place. We return to a technical discussion of inferential leverage towards the end of our presentation of the model.

Data Collection

To obtain our data, we started with a universe of tweets from the early stage of the Arab Uprising, December 2010 to April 1st, 2011 that all matched the search keywords “Cairo”, “Alexandria” and “Tunis” in the user self-reported location field in Twitter.¹ While this dataset comprised 11 million tweets, it nonetheless did not capture all the important or influential users because it is common for Twitter users either to not report their location or to list a location that is not geographic. To identify influential users not in this sample, we parsed the tweets in order to find those users that had received at least a thousand mentions during that time period. In this way, even if an influential user was not in the original sample, we were able to locate most popular Twitter users in the country by analyzing the content of the tweet database.

We curated the resulting list of elite users, both by removing accounts that were later abandoned and by adding in accounts that were created later in time. We removed accounts that were solely focused on media aggregation or other types of non-partisan content. For additions, we focused in particular on identifying influential politicians, such as those associated with the Egyptian Muslim Brotherhood and other political parties. Some of these additions were found by examining the follower networks of elites we could identify (i.e., who did each elite choose to follow). The final sample amounted to 148 Twitter users, 56 from Tunisia and 92 from Egypt (the full list is in Appendix F).²

Within this finished list of elite users, we collected full tweet histories from March 31st, 2013 to December 31st, 2013. We then filtered out each user’s retweets, reducing the number of tweets

1. The tweets were purchased through the Gnip corporation (now owned by Twitter). The names of cities were chosen to permit within-country inference, though in this paper we aggregate the data to the country level.

2. The larger number from Egypt reflects the much larger Twitter-sphere in the country and hence the need to obtain a broader sample of users. While there may be users whom we did not find through these methods, we believe the lists to be broadly representative of the different ideological groups we are interested in.

from 1.7 million to 1.2 million. Using Twitter’s REST API and the R package `rtweet` (Kearney 2019), the retweets of all of these 1.2 million original tweets was then downloaded as a list of user IDs for each user per day, for a total of 275 days. The use of Twitter’s publicly available REST API represents a limitation in the data collection because only 100 re-tweets of a given retweet may be downloaded; however, this limit was rarely reached in practice because very few of the tweets in question had more than 100 retweets. Even Morsi, who had more than a million Twitter followers, averaged only a few hundred re-tweets per tweet sent during 2013.³ We refer the reader to Appendix E for an overview of descriptive statistics related to the data.

To perform our analysis, we needed to code each of these elite users along two axes: Islamism and secularism and pro- and anti-democratic attitudes. We describe in Appendix F the procedures which we used to do so, including both human coding and the employment of separate IRT models to account for difficult-to-code users.

As a result, the final database comprises a set of elites i and citizens j in which the outcome is the number of times that j retweeted i for each 24-hour period in the sample. We removed all citizens who did not re-tweet at least five different elites and at least two different elite groups during the entire time period, resulting in a final dataset of 119,470 citizen-elite interactions with a total of 6,134 unique citizen Twitter users, or an average of 19 retweets per user of elites during the sample period.

To fully capture the nuance in our data, we expanded this dataset by including all absent interactions; i.e., for every 24-hour period, we include zeroes for all elite tweets which a citizen did not retweet. This dramatically expands the volume of data to 11,807,950 observed and possible

3. While these retweet counts are small relative to many political accounts today, it should be noted that this data collection took place before it was commonplace for collections of Twitter bots to immediately retweet influential accounts. In addition, simple amplification of accounts will not affect model estimates as it is only the relative change in the proportions of retweets which can affect ideal point scores.

Figure 1: Count of Retweets by Ideological Group Over Time

interactions. Including zeroes is important so that we do not assume that each citizen had an equal chance of retweeting all the elite Twitter accounts in the sample, as we explain in our modeling strategy.

Figure 1 shows the number of retweets of elite users aggregated by sectarian affiliation on each day in the sample. As can be seen, Twitter activity varies tremendously over time, with prominent spikes around July when the Egyptian military overthrew the Muslim Brotherhood president Mo-hamed Morsi. It is the job of the measurement model we explicate in the next section to separate mere heightened Twitter activity from true group polarization in these retweet data.

Modeling Ideological Polarization Over Time

Incorporating measurement uncertainty is very important for Twitter analysis, and has been approached in a variety of ways by existing studies. Generally speaking, scholars have applied some kind of model or aggregation algorithm to Twitter data before running statistical models, such as

sentiment coding (Jamal et al. 2015; Siegel et al. 2018; Weber, Garimella, and Batayneh 2013), network statistics (Freelon, Lynch, and Aday 2015), calculating the diffusion of known hashtags or keywords (Steinert-Threlkeld 2017; Metzger and Tucker 2017; Driscoll and Steinert-Threlkeld 2020) or the creation of very large corpora coded by human beings (Mitts 2019; Siegel et al. 2018). The study of ideological polarization requires an assumption that the aggregation used accurately reflects the underlying social process that the analyst wishes to identify.

To address these concerns, we focus on a single type of Twitter-based behavior—retweets—and design a model that combines measurement with direct statistical analysis of group polarization. Thus, our model need make only two basic assumptions: 1) retweets are a noisy signal of underlying ideological agreement and 2) our codings of elite Twitter users contain at least some information about the concept we want to measure. The first assumption has been shown to be valid through analyses of the political content of retweet networks (Conover et al. 2011; Stieglitz and Dang-Xuan 2012; Barberá et al. 2015; Halberstam and Knight 2016) and surveys of Twitter users (Metaxas et al. 2015).⁴ Barberá (2015) showed quite convincingly that Twitter follower networks, which are plausibly a less active signal of agreement, are highly predictive of DW-NOMINATE scores for U.S. Congresspeople. As we describe later, our model implicitly accounts for follower networks because the decision to follow users is necessarily correlated with one’s ability to retweet them.

For the second assumption, we try to be as transparent as possible about our coding of elite users, as discussed in Appendix F. Because we employ a joint measurement and inference model, we need not claim that we have *true* or *error-free* measures of group affiliation as our residual measurement uncertainty will propagate through our estimates. If our groups prove to be irrelevant

4. We note that we do not need to take a position on the debate in the literature over whether Twitter encourages echo chambers. Retweet networks may result in echo chambers or more diverse groups; our model allows for either to occur empirically.

to retweet behavior, then our estimates will deflate to zero or have very large uncertainty intervals (or both).

To create this new model, we combine ideas from two distinct literatures in statistics: time-series econometrics for inference, in particular vector auto-regression, and item-response theory for measurement. We are the first to combine item-response theory with multivariate time series estimation, though prior research has employed IRT with univariate time series methods (Martin and Quinn 2002; Kropko 2013). Our employment of multivariate stationary time series techniques is necessary to capture the transnational element in ideological polarization. We want to know whether an increase in group polarization in Egypt causes changes in group polarization in Tunisia and vice versa. For that reason, we use vector autoregression (VAR) as our inference engine for measuring transnational diffusion. The VAR method has important similarities to the spatial lag model used to measure geographical diffusion and interdependence (Franzese Jr. and Hays 2007). Our main innovation in this paper is to jointly fit a VAR and an IRT model so that the VAR can fully incorporate the measurement error present in the Twitter data.

The primary purpose in employing a VAR as the inferential method is that it enables us to track the endogenous feedback effects of the ideal points of transnational groups on each other over time. To set up this model, we start with two time series that represent the latent ideal points of different religious or political groups: $y_{c,g,t}$ and $y'_{c',g,t}$. These series are observed at discrete time units t and each belongs to the same religious affiliation $g_1 \in \{Secularist, Islamist\}$ or political affiliation $g_2 \in \{Democratic, Authoritarian\}$ but different countries $c \in \{Tunisia, Egypt\}$. While we have two dimensions in our model (religion and democracy), we suppress the dimensional subscript and present the model in terms of the first dimension for simplicity (i.e., only one group g_1). In a VAR framework, we can use the following equation to signify two series (y and y') of the same group g but we separate countries (c and c') through their prior period lags. The parameters $\beta_{c,g,I}$ and $\beta_{c',g,E}$ control the relative influence of prior period lags for internal I influence (prior lag

of same series) and external E influence (prior lag of group in other country):

$$y_{c,g,t} = \gamma_{c,g} + \beta_{c,g,I}y_{c,g,t-1} + \beta_{c',g,E}y'_{c',g,t-1} + \beta_{c,g,x}X_t + \epsilon_{c,g,t} \quad (1)$$

$$y'_{c',g,t} = \gamma_{c',g} + \beta_{c',g,I}y'_{c',g,t-1} + \beta_{c,g,E}y_{c,g,t-1} + \beta_{c',g,x}X_t + \epsilon_{c',g,t} \quad (2)$$

To make the model stochastic, we include $\epsilon_{c,g,t}$ and $\epsilon_{c',g,t}$ as white noise (stationary) errors assuming that $\beta_{c,g,I}$ and $\beta_{c,g,E}$ meet the VAR stability conditions (Zivot and Wang 2006, 386-387). If these parameters meet the stability conditions, the latent ideal points will over time return to their long-run equilibrium value $\gamma_{c,g}$ (i.e., the intercept). Substantively, these parameters provide estimates of how quickly a time series will return to its long-term mean given an exogenous shock ($\beta_{c,g,I}$) and the strength of influence of the other time series ($\beta_{c,g,E}$).

We include an additional parameter in each series, $\beta_{g,c,x}$. This parameter does not vary over time, but rather represents the effect of the exogenous event X_t , which equals 1 after the coup against Morsi and 0 before the coup. As such, we can use $\beta_{g,c,x}$ as a direct measure of the long-term polarizing effect of the coup on each of the series. A null hypothesis of no effect of the coup would be the case in which $\beta_{g,c,x} = 0$.

Given that we have two groups and two countries, we have two sets each of ideal points series $y_{c,g,t}$ and $y'_{c',g,t}$ with two dimensions, which comprises a seven time-series system (Tunisian and Egyptian secularists, Islamists, democrats and authoritarians).⁵ While we could pair each series with every other series, we instead chose to pair each ideological group only with its ideological allies in the other country. We impose this restriction because we aim to identify the effect of transnational polarization, and also because the within-country groups are separately related through the IRT model that we explicate below. In other words, the within-country groups are

5. As mentioned elsewhere, we do not have Egyptian Islamists who are also authoritarian.

implicitly scaled relative to each other; we just do not add parameters that would measure the influence of prior period lags within countries across groups. Again, this is primarily for reasons of model parsimony.

However, in order to calculate all of these effects, we need to obtain the ideal point series themselves (y_t and y'_t) from relatively noisy Twitter data. For that reason, we turn to item-response theory (IRT).

We illustrate the model in Figure 1 in Appendix C. The figure shows a random assortment of citizens j and elites i in a two-dimensional latent space, with the democracy-authoritarianism scale on the y axis and Islamism-secularism on the x axis. At time $t = 0$, the elites i are in an equilibrium that corresponds to low group polarization as expressed in our theory. The elites i are located close to a middle group of citizens j who could be thought of as moderates. The counts of retweets are shown by directional ties between nodes, with the shade of the tie indicating the number of retweets sent by j to i in a given time point.⁶ From $t = 0$ to $t = 1$ an exogenous shock occurs, forcing the elites i , who are colored to represent different groups, to move away from the moderates and toward the extremes. This event reduces retweeting by moderate citizens and increases retweeting by extremist citizens. In our model, this would represent a positive (negative) value for the coup parameter $\beta_{g,c,x}$. It should be noted that it is the elites who are polarizing in the model; the citizens' ideal points are by comparison fixed over time. It is possible, of course, to consider the endogenous relocation of citizens due to group polarization, but such an extension is beyond the scope of this paper.

To be clear, both citizens and elites are taking actions that co-determine their position in the latent space. Elites release tweets, which can be thought of as being more or less polarizing in terms of their content. Citizens, whose group identities are relatively fixed over this short time

6. We thank an anonymous reviewer for providing the inspiration for this figure.

frame (approximately 9 months), retweet the elites' tweets depending on how close the tweet is to their preferred level of group polarization. Moderate users will prefer to retweet statements that are less polarizing, while more partisan users will retweet statements that are more polarizing. Because it is elites' ideal points that vary over time, the model assumes that elites are strategically updating their ideal points in response to over-time trends and external events, which is then picked up as varying retweet patterns among like-minded Twitter users. We believe this model to be appropriate to the domain under study as polarizing events such as coups are primarily occurring at elite levels.

To construct the latent inter-group distance time series, we employ as our base specification the standard 2-PL IRT model that can be used to estimate the canonical ideal point model (Clinton, Jackman, and Rivers 2004). Formally, we use this model to predict the mean of retweet counts $R_{c,g,t,j}$:

$$R_{c,g,t,j} \sim \text{Pois}(\delta_j \alpha_{c,g,t} - \beta_j) \quad (3)$$

In this model, $\alpha_{c,g,t}$ represent the latent ideal points of all the elites in each ideological group-country combination at each time point t , while δ_j represents how strongly ideological citizen j 's retweet pattern is and β_j is a citizen-specific intercept. We know from Clinton, Jackman, and Rivers (2004) that we can interpret this model as citizens choosing to retweet an elite if and only if that elite's ideal point in a latent space is closer to the citizen's ideal point in the latent space than the ideal point of any other elite.

In order to estimate this model, we situate equation 3 in a Bayesian framework in which we define θ as the full set of parameters we can estimate in (3), and we want to know the most likely values of θ conditional on the observed data $R_{c,g,t,j}$:

$$p(\theta | R_{c,g,t,j}) \propto p(\theta) p(R_{c,g,t,j} | \theta) \quad (4)$$

Using this standard form of Bayesian inference, equation 3 becomes the likelihood $p(Y_{c,g,j,t}|\theta)$, and we can then look at endogenous relationships between ideal point parameters via the priors of these parameters, $p(\theta)$. In particular, building on Martin and Quinn (2002) and Kropko (2013), we can model the vector auto-regression between the ideological groups g via priors on the ideal points $\alpha_{c,g,t}$:

$$\alpha_{c,g,t} \sim N(\gamma_{c,g} + \beta_{c,g,I}\alpha_{c,g,t-1} + \beta_{c,g,E}\alpha_{c',g,t-1} + \beta_{g,c,x}X, \sigma_{c,g}) \quad (5)$$

Equation 5 shows how any one elite group $\alpha_{c,g,t}$'s latent score in time t is a function of its prior time period latent score, $\beta_{c,g,I}\alpha_{c,g,t-1}$, and the latent score of the same group g but opposite country c' in the previous time period, $\beta_{c,g,E}\alpha_{c',g,t-1}$. As can be seen relative to equation 1, equation 5 substitutes the observed time series y_t and y'_t with the latent ideal scores $\alpha_{c,g,t}$, but otherwise has the same parameters $\beta_{c,g,I}$ and $\beta_{c,g,E}$. In other words, we use the IRT model to construct the time series by estimating the latent positions of the elite actors, but we are also able to directly estimate parameters of interest even with this measurement uncertainty. Because these priors are multiplied with the likelihood $p(R_{c,g,t,j}|\theta)$, we can then estimate a full joint density of both the IRT model and the vector autoregression between latent ideological positions so that uncertainty in both models is appropriately captured.

There are two other notable features of equation 5. First, we include an exogenous regressor $\beta_{g,c,x}X$. $\beta_{g,c,x}X$ represents a binary vector that equals 0 before a polarizing event occurs, and 1 afterward, so that we can directly measure the effect of polarizing events on the ideal points $\alpha_{c,g,t}$. We summarize these important parameters and their interpretation in Table 1.

To achieve identification, we constrain one of the intercepts $\gamma_{c,g}$ for one group in each dimension to be equal to +1 and -1 (see Bafumi et al. (2005) for a full discussion of ideal point identification). In addition, we fix one of the variance parameters $\sigma_{c,g}$ to .1 to identify the scale of the ideal points

because the addition of the over-time variation adds another dimension of potential multi-modality in the model’s joint posterior distribution.

Table 1: Important Parameters of Interest in IRT-VAR Model

Parameter	Meaning
$\beta_{c,g,x}$	Effect of the coup X on a given country-group’s ideal points. Exogenous to time process.
$\beta_{c,g,I}$	Over-time trend of a given country-group’s ideal points (auto-regressive parameter).
$\beta_{c,g,E}$	Over-time trend of the ideological ally group in a foreign country (auto-regressive parameter).

At this point, we have defined the IRT-VAR model that allows us to make time-series inferences on the over-time changes in the elite ideal points $\alpha_{c,g,t}$. However, this model is only defined over the retweet counts in which we have observed a citizen j retweet one of the elites in a specific time period t . As mentioned in the previous section, we expand our observed data to include all the times that each citizen j *does not* retweet an elite in each time period, or $R_{c,g,j,t} = 0$. If we simply include those zeroes in our likelihood $L(R_{c,g,j,t}|\theta)$ as additional data, we will be making the very strong assumption that each citizen j looked at every tweet from every elite in time period t and decided whether or not to retweet each tweet. In fact, that assumption may not hold for any of the citizens in our data except for unusually thorough citizens who want to have all influential Twitter users in their feed. As a result, we are concerned about a form of selection bias in which citizens are only exposed to tweets from those elites who are ideologically proximate to them, either because 1) Twitter’s recommendation algorithm suggests that they follow elites who are ideologically proximate or 2) the citizen prefers a network full of ideological allies or 3) both of these factors in interaction.

For these reasons, we need a separate likelihood for the case when $R_{c,g,j,t} = 0$. To do so, we incorporate the missing-data mechanism employed by Kubinec (2019), in which a hurdle model is used to account for missing data in an ideal point model when missingness is plausibly a function

of the value of a person's ideal points. Given that this missingness pattern is very likely present in our data for the reasons previously described, we define a new likelihood $L(Y_{c,g,t,j} = 0|\lambda)$ conditional on a different set of parameters λ :

$$L(R_{c,g,j,t} = 0|\lambda) = \prod_{C=1}^c \prod_{G=1}^g \prod_{J=1}^j \prod_{T=1}^t \text{logit}^{-1}(\delta_{Aj} \alpha_{c,g,t} - \eta_{Aj}) \quad (6)$$

What should be noted about this model is that we now have a Bernoulli-distributed random variable $Y_{c,jgt} \in \{0, 1\}$, and so we predict this probability using a logit link function of the parameters in λ . These parameters represent a separate IRT equation with the same elite ideal points $\alpha_{c,g,t}$ but separate discrimination parameters for the citizens δ_{Aj} . These separate set of citizen parameters represent a citizen's latent willingness to look at tweets from across ideological spectrum that is independent of that citizen's own individual ideal point, or what could be thought of as that citizen's desire to be informed of tweets from different points of view. We include citizen intercepts η_{Aj} that represent missingness that is ignorable, which will occur if $\delta_{Aj} = 0$ and the probability of a citizen seeing a tweet is equal to that citizen's average number of tweets per group for the sample period. Ignorable missingness can occur if, for example, a citizen does not see tweets because they are working or on vacation. Importantly, in either case, the elite ideal points $\alpha_{c,g,t}$ that are our focus of inference will adjust to the uncertainty in this first-stage selection model. Our revised Bayesian model can then be written as follows in terms of the joint distribution of $R_{c,jgt} = 0$, $R_{c,jgt} \neq 0$, λ and θ :

$$p(\theta, \lambda | R_{c,jgt} \neq 0, R_{c,jgt} = 0) \propto p(\theta)p(\lambda). \quad (7)$$

$$[L(R_{c,jgt} = 0|\lambda) + (1 - L(R_{c,jgt} = 0|\lambda))L(R_{c,jgt} \neq 0|\theta)] \quad (8)$$

To finish our model specification, we include here our prior distributions for all other parameters

in the model:

$$\eta_{Aj} \sim N(0, 3) \quad (9)$$

$$\delta_j \sim N(0, 3) \quad (10)$$

$$\delta_{Aj} \sim N(0, 3) \quad (11)$$

$$\gamma_{cg} \sim N(0, 3) \quad (12)$$

$$\beta_{c,g,x} \sim N(0, 5) \quad (13)$$

$$\beta_{c,g,I} \sim N(0, 1) \quad (14)$$

$$\beta_{c,g,E} \sim N(0, 1) \quad (15)$$

$$\sigma_{c,g} \sim E(1) \quad (16)$$

Crucially, in addition to the summary estimate of $\beta_{g,c,x}$ for the exogenous covariate, we can also use the values of $\beta_{c,g,I}$ and $\beta_{c,g,E}$ to calculate impulse-response functions (IRFs) for a shock to the elite group's ideal points coming from the group's own time series, or the indirect effect coming from a shock to a different group's time series. We can express this mathematically as the derivative of an exogenous shock $\beta_{c,g,x}$ with respect to the value of the ideal point $\alpha_{c,g,t}$ at time points after the shock from $t + n, n \in \{1, 2, \dots, 10\}$ while incorporating over-time domestic trends $\partial\beta_{c,g,I}$:

$$\frac{\partial\alpha_{c,g,t+n}}{\partial\beta_{c,g,t+n}\partial\beta_{c,g,I}} \quad (17)$$

This IRF essentially measures the decaying (if the time series is stable) average effect of a shock to the latent ideal points over time, and provides a straightforward measure of the effect of our explanatory variable on the outcome over time. We can also use this same framework to calculate another important IRF of interest, which is how the effect of the coup on the transnational group will affect the domestic group's ideal points given external influence $\partial\beta_{c,g,E}$:

$$\frac{\partial \alpha_{c,g,t+n}}{\partial \beta_{c',g,x} \partial \beta_{c,g,E}} \quad (18)$$

Each of these effects incorporates the endogenous nature of these processes, allowing for heightened polarization in prior time periods to influence the present. To summarize the model, then, we can match these estimates to our hypotheses to provide very specific null hypothesis tests of our arguments that, as we have noted, incorporate our measurement uncertainty in using Twitter data. We show how this model relates to our hypotheses in Table 2.

Table 2: Hypotheses and Associated Tests from our Model

Hypothesis	Definition	Expected Outcome
H1	After the coup against Mohamed Morsi, the difference in latent ideological positions between Islamists and secularists and democrats and authoritarians will diverge in direct reaction to the coup (direct effect).	$\frac{\partial \alpha_{c,g,t+n}}{\partial \beta_{c',g,x} \partial \beta_{c,g,I}} \neq 0$ for all $c \in C$ and $g \in G$
H2	After the coup against Morsi, the difference in latent ideological positions between Islamists and secularists and democrats and authoritarians will diverge from each group's reaction to their ideological allies' shift in latent ideological position (indirect effect).	$\frac{\partial \alpha_{c,g,t+n}}{\partial \beta_{c',g,x} \partial \beta_{c,g,E}} \neq 0$ for all $c \in C$ and $g \in G$

We return now to our discussion of the coup against Morsi and its exogeneity with reference to the model parameters. We note that we can only infer exogeneity from the timing of the coup, i.e., the particular day that the military chose to surround President Morsi's palace and remove him from office. The coup itself was not a surprise in that large anti-Morsi protests erupted in several cities in late June, answered by large pro-Morsi counter-demonstrations. Thus we can assume that the timing of the coup is exogenous. Furthermore, actions taken after the coup by both groups, including demonstrations and various political statements, are not exogenous to polarization. As such, we limit our claim of any causal identification of the coup parameter $\beta_{c,g,x}$ to several days following the actual coup date. As we show in the results section, given the decay in the over-time

effects measured in days, we can make the qualified inference that this movement was likely due to the coup, and not to some other factor that could explain both polarization and the coup. We note as well that the model explicitly incorporates endogenous trends in polarization over time, and so the mere fact that Egypt and Tunisia were relatively polarized prior to the coup does not jeopardize inference. If the coup did not have any independent effect on group polarization, its estimate would be zero as residual polarization would be picked up by the endogenous parameters in the model. Put another way, we are trying to explain the timing of the sharp increase in polarization on and immediately after July 3.

Second, we note that it is only the coup parameter $\beta_{c,g,x}$ that can be plausibly interpreted as *causally* identified. The other parameters that measure direct and indirect influence, $\beta_{c,g,I}$ and $\beta_{c,g,E}$, are not causally identified given the exogeneity of the coup's timing. The reason why we do not infer causal identification is because the model measures these channels of polarization, but does not assign a causal interpretation. In other words, when we refer to a "missing counterfactual", we mean in terms of measurement, i.e., being able to track transnational group polarization over time and separate domestic versus international forms of polarization. The particular mechanisms and variables underlying these processes are not identified in the model, such as the types of strategic communications sent by groups as they adjust to new equilibria in group salience. Our model provides insight into how we can measure transnational polarization. It does not provide inferences as to how we can manipulate these different channels of transmission, although we hope further research may provide more evidence on these questions.

Model Results

Estimating this model using full Bayesian inference would be computationally prohibitive given the volume of the data and the tens of thousands of model parameters. One possibility is to employ

variational inference, which simplifies the joint posterior into factorized distributions that are much easier to sample. However, variational inference can understate uncertainty in its estimates, and the approximation can fail to adequately describe the true posterior (Grimmer 2011). Instead we employed a novel technique recently developed by the Stan team for the parallelization of Markov chains. We used code that parallelized the gradient calculations necessary for each iteration of a Hamiltonian Markov sampler and are the primary computational burden. Employing a cluster computer system with 700 cores, we were able to estimate a converged chain within 48 hours.⁷ Based on conversations with the Stan team in August 2019, we believe that this is the largest such model estimated to date. Our full Bayesian estimation produced 6,134 each of discrimination δ_j for all of the citizens in the model, but we will focus on the four group parameters that varied over time, $\alpha_{c,g,t}$, with one for each country-ideological pairing: Tunisian and Egyptian Islamists-secularists and democrats-authoritarians, as these are our primary focus for inference.

In Appendix D we show the estimated ideal points at each time point for the seven ideological groups as a scatter plot with both dimensions. To examine the overall time trends, we plot each set of dimensions and groups separately in Figure 2 along with vertical lines showing when the following events occurred: 1) the military coup against Morsi, and 2) the assassination of the secular leftist Tunisian politician Mohamed Brahmi. Increasing group polarization is evidenced by increasing distance between the group-level ideal points, while decreased distance is evidence of decreasing group polarization. It should be noted that only Morsi’s coup was explicitly parameterized in the model; the spikes in polarization occurring around the Brahmi’s assassination emerged endogenously from the model. The confidence intervals on the chart reflects the 5%-95% quantiles

7. To optimize computer time, we ran each chain for a total of 550 iterations, 300 of which were discarded as warmup iterations. While this number of iterations is smaller than other samplers, such as Gibbs samplers, it is sufficient for Hamiltonian Monte Carlo to reach convergence and exceeds the recommendations of the Stan manual (minimum 100 warmup iterations). All parameter estimates had Rhat values 1.1 or lower, and more than 99% of parameters had Rhat values lower than 1.05.

of the empirical posterior, hereafter referred to as the high-posterior density (HPD) interval.

From this chart we can make useful descriptive inferences. The average position over time of the time series gives us a sense of the general level of polarization in the data during this time period. First, the sectarian groups are much closer to each other than to the opposing group of the same country. These average locations would suggest that there is substantial ideological similarity in these groups as expressed in their retweet patterns. Second, polarization following the anti-Morsi coup is easily evident on the chart without need to examine the model's coefficients. Islamists and secularists both diverged after the coup in opposite directions. Interestingly, while Islamists moderated somewhat over time, they remained closer than they were before the coup, while secularists start to diverge from each other after the coup. While some of these findings, such as the secularist divergence, invite speculation, the plot as a whole shows how it can capture many of the subtleties of this form of polarization graphically. Furthermore, the method is able to identify events that were particularly polarizing, especially the assassination of the secularist Brahmi, without prior parameterization.

It is worthwhile to compare Figure 2 with Figure 1 that show raw counts of retweets for these same ideological groups. While Figure 1 had spikes around notable events like the anti-Morsi coup, the estimated ideal points in Figure 2 generally do not share these spikes and certainly not in the same proportions. The reason for this disparity is the work of the IRT-VAR measurement model, which only picks up changing proportions of retweets among users in terms of the ideological groups of interest. This disparity between raw data and the measured ideal points shows how the model is separating the chaff from the wheat by identifying salient trends in the data. For example, simple amplification of individual users, as happens due to Twitter "bots", will not affect the results of the model, as it is the relative weight of *ideological* retweeting between groups that is being estimated.

Figure 2: Over-Time Trends for Ideal Points of Secularists and Islamists (First Dimension)

We can similarly show the trends over time for our pro-democracy and anti-democracy dimension in Figure 3. These results are remarkably different than the trends for the religious dimension. First, we would note that the groups are more clustered in terms of their national versus sectarian identity. In other words, Tunisian and Egyptian democrats are more “democratic” than their country counter-parts, but Tunisians on the whole are more democratic than Egyptians. This scaling artifact should not be interpreted literally; rather, the retweet patterns along this dimension are more distinct geographically than retweet patterns for the religious dimension. Interestingly, we do

not see as much movement in general for Tunisians along this axis except for a pronounced spike following Brahmi's assassination. Egyptians, on the other hand, become more pro-democratic in their retweeting patterns over time.

Figure 3: Over-Time Trends for Ideal Points of Secularists and Islamists (First Dimension)

Again, we need to be careful in interpreting these charts, but the over-time trends would suggest that both sides of the debate over democracy in Egypt were trying to reclaim the democracy space following the coup. This movement is likely due to the need to claim legitimacy arising from

democratic discourse even as (or because of) democratic norms are violated. Indeed, even though the military regime in Egypt is seen as one of its most brutal in its history, the regime has faithfully implemented elections for the legislature and the executive, and has also held referenda. As such, promoting pro-democratic discourse may have become paradoxically *more* important following the coup. We would note, however, that even as both groups in Egypt exhibited a pro-democratic shift, the net distance between them expressed in retweet patterns increased following the coup.

Following this descriptive analysis of the results, we turn to our inference on our parameters measuring the effect of the anti-Morsi coup. To do so, we present an interpretation of the marginal effect of the coup on retweet counts in the first dimension (religion) in Figure 4 and the second dimension (democracy) in Figure 5. These marginal effects are weighted by the discrimination values, or the estimated ideological affiliation of the citizens in this case. These marginal effects were created by averaging the exponentiated effect of the coup $\beta_{c,g,x}$ on the retweet count $R_{c,g,t,j}$ over all of the positive and negative citizen discrimination parameters δ_j . These calculations are performed separately for each posterior draw $s \in S$ to capture uncertainty in the estimate:

$$\frac{\partial \hat{R}_{c,g,t,j}}{\partial \beta_{c,g,x} | \delta_j > 0} = \frac{\sum_{s=1}^S \frac{\sum_{j=1}^{J_{\delta_j > 0}} e^{\beta_{c,g,x,s} \delta_{j,s}}}{J}}{S} \quad (19)$$

$$\frac{\partial \hat{R}_{c,g,t,j}}{\partial \beta_{c,g,x} | \delta_j < 0} = \frac{\sum_{s=1}^S \frac{\sum_{j=1}^{J_{\delta_j < 0}} e^{\beta_{c,g,x,s} \delta_{j,s}}}{J}}{S} \quad (20)$$

These marginal effects are a way of interpreting the effect of the coup in terms of the observed retweet counts $R_{c,g,t,j}$ that takes into account the total level of ideological polarization among Twitter users. Given that the outcome is modeled using the Poisson distribution, the exponentiation of the estimate allows us to interpret the coefficient as percentage change in retweet counts. These percentages represent the increase or decrease in retweets coming from a typical Islamist or secu-

larist user. In other words, as a result of weighting by the discrimination parameters, the estimate is *weighted by the underlying level of polarization among citizens*. For this reason, these ideal point marginal effects are a way to capture our full measurement uncertainty in the latent scale while providing a digestible “real world” interpretation. As can be seen in the Figures 4 and 5, there are two sets of marginal effects, one set for each end of the latent scale. An alternative way of phrasing the coefficients would be as the effect of the coup in encouraging retweets from ideologically like-minded and ideologically opposed citizens.

As can be seen in Figure 4, the coup had very strong effects on the retweet patterns of Islamist users. Elite Islamists in Egypt after the coup received approximately 50% more retweets from Islamists and 25% fewer retweets from secularists. These effects are not identical because there are weighted by the total level of underlying polarization, or ideological identification, of these users. We see that the coup had the second largest effect on polarization among secularists in Egypt with similar increases in retweets from ideologically similar users. The effects are weaker for Tunisian groups, as might be expected given that the coup occurred in Egypt. In addition, the effects are much stronger for Tunisian Islamists than for secularists. Regardless of size, all of the effects point in a polarizing direction, e.g., more retweets coming from ideologically similar users and fewer from ideologically dissimilar users.

As might be expected from a perusal of the over-time trends for the second dimension, Figure 5 shows much weaker effects of the coup on ideological (de-)polarization for the democratic-authoritarian axis. First, we see that the effects of the coup are no longer uniformly polarizing. As could be seen visually in Figure 3, both Egyptian groups received more retweets from pro-democratic users and decreased retweets from anti-democratic users. The effects of the coup on Tunisian groups, on the other hand, were too weak to be statistically distinguishable. The model provides only limited evidence that authoritarian Tunisians became more authoritarian and pro-democratic Tunisians became modestly more pro-democratic, but the opposite conclusion is still

Figure 4: Marginal Effects of Coup Effect $\beta_{c,g,x}$ on Retweet Counts Weighted by Citizens' Religious Affiliation

Figure 5: Marginal Effects of Coup Effect $\beta_{c,g,x}$ on Retweet Counts Weighted by Citizens' Pro-Democratic Inclinations

plausible. On the whole, these figures reveal that most of the polarization occurring as a result of the coup is along the religious identity axis, although there is interesting tandem movement in the pro-democracy direction for Egyptian groups as we described earlier based on visual analysis.

Finally, we can examine whether or not there was additional influence on the ideological locations of groups due to *transnational* influence. To do so, we calculate the impulse-response functions for a ten-day window following the coup. As explained previously, these impulse response functions (IRFs) show the average over-time decay of the effect of the coup on the ideal points of different groups. By altering the parameters of the simulation used to calculate these effects, we can separately identify the *direct* effect of the coup on a group and the *indirect* effect of the coup on a group's transnational allies transmitting across borders. Doing so is straightforward as it involves setting one of the autoregressive lag coefficients $\beta_{c,g,t,I}$ for internal influence or $\beta_{c,g,t,E}$ for external influence equal to zero. The results of these direct (light blue) and indirect (dark blue) effects for the religious dimension are shown in Figure 6, while the same IRFs for the democracy dimension are shown in Figure 7. The scale in these figures is taken from movement in the underlying ideal point scale.

What is immediately apparent from these figures is that for three out of the four groups, the indirect and direct effects are mirror images of each other. Structurally, the Egyptian Islamists show a different type of over-time autocorrelation as their ideal points oscillate back and forth over time. This behavior is a result of a negative autoregressive parameter $\beta_{c,g,t,I}$ that allows the series to oscillate. It is difficult to ascribe a clear meaning to this oscillation, although it should be noted that the Egyptian Islamists exhibit the largest amount of over-time change (i.e., variance) relative to other groups. It could be a sign of the stress that the group is under during this time.

However, whether the time series exhibit oscillation or stable over-time decay, the indirect effects counter-balance the direct effects. What this result implies is that these ideological groups

Figure 6: Indirect and Direct Effects of Coup Over Time on Islamists and Secularists

were *de-polarized* by the polarization occurring among their transnational allies. The effects are strongest for the Egyptian Islamists. As their counterparts in Tunisia polarized due to the coup, Egyptian Islamists de-polarized. This inference is made possible because the model permits us to examine crucial unobserved counterfactuals: if the coup had occurred but the Tunisian Islamists had not existed, what would have been the effect of the coup on Egyptian Islamists? The answer is that Egyptian Islamists would have ended up *even more* polarized than they did in fact become.

Figure 7: Indirect and Direct Effects of Coup Over Time on Democrats and Authoritarians

The pattern is similar for secularists in both countries, though the effects are relatively weaker.

Fascinatingly, Tunisian Islamists show little indirect influence coming from their Egyptian ideological allies. This minimal influence occurred because the indirect influence parameter for the Tunisian Islamists is estimated very close to zero with a posterior mean of 0.017 (HPD -0.088, 0.112). While the Tunisian Islamists reacted to the *coup* very strongly, they did not react at all to the polarization occurring between their transnational allies and secularist adversaries. Again, it is

the model that permits us to test this important counterfactual in a highly endogenous system: how would the Tunisian Islamists have reacted if the Egyptian Islamists had not existed? As it turns out, not very differently at all.

Finally we turn to the similar effects for the democracy dimension in Figure 7. Just as the effect of the coup was weaker for this dimension, so the indirect and direct simulations show less information. The only pronounced indirect effect is for Tunisian democrats, who became slightly more authoritarian due to the movement of their ideological allies. This finding suggests that there is indeed some kind of strategic positioning at work in democratic discourse among democrats and authoritarians following the coup, with the definition of democracy itself a contested value.

Discussion

Explanation of the exact trajectory of the ideal points over time will require further specification of the sometimes conflicting determinants of group polarization that we present in this paper. For example, in Tunisia, Islamists were under considerable pressure to dampen ideological conflict during this time period as they faced rising social unrest due to Islamist radical violence (McCarthy 2016). For that reason, after the coup in Egypt they may have feared supporting their co-religionists too publicly, lest they suffer a similar fate within their own country. This kind of suppression, or what we call de-polarization, may explain why Egyptian Islamists de-polarized in response to witnessing polarization of Tunisian Islamists and secularists.

An important question stemming from our results is whether and to what extent the diffusion of sectarianism is driven by exposure to Twitter as a medium versus other technologies and forms of social communication. Our results do not require Twitter to be the medium of polarization as we only use Twitter as a measure of underlying group polarization. As we described in our in-

formal theory, group polarization is driven by information about polarizing events, which could travel through offline or online networks. However, as we show in Appendix E, there is substantial overlap across countries among the Twitter followers of elite users, and Islamist elites such as Rached Ghannouchi and the Muslim Brotherhood's official account have the highest number of foreign followers, suggesting that Islamist networks are more integrated across territorial boundaries. These descriptive statistics, combined with the results of our ideal point analysis, suggest that Twitter itself is likely a source of group polarization even as it is a useful tool for measuring it by permitting transnational groups to become more responsive. Nonetheless, fully substantiating the share of group polarization attributable to Twitter would require broadening the analysis to account for other types of media, a potentially fruitful area for future research.

Conclusion

For all of the theoretical and empirical progress in the study of ideological diffusion, lacking have been studies demonstrating the spread of contention across national boundaries that account for pre-existing contention in the receiving country. In this paper we address that gap by advancing a method of estimating the latent positions of ideological groups in Egypt and Tunisia during the tumultuous period of the Arab Uprising. We exploit Twitter's widespread usage to provide inference on a crucial counterfactual: how ideological groups would react to domestic polarization and transnational diffusion of polarization separately.

The precision of the hypothesis tests we are able to implement in this paper enables us to identify the direct and indirect effects of group polarization. Being able to separate these different components of the feedback process allows us to substantiate the major elements of the theory, and to support the central point of our paper that transnational linkages among ideological groups can endogenously heighten or dampen polarization independent of what is occurring within each

group's country. While we are not the first to document such linkages, we are the first to identify this kind of transnational ideological polarization in a way that incorporates our uncertainty in measuring latent social cleavages.

References

- Bafumi, Joseph, Andrew Gelman, David K. Park, and Noah Kaplan. 2005. "Practical Issues in Implementing and Understanding Bayesian Ideal Point Estimation." *Political Analysis* 13 (2): 171–187.
- Barberá, Pablo. 2015. "Birds of the Same Feather Tweet Together: Bayesian Ideal Point Estimation Using Twitter Data." *Political Analysis* 23:76–91.
- Barberá, Pablo, John T. Tost, Jonathan Nagler, Joshua A. Tucker, and Richard Bonneau. 2015. "Tweeting from Left to Right: Is Online Political Communication More Than an Echo Chamber?" *Association for Psychological Science* 26 (10): 1531–1542.
- Clinton, Joshua, Simon Jackman, and Douglas Rivers. 2004. "The Statistical Analysis of Rollcall Data." *American Political Science Review* 98 (2): 355–370.
- Conover, M. D., J. Ratkiewicz, M. Francisco, B. Goncalves, A. Flammini, and F. Menczer. 2011. "Political Polarization on Twitter," *Association for the Advancement of Artificial Intelligence*.
- Dempsey, Judy. 2013. "Countdown to a Coup d'Etat in Egypt?" *Carnegie Europe*.
- Driscoll, Jesse, and Zachary C Steinert-Threlkeld. 2020. "Social media and Russian territorial irredentism: some facts and a conjecture." *Post-Soviet Affairs* 36 (2): 101–121.

- Egypt and the Coup: Inside the 11 Days that Toppled Morsi*. 2018. Working Paper. Middle East Eye, July 4. <https://www.middleeasteye.net/news/egypt-and-coup-inside-11-days-toppled-morsi>.
- Franzese Jr., Robert J., and Jude C. Hays. 2007. "Spatial Econometric Models of Cross-Sectional Interdependence in Political Science Panel and Time-Series-Cross-Section Data." *Political Analysis* 15:140–164.
- Freelon, Deen, Marc Lynch, and Sean Aday. 2015. "Online Fragmentation in Wartime: A Longitudinal Analysis of Tweets about Syria, 2011-2013." *The Annals of the American Academy of Political and Social Science* 659:166–179.
- Gerling, Lena. 2017. "Riots and the Window of Opportunity for Coup Plotters: Evidence on the Link between Urban Protests and Coups d'État." *Center for Interdisciplinary Economics*.
- Grimmer, Justin. 2011. "An Introduction to Bayesian Inference via Variational Approximation." *Political Analysis* 19 (1): 32–47.
- Hafer, Catherine, and Dimitri Landa. 2006. "Deliberation and Social Polarization." *SSRN*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=887634.
- Halberstam, Yosh, and Brian Knight. 2016. "Homophily, group size, and the diffusion of political information in social networks: Evidence from Twitter." *Journal of Public Economics*, no. 143: 73–88.
- Jamal, Amaney A., Robert O. Keohane, David Romney, and Dustin Tingley. 2015. "Anti-Americanism and Anti-Interventionism in Arabic Twitter Discourses." *Perspectives on Politics* 13 (1): 55–73.
- Kearney, Michael W. 2019. "rtweet: Collecting and analyzing Twitter data." *Journal of Open Source Software* 4 (4): 1829. doi:10.21105/joss.01829.

- Kropko, Jonathan. 2013. *Measuring the Dynamics of Political Power: a Time-Series IRT Model*. Presentation. Political Methodology Society.
- Kubinec, Robert. 2019. "Generalized Ideal Point Models for Time-Varying and Missing-Data Inference." *Open Science Foundation Preprints*. doi:10.31219/osf.io/8j2bt.
- Luttwak, Edward. 1968. *Coup d'Etat: A Practical Handbook*. Harvard University Press.
- Martin, Andrew D., and Kevin M. Quinn. 2002. "Dynamic Ideal Point Estimation via Markov Chain Monte Carlo for the U.S. Supreme Court, 1953-1999." *Political Analysis* 10 (2): 134–153.
- Mason, Lilliana. 2014. "'I Disrespectfully Agree': The Differential Effects of Partisan Sorting on Social and Issue Polarization." *American Journal of Political Science* 59 (1): 128–145.
- McCarthy, Rory. 2016. "The Tunisian Uprising, Ennahda and the Revival of an Arab-Islamic Identity." In *Political Identities and Popular Uprisings in the Middle East*, edited by Shabnam J. Holliday and Philip Leech, 157–176.
- Metaxas, Panagiotis, Eni Mustafaraj, Kily Wong, Laura Zeng, Megan O'Keefe, and Samantha Finn. 2015. "What Do Retweets Indicate? Results from User Survey and Meta-Review of Research." In *Ninth International AAAI Conference on Web and Social Media*.
- Metzger, Megan MacDuffee, and Joshua Tucker. 2017. "Social Media and EuroMaidan: A Review Essay." *Slavic Review* 76 (1): 169–191.
- Meyersson, Erik. 2016. "Political man on horseback: coups and development." *Stockholm Institute for Transition Economics (SITE)*.
- Mitts, Tamar. 2019. "From Isolation to Radicalization: Anti-Muslim Hostility and Support for ISIS in the West." *American Political Science Review*, 113th ser., 1:173–194.

- Moulaert, Frank, and Farid Sekia. 2003. "Territorial Innovation Models: A Critical Survey." *Regional Studies* 37 (3): 289–302.
- Siegel, Alexandra A., Joshua A. Tucker, Jonathan Nagler, and Richard Bonneau. 2018. *Tweeting Beyond Tahrir: Ideological Diversity and Political Intolerance in Egyptian Twitter Networks*. Working Paper. New York University, March.
- Singh, Naunihal. 2014. *Seizing Power: The Strategic Logic of Military Coups*. Johns Hopkins University Press.
- Steinert-Threlkeld, Zachary C. 2017. "Spontaneous Collective Action: Peripheral Mobilization During the Arab Spring." *American Political Science Review* 111 (2): 379–403.
- Stieglitz, Stefan, and Linh Dang-Xuan. 2012. "Political Communication and Influence through Microblogging—An Empirical Analysis of Sentiment in Twitter Messages and Retweet Behavior." In *2012 45th Hawaii International Conference on System Sciences*.
- Sunstein, Cass R. 2002. "The Law of Group Polarization." *The Journal of Political Philosophy* 10 (2): 175–195.
- Weber, Ingmar, Venkata R. Kiran Garimella, and Alaa Batayneh. 2013. "Secular vs. Islamist Polarization in Egypt on Twitter." In *2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)*. August 25.
- Zivot, Eric, and Jiahui Wang. 2006. *Modeling Financial Time Series with S-PLUS*. Springer.