Online Appendix to:
"Do Islamic State's Deadly Attacks Demotivate, Deter, or Mobilize Supporters?"

Contents

A Formal modelling of the response to an IS attack

The purpose of this section is to analytically identify the sources of the change in the number of the followers of the reported accounts as a response to the IS attacks. As already proposed, the observed number of the followers consists of the group of leaders, as well as actively and passively involved people who are not in a leadership position; furthermore, some followers are just "observers": terrorism fighters, analysts or journalists. Hence, we model the mechanism of transition between those categories and non-followers. In this section, we provide a formal justification for our main hypothesis and show that finding a negative effect of the IS attack on the number of Twitter followers indicates that the de-mobilization effect of the attacks is stronger than the mobilization effect across ground boundaries.

Importantly, the following model contains five categories while empirically we can discriminate only between followers and non-followers. We use here five categories rather than three because we unpack insiders of the organization in tree categories of intensity: leaders, active, and passive members. Having a richer classification in the formal model than in the data is crucial as it enables to underline three key points. First of all, this model proposes a simple way to look into the general dynamics of the online extremist mobilization on Twitter. Second, as a result of the first point and crucially for our main argument, it allows showing explicitly that the existence of the hidden within-group mobilization dynamics does not contradict our hypotheses about the mobilization dynamics given our data. Third, the model explicitly links the broader formal setting to our causal empirical claims, providing additional evidence for them.

A.1 Individual transition model

Based on the classification in the previous section, here we describe the individual likelihoods of the transition between the categories. The stochastic matrix (Table A.1) formalizes our substantive assumptions about the types of the followers and their likely responses to a terrorist attack by IS. We design the movements across the categories as a Markov process with no memory - only the current state affects the possible dynamics. Clearly, from the substantive perspective, the history influences the current position as well. However, since all that information is embedded in the present state the lagged variables are omitted. The model has five individual states, two of which are fixed-points: "leader" and "observer." Meanwhile, three other states enable movements up or down one category regarding the individual mobilization.

Table A.1: Individual category transition stochastic model

	leader group	active	passive	non-followers	observers
leader group	1	0	0	0	0
active	α_0	α_1	α_2	0	0
passive	0	β_0	β_1	β_2	0
non-followers	0	0	γ_0	γ_1	γ_2
observers	0	0	0	0	1

Each row shows the non-zero probabilities of the member of a particular group to transition to another category or to stay in the same category: $\alpha_0 + \alpha_1 + \alpha_2 = 1$, $\beta_0 + \beta_1 + \beta_2 = 1$, and $\gamma_0 + \gamma_1 + \gamma_2 = 1$. First, we assume that the current IS leaders do not respond in any way to the attack. That is why the only non-zero value on the first row is on the diagonal. Contrary to them, the active supporters might become more or less mobilized as a response to the attack, or they might keep the same level of interest. Most importantly, we assume that they will be still among the detected followers even if their support becomes milder. The categories providing the variation in the observed numbers of the followers are passive supporters and former non-followers. Talking of the passive supporters, some of them might radicalize. Meanwhile, some may dislike what they see or become afraid of being tracked as connected to IS and stop following IS. The observers, who are most likely IS enemy fighters or journalists following IS, will not change their behavior as a response to the attack. Finally, some of the non-followers might become passive supporters as the terrorist attack makes IS more visible; this would imply a broadening of the base of supporters of the organization—what we call the effect of mobilizing outsiders.

A.2 Group comparative statistics

Let's denote: $x = (x_1, x_2, x_3, x_4, x_5) =$ (leader group, active, passive, non-followers, observers). Our data enables us to observe the total number of the followers of the reported accounts. This is the exact upper bound of the actual total number of the people in categories 1-3 and 5, since some reported Twitter accounts may have common users following them. We assume that the correlation between the observed upper bound and the total number of the distinct followers is approximately the same before and after the attack. Hence, if we observe the decrease in the exact upper bound after the attack, it indicates a decline in the total number in categories 1-3 and 5.

To sum up, the total observed number of the followers as: $x = x_1 + x_2 + x_3 + x_5$, where x_i denotes the number of the people in group i (x_4 are the non-followers before a terrorist attack). We need to estimate $E(x^*)$ the expected number of the observed followers after the attack based on Table A.1. Given that every follower is independent of one another, the estimates for the categories after the attack are:

$$E(x^*) = A^{'}x^{'} = \begin{pmatrix} x_1 + \alpha_0 x_2 \\ \alpha_1 x_2 + \beta_0 x_3 \\ \alpha_2 x_2 + \beta_1 x_3 + \gamma_0 x_4 \\ \beta_2 x_3 + \gamma_1 x_4 \\ \gamma_2 x_4 + x_5 \end{pmatrix}^{'}$$

Summing up, what we observe after an attack:

$$x_1^* + x_2^* + x_3^* + x_5^* = x_1 + x_2 + (\beta_0 + \beta_1)x_3 + (\gamma_0 + \gamma_2)x_4 + x_5$$

Hence, the observed change is:

$$\sum_{i=1,2,3,5} x^* - \sum_{i=1,2,3,5} x = \delta = -\beta_2 x_3 + (\gamma_0 + \gamma_2)x_4 \tag{1}$$

where $\beta_2 x_3$ is *the de-mobilization effect*, $\gamma_0 x_4$ is *the effect of mobilizing outsiders*.

Hence:

$$\delta < 0 \implies \{\gamma_2 x_4 > 0\} \implies \beta_2 x_3 > \gamma_0 x_4 \tag{2}$$

Importantly, because of $\gamma_2 x_4 \neq 0$ if $\delta > 0$ we are not able to evaluate the relation between *the* *external de-mobilization effect*, $\beta_2 x_3$, and *the effect of mobilizing outsiders*, $\gamma_0 x_4$.

Three possible empirical outcomes conclude this formal section. If we observe an *increase* in the total number of followers of the reported accounts, it implies that the combination of the mobilization and the attention effect dominates the de-mobilization effect ($\gamma_0 x_4 + \gamma_2 x_4 > \beta_2 x_3$). If the number of followers is not altered, then it must be the case that the de-mobilization effect cancels the other two effects ($\gamma_0 x_4 + \gamma_2 x_4 = \beta_2 x_3$) and we still could claim that $\beta_2 x_3 > \gamma_0 x_4$. Finally, if we observe a *decrease* in δ, the de-mobilization effect of the attacks dominates their mobilization effect ($\gamma_0 x_0 < \beta_2 x_3$).

B The Process of Account Suspensions

Twitter suspends an account when there is sufficient evidence that it is a bot, it has been hacked or compromised, or spreads inappropriate content such as child pornography, right-wing or religious extremist content. While Twitter suspends all of such accounts, the process is far from automatic.

The timeline for account suspension as a result of Anonymous' flagging goes, at least, through four stages: (1) identification; (2) reporting; (3) checking; and, finally, (4) suspension. In the first stage, Anonymous mobilizes volunteers to *identify* Islamic radical content by searching through tweet hashtags associated with radical content. Once a suspicious account is identified, a volunteer informs one of the core Anonymous activists (generally nicknamed "CtrlSec"). Second, one of the core activists reviews the suspicious account and, if appropriate, *reports* it to Twitter for review. Twitter surveilance team manually *checks* each reported account. Finally, the Twitter surveillance team decides whether to suspend it or not.

From the information in our dataset, we can evaluate the time that it takes for the entire process to complete. For this, we calculate the difference between the date of the first time an account has been reported to Twitter and the suspension date. The average time of the suspension is 27.92 days after the first time of the Anonymous' report.

C The Sample of Reported Accounts

This online Appendix provides a brief overview of some characteristics of sample accounts and a discussion of potential reporting errors.

C.1 Examples of Reported Accounts

To begin with, Figure C.1 shows a snapshot of a sample of *Anonymous*' reported accounts. These snapshots allow us to glimpse how these accounts look like. Images with radical and even violent content, as well as extremist religious messages, generally explain why Anonymous activists report these accounts to Twitter.

C.2 The Reported Accounts

The previous section provides anecdotal evidence that the accounts reported by Anonymous seem to be related to Islamic-related radical content. This subsection moves beyond these anecdote to report a more systematic evaluation of the reported accounts.

Specifically, a working paper authored by some members of our research team ('reference omitted to maintain anonymity') provides further descriptive information about the IS-related accounts reported by Anonymous used in our empirical examination.[24]

In one of the empirical analysis of the working paper, the authors use semantic topic analysis based on the profile descriptions of the same accounts we use in this paper. Their results reveal that their profile information center around five topics: (1) Islamic cosmology (e.g., dunya / god / peopl / follow / syria); (2) martyrdom (e.g., life / fight / death / martyr / fear / heaven / libya); (3) piety (e.g., land / soldier / back / show / call / moham / eye); (4) jihad (e.g., islam / jihad / lord / caliph / global / levant / blood / servant); (5) religious blessings (e.g., free / truth / good / merci / love / heart / bless / success / forget). Among the reported accounts, martyrdom is the most popular topic, followed by

[24]The contribution of the working paper is mostly methodological, as it aims to be a review of some major families of ML models with a detailed presentation on Boosted Decision Tree (BDT) models. The working paper illustrate the utility of BDT models for political science by implementing an accuracy models using our dataset of IS-related accounts. The results from this working paper are used in this appendix to provide a thicker description of the sample of accounts. However, we should note that there is effectively little overlap between our main manuscript and this methodological working paper.

Figure C.1: Sample of Reported Accounts

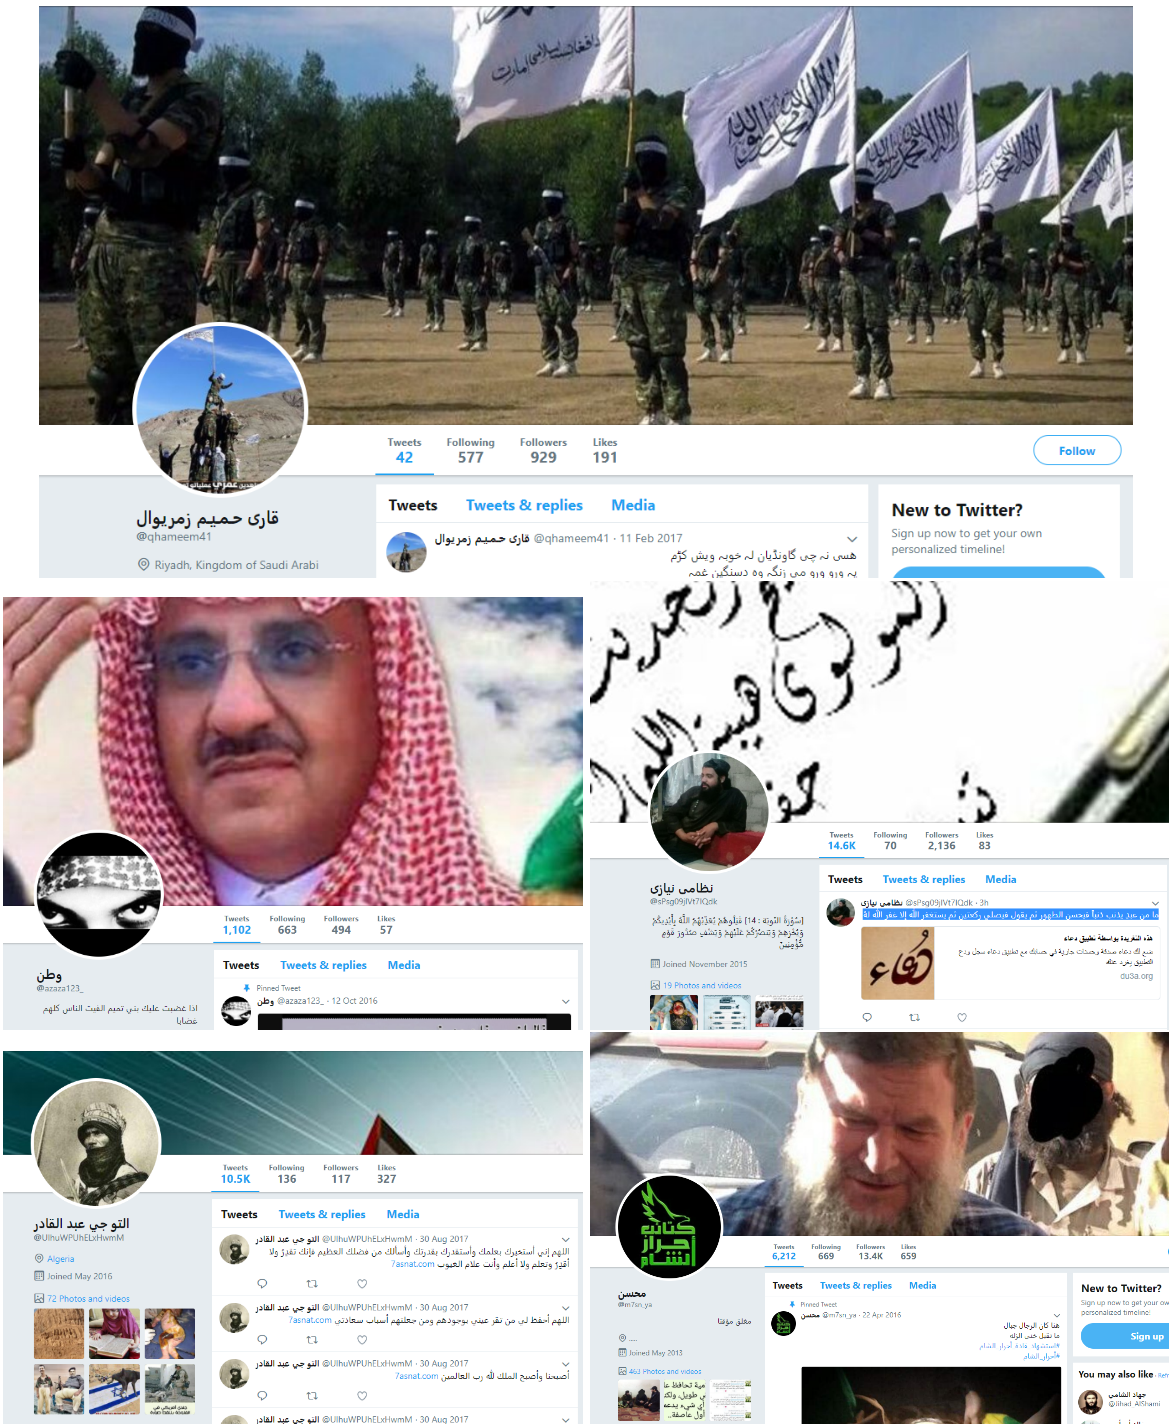

piety, jihad, Islamic cosmology, and religious blessings.

Moreover, the authors also provide an evaluation of the informational value of Anonymous reports; that is, how many false positive might there be in our sample? To estimate this, they note that Twitter suspends (1) bots; (2) hacked or compromised accounts; and, (3) abusive accounts, which are those that spread child pornography or right-wing and religious extremism. As we have seen in the semantic topic modelling, an ultimate suspension of an account indicates that an account certainly had Islamic-related extreme content.

Finally, the authors implement a Boosted Decision Tree (BDT) model, the gold standard of Machine Learning models, and estimate the accuracy rate of Anonymous activists in tagging Islamic extremism. Overall, their model predicts that 87% of the observations in the test dataset are affiliated with Islamist extremism. Therefore, we should conclude that there is some misreporting and false positive in our sample, yet there is suggestive evidence to claim that the vast majority of the accounts in our sample were accurately reported by Anonymous as associated with extreme Islamic activity.

C.3 The Geographic Distribution of Reported Accounts

Another approach to evaluate the face validity of our sample of IS-related accounts is to observe their geographic distribution by country. If the accounts are really related to IS, we should expect the overrepresentation of accounts from countries with a strong support base for IS. As we discuss in greater detail in the online Appendix I, the *location* of each Twitter account is self-reported and voluntary. Table C.2 shows the ratio of the share of the detected IS-related accounts in each country over the number of internet users. The five countries with most accounts reported by IS per online user are Somalia, Iraq, Syria, Libya, and Afghanistan. They are countries with a well-known strong and influential base of IS supporters. In addition, the geographic distribution of accounts also alleviate concerns that the sample would simply include Muslim accounts. We can observe that the sample includes few accounts from Muslim countries with little links to IS such as Iran, UAE, or Indonesia.

Figure C.2: Distribution of IS-related Accounts per online user by Country

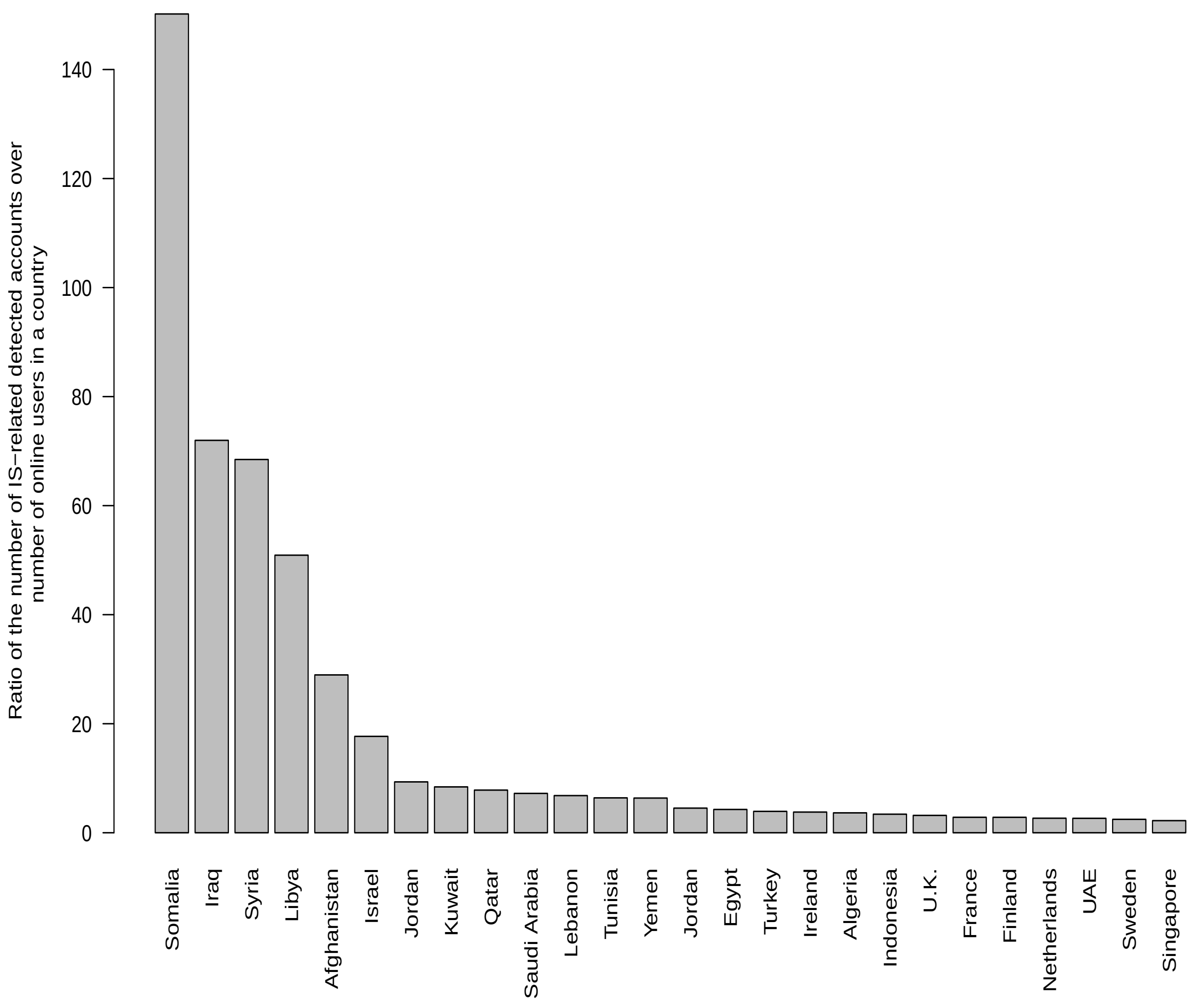

D Terrorist Data: List of Attacks and Number of Victims

Table D.1: List of IS Attacks in the Analysis

Terrorist Attack	Date	Location	Victims	References
Istanbul Explosion	March 19	Istanbul, Turkey	5	Dearden (2016); Tattersall and Yackley (2016)
Haqlaniyah	March 21	Haqlaniyah, Iraq	24	Acosta (2016)
Albu Obaid	March 21	Albu Obaid, Iraq	3	Acosta (2016)
Brussels Bombings	March 22	Brussels, Belgium	36	“Brussels Explosions” (2016)
Yemen Bombings I	March 25	Aden, Yemen	26	“Yemen Bombings” (2016)
Iraq Stadium	March 26	Al-Asriya, Iraq	41	“Iraq buries young” (2016); “Iraq Violence” (2016)
Workers’ Attack	March 29	Baghdad, Iraq	7	Acosta (2016)
Kurdish Policemen	March 31	Makhmour, Iraq	3	Acosta (2016)
Security Checkpoints	April 4	Mishahda, Iraq	10	Acosta (2016)
Commercial Center	April 4	Basra, Iraq	19	Acosta (2016)
Attack on Army Recruits	April 12	Aden, Yemen	5	Acosta (2016)
Mosque	April 23	Baghdad, Iraq	13	Acosta (2016)
Binnish Bombing	April 23	Binnish, Syria	4	Acosta (2016)
Military Checkpoint	April 24	Baghdad, Iraq	14	Acosta (2016)
Commercial Area	April 25	Baghdad, Iraq	12	Acosta (2016)
Military Checkpoint	April 25	Zeinab, Syria	8	Acosta (2016)
Attack on LGBT activists	April 25	Dhaka, Bangladesh	2	“Editor Hacked” (2016)
Baghdad Bombing I	April 30	Baghdad, Iraq	38	Adel (2016*b*)
Qamishli Attack I	April 30	Qamishli, Syria	5	Acosta (2016)
Baghdad Bombing II	April 30	Baghdad, Iraq	21	Acosta (2016)

Continued on next page

Table D.1 – continued from previous page

Terrorist Attack	Date	Location	Victims	References
Samawa Twin Explosion	May 1	Samawa, Iraq	33	"Rare IS Bombings" (2016)
Police Attack	May 1	Ganziantep, Turkey	2	Acosta (2016)
Attack against pilgrims	May 2	Baghdad, Iraq	18	Acosta (2016)
Attack on funeral	May 8	Baghdad, Iraq	5	Acosta (2016)
Attack on restaurants	May 9	Baquba, Iraq	13	Acosta (2016)
Attack against guards	May 11	Tataouine, Tunisia	4	Acosta (2016)
Baghdad bombings III	May 11	Baghdad, Iraq	110	"IS Kills Dozens" (2016)
Naval checkpoint	May 12	Mukalla, Yemen	15	Acosta (2016)
Real Madrid massacre I	May 13	Balad, Iraq	16	Stephen (2016)
Police station	May 15	Mukalla, Yemen	40	Acosta (2016)
Natural gas plant	May 15	Taji, Iraq	14	Acosta (2016)
Attack against Kurdish	May 21	Rojava, Syria	5	Acosta (2016)
Qamishli attack II	May 22	Qamishli, Syria	5	Acosta (2016)
Yemen bombings II	May 23	Aden, Yemen	45	"ISIL Blamed" (2016)
Jableh massacre	May 23	Jableh, Syria	148	Acosta (2016) "Bombs kill" (2016)
Azzawi attack	May 27	Azzawi, Syria	4	Acosta (2016)
Cafe Attack	May 29	Muqdadiya, Iraq	7	Acosta (2016)
Real Madrid massacre II	May 29	Balad, Iraq	12	Couzens (2016)
Mosque attack	June 2	Latakia, Syria	3	Acosta (2016)
Attack against militias	June 2	Jalalabad, Afghanistan	2	Acosta (2016)
Military checkpoint	June 4	Tarmiyah, Iraq	8	Acosta (2016)
Aktobe shootings	June 5/8	Aktobe, Kazakhstan	7	"Police Arrest" (2016); Dubnov (2016)

Continued on next page

Table D.1 – continued from previous page

Terrorist Attack	Date	Location	Victims	References
Church attack	June 6	Hah, Turkey	5	Acosta (2016)
Commercial area	June 9	Baghdad, Iraq	19	Acosta (2016)
Checkpoint attack	June 9	Taji, Iraq	12	Acosta (2016)
Shi'a shrine	June 11	Zeinab, Syria	20	Acosta (2016)
Hospital attack	June 12	Sirte, Lybia	3	Acosta (2016)
Iraqi troops	June 13	Ramadi, Iraq	5	Acosta (2016)
Magnanville stabbing	June 13	Magnanville, France	2	"French Jihadist" (2016)
Military checkpoint	June 15	Yusufiyah, Iraq	7	Acosta (2016)
Security forces	June 16	Abu Grein, Libya	11	Acosta (2016)
Memorial attack	June 19	Qamishli, Syria	3	Acosta (2016)
Iraqi Troops	June 20	Taji, Iraq	4	Acosta (2016)
Security forces	June 21	Ruqban, Jordan	6	Acosta (2016)
Qaa bombings	June 27	Qaa, Lebanon	5	"Lebanon" (2016)
Compound	June 28	Mukalla, Yemen	42	Acosta (2016)
				Almasmari and Sterling (2016)
Ataturk airport	June 28	Istanbul, Turkey	44	"Airport Attack" (2016)
Local office	June 29	Tel Abyad, Syria	5	(2016)
Dhaka attack	July 1	Dhaka, Bangladesh	21	Hanna et al. (2016)
Shi'a mosque	July 1	Imam Ahmad, Iraq	2	Acosta (2016)
Karrada bombings	July 3	Baghdad, Iraq	341	Adel (2016*a*)
Saudi Arabia bombings	July 4	Saudi Arabia	4	Robertson et al. (2016)
Bakery attack	July 5	Hasakah, Syria	16	Acosta (2016)
Army post	July 6	Benghazi, Libya	12	Acosta (2016)
Police attack	July 6	Balad, Iraq	40	Acosta (2016)

Continued on next page

Table D.1 – continued from previous page

Terrorist Attack	Date	Location	Victims	References
Market	July 12	Baghdad, Iraq	12	Acosta (2016)
Police checkpoint	July 12	Rashidiya, Iraq	9	Acosta (2016)
Police checkpoint	July 13	Baghdad, Iraq	7	Acosta (2016)
Tribal leader	July 13	Baquba, Iraq	2	Acosta (2016)
Truck attack	July 14	Nice, France	86	"Truck Attack" (2016)
Soldiers	July 20	Aden, Yemen	5	Acosta (2016)

Note on inclusion criterion: A borderline case is the attack on the *Purse club* in Orlando, on June 14, 2016. Though some initial information linked it to Islamic State, Barack Obama addressed the nation shortly after the attack by stating that the Orlando attack had been an act of "homegrown terrorism" carried out by legally purchased firearms ("Orlando Shooting" (2016)). This contrasts to François Hollande statement immediately after the truck attack on Nice, who strongly linked it to ISIS by stating that "all of France is being menaced by Islamic fundamentalist terrorism". Because our treatment effect should be evaluated in the short-term with the available information at that time, we choose to code the Nice attack as an IS attack, and the Orlando attack as a "home-grown" hate crime (Mestre, Revault d'Allonnes and Bissuel, 2016).

Note on measurement error: Figures on the number of deaths are approximate because they may vary depending on the news source.

E Google Trends and the Cumulative Lagged Measure of Deaths

Figure E.1 shows the cumulative death parameter around the Brussels bombings on March 22, 2016, with r=0.5, as having values of 0, 36, 24, 16, 10.67, and so on, for the days between 21th to 25th March, respectively. This discount pattern parallels the changes in the keywords "Brussels bombing" reported by Google Trends. On the date of the occurrence, Google trends reports a value of 100 (its standardized base value), a decrease to 7 one week later, and a further decrease to 0 two weeks later. Similarly, if we applied a discount factor of, for example, 50% to an event with 100 deaths—as the base in Google trends for comparison—then the trend would be 100 in the same day, 5.85 a week later, and 0.34 two weeks later. The correlation of Google attention to the event and our cumulative value is above 0.90 within a one-month window around the event. Other attacks in the sample do not differ in the evolution of their attention over time, but some of them had obviously less overall attention. It is worth noting that choosing a greater discount parameter does not alter any of the results presented in the paper.

Figure E.1: Google Trends and the Cumulative Lagged Measure of Deaths: Brussels Bombings (keywords "Brussels Bombing"

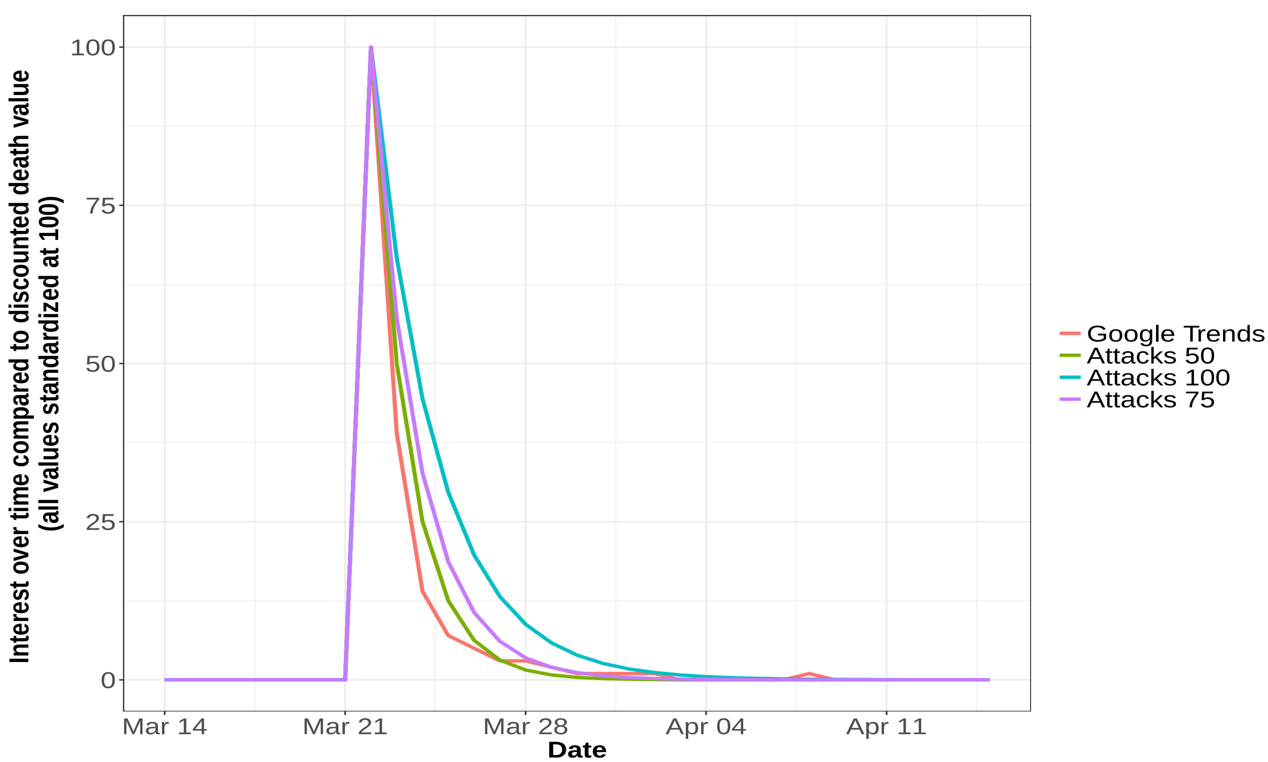

Figure E.2 shows the cumulative death parameter throughout the entire period together with the variation in Google Trends with respect to two sets of keywords "Terrorist Attacks" and "Terrorist Attacks Islamic State." The Google Trends has two important peaks: the Brussels bombing in late March and the Nice truck attack in late July. More importantly, we can observe that the shape of attention to the attacks follows quite closely the cumulative death parameter with the chosen discount rates.

Figure E.2: Google Trends and the Cumulative Lagged Measure of Deaths in Europe and the US: Entire Period

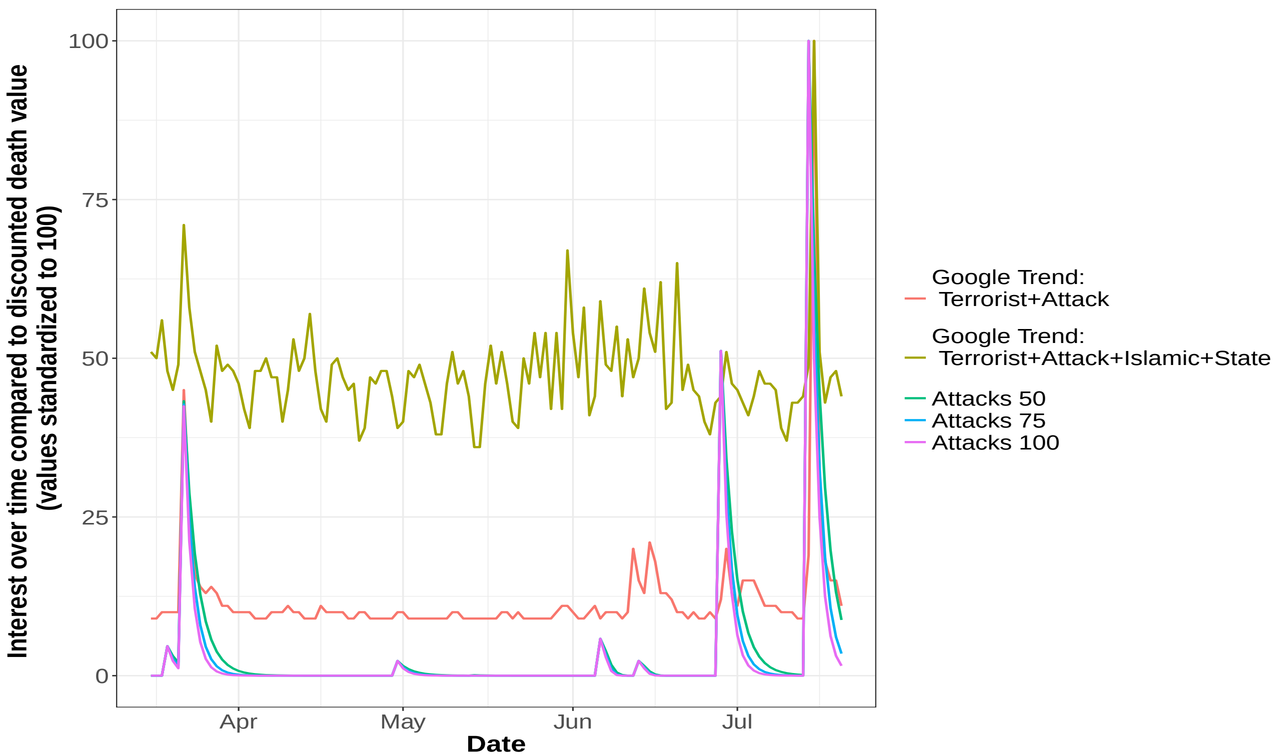

F The Long-term Effects of Attacks

As we show in the online Appendix E, we choose the discount rates for the cumulative death parameter that better captures the shape of people's and media attention, as measured by the trend in Google searches over time. Yet, another utility of the discount rates is that they allow us to explore the rate of decay of the effect of attacks on followers. For this, we generate death cumulative parameter by using extremely high and low discount rates, and explore their pattern.

Figure F.1 shows the trend of a cumulative death parameter for an attack of 100 deaths. These trends illustrate the variation in the lingering effects of attacks after several periods from using distinct discount rates for the attacks.

Figure F.1: Cumulative Death Parameter Across Several Discount Rates

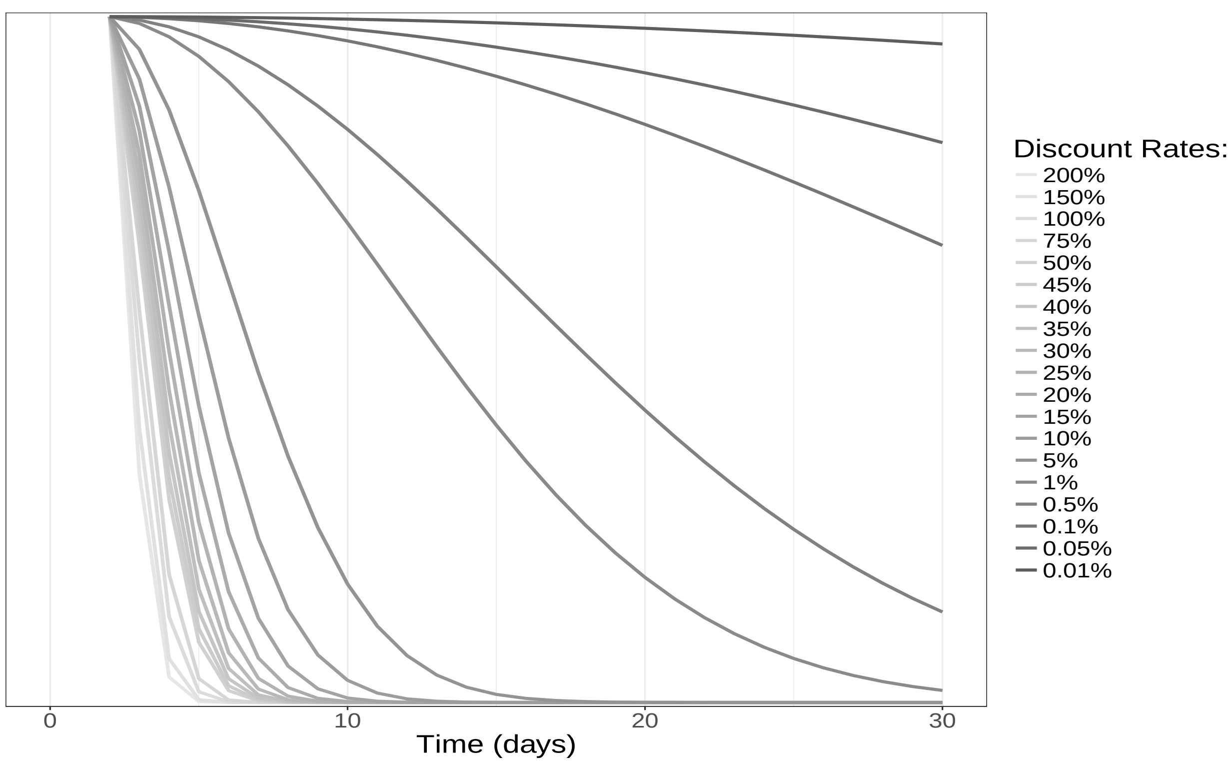

Figure F.2 reports the effect of terrorist attacks across several discount rates. As we can observe, the coefficients presented in the main body of the text are not sensitive to the choice of a specific discount rate. In addition, they also give us information about the long-term effect of attacks. In this regard, we can observe that the effect of terrorist attacks in European soil diminishes as the discount rate increases; that is, an attack's effect lasts across more periods. Specifically, the size of the effect halves at discount rates of 5% or lower. In terms of the discount rates, this means that when we model the effect of attacks to vanish in 15 days or longer (5% or lower in Figure F.1), the reported effect significantly decreases, although it remains significant at the 99% confidence level. This suggests that the effect of attacks extend to periods over 15 days, although the magnitude significantly decreases over time. By contrast, we do not observe a diminishing trend in longer periods in the attacks outside Europe.

Figure F.2: The Effect of Terrorist Attacks Across Several Discount Rates

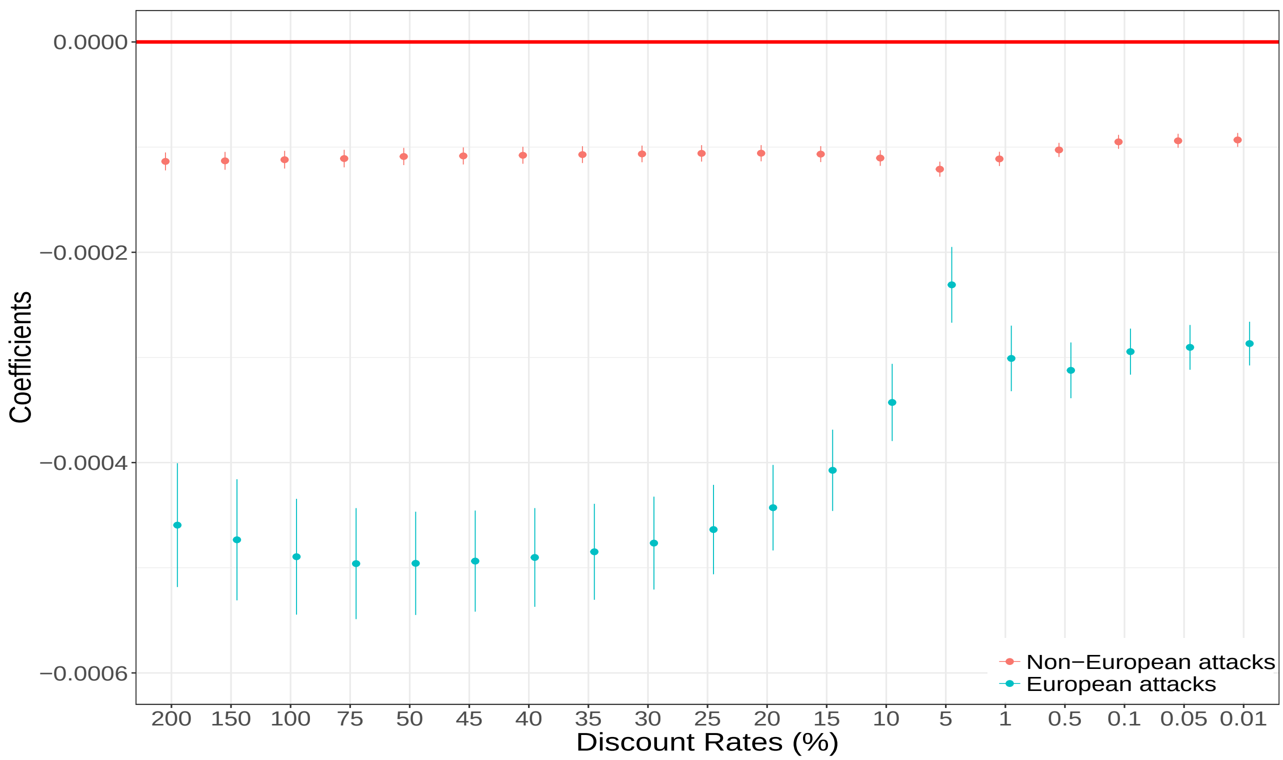

G Sensitivity Analysis to Aggregation Level of Time

In this Appendix, we acknowledge that "day" as the unit of analysis is an arbitrary choice. Note, however, that day is the smallest unit of time that is available in the dataset. Thus, we show our main findings in the main body of the text with data at a one-day intervals. We test here whether any of our main findings is altered by using 1-day intervals rather than more aggregate levels of time. Figure G.1 shows the coefficients for both European and non-European attacks on the number of followers to IS-related Twitter accounts. The horizontal red line indicates the point of no-effect. We can observe that the magnitude of the effect, if anything, increases when we use greater levels of aggregation, although this comes at a cost of greater variance in our estimates. Therefore, the main results remain unaltered at different levels of time aggregation.

Figure G.1: Sensitivity of Coefficients to the Level of Aggregation of Time

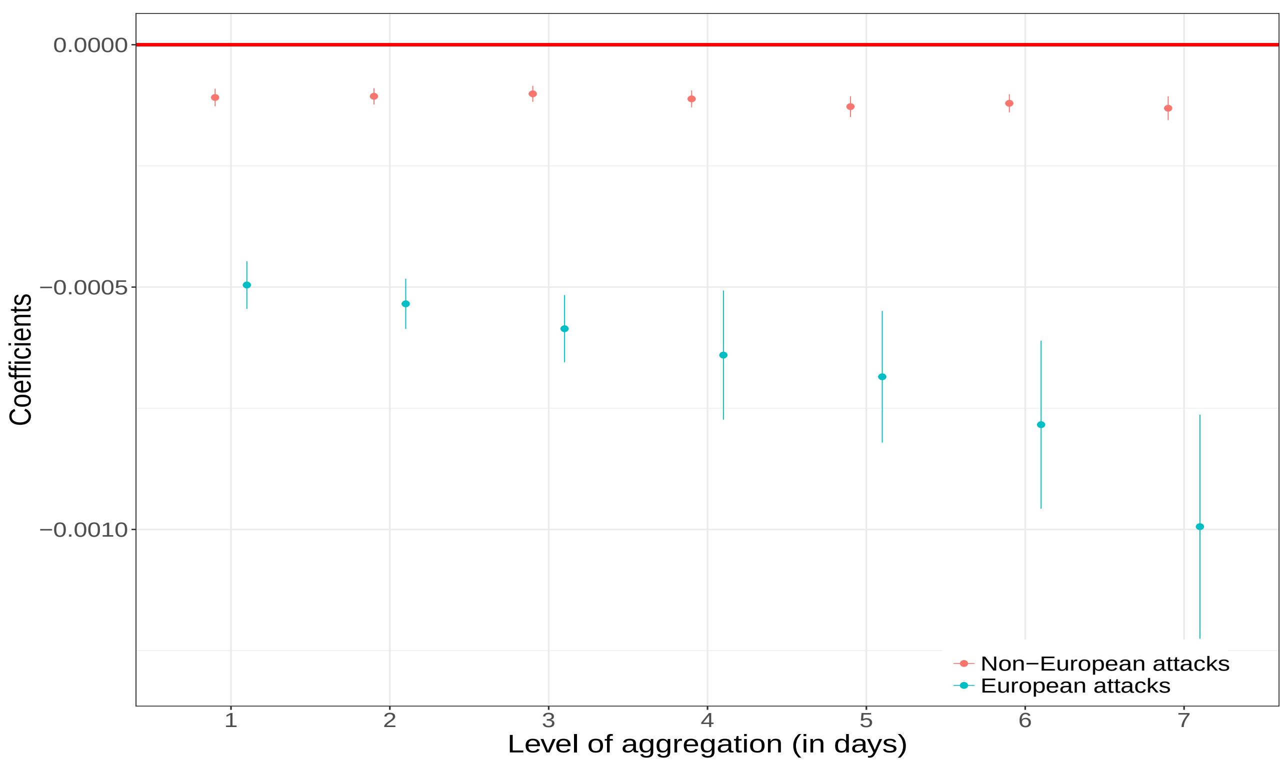

H Endogeneity Test: Placebo Timing of Attacks

In this Appendix, we develop a placebo test to check the robustness of our models to several sources of endogenity. An adequate placebo test when dealing with treatment effects that are based on time discontinuities is to test that a similar effect would not be seen if alternative timings are examined. If the effect we see is severely biased due to endogeneity or a result of other phenomena but the attacks, we should expect the same relationship to emerge when picking other dates.

To test this proposition, we generate 500 datasets with simulated attacks with randomly assigned timings. For each dataset, we run the fixed effects model in Table 2 with a discount rate of 50 and store each coefficient. Across all simulated datasets, we ensure that the mean and variance of the cumulative death rate is the same as in the original data. This allows us to directly compare the magnitude of coefficients between the observed and the simulated datasets.

Figure H.1 shows these tests. The histograms show the density of the 500 coefficients for attacks at random times (in black). The red line shows the observed coefficient when using the actual timing of attacks. We can observe that the observed coefficient in the models is clearly stronger than it would be given a random allocation of the attacks. This means that these relationships are not empirical artifacts as they require the death cumulative rate to reflect the actual timing of attacks to be able to yield a significant effect on the number of followers of IS-related accounts.

Figure H.1: Placebo Test: Coefficients of Random Timing of Attacks on Followers

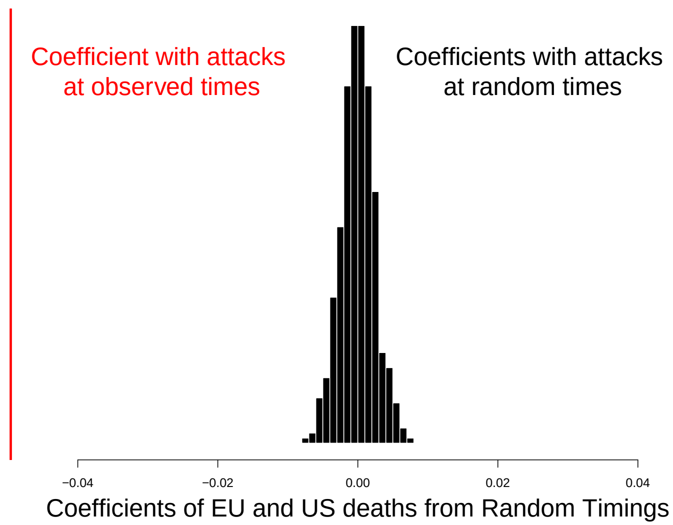

(a) Coefficients for attacks in European soil:
Observed (red) and simulated (black) Datasets

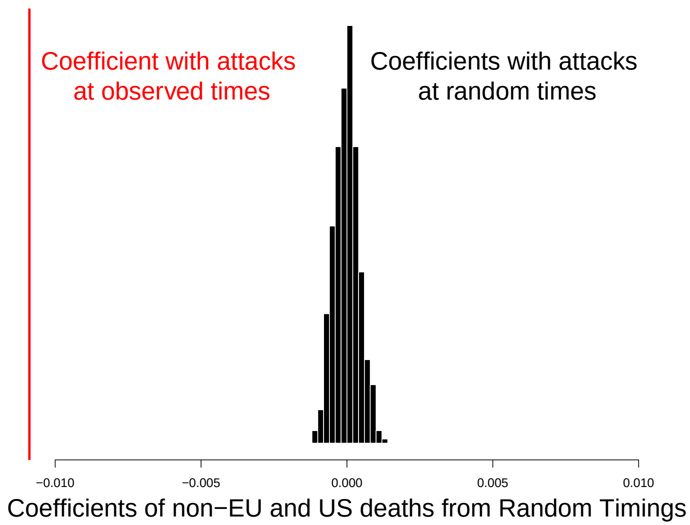

(b) Coefficients for attacks in non-European soil:
Observed (red) and simulated (black) Datasets

I Predicting National Material Capabilities (NMC)

In the last section of our paper, we look at how National Material Capabilities (NMC) of the country in which the account is located moderates the effect of the terrorist attacks on the number of followers of likely IS-related accounts. The country of a Twitter account defines its relative NMC. However, not all accounts have information enabling to determine their geographic origin. Therefore, we implement an imputation process to generate NMCs for those accounts from which we do not have sufficient information to determine their geographic location.

I.1 Inferring the country of origin

The profile information of each Twitter account has the feature *location*. This field is self-reported and voluntary. Therefore, many users leave this field empty (9580 of 13,300 or 72% of our sample). However, 1615 observations (43% of the non-empty or 12% of our sample) have sufficient information to infer their country. Table I.1 shows the shares of the countries that we can observe. They constitute the observations in our training dataset.

I.2 Matching accounts to the NMC: Clusterization of the countries

The distribution in our training set is unbalanced relative to the countries: for some countries we do not have enough information to reliably train a machine learning model to predict the country of account's origin. Therefore, we cannot predict the specific country of all the accounts that do not report their country of origin. Yet, we can reliably predict the account's characteristic of interest, the NMC of the account's country. This is a continuous attribute with much less variation. As Figure I.1 shows, the distribution is not uniform and proposes to take its clustered structure into account. We use the K-means algorithm to assign all countries based on their NMC to 5 clusters with the incremental labels.

Table I.1: Shares of the Twitter accounts by country (training dataset)

Country	Share	Country	Share	Country	Share
Syria	11.19%	Russia	0.98%	Maldives	0.43%
Iraq	10.46%	Spain	0.98%	Mexico	0.43%
United States	6.73%	Lebanon	0.92%	Philippines	0.43%
United Kingdom	5.69%	Algeria	0.86%	China	0.37%
Indonesia	5.38%	Malaysia	0.86%	Ukraine	0.37%
Turkey	5.38%	Somalia	0.86%	Nigeria	0.31%
France	4.71%	Brazil	0.86%	Romania	0.31%
Saudi Arabia	4.46%	Japan	0.80%	Singapore	0.31%
Egypt	3.91%	Kuwait	0.80%	Argentina	0.24%
Israel	3.12%	Morocco	0.80%	Bosnia and Herzegovina	0.24%
Germany	2.08%	Australia	0.73%	New Zealand	0.24%
Libya	2.02%	Italy	0.73%	Norway	0.24%
Afghanistan	1.96%	Sweden	0.67%	Austria	0.18%
India	1.83%	United Arab Emirates	0.67%	Bahrain	0.18%
Canada	1.59%	Belgium	0.55%	Denmark	0.18%
Pakistan	1.28%	Switzerland	0.55%	Ethiopia	0.18%
Yemen	1.28%	Qatar	0.49%	Iran	0.18%
Netherlands	1.16%	Bangladesh	0.43%	Oman	0.18%
Tunisia	1.04%	Finland	0.43%	Other	4.40%
Jordan	0.98%	Ireland	0.43%		

Figure I.1: National Material Capabilities: Ordered Accounts

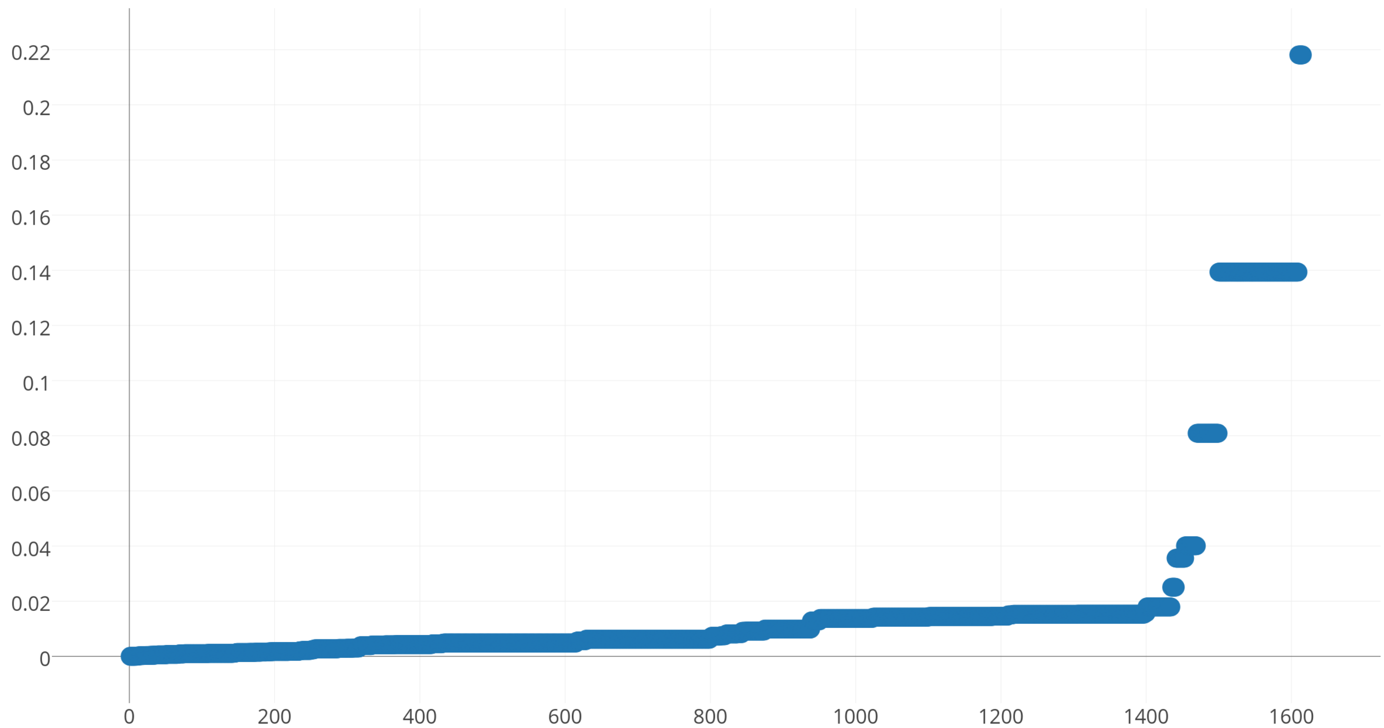

I.3 Training the machine learning model

The classification algorithms are well-developed in machine learning and, importantly for us, have very clear evaluation measures. In particular, we can explicitly see how many values in the data-set are predicted correctly. That is why we are making use of the clustered structure of NMCs. We train the ordinal regression model using *the Two-Class Boosted Decision Tree* (Elith, Leathwick and Hastie, 2008) based on the features that we have for all accounts in our dataset: *language*, *time-zone*, *number of friends*, and *number of favorites*, and *number of statuses*[25]. Figure I.2 shows the results from 10-fold cross-validation from the trained model for the test sample. In addition, we can see that the prediction model has a very good fit. Hence, we apply the trained model to obtain NMC results for the rest of the sample. Overall, this procedure generates NMC scores for all accounts in our sample. At one extreme, those accounts located in a weak country have a NMC score of 0. At the other extreme, those accounts located in countries with strong material capabilities have a score of 4.

Table I.2: Cross-validation: Ordinal Regression of the Test Sample via the Two-Class Boosted Decision Tree

Fold Number	Observations	Error
0	162	0.17
1	161	0.16
2	161	0.14
3	162	0.12
4	162	0.13
5	161	0.15
6	162	0.14
7	161	0.14
8	161	0.12
9	162	0.14
Mean	1615	0.14
Standard Deviation	1615	0.02

[25]We exclude the number of followers from the features used to train the model since we use it as a dependent variable further in our analysis.

References

Acosta, Benjamin. 2016. "Dying for survival Why militant organizations continue to conduct suicide attacks." *Journal of Peace Research* .

Adel, Loaa. 2016*a*. "Baghdad bombing: Scores of people dead, dozens injured." *Iraqi News, 3 July* .
URL: *http://www.iraqinews.com/iraq-war/baghdad-bombing-scores-of-people-dead-dozens*

Adel, Loaa. 2016*b*. "Samawah explosion casualties jump to 38 dead and 86 wounded." *Iraqi News, 1 May* .
URL: *http://www.iraqinews.com/iraq-war/samawah-explosion-casualties-jump-38-dead-86*

Adnan, Ghassan and Asa Fitch. 2016. "Islamic State Bombings Kill Dozens in Baghdad." *Wall Street Journal, 11 May* .
URL: *http://www.wsj.com/articles/at-least-29-killed-in-baghdad-car-bombing-14629605*

Almasmari, Hakim and Joe Sterling. 2016. "More than 40 killed in Yemen suicide attacks.".
URL: *https://edition.cnn.com/2016/06/27/middleeast/yemen-suicide-car-bombs/index.html*

Brussels explosions: What we know about airport and metro attacks. 2016. *BBC, 9 April* .
URL: *http://www.bbc.com/news/world-europe-35869985*

Couzens, Gerard. 2016. "ISIS gunmen kill at least 12 Real Madrid fans as they watch Champions League Final in fresh attack on the club's supporters in Iraq." *Daily Mail, 31 May* .
URL: *http://www.dailymail.co.uk/news/article-3615485/ISIS-gunmen-kill-12-Real-Madrid-fans-watch-Champions-League-Final-fresh-attack-club-html*

Davison, John. 2016. "Bombs kill nearly 150 in Syrian government-held cities: monitor." *Reuters, 23 May* .
URL: *goo.gl/i5v9CW*

Dearden, Lizzie. 2016. "Istanbul explosion: At least five dead as tourist shopping area Istiklal Street hit by suicide bombing in Turkey." *The Independent, 19 March* .

URL: *http://www.independent.co.uk/news/world/europe/istanbul-explosion-injuries-reported-as-tourist-shopping-area-hit-by-blast-in-turkis html*

Dubnov, Arkady. 2016. "Nazarbayev on the Rocks: 'Terror,' 'Coup' Claims Shake Kazakhstan." *Moscow Times, 8 June* .

URL: *https://themoscowtimes.com/articles/nazarbayev-on-the-rocks-terror-coup-claims*

Editor of Bangladesh Gay Magazine Hacked to Death in His Home. 2016. *Wall Street Journal, 25 April* .

URL: *http://www.wsj.com/articles/two-men-including-usaid-employee-killed-in-banglad*

Elith, Jane, John R Leathwick and Trevor Hastie. 2008. "A working guide to boosted regression trees." *Journal of Animal Ecology* 77(4):802–813.

French jihadist police killer 'obeyed Islamic State call. 2016. *BBC, 14 June* .

URL: *http://www.bbc.co.uk/news/world-europe-36530710*

Hanna, Jason, Ralph Ellis, Saeed Ahmed and Emanuella Grinberg. 2016. "Terror attack in Bangladesh: The victims." *CNN, 6 July* .

URL: *http://www.cnn.com/2016/07/02/asia/bangladesh-dhaka-attack-victims/*

Iraq buries young victims of football pitch bomb carnage. 2016. *Agency France-Presse, 25 March* .

URL: *https://www.yahoo.com/news/suicide-bomber-kills-25-south-baghdad-police-172305 html?ref=gs*

Iraq violence: 'IS suicide attack' kills 29 in football match. 2016. *BBC, 25 March* .

URL: *http://www.bbc.com/news/world-middle-east-35901358*

ISIL blamed for deadly blasts in Yemen's Aden. 2016. *Al-Jazeera, 24 May* .

URL: *http://www.aljazeera.com/news/2016/05/dozens-dead-car-bomb-attack-yemen-16052
html*

Istanbul airport attack toll rises to 45 as child dies. 2016. *Agency France-Presse, 2 July* .

URL: *http://www.straitstimes.com/world/middle-east/istanbul-airport-attack-toll-rises-to-45-as-child-dies*

Lebanon: Second wave of suicide bombings hits Qaa. 2016. *Al Jazeera and agencies, 27 June* .

URL: *http://www.aljazeera.com/news/2016/06/deadly-suicide-bombing-rocks-village-ea
html*

Mestre, Abel, David Revault d'Allonnes and Bertrand Bissuel. 2016. "Attentat de Nice : les responsables politiques, entre motion et colre." *Le Monde, July 15, 2016* . [Online; accessed on 10/22/2016].

URL: *goo.gl/UZslZf*

Nice truck attack: French police arrest eight more suspects. 2016. *The Guardian, 20 September* .

URL: *https://www.theguardian.com/world/2016/sep/20/nice-truck-attack-french-police-arrest-eight-new-suspects*

Orlando shooting: What motivated a killer? 2016. *CNN, 14 June* .

URL: *http://www.cnn.com/2016/06/13/us/orlando-nightclub-shooting/*

Police arrest, kill radical Islamist gang, after attack on military base in Kazakhstan. 2016. *Russia Today, 5 June* .

URL: *https://www.rt.com/news/345468-kazakhstan-shooting-hostages-attacks/*

Rare IS group bombings in southern Iraq kill 33. 2016. *Agency France-Presse, 1 May* .

URL: *https://www.yahoo.com/news/twin-bombings-southern-iraq-kill-33-113911468.html*

Robertson, Nick, Mohammed Tawfeeq and Joe Sterling. 2016. "Suicide attackers launch wave of strikes in Saudi Arabia." *CNN, 5 July* .

URL: *http://edition.cnn.com/2016/07/04/middleeast/saudi-arabia-jeddah-bomb/*

Stephen, Jones. 2016. "ISIS kill 16 Real Madrid fans in massacre at supporters club headquarters." *Mirror, 13 May* .
URL: *http://www.mirror.co.uk/news/world-news/isis-kill-16-real-madrid-7962654*

Tattersall, Nick and Ayla Yackley. 2016. "Suicide bomber kills four, wounds 36 in Istanbul shopping district." *Reuters, 20 March* .
URL: *http://www.reuters.com/article/us-turkey-blast-idUSKCN0WL0D5*

Yemen bombings claimed by Islamic State kill at least 26. 2016. *Reuters, 25 March* .
URL: *http://www.reuters.com/article/us-yemen-security-idUSKCN0WR1EZ*