Supporting Information For Over Pressure: Grassroots-Driven Transformation of (Militant) Organizations

February 13, 2023

Contents

1 Introduction

This appendix presents supporting information for the quantitative analysis sections of Over Pressure: Grassroots-Driven Transformation of (Militant) Organizations. The analysis covers the data selection, processing, and results for a random forest, support vector machine, and t-Distributed Stochastic Neighbor Embedding (t-SNE) clustering of news articles about violence in Yemen attributed to AQAP, Ansar al-Shariah, and the Houthi militant groups. It also describes the data, analysis, and results for a Structural Topic Model of communications from AQAP, Ansar al-Shariah, and as-Sahab.

2 Media Texts and Processing

News stories originated in the ICEWS database and were selected by first querying the database for stories about events located in Yemen. This resulted in 47,385 stories ranging from January 15, 1991 through January 4, 2015. I selected only "violent" events, defined as events that fall into one of the following ICEWS event types: "Threaten with military force," "Use unconventional violence," "Violate ceasefire," "Use as human shield," "Threaten," "Occupy Territory," "Physically assault," "Mobilize or increase armed forces," "Engage in violent protest for leadership change," "Engage in mass killing," "Conduct suicide, car, or other non-military bombing," "Carry out suicide bombing," "Attempt to assassinate," "Assassinate," "Abduct, hijack, or take hostage," "Fight with small arms and light weapons," or "Fight with artillery and tanks."

This resulted in 10,818 stories, of which I took a random sample of 1772 stories. ICEWS codes for event date, event type, source actor, and target actor. However, the source and target actor codes typically characterize the actor by their role, such as "Armed Rebel" or "Militant," rather than by group affiliation.

In order to generate data on how groups operate, I re-coded the reports to include a variable

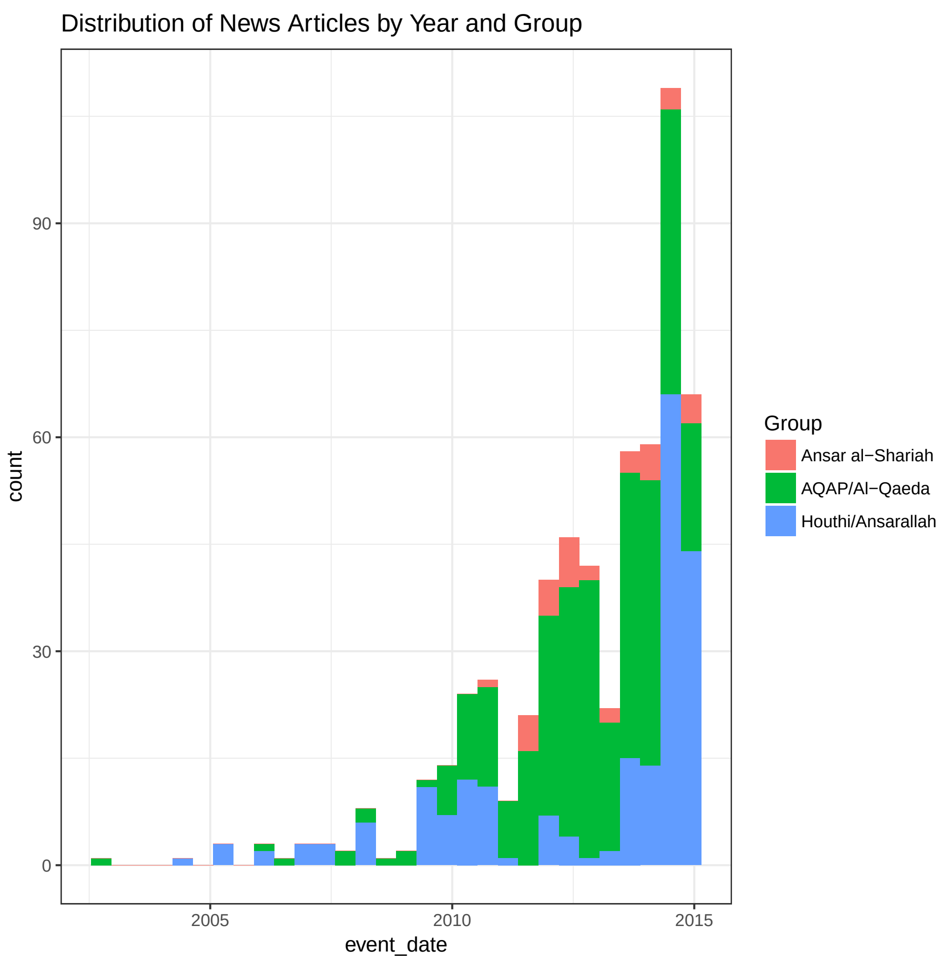

Figure 1: Distribution of News Articles

for group or movement affiliation. I first sent the sample to Amazon's Mechanical Turk platform, asking the workers to categorize the stories as relating to an action carried out by Ansar al-Shariah, AQAP, Houthi/Ansarallah, Yemeni Government, Tribal Uprising, Other, Multiple Actors, or Unknown. I kept the tags for the 283 stories that both coders agreed on, and hand-coded the remaining 1489 stories. The temporal distribution of the news articles and actor labels can be seen in Figure 1. I then further subset the data to keep only the stories tagged as describing a violent event carried out by one of the three militias of interest. This produced the final 720-story corpus of news events.

For each of the development, validation, and test sets, I used the tm() package to tokenize the words in each story and remove numbers, standard English stopwords, whitespace, and stray HTML markup. I additionally removed a custom list of stopwords that strongly signal the group, such as variations on the group name and signifiers of sectarian identity. These custom stopwords are comprised of:

Words that signal AQAP: "qaeda," "alqaida," "alqaeda," "qaida"

Words that signaled Houthis: "houthi," "huthi," "houthis," "zaidi," "alhouthi"

Terms that signaled Ansar al-Shariah: "ansar," "sharia," "alsharia";[1]

Terms that suggest an al-Qaeda affiliation: "laden," "osama"; words often used to summarize location of action for one of the groups: "peninsula," "northern," "southern," "arabian" "yemen[-]based"

And, finally, words that denote sectarian identity: "sunni," "shia," "shiite."[2]

Word frequency was normalized via term frequency-inverse document frequency (tf-idf), producing a pair of tf-idf matrices, from which I took the intersection of features (i.e. words). This generated a set of 2,222 "features" for classification in the texts; which reduced the available terms significantly but was necessary to test models across the training, validation, and test sets.

[1] The robustness models also remove "alshariah" with little change in results.

[2] I did not remove areas of operation from the texts as the goal of the classifiers was to seek discussion of operational differences. Locations of operation are substantively meaningful.

	Absolute Frequency	Proportion of Documents
Ansar al-Shariah	27	5.8%
AQAP/Al-Qaeda	260	56.3%
Houthi/Ansarallah	174	37.7%
Total	461	99.8%

Table 1: Distribution of group labels in "development" set

I reattached metadata to each of the term document matrices. Metadata included group label, date, and whether the story was coded by Mechanical Turk workers.

The distribution of group labels in the development set can be seen in Table 1 with the corresponding distribution from the validation set in Table 2.[3]

	Absolute Frequency	Proportion of Documents
Ansar al-Shariah	9	7.8%
AQAP/Al-Qaeda	67	58.2%
Houthi/Ansarallah	39	33.9%
Total	115	99.9%

Table 2: Distribution of group labels in "validation" set

3 Machine Learning Classifiers

The supervised machine learning techniques used in the paper provide a strategy to adjudicate between the counterfactuals introduced in the theory and qualitative sections. In particular, the clustering analysis indicates that international and local journalists writing about events in Yemen use similar terms when describing the activities of AQAP and Ansar al-Shariah. This suggests that AQAP has been unable to maintain a local spin-off with a distinctive operational profile. However, one significant caveat is that these techniques are unable to distinguish between AQAP acting like Ansar al-Shariah, Ansar al-Shariah acting like AQAP or journalists conflating Sunni insurgent groups. Distinguishing between the three possibilities is important to assess the theoretical expectation that an inflow of re-

[3]Deviations from 100% in the relative frequency sums is due to rounding.

cruits should pressure AQAP's leadership to adopt a local emphasis. The topic modeling section addresses concerns about the direction of convergence. One natural counterfactual in which an influx of local fighters is followed by behavioral convergence of AQAP and Ansar al-Shariah but which does not follow the mechanism hypothesized by the bottom-up transformation theory could be that AQAP's socialization has been so successful that the group has changed the preferences of the communities in which they operate. In this scenario, the local Ansar al-Shariah should gain a greater international focus as local actors are socialized into the transnational jihadi ideology.

The following section provides technical details about the implementation of the t-SNE visualization and the SVM and Random Forest classifiers.

3.1 t-SNE Hyperparameter Selection

Figure 2 presents t-SNE clustering for all stories published between 2011 and 2014, presented according to year of publication. These dates provide a snapshot of writing about each of AQAP (triangle), Ansar al-Shariah (circle), and the Houthis (square), and provide a high-level visualization of the separation or convergence among the words used to describe each of the three groups. The yearly clustering displayed in 2 features one point per story, and indicates that across the time period, stories about the Houthi insurgency appear to be systematically different from stories about the two Sunni groups, and is suggestive of a pattern in which Ansar al-Shariah stories become progressively more similar to AQAP stories from 2012 through 2014.

The visualization presented here was generated by running the t-SNE algorithm on for 5,000 iterations on the pooled data. The perplexity hyperparameter presented below was selected after grid sweeping from 5-50, at intervals of 5. Sweeping the perplexity hyperparameter changes the exact outcome, as expected from a probabilistic approach to summarizing structure in complex high dimensional data, the conclusions are broadly consistent across the

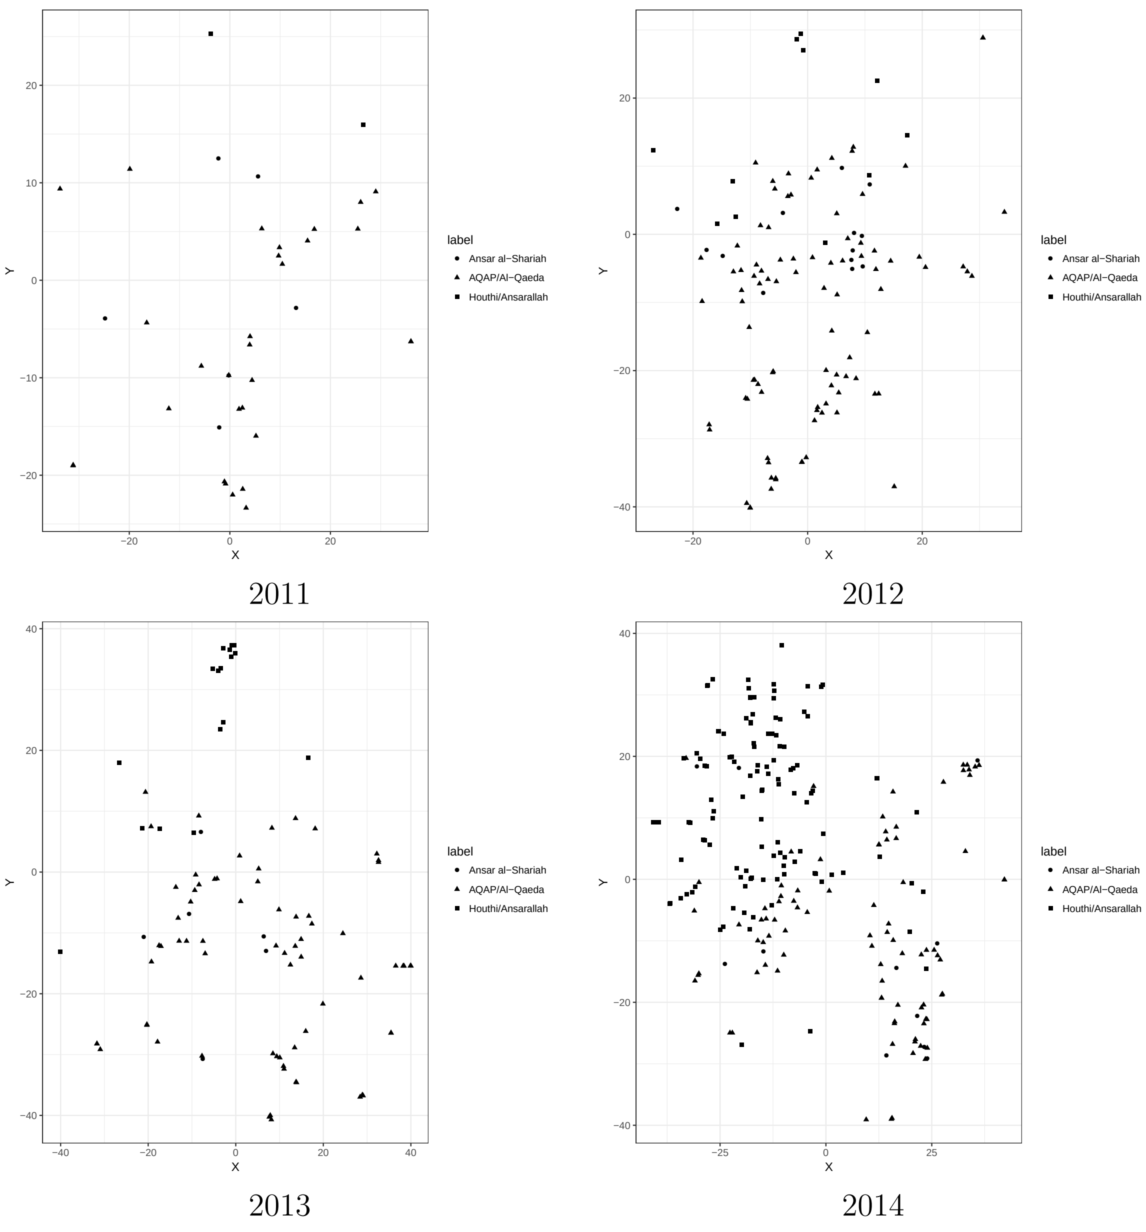

Figure 2: tSNE Clustering, All Stories 2011-2014

specifications. To address concerns that the observed clustering is random noise or driven by a specific initialization, clustering was carried out in parallel on two different machines using the same specifications but different starting points. The results were broadly similar, namely clear separation between the Sunni and Houthi stories but lack of clear separation among stories about AQAP and Ansar al-Shariah. As the diffusion and relative positioning of clusters generated using t-SNE are not inherently meaningful, comparison across the runs is only impressionistic and presenting averaged results would not be meaningful.

Figure 2 presents t-SNE clustering for AQAP and Ansar al-Shariah stories published between 2011 and 2014, presented according to year of publication. These dates provide a snapshot of writing about each of AQAP (green), Ansar al-Shariah (red), and the Houthis (blue), and provide a high-level visualization of the separation or convergence among the words used to describe each of the three groups. The yearly clustering displayed in Figure 2 features one point per story, and indicates that across the time period, stories about the Houthi insurgency appear to be systematically different from stories about the two Sunni groups, and is suggestive of a pattern in which from 2012 through 2014, the Ansar al-Shariah stories become progressively more similar to AQAP stories. Figure 3 focuses only on whether or not the t-SNE visualization differentiates between Ansar al-Shariah and AQAP stories. As compared to the separation of the Houthi stories in the corpus, the two Sunni groups demonstrate no apparent separation. This implies that stories about the two groups are much more similar than are stories about AQAP and the Houthi insurgency.

3.2 Random Forest Parameter Selection

The random forest classifier used the randomforest() method from the randomForest R package. The specification used the development data as training data, and the validation data as a test set, with story label as the classifier to predict. The confusion matrix for the

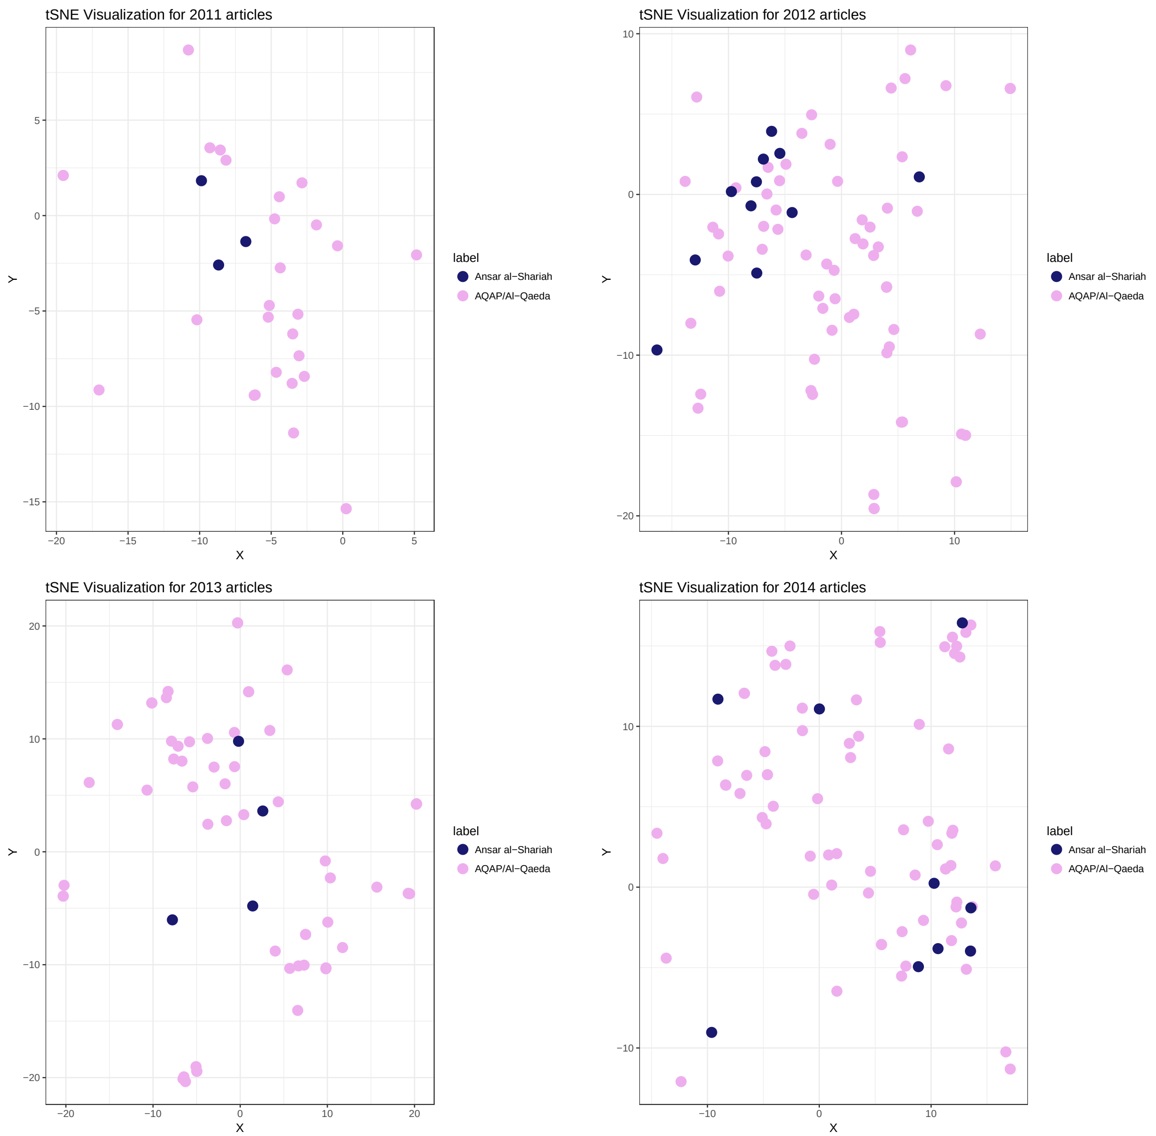

Figure 3: tSNE Visualization Sunni Groups, 2011-2014

training data is featured in the full text, while the confusion matrix for the training data appears below in Figure 3, followed by the classification results on the test data, in Figure 4

.

The model grew 500 trees, from which is generated a proximity measure for each document. The variable importance plot was generated after extracting the variable importance measure and plotted using the varImpPlot() method native to randomForest.

	Ansar al-Shariah	AQAP/Al-Qaeda	Houthi/Ansarallah	Class Error
Ansar al-Shariah	0.00	24.00	2.00	1.00
AQAP/Al-Qaeda	4.00	240.00	3.00	0.03
Houthi/Ansarallah	0.00	7.00	152.00	0.04

Table 3: Random Forest Confusion Matrix, Training Data

	Ansar al-Shariah	AQAP/Al-Qaeda	Houthi/Ansarallah	Class Error
Ansar al-Shariah	0.00	6.00	3.00	1.00
AQAP/Al-Qaeda	0.00	48.00	19.00	0.28
Houthi/Ansarallah	0.00	14.00	25.00	0.36

Table 4: Random Forest Confusion Matrix, Test Data

A second version of the random forest model estimated the model only on Ansar al-Shariah and AQAP stories. As with the full data, this model failed to predict any stories for the Ansar al-Shariah label. However, it did predict two AQAP-labeled stories as being likely Ansar al-Shariah stories. These stories were an August 3, 2014 story about a clash between al-Qaeda militants and police forces and a December 17, 2014 story about a car bomb carried out against Yemeni police officers. Notably, both stories were one of the 11 stories from Xinhua News that used the names AQAP and Ansar al-Shariah interchangeably.

One advantage of the random forest approach to classifying text is that the features used for classification are also terms in the document, which provide interpretable insights into what words drive separation among stories in the data. The fifteen most important words for the random forest classification—after removing stopwords that describe or name the active group— suggest that the reason for the clean split across the Sunni and Shia movements

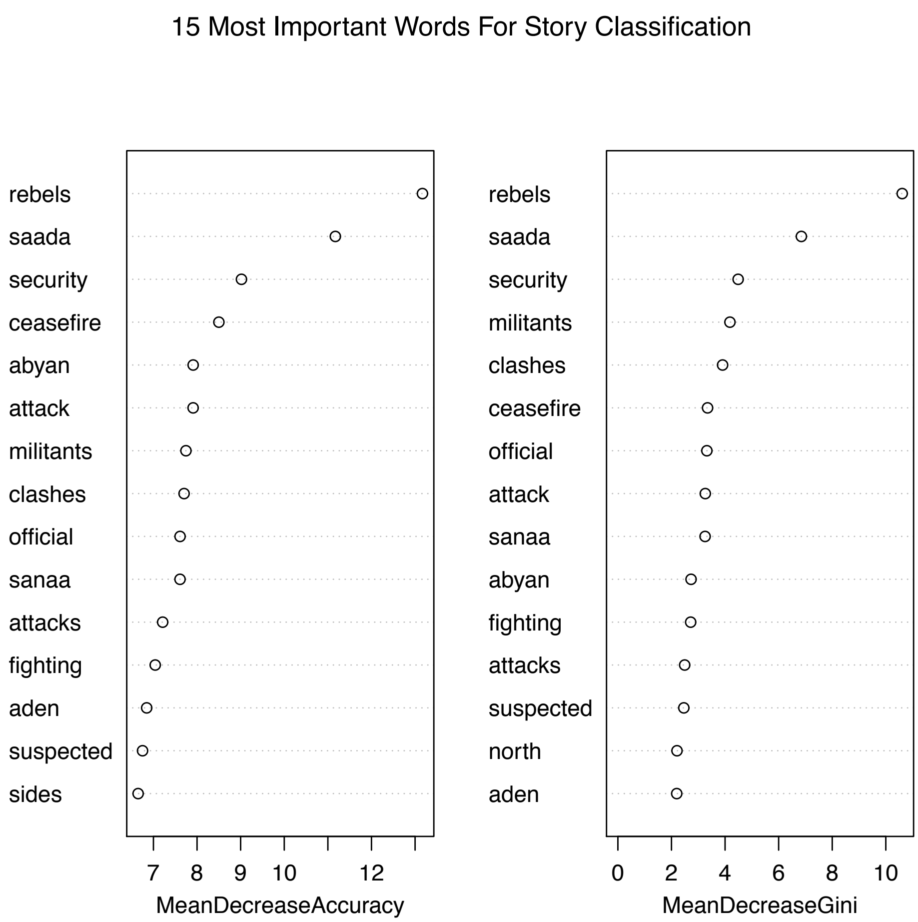

Figure 4: Important Words for Classification

lies in framing. The list of these terms, ranked in order of their importance for decreasing the classifier accuracy (left) and importance for decreasing homogeneity in the final nodes of the classification (right), can be seen in Figure 4. The importance of "rebels" as the top term for both accuracy and Gini coefficient is revealing: indeed in the texts, Houthis are consistently described as "rebels" while the Sunni fighters are frequently presented as "militants." After the rebel/militant split, words that describe the location of operations and military occupations are, unsurprisingly, important classifiers: Aden, Saada and Sanaa are regions associated with Houthi territorial gains, while Abyan is more closely linked to AQAP and Ansar al-Shariah activities.

The above results present a random forest model estimated over the entire dataset, but a cursory examination indicates that the data is extremely unbalanced. Notably, there are almost 10 times more articles describing AQAP events than there are articles describing Ansar al-Shariah activity. Because the data is so imbalanced, we should not be surprised that the model has a difficult time separating Ansar al-Shariah articles. We can balance this by taking a stratified sample, balancing the frequency of Ansar al-Shariah, AQAP, and Houthi articles:

3.3 Principal Component Analysis

Another view of the separation among the stories is shown in Figure 5, which uses a principale component analysis (PCA) to plot proximity of stories, as measured by the proportion of times that individual stories are in the same terminal node.[4] This reaffirms the takeaway from the confusion matrix: Houthi stories are distinct from AQAP stories, but Ansar al-Shariah stories contain enough words in common with AQAP stories that the two are difficult to distinguish via the random forest's iterated decision trees.

The principalecomponent analysis presentation was done using the extract_proximity() and

[4] Jones and Linder 2016.

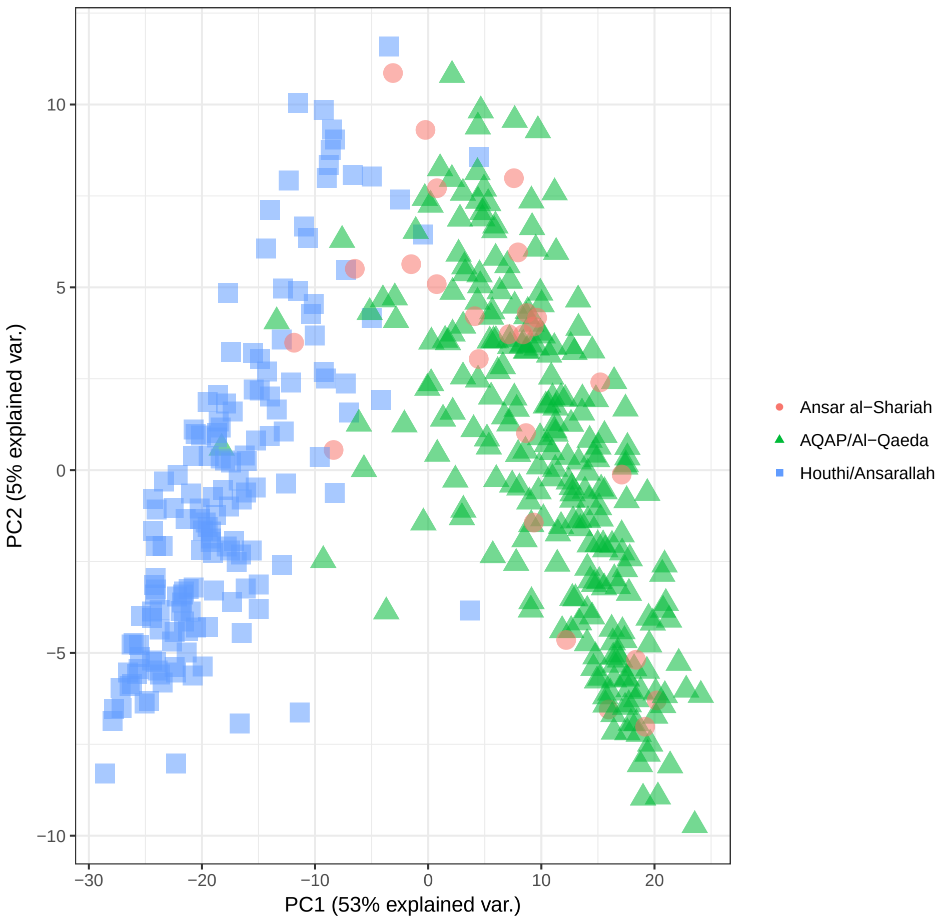

Figure 5: PCA Visualization of Group Classification

prcomp() methods from the edarf R package. The proximity extraction categorizes document proximity as the proportion of times that two observations are in the same terminal node across each estimated tree. This generates an NxN matrix of location similarities. This proximity is visualized using a principal components analysis of the proximity matrix, implemented in the edarf package.

3.4 SVM Specification

The support vector machine classification was developed using the ksvm() method from the kernlab R package. I used the "development" data as a training set and the "validation" data as the test set, with "label" as the classification attribute to predict and the words from each term document matrix as the features used in the classification.

The kernel specified was the radial basis function (RBF) Gaussian kernel via kernlab's "rbfdot" implementation. The RBF kernel is defined as:$K(x, x\prime)$ =exp$(-\frac{||x-x\prime||^2}{2\sigma^2})$. This kernel is based on squared Euclidean distance between feature vectors for each document, and thus is amenable to interpretation as distance between the documents. The svm carried out a 3-fold cross-validation on the test set, and used the resulting model to predict category assignment in the test (validation) set. In the analysis presented in the paper, each class was given a weight of 1. The paper reports predicted probability of label assignment.

Given the imbalance in the data, I developed two additional versions of the SVM model, subsetting the development-validation data subset to include only AQAP and Ansar al-Shariah stories. One of these replicated the original support vector machine with an equal weighting on each possible class, and one of which assigned class weights to the labels in the 90-10% distribution of AQAP and Ansar al-Shariah stories. Results of these two models were similar to the model presented in the paper: given the words in the news stories as features, the SVM does not distinguish Ansar al-Shariah stories from AQAP stories.

Figure 6 provides a closer look at whether the support vector machine's confidence in whether to assign stories to Ansar al-Shariah or AQAP changes over time. Ideally, support for the transformation theory would indicate the SVM assigning increasing weight to the Ansar al-Shariah label for news stories about AQAP activities as AQAP members push the group to engage with local concerns. The plot focuses on stories with a true label of AQAP and depicts the predicted probability that an article would be assigned to the label of Ansar al-Shariah given that it is a story about AQAP actions (red) as well as the SVM's confidence in the classification for AQAP (blue) for the 144 stories in the test set. As the SVM consistently predicts all AQAP and Ansar al-Shariah stories for AQAP, the expected probability of assignment to Ansar al-Shariah remains constant at approximately $p = .15$. The classification predictions provide mixed support for the expectation that an influx of local members should increase the difficulty of assigning group labels: although the SVM's confidence in predicting the AQAP label becomes more variable over time, the predicted probability of assignment to Ansar al-Shariah remains constant over the time period

4 Structural Topic Model

The three STM models are based on a corpus of 1353 documents, spanning October 25, 2005 through September 21, 2016. Approximately 500 documents are associated with as-Sahab. A histogram of the distribution of can be seen in Figure 7.

4.1 Document Pre-processing

I focus on media released online to jihadi media platforms and outlets. Preprocessing removed words that occurred in fewer than two or more 70% of the documents in the corpus.[5]

Models One, Two, and Three each use time as a covariate. The time variable is expressed

[5]In the AQAP corpus, there was no change to the number of tokens in the corpus for an upper bound threshold between 70-95%. I evaluated coherence and exclusivity at an upper threshold of 50%, but did not find results that would suggest either a coherence or exclusivity benefit from the additional reduction in corpus size.

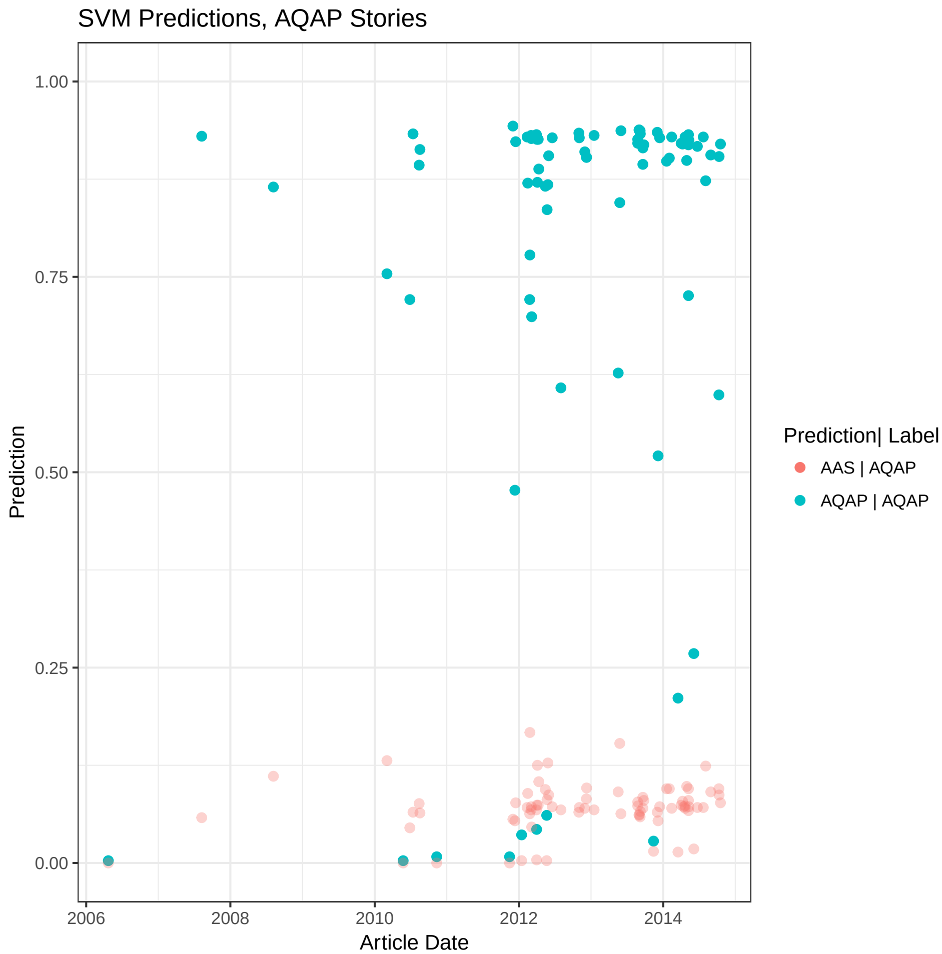

Figure 6: SVM Predictions Over Time For AQAP Stories

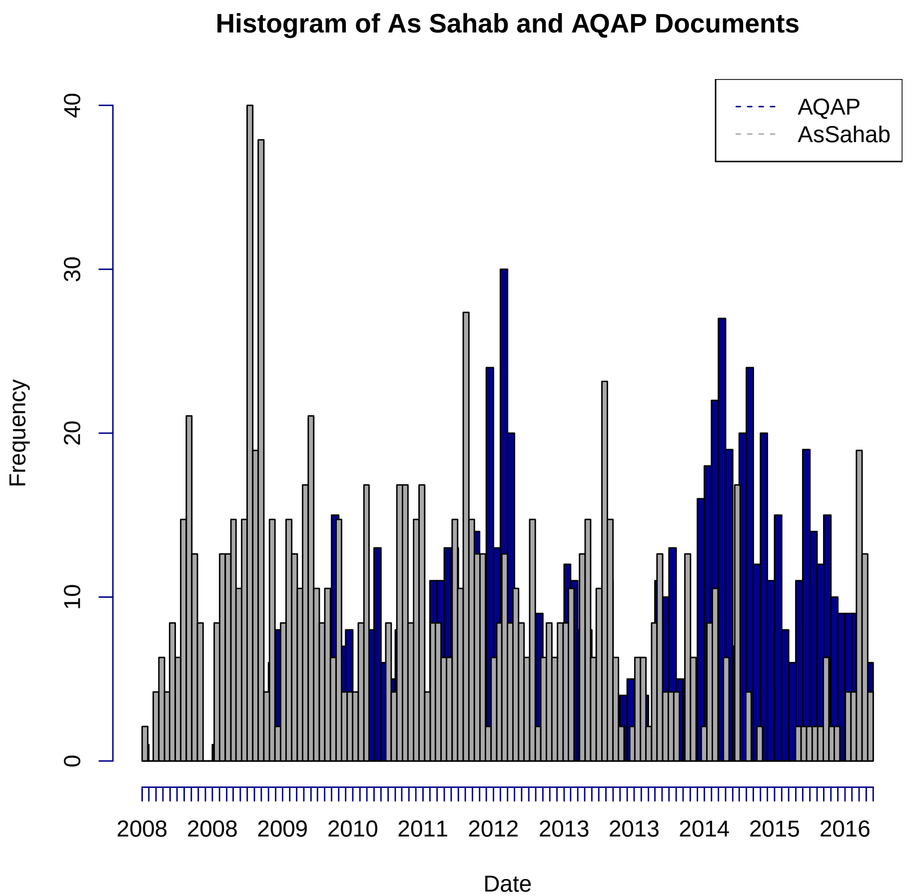

Figure 7: Distribution of AQAP and AQC Materials

in the data as a running counter of days from the oldest document in each corpus. Thus, the date of the oldest document is given as 1 and the "day" of each subsequent document is modeled as the number of days between the first date and the date of the individual document. These dates are linked to the translation date rather than release date as the former can be accurately pinpointed for each document in the corpus. For the vast majority of the corpus, the translation date closely coincides with the date that the document was released to online jihadi media outlets. The precision of release dates contained in the original Arabic text can vary according to type of document: communiques are typically dated to a specific day, while strategy documents or promotional magazine can be dated with a day, a month, or even a season. Thus, for consistency, the date covariate is linked to the translation date.

4.2 Model Selection

Topic models rely on the user to prespecify a number of topics for the algorithm to search for. However, this parameter fundamentally influences the themes that will be identified in the documents. For models one and two, I selected the number of topics by doing a sweep of model specifications with 10 to 30 topics. I selected a topic number that performed best on both semantic coherence and exclusivity.[6] After this process, a model with 18 topics appeared to present the greatest gains to semantic coherence without trading off exclusivity. Moreover, the 18-topic model identified topics that were particularly substantively coherent.

For the joint model, after comparing the semantic coherence-exclusivity trade-off for models across a sweep from 10 to 40 topics, I set the number of topics to discover at 34. The increase of topics reflects the expectation that the two organizations are already rhetorically distinct, and so the joint corpus should require more topics. Specifically choosing an output that

[6]Ideally, the selected number of topics would have relatively high exclusivity and semantic coherence. I often faced a trade-off between the two. When determining the trade-off, I prioritized semantic coherence over exclusivity. The exclusivity bands were, overall, narrow while coherence varied substantially.

doubled the number of topics slightly penalized semantic coherence over a model with fewer topics, but allowed for a more precise comparison of topics between each group.

As topic models are, by nature, non-deterministic, each implementation of a given model will produce slightly different results. Thus, after selecting the number of topics for the STM to identify, I ran each model specification ten times to create a range of possible output models for analysis. I compared the average semantic coherence and exclusivity for each of the models. For each of the three models below, I found that the averages within each ten-model set were nearly identical. To avoid biasing my results by selectively choosing the output that best confirms my theoretical expectations, I chose which specific models to analyze by maximizing average coherence and exclusivity metrics. As no model clearly dominated the coherence-exclusivity trade-off, I assigned a relatively stronger weighting to semantic coherence when selecting a specific iteration to present. I then selected a model to present before qualitatively evaluating any of the topics. This decision was intended to avoid bias in choosing how to prioritize coherence gains against exclusivity losses.[7]

Finally, after selecting which model to present, I evaluated the remaining models to ensure that the output was consistent across the set of ten results for each model. In particular, I verified general agreement on the thematic content identified across the runs.

Figure 9 depicts the expected proportion of the transnational jihadi topics presented according to time.[8] The y-axis represents the expected proportion of each document dedicated to each topic. For context, I added four vertical lines marking important dates identified above. From left to right, the lines represent: the al-Majalah airstrike on December 17, 2009; the start of the Yemeni Revolution on January 27, 2011; the death of Usama bin Laden on May 2, 2011; the end of the first Obama Administration on January 19, 2012; and the Houthi takeover of Sanaa on September 21, 2014.

[7]A plot of average semantic coherence and exclusivity scores is available upon request.

[8]For interpretability, the x-axis is labeled by year. The model was estimated according to the number of days from the start of the document corpus.

4.3 Local Conflict Topics

The “Local Conflict” cluster is comprised of four STM-identified topics. The topics are unified by a shared focus on people and places local to Yemen, as well as tactical terms that suggest military operations.

Within the cluster, specific topics are internally unified by their focus. One topic, which I have named ”Southern Yemen” is characterized by geographic references to Abyan Governorate, which was held by Islamic militants in 2012. Another, the ”Local Targets” topic has a geographic focus on Hadramawt Governorate in Central Yemen and makes frequent references to specific operational behaviors and targets. The “Houthis” local topic combines frequent temporal references with military terms, a combination that is indicative of claims of military operations.

Figure 8 shows temporal trends in the expected proportion of the corpus dedicated to each of the four topics in the “Local Topics” cluster. Notably, the “Houthi” and “Southern Yemen” topics closely track then-current events. Each topic rises in prevalence in the corpus corresponding to the dates of the respective military offensives. Moreover, that the unsupervised topic model “discovered” the military offensives provides a useful source of external validity for the 18-topic model.

4.4 Transnational Jihadi Topics

The first topic is dedicated to naming groups that the jihadi worldview considers global enemies of Islam. I term this the “Clash of Civilization” topic, as the FREX words reflect a pervasive jihadi doctrinal focus on fighting a perceived global alliance of Jews and Christians who are attempting to subjugate Muslims. The “Clash of Civilizations” topic began to decline after about 2009, an important benchmark, as the year saw a number of high-profile drone strikes that caused widespread resentment. This decline may speak to the bottom-up

Figure 8: Changes over time to attention dedicated to local topics

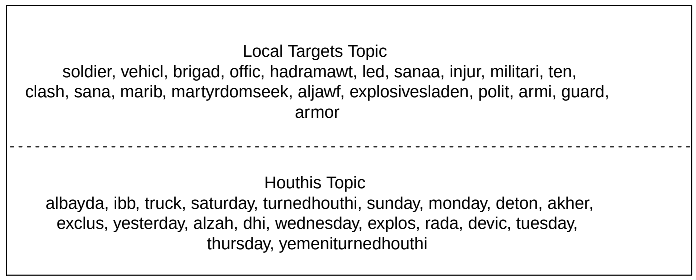

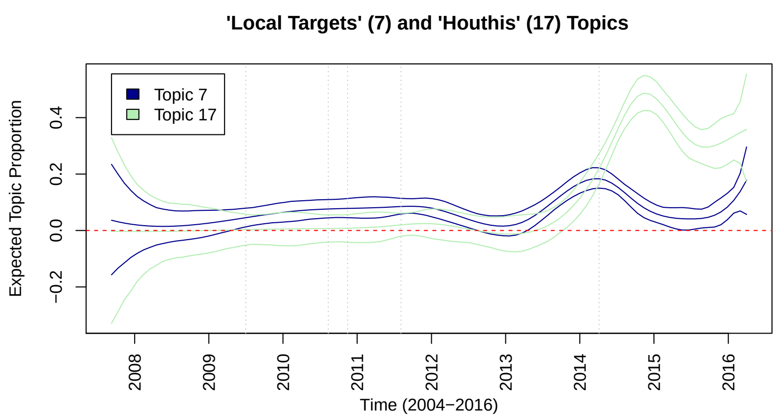

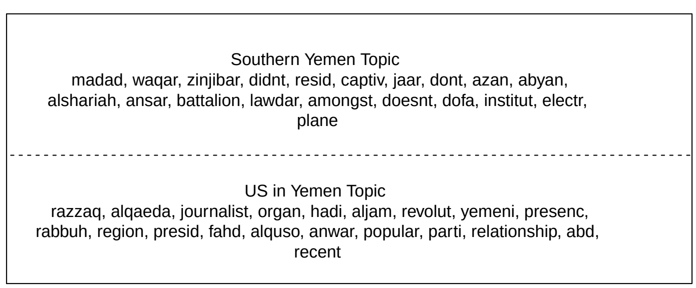

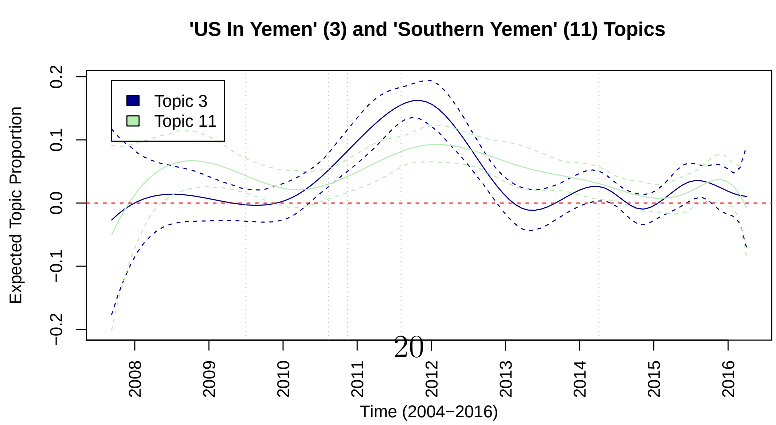

transformation of interest described above: if localizing rhetoric was driven by top-down marketing decisions, and the group turned away from global jihadi branding to capitalize on domestic frustrations, we might expect to see a sharper decline in the topic during 2009.

The second transnational jihadi topic, topic 15, centers on words that indicate an attempt to sooth conflicts among and between jihadi communities and other Muslim groups. Thus, I label this topic "Jihadi Factionalism." FREX words for the topic feature words used when attempting to recapture ideological legitimacy.

The third topic presented in Figure 9, the "Jihadi Revolution" topic, is centered around concepts used to incite for overthrow of secular governments and implementation of an Islamic theocracy. Such revolutionary rhetoric is central to the transnational jihadi view of themselves as a vanguard of social and political revolution. The topic declines briefly after the first inflection point, then rises from mid-2010 through mid-2013 before taking a more dramatic downturn at the second point. One reason why the "Jihadi Revolution" topic may not reflect the theorized "transnational" topic decline may be that the 2011 Yemeni Revolution increased the group's interest in presenting itself as a viable alternative to the Yemeni state.

The fourth transnational topic in Figure 9 is topic 2, labeled "Demise of al-Saud." The topic is focused on Saudi-centric themes, and references to Saudi officials and the state's repression of jihadi dissidents. The topic is largely stable at an expected prevalence of approximately 5% throughout the time period. I code this topic as a "transnational" topic because the Saudi-lead intervention in the Yemeni civil war, "Operation Decisive Storm," was launched only in 2015 and so for the majority of the dates analyzed, Saudi Arabia was an external target.

Comparison to local spin-off

Finally, I conclude by comparing reporting of AQAP against it's own local spin-off organization, Ansar al-Shariah (Supporters of the Shariah). Ansar al-Shariah was established in

Figure 9: Changes over time to attention given to global jihadi topics

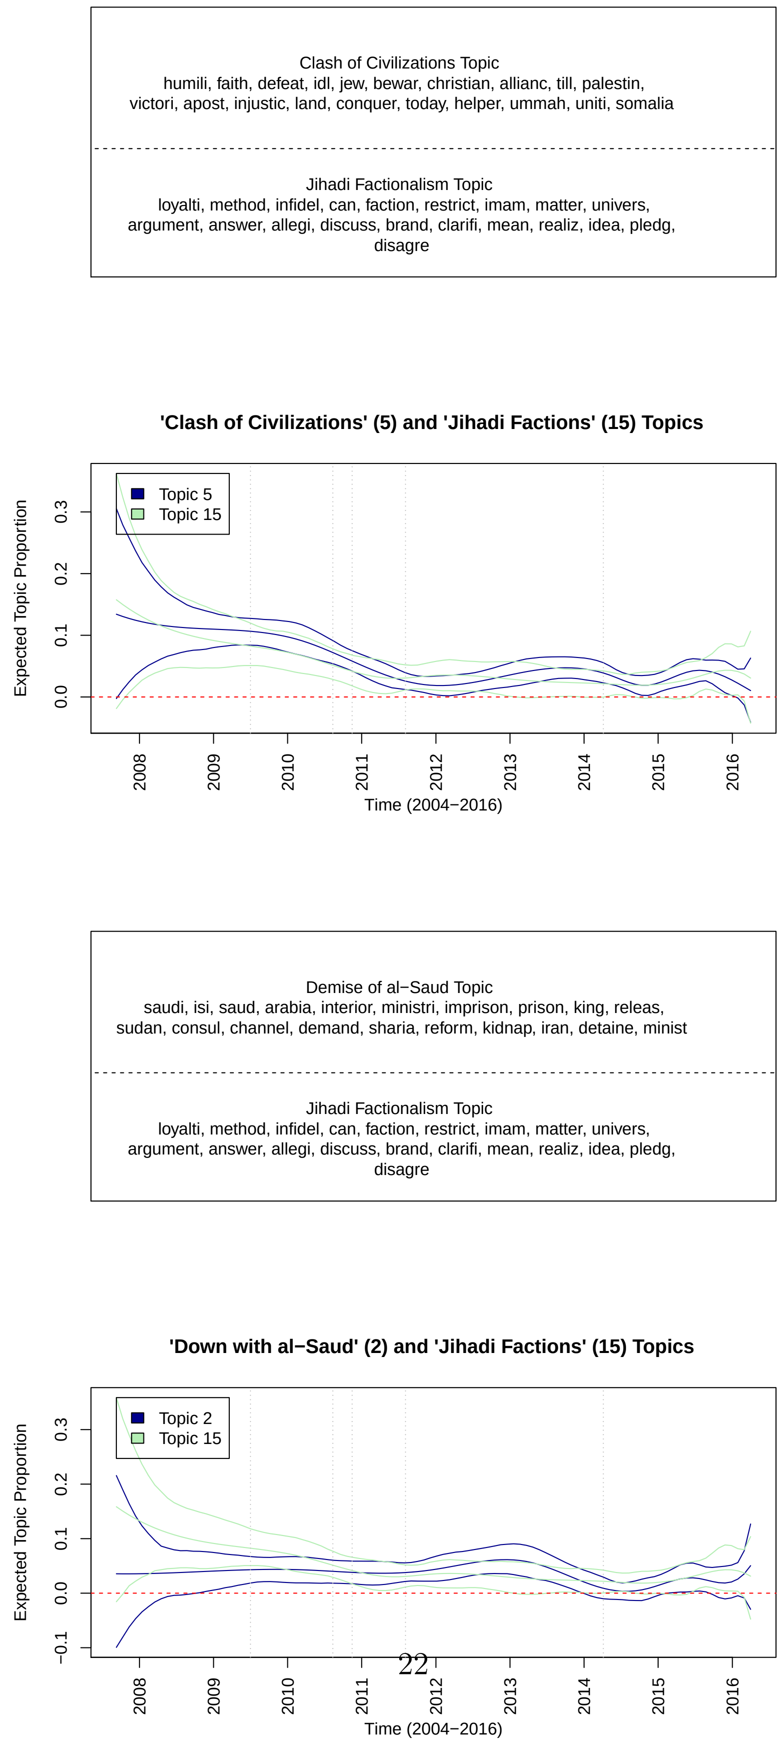

2011 as an arms-length local wing that could focus on domestic grievances and administration rather than AQAP's transnational mission and which would be free of negative local sentiment associated with the al-Qaeda brand.[9] Although quickly identified as an alias for AQAP, having two different brands provides a reference point. Under the Ansar al-Shariah name, AQAP could strike a more parochial message, exploit local grievances, and avoid the encumbrances of the al-Qaeda brand.[10] In keeping with the expectation that local recruits would be primarily invested in the local conflict, many of these fighters "have deployed exclusively for an insurgency against the Yemeni government".[11]

The transformation mechanism outlined throughout this manuscript predicts that AQAP has become increasingly constrained by growing local preferences within their rank-and-file. If this is the case, despite AQAP's attempt to create a local spin-off, their base could be expected to exert internal pressure to become more locally involved. The actions of AQAP should become similar to those of Ansar al-Shariah. Conversely, if AQAP's leaders are not experiencing internal pressures to accommodate local preferences and interests, AQAP should be expected to implement their leaders' stated preferences to clearly differentiate the globally-branded AQAP from the locally-branded Ansar al-Shariah.

An ideal quantitative test of the theory would draw on micro-level recruitment and operations data that can identify actors, tactics, and strategic priorities. Absent this systematic feature-rich organizational data, I treat news texts as a source of data that encodes observed behavior of local actors and the views of regional experts.[12]

To tighten the focus on whether AQAP is distinguishable from Ansar al-Shariah, I evaluated a sample of 566 articles from the ICEWS media database that reported on events attributed

[9]International Crisis Group 2017.
[10]Swift 2012.
[11]Human Rights Watch 2013, p. 14.
[12]Existing sources of conflict event data emphasize the accuracy of event counts over dense metadata. Although the emphasis on event de-duplification is important, exiting sources of conflict event data often do not feature granular attribution of activities to specific conflict actors. Because this project is interested in whether groups behave in a more or less similar way than absolute activity levels, my data do not need the same level of attention to deduplification of reports as most automatically-generated event data.

to al-Qaeda in the Arabian Peninsula, Ansar al-Shariah, and the Houthi militias. I developed this sample by first creating a corpus of media reports of violent activity from Yemen for 2009-mid 2015 from the ICEWS database.[13] This generated 10,818 stories, covering November 1993 through January 2015. I randomly sampled 1,772 stories of violent activities and hand-coded the primary violent-producing actor featured in the selected articles.[14] From this set 566 articles described violence attributed to AQAP, Ansar al-Shariah, or the Houthi insurgency. I then randomly divided the stories into training (67%) and test sets (33%), yielding a training corpus with 432 articles and a test corpus with 144 articles. Both document sets spanned October 30, 2002 through January 3, 2015. Each tagged story was converted into a tokenized bag of words, normalized via term frequency-inverse document frequency (tf-idf). This produced a 432 x 2,222 matrix of tokens common to both sets. I then applied a suit of machine learning classification algorithms— random forest classification, PCA decomposition, support vector machine classification, and tSNE visualization— to determine whether the algorithms could systematically differentiate among stories describing violent activities ascribed to AQAP from violent activities ascribed to Ansar al-Shariah. The classifiers failed to consistently separate stories about the two sets of actors. Finally, I used the same set of algorithms to ensure that the classifiers could differentiate both AQAP and Ansar al-Shariah from their sectarian rival, the Houthi insurgency. The full results of this analysis are accessible in the Supplemental Appendix.

The analysis suggests that local observers generally characterized AQAP and Ansar al-Shariah using similar terminology and framing. Despite investing in Ansar al-Shariah as the local face of the movement, AQAP's communiques became progressively more local in theme.

[13] Boschee et al. 2015.

[14] Approximately 15% of the data was coded by workers on Mechanical Turk, the rest was coded by the author.

5 Comparison between AQAP and al-Qaeda Central Messaging

This model addresses the counterfactual that regional and global developments may account for changes in AQAP messaging, independent of any changes in membership base. To account for the possibility that observed shifts in messaging were driven by top-down directions from al-Qaeda leadership or general trends in the jihadi environment, the model contrasts the rhetoric of AQAP with that featured in documents released by As-Sahab, a media production house closely associated with al-Qaeda's senior leadership.

AQAP's changing rhetorical style is presented alongside that of As-Sahab to establish that observed changes in AQAP rhetoric are not driven by an underlying pan-jihadi trend. This model identified themes in a corpus of 1375 documents, of which 875 were AQAP communiques and 500 were releases from As-Sahab. The model, estimated for 34 topics, predicted topic prevalence as a function of time interacted with an indicator for whether the document was authored by AQAP or As-Sahab.[15]

The patterns are generally intuitive: overall AQAP is more likely to talk about Yemen-related topics while topics that discuss other battlegrounds and targets for revolution are more associated with As-Sahab.[16] However, the divergence in thematic prevalence of pan-jihadi topics between releases issued by As-Sahab and AQAP indicate that AQAP's increasing Yemen focus was not indicative of a localizing turn lead or directed by al-Qaeda Central, via their As-Sahab mouthpiece. Two such themes are highlighted below. The two highlighted themes were chosen for their popularity among globally-minded jihadis. The first is characterized by an explicitly transnational list of FREX words: it features countries with active jihadi battlegrounds as well as references to what jihadis perceive as an American-Israeli alliance against Muslims around the world. The second topic is a relatively abstract transnational

[15] Time is included in this model as a linear function for computational tractability.

[16] A full analysis is discussed in the Technical Appendix

theme that attempts to mobilize jihadi supporters by referencing vulnerable demographics of Muslims, such as women and children, experiencing hardship and travails.

Due to the complexity of a 34-topic model, the section briefly summarizes the results of the STM model from the joint AQAP and As-Sahab corpus. Twenty-nine of the substantively interesting topics from the model are featured in Figure 10.[17]

For ease of interpretation, the topics are clustered into four general categories: those primarily relating to Yemen, topics with strong religious overtones, topics that suggest engagement with global jihadi issues, and topics that address specific countries and battlefields other than Yemen. This clustering was done on the unsupervised STM output using the author's substantive expertise. Within each cluster, topics are summarized according to the top FREX words.

Rather than present expected topic proportions, which are difficult to process visually for many topics, Figure 10 presents the topics according to their likelihood of being associated with either AQAP or As-Sahab. Positive values indicate a stronger association with AQAP's corpus, while negative values indicate topics more closely associated with As-Sahab. The point estimates of the difference in topic prevalence are presented along with 95% confidence intervals. Unsurprisingly, topics associated with battlegrounds other than Yemen, including Afghanistan, the Indian subcontinent, Libya, Somalia, and Syria, are all statistically more likely to have occurred in the As-Sahab corpus. Conversely, AQAP is significantly more likely to use terms that refer to specific events and locations in Yemen.

The two topics highlighted in the main text, named as "Crusader and Zionists" and "Defending the Weak" both occur in the "Global Jihadi Topics" cluster. Topics in this cluster are, on the whole, not statistically associated with either media group or the other, and thus group-level differences in attention to them over time are unlikely to be an artifact of

[17] Of the five topics excluded from the summary, four were associated with editing and document preparation and the fifth consists of declarations of defiance and intention. FREX words for this last topic include: "know," "think," "can," "see," "thing," "happen," and "now".

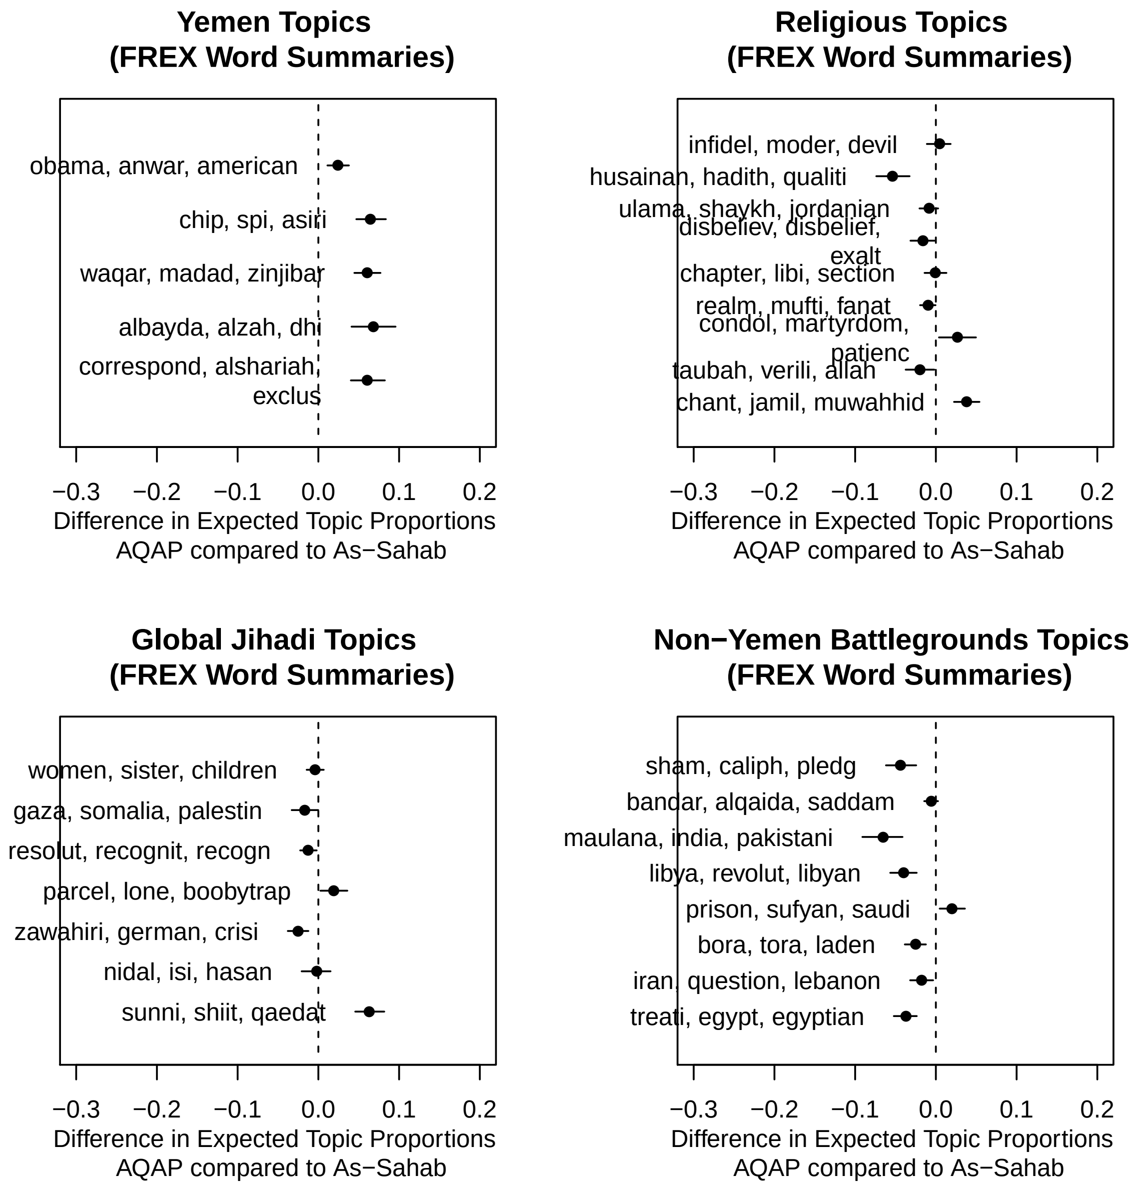

Figure 10: Summary of 34-Topic AQAP and As-Sahab Model

frequency of messages from one of the two organizations. Indeed, with a point estimate near zero and confidence intervals crossing the null line, the across-time group association of the "Defending the Weak" topic, summarized with the FREX terms "women, sister, children" is statistically indistinguishable for AQAP and As-Sahab. When disaggregating the topic prevalence across the dates of the corpus, the topic is statistically more likely to feature in later As-Sahab propaganda. Similarly, the group-association point estimates for the "Crusader and Zionists" topic is featured directly under that of the "Defending the Weak" topic. It is slightly associated with As-Sahab in the plot in Figure 10, but significantly more likely to feature in media from As-Sahab. Taken together, these results suggest that the groups started out the time period with relatively similar attention to these topics, but that over time AQAP has directed their propaganda attention completely away from this topic as they became more preoccupied with the local Yemeni conflict.

Figure 11 highlights two topic outcomes from the model. The first topic is analogous to both of the "Clash of Civilizations" topics above, here titled "Crusaders and Zionists."[18] This topic is notable for having a strong transnational jihadi focus; exactly the type of subject that AQAP should cease to discuss if their domestic recruitment is driving a local focus. Indeed we see that although al-Qaeda Central's rhetoric does increasingly feature this topic, AQAP becomes progressively less and less inclined to use words associated with the topic. Interestingly, AQAP's move away from the "Crusaders and Zionists" topic occurs even though AQAP documents were originally more likely to feature the topic than were as-Sahab documents.

Similarly, after 2011, as-Sahab documents become more likely to discuss a topic addressing alleged injustices against vulnerable populations such as women and children. The topic, which I name "Defending the Weak" expresses indignation about alleged crimes against Muslim women and children and is a pervasive theme of transnational jihadi rhetoric. Throughout

[18] As a different model, the topics are unique to this model although generally similar to the themes from Models One and Two.

the time period covered by the corpus, the topic becomes less common in AQAP documents and more prevalent in as-Sahab releases.

6 Cases and Selection Criteria

The theoretical contribution of this manuscript outlines a process by which mission drift in the militant group occurs after an expansion into a demographic that has different preferences over strategic outcomes or differences in tactics that has important knock-on effects (notably extractive/non-extractive preferences). This is a difficult research design, as information about recruitment rates as well as internal managerial capacity and dynamics are not always systematically available.

An ideal research design would include compiling systematic information about the growth and management patterns of the universe of cases of 20th-century militant groups. However, given limitations on the availability of information, it is difficult o obtain precise growth and internal dynamics information about the universe of militant groups (do I have a number for how many to expect?).

Based on my knowledge of managerial literature produced by al-Qaeda and affiliates, I take the AQAP case as a theory-building case. I observed with both my own reading and in an analysis of al-Qaeda's management a consistent complaint from dissenters that the al-Qaeda decision-makers focused on growth and expansion rather than selection (eg: Brown 2007).

To identify generality and flesh out the accommodation mechanism, I looked for other cases that shared a period of extremely rapid expansion. I identified potential cases via a combination of surveying case experts (i.e. asking whether they have seen an expansion-accommodation dynamic at work) and searching academic and practitioner literature via Google Scholar, ReliefWeb, and the CTC Sentinel. Search terms included:

- Rapid + expansion OR recruitment OR growth + militant

Figure 11: AQAP and as-Sahab Divergence

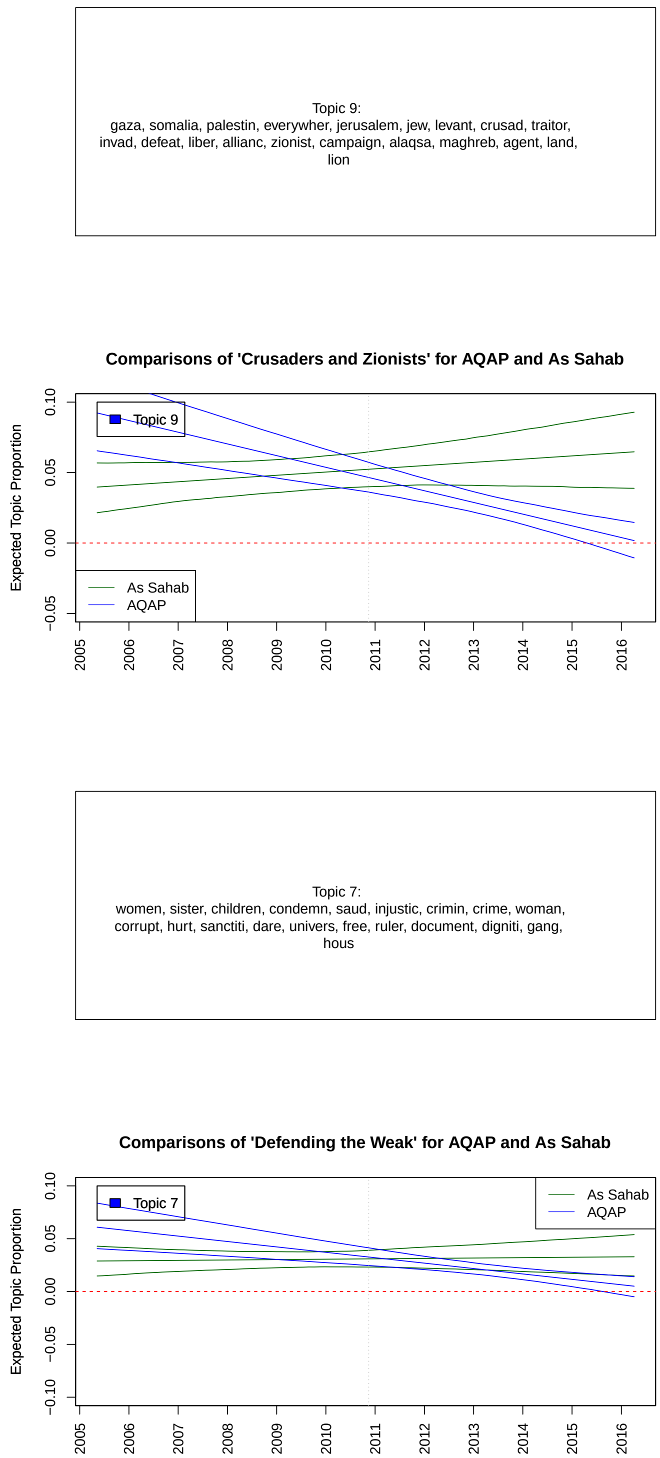

- flood + recruits + militant + groups
- influx + recruits + militant

This process generated a handful of cases of rapid growth, from which I selected to look for evidence of accommodation pressure. Specifically, I searched each result for passages that indicated specific periods of rapid intake of new members or personnel. The identified cases are listed below, along with specific passages where applicable:

Case	**Region**	**Specific Passage, if Available**	**Years**	**Citations**
al-Qaeda	Middle East	See, in particular, Walid al-Masri's criticism of recruitment as quality over quantity (Brown 2007)	1990s	Brown 2007; Fishman et al. 2008; Moghadam and Fishman 2010
Ethnic militias in Nigeria	West Africa	"The roots of the conflict lie in climate-induced degradation of pasture and increasing violence in the country's far north, owing to the rapid growth of ethnic militias, which have forced herders south"Hussona 2021	2010s -	Hussona 2021
Islamist militancy (Sahel)	North Africa	"The Sahel has experienced the most rapid increase in militant Islamist group activity of any region in Africa in recent years" Le Roux 2019, p. 1	2010 -	Le Roux 2019

PLO	Middle East	"sudden influx [that] was well beyond the factions' absorptive capacity"Szekely 2017, p. 85	1956	Brynen 2019; Sayigh 1997; Szekely 2017
Kosovo Liberation Army	Europe	"In 1996, the KLA could muster a few hundred fighters. By the end of 1998 it had over 17,000 men." The huge influx identified as linked to the Jashari Massacre (Perritt 2010, p. 100)	1990s	Koktsidis and Dam 2008; Perritt 2010
Provisional IRA		"After the massive influx of volunteers who were motivated by more temporal concerns such as revenge and anger, the PIRA took on a very different character"Kenny 2010, p. 551	1970s	Kenny 2010
Cumann na mBan (women's wing of IRA)	Europe	"The women who joined the Republican Movement in the 1960s or earlier were attracted to Republicanism by political ideology and biographical continuity, while the majority of the activists recruited after 1969 joined to fight British soldiers." Reinisch 2016, pp. 159–160	1969	Reinisch 2016

Croatian Separatists	Europe	a large influx of recruits were drawn from Croatian workers migrating to West Germany to fill a labour shortage in the 1960s, they came from underdeveloped areas that were a hotbed of Croatian nationalism tokic2009diaspora	1960s	Tokić 2009
Saudi Arabian Army	Middle East	In November 1948, pay was raised which resulted in an "influx of recruits to al-Taif." Cronin 2013, p. 33 To the magnitude of over 1,000 over a few months.	1950	Cronin 2013
Jebha	East Africa	Huge influx of Christians trying to escape repression. The organization grew from 2,000 to 10,000 with a similar influx happening the year before.	1975	Woldemariam 2018
FSLN				Mosinger 2019

Naga Groups	Southeast Asia	"Following the 1997 cease-fires, however, Naga groups saw the biggest influx of recruits in their long histories. The Naga factions more than doubled their ranks during the first few years of cease-fire (to about 10,000), and they have remained that size over the two decades since."Hanson 2021, p. 820	1997	Hanson 2021
European Islamist Militants	SE Asia	"an influx of new Western recruits into the tribal areas since mid-2006" McConnell (2008) in Cruickshank 2009, p. 1	2000s	Cruickshank 2009
Muslim Brotherhood Syria	Middle East	"[leadership centralized power] by drawing on a new influx of fighters "recruited just out of the mosques, universities and even high schools." Lefèvre 2013, p. 106 in Mosinger 2019, p. 973		Lefèvre 2013

Prosper Network	Europe	“Prosper grew quickly...This rapid expansion led Prosper to become the primary target of the Gestapo in Paris.... However, Prosper’s rapid growth came at the expense of careful security measures.” Boutton and Dolan 2021	1940s	Boutton and Dolan 2021

The identified cases are strongly skewed towards organizations based in European and the Middle Eastern, particularly Islamist groups. Potentially this is driven by rhetorical idiosyncrasies: I build the list of cases by pinpointing a specific descriptive trope — the “influx” or “flood” of recruits—which may be relatively more present in the English-language reporting during a time and place (late 2010s) that has also seen more attention to the dynamics of conflicts in the Middle East.