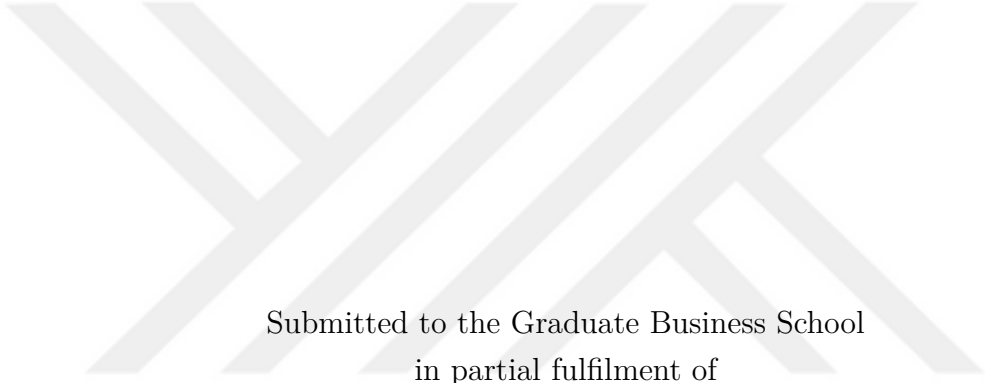


**WHO RALLIES AROUND THE FLAG? ANALYZING THE IMPACT
OF FOREIGN INTERVENTIONS ON NATIONS' POLITICAL
STANCE USING SOCIAL MEDIA DATA**

by
VAHID GHAFOURI



Submitted to the Graduate Business School
in partial fulfilment of
the requirements for the degree of Master of Science in Business Analytics

Sabanci University
June 2020


WHO RALLIES AROUND THE FLAG?ANALYZING THE IMPACT OF FOREIGN
INTERVENTIONS ON NATIONS' POLITICAL-STANCE USING SOCIAL MEDIA
DATA

Approved by:

Prof. Nihat KASAP
(Thesis Supervisor)



Assoc. Prof. Ayşe
KOCABIYIKOĞLU



Assoc. Prof. Enes ERYARSOY



Approval Date: June 29, 2020



Vahid Ghafouri 2020 ©

All Rights Reserved

ABSTRACT

WHO RALLIES AROUND THE FLAG? ANALYZING THE IMPACT OF FOREIGN INTERVENTIONS ON NATIONS' POLITICAL STANCE USING SOCIAL MEDIA DATA

VAHID GHAFOURI

Business Analytics M.Sc. Thesis, June 2020

Thesis Supervisor: Prof. Nihat Kasap

Thesis Co-advisor: Assist. Prof. Babak RezaeeDaryakenari

Keywords: foreign intervention, Iran, US, NLP, Twitter

There is a common-sense that in times of foreign interventions, the country's political actors are likely to set aside their differences and support the state or the government for a period of time as a temporary reaction to that foreign intervention. This study focuses on the specific case of Iran-U.S. conflicts to investigate the effects of events such as U.S. sanctions and military interventions on political discourse among Iranian influencers on Twitter. The quantitative approach in this study utilizes Classical Natural Language Processing to measure Iranian tweeps' sentiment towards the state across the time. We have grouped Iranian Twitter Influencers by their political affiliations to analyze which political affiliations in Iran (e.g. Conservatives, Reformists, . . .) are more likely to rally around the flag in correspondence to foreign interventions and what categories of foreign interventions have more potential in stimulating them for such reactions.

ÖZET

BAYRAK ETRAFINDA KİMLER TOPLANIR? MİLLETLERİN SİYASİ TUTUMLARI ÜZERİNDEKİ DIŞ MÜDAHALELERİN ETKİLERİNİ SOSYAL MEDYA VERİLERİYLE ANALİZ ETMEK

VAHİD GHAFOURİ

İş Analitiği Yüksek Lisans Tezi, Haziran 2020

Tez Danışmanı: Prof. Dr. Nihat Kasap

Tez Danışmanı: Dr. Öğr. Üyesi Babak RezaeeDaryakenari

Anahtar Sözcükler: dış müdahale, İran, ABD, doğal dil işleme, Twitter

Dış müdahaleler sırasında ülkenin siyasi aktörlerinin, farklılıklarını bir kenara bırakması ve bu dış müdahaleye geçici bir tepki olarak bir süre devleti veya hükümeti desteklemesi, akli selime uygun bir harekettir. Bu çalışma, ABD yaptırımları ve askeri müdahaleleri gibi olayların, İranlı etkileyicilerin Twitter'daki politik söylemleri üzerindeki etkilerini araştırmak için, İran-ABD uyuşmazlığının özel durumuna odaklanmaktadır. Bu çalışmadaki nicel yaklaşım, İranlı twitter arkadaşlarının, zaman içinde devlete karşı duyarlılığını ölçmek için Klasik Doğal Dil İşleme kullanmaktadır. İran'daki hangi siyasi bağlantıların (örneğin Muhafazakarlar, Reformcular, ...) dış müdahalelere karşı bayrak etrafında toplanma olasılığının olduğunu ve hangi çeşit dış müdahalelerin bu tür reaksiyonu uyarmada daha fazla potansiyelinin olduğunu analiz etmek için, İranlı Twitter etkileyicilerini politik bağlantılarına göre grupladık.

ACKNOWLEDGEMENTS

This thesis was supported by the help and advice of many. Firstly, I would like to deeply thank Prof. Kasap and Dr. RezaeeDaryakenari for their supervision and insight.

I'm also profoundly grateful to my father, step-mom, and brother for their inspiration and consult in critical moments. It is crystal clear to me that none of these would have been nearly possible if they weren't there for me.

I would also like to extend my sincerest gratitude to Pouya Yousefi, a friend of mine who helped me with some technical issues in LaTeX, without having any personal gain in this project. Such altruistic people are rare and precious.

Since the literature of the thesis was completed during the COVID-19 pandemic, it is notable to admire the Sabanci University's administration for their intensive efforts in managing the lock-down situation inside the campus.

Last but not least, I sincerely appreciate the Sabanci Foundation for granting me the scholarship and opportunity for my M.Sc. studies. I wish I'd be able to pay back their kindness in the future.

TABLE OF CONTENTS

LIST OF TABLES	ix
LIST OF FIGURES	x
LIST OF EQUATIONS	xiii
LIST OF ABBREVIATONS	xiv
1 INTRODUCTION	1
2 LITERATURE REVIEW	3
2.1 The Concept of 'Rally Round the Flag'	3
2.2 Methods of Inferring Social-Media Users' Ideology	5
3 DATA-SET	7
3.1 Tweeps	7
3.2 Tweets	9
3.3 Text Preprocessing	10
3.4 Stratified Sampling on the Events	10
4 METHODOLOGY	12
4.1 Sentiment Analysis	12
4.1.1 Hazm Library	12
4.1.2 TF-IDF	14
4.1.3 Classification	15
4.1.4 The Challenge of Simultaneous Oversampling & Cross- Validation	15
4.1.5 Grid Search	16
4.1.6 Accuracy on the Test-Set	18
4.2 Hypothesis	20
5 RESULTS	21
5.1 Bubble-Charts	21

5.1.1	Aggregated Bubble-Charts	26
5.2	Average of Likes	28
5.3	Regression Analysis	31
6	CONCLUSION & FUTURE WORK	35
	BIBLIOGRAPHY	38
	Appendices	40
A	GRID-SEARCH PLOTS	41
B	CONFUSION MATRICES	43
C	BUBBLE CHARTS	46
D	AGGREGATED BUBBLE CHARTS	51

LIST OF TABLES

Table 3.1 Stratified-Sampling's scheme	11
Table 4.1 Accuracy-Measures of each classifier (1st Approach)	18
Table 4.2 Confusion-Matrix for Logistic Regression on the whole test- data (1st approach)	18
Table 4.3 Accuracy of the 1st approach on each affiliation Separately . .	19
Table 4.4 Accuracy of the 2nd approach on each affiliation separately . .	19
Table B.1 Naive-Bayes Classifier on the whole test-set (1st approach) . . .	43
Table B.2 Random Forest Classifier on the whole test-set (1st approach) .	43
Table B.3 Logit on the Conservatives (1st approach)	43
Table B.4 Logit on the reformists (1st approach)	44
Table B.5 Logit on the Transitionists (1st approach)	44
Table B.6 Logit on the Non-Politicals (1st approach)	44
Table B.7 Logit on the Unclears (1st approach)	44
Table B.8 Logit on the Overthrowers (1st approach)	44
Table B.9 Logit(C=2) on the Conservatives (2nd approach)	44
Table B.10 Logit(C=0.4) on the Reformists (2nd approach)	45
Table B.11 Logit(C=0.64) on the Transitionists (2nd approach)	45
Table B.12 Logit(C=0.2) on the Non-Politicals (2nd approach)	45
Table B.13 Logit(C=1) on the Unclears (2nd approach)	45
Table B.14 Logit(C=5) on the Overthrowers (2nd approach)	45

LIST OF FIGURES

Figure 3.1	Proportion of Affiliations Among the Tweeps	8
Figure 3.2	Proportion of Twitter page types	8
Figure 4.1	Grid-Search for tuning the best regularization-parameter on Logit	17
Figure 4.2	Grid-Search for tuning the best regularization-parameter on Naive-Bayes	17
Figure 5.1	The week before U.S. pulled out of the JCPOA	23
Figure 5.2	The week U.S. pulled out of the JCPOA (we anticipate the rally here)	23
Figure 5.3	The week after U.S. pulled Out of the JCPOA	23
Figure 5.4	The week before Iranian seizure of the British oil-tanker	24
Figure 5.5	The week of Iranian seizure of the British oil-tanker	24
Figure 5.6	The week after Iranian seizure of the British oil-tanker	24
Figure 5.7	The week before Iranian shoot-down of the U.S. drone	25
Figure 5.8	The week of Iranian shoot-down of the U.S. drone	25
Figure 5.9	The week after Iranian shoot-down of the U.S. drone	25
Figure 5.10	The week before U.S. pulled out of the JCPOA	27
Figure 5.11	The week U.S. pulled out of the JCPOA (we anticipate the rally here)	27
Figure 5.12	The week after U.S. pulled out of the JCPOA	27
Figure 5.13	Average of likes - U.S. pulling out of JCPOA	29
Figure 5.14	Average of likes - U.S. designation of IRGC as an FTO	29
Figure 5.15	Average of likes - Trump’s tweet about the ‘End of Iran’	29
Figure 5.16	Average of likes - Iranian shoot-down of American drone	30
Figure 5.17	Average of likes - UK’s seizure of Iranian-flagged ship in Gibraltar	30
Figure 5.18	Average of likes - Iranian’s seizure of the British oil-tanker	30
Figure 5.19	Average of likes - U.S. imposition of sanction on Javad Zarif	31
Figure 5.20	Regression Analysis - U.S. pulling out of JCPOA	32
Figure 5.21	Regression Analysis - U.S. designation of IRGC as an FTO	32

Figure 5.22 Regression Analysis - Trump's tweet about the 'End of Iran' .	32
Figure 5.23 Regression Analysis - Iranian shoot-down of American drone .	33
Figure 5.24 Regression Analysis - UK's seizure of Iranian-flagged ship in Gibraltar	33
Figure 5.25 Regression Analysis - Iranian's seizure of the British oil-tanker	33
Figure 5.26 Regression Analysis - U.S. imposition of sanction on Javad Zarif	34
Figure A.1 Grid-Search for tuning the best regularization-parameter for Conservatives	41
Figure A.2 Grid-Search for tuning the best regularization-parameter for Reformists	41
Figure A.3 Grid-Search for tuning the best regularization-parameter for Transitionists	42
Figure A.4 Grid-Search for tuning the best regularization-parameter for Non-Politicals	42
Figure A.5 Grid-Search for tuning the best regularization-parameter for Overthrowers	42
Figure C.1 The Week Before U.S. designated JCPOA as an FTO	46
Figure C.2 The week U.S. designated JCPOA as an FTO	46
Figure C.3 The week after U.S. designated IRGC as an FTO	47
Figure C.4 The week before Trump's tweet about 'the End of Iran'	48
Figure C.5 The week of Trump's tweet about 'the End of Iran'	48
Figure C.6 The week after Trump's tweet about 'the End of Iran'	48
Figure C.7 The week before UK seizure of the Iranian oil-tanker	49
Figure C.8 The week of UK seizure of the Iranian oil-tanker	49
Figure C.9 The week after UK seizure of the Iranian oil-tanker	49
Figure C.10 The week before U.S. imposition of sanctions on Javad Zarif	50
Figure C.11 The week of U.S. imposition of sanctions on Javad Zarif	50
Figure C.12 The week after U.S. imposition of sanctions on Javad Zarif	50
Figure D.1 The Week Before U.S. Pulled Out of the JCPOA	51
Figure D.2 The week U.S. pulled out of the JCPOA (we anticipate the rally here)	51
Figure D.3 The week after U.S. pulled out of the JCPOA	52
Figure D.4 The week before U.S. designated IRGC as an FTO	53
Figure D.5 The week U.S. designated IRGC as an FTO	53
Figure D.6 The week after U.S. designated IRGC as an FTO	53
Figure D.7 The week before Trump's tweet about 'the End of Iran'	54
Figure D.8 The week of Trump's tweet about 'the End of Iran'	54

Figure D.9	The week after Trump’s tweet about ‘the End of Iran’	54
Figure D.10	The week before Iranian shoot-down of the U.S. drone	55
Figure D.11	The week of Iranian shoot-down of the U.S. drone	55
Figure D.12	The week after Iranian shoot-down of the U.S. drone	55
Figure D.13	The week before UK seizure of the Iranian oil-tanker	56
Figure D.14	The week of UK seizure of the Iranian oil-tanker	56
Figure D.15	The week after UK seizure of the Iranian oil-tanker	56
Figure D.16	The week before Iranian seizure of the British oil-tanker	57
Figure D.17	The week of Iranian seizure of the British oil-tanker	57
Figure D.18	The week after Iranian seizure of the British oil-tanker	57
Figure D.19	The week before U.S. imposition of sanctions on Javad Zarif	58
Figure D.20	The week of U.S. imposition of sanctions on Javad Zarif	58
Figure D.21	The week after U.S. imposition of sanctions on Javad Zarif	58

LIST OF EQUATIONS

4.1 TF-IDF	15
4.2 C Parameter in Logistic Regression	17
4.3 Each Tweep’s Sentiment Toward the State in Each Week	20
4.4 The Change of Each Tweep’s Sentiment Toward the State in the Week Comparison to Its Previous Week	20
4.5 Hypothesis	20
5.1 Regression equation	31

LIST OF ABBREVIATIONS

FTO Foreign Terrorist Organization	x, xi, 1, 29, 32, 34, 35, 46, 53
IRGC Islamic Revolutionary Guard Corps.....	x, xi, 1, 29, 32, 34, 35, 53
IRGC-QF Islamic Revolutionary Guard Corps - Quds Force.....	1
JCPOA Joint Comprehensive Plan of Action .	x, xi, 23, 27, 29, 32, 34, 35, 46, 51, 52
NLP Natural Language Processing.....	2
NLTK Natural Language Toolkit	12
P5+1 UN Security Council’s five permanent members; namely China, France, Russia, the United Kingdom, and the United States; plus Germany	35
TF-IDF term frequency–inverse document frequency	vii, 14
VOA-PNN Voice of America Persian News Network	1

1. INTRODUCTION

On May 08, 2018, the U.S. president, Donald Trump, declared its withdrawal from the Joint Comprehensive Plan of Action, also known as the Nuclear Agreement between Iran and P5+1, calling it “rotten and decaying” raising concerns about not only the sanctions’ impact on the Iranian people and global market but also the possibility of a tension outbreak in the classical tense relationships between U.S. and Iran. One year later, on April 08, 2019, the White House designated Iran’s Islamic Revolutionary Guards Corps (IRGC) as a foreign terrorist organization (FTO) which makes it the first time Washington has officially used this term for a foreign country’s military. Meanwhile, the use of hostile language among US-Iran’s officials on social media mounted to its peak when Trump threatened the “End of Iran” as a response to a missile attack on the US embassy in Baghdad which Iranian-backed forces were accused of. Later on, events such as the Iranian shoot-down of a U.S. drone and the UK’s seizure of Iranian-flagged ship in Gibraltar fueled the tensions in a limited scope. Furthermore, when the U.S. imposed sanctions on MohammadJavad Zarif, Iranian foreign minister, it occurred to many as if the gates of diplomacy between the two countries is going to be totally closed for a while.

US-Iran tensions raised to its peak, and evolved into a direct conflict after U.S. military forces allegedly assassinated Gen. Qasem Soleimani, the commander of Iran’s Islamic Revolutionary Guard Corps-Quds Force (IRGC-QF) and one of Iran’s most important military commanders, in an airstrike in Baghdad on January 2, 2020.

Although Iran had been facing intense economic protests just a few months prior to the U.S. strike, a surprisingly vast number of people in Iran and Iraq came to the streets mourning the death of Qassem Soleimani and Abu-Mahdi Al-Muhandis. Many of the conventional opposition figures (e.g. foreign minister of the previous regime, Ardashir Zahedi; the former executive editor of VOA-PNN, Mohammad Manzarpour) condemned the attack supporting the Iranian state’s position. This has risen an interesting case-study of the phenomenon called ‘Rally Round the Flag’.

There is abundant literature focusing on the impact of foreign interventions on democracy and human rights. For instance, Peksen (2011) argues that the chances of extra-judicial killing, disappearance, political imprisonment, and torture increase under supportive and neutral interventions, and hostile interventions merely increases the probability of political imprisonment. However, due to the recent events happening in Iran, we found our interest in analyzing the impact of foreign intervention on a country's internal political discourse and public sentiment towards the state.

More specifically, we are curious to see that in reaction to what types of foreign intervention cases, what certain political affiliations and camps are likely to set aside their differences and support the state or the government for at least a period of time.

Our work is a Social-Media-Analysis which mainly focuses on observing the 'Rally Round the Flag' phenomenon on Twitter by looking into the variations in behaviour of the accounts of Iranian Influencers. Despite some of the previous who focus the network analysis of tweeps, we use a content-based approach that looks into the sentiments of the tweets generated by the tweeps and tries to measure its variations across the time. Using a semi-supervised NLP algorithm, we show that the average of positive sentiments among Persian Twitter influencers, increases in three out of seven cases of foreign conflict, which we interpret as empirical proofs for 'rally round the flag' effect in our case-study. Moreover, we detect the effect in tweeps from each political affiliation separately to provide further intuition on the groups who are more prone to the phenomenon.

2. LITERATURE REVIEW

In this chapter, we will initially review the previous literature on the concept of 'rally round the flag' effect. Then we will look into the methods used for ideology measuring in social-media and the reason we decided to use a content-based approach rather than the more prevalent network-based methods.

2.1 The Concept of 'Rally Round the Flag'

The 'Rally Round the Flag' is a term suggested by John Muller for discussing the temporary effect of popular support for the President of the United States during periods of international crisis or war (Goldstein, Pevehouse & Sernau, 2008). It is a common belief that in times of major national conflicts American people are likely to set aside their disagreements with the incumbent presidents' policies or performance in the office to demonstrate a united front to the international community (Baker & Oneal, 2001).

Mueller (1970) proposed that the 'Rally Round the Flag' effect would be emerging as a response to an event with three qualities:

1. "Is international"
2. "Involves the United States and particularly the President directly"
3. "Specific, dramatic, and sharply focused"

In addition, Mueller split the rallies into five categories:

1. "Sudden US military intervention" (e.g., Korean War, Bay of Pigs Invasion)"
2. "Major diplomatic actions" (e.g., Truman Doctrine)

3. "Dramatic technological developments" (e.g., Sputnik)
4. "US-Soviet summit meetings" (e.g., Potsdam Conference)
5. "Major military developments in ongoing wars" (e.g., Tet Offensive)

Although Mueller's three-part definition of the rally effect has been widely accepted by other scholars, his five-category list of rally-inducing events is thought to be old-fashioned by modern political scientists due to their heavy reliance on Cold War events (Hetherington & Nelson, 2003).

By extending the historical time-frame and embracing the Nixon, Ford, Carter, and Reagan administrations, Richard Brody (1991) found that events in the latter two of Mueller's five categories namely, major military developments in ongoing wars and U.S.-Soviet summits were as likely to be causing drops in presidential approval as resulting in boosts (Hetherington & Nelson, 2003).

Although most of the previous literature on the 'rally round the flag' effect has been focusing on the U.S. administration as their case-study, Wood (2008) investigates the 'rally round the flag' effect alongside democracy and human-rights' conditions in countries which have been facing U.S. or UN sanctions, arguing that economic sanctions impose political, social, and physical difficulty on civilians by pushing incumbents toward imposing further repression. However, it still does not offer empirical support at the sub-national level. The recent US-Iran conflicts have drawn lots of attention to the potential of a sub-national empirical case-study on Iran's political atmosphere.

The vast number of Iranian people coming to streets as a protest to US assassination of general Qassem Soleimani and their surprisingly massive amount of participation in his funeral, which took place just a few months after the oil-price protests, demonstrates Iran as a good case-study around the topic of 'Rally Round the Flag'. In this research, we apply quantitative methods on Twitter to explore the rally round the flag in Iran's case for seven Iranian foreign conflict cases prior to the assassination of Qassem Soleimani. Unfortunately, due to the early time-frame of our research, the specific case of Qassem Soleimani is not included in our analysis and left for future works.

2.2 Methods of Inferring Social-Media Users' Ideology

In this study, we will review the previous literature on methods of ideology measurement and try to explain the factors which encouraged us to choose a content-based approach for our task:

In order to measure variations of sentiment toward state among twitter users across the time, we initially need to come up with an assessment method of a tweep's ideology using social media. The literature on methods of inferring social-media users' attributes based on their network positions is limited, though prosperous (Barberá, Jost, Nagler, Tucker & Bonneau, 2015). According to Barberá et al. (2015) although it is said that one's ideology could be an estimator of his/her social and political behavior (Jost, 2006), very limited work has been done to measure ideology by utilization of social-media data (e.g. Back, Stopfer, Vazire, Gaddis, Schmukle, Egloff & Gosling, 2010; Conover, Ratkiewicz, Francisco, Gonçalves, Menczer & Flammini, 2011; Kosinski, Stillwell & Graepel, 2013).

For instance, Weber, Garimella & Batayneh (2013) is an analysis run on Egyptian tweeps to study the phenomenon of secular vs. Islamist polarization in twitter considers the following as key findings while measuring tweeps' partisan group:

- **Retweeting signifies endorsement.** One way is to utilize the simple retweet data to check whether the tweep belongs to an Islamist or secular partisan group.
- **Similar user sets.** On average, the personal traits of tweeps in both parties follow akin distribution; not only in terms of activity rate but also gender distribution and demographics.
- **Polarized hashtags can be identified.** One could assign a polarity score to hashtags with the aim of recognizing the politicized topics automatically.
- **Tension over time.** Following the polarization and “distance from the center” in the set of all hashtags across time (i.e., how much a hashtag is solely adopted by one political affiliation or conversely how diverse it is among the affiliations) could provide us with an overall understanding of political conflict which could not be identified by simple measures of volume or per-user polarity trends.
- **Community structure in retweeter graph.** It is possible to observe and measure a tendency for people with similar affiliations to connect to each other

in their retweet network (Weber et al., 2013).

Looking into individual tweeps' network positions, could be useful for tasks who focus mainly on density of interactions among various tweeps' with different ideologies (e.g. Weber et al., 2013) or the ones who aim to approximately measure the similarity of tweeps' ideology or exposure of a certain group of certain tweeps to some other tweeps' ideas in a network (Larson, Nagler, Ronen & Tucker, 2019).

However, the deficiency of this method for our specific task is that it does not assess the aspect of which two tweeps may have moved towards each other. Two users might retweet from each other at some point but it looks problematic to identify that in what sense they have found a common ground in their discourse which has motivated them to retweet from one-another. For example, a pro-regime and an anti-regime feminist might retweet a tweet about the "metoo" case from each other which shouldn't be interpreted as their agreement on supporting the state. Indeed, in our task, the content of the tweet plays a much more crucial role.

Since we want to detect the Rally-Round-the-Flag sentiment in each tweet over time, we need to know if that similarity of opinion means that both tweeps are supporting the state, or is it another issue they have both agreed on. Thus, we had to try a new approach that focuses on the content of tweets rather than network positions of the tweeps.

3. DATA-SET

In recent years, Twitter has become one of the main sources of political news and interactions among citizens and political authorities. Thus, it could somewhat be claimed that Iranian twitter data could reflect the variations in Iran’s political atmosphere (Hajizadegan, 2019). Accordingly, our dataset contains of two sections, one of which embraces the set of influencer tweeks and the other one contains the set of tweets generated by those tweeks. In this chapter we will cover the structure of these two datasets.

3.1 Tweeks

We used a list of Iranian Twitter influencers gathered by Hajizadegan & Jalaeipour (2018). In that study, the term ‘influencer’ was attributed to the tweeks who satisfied the following criteria:

1. Iranian tweeks are of their main audiences
2. Possess more than 10000 followers
3. The number of their followers are above two times more than their followings

A dataset of 1765 Iranian twitter influencers was gathered at the end of the day. The dataset also contained 33 columns representing different attributes of each tweek. The ones we utilized for our tasks were as follows:

- **Political Affiliation:** The dataset had divided the tweeks into six categories of political affiliation. In a rough descending sorting based on their support for the state, they are tagged as Conservatives, Reformists, Transitionists, Non-Politicals, Overthrowers (literal translation of the Persian word ‘barandaz’, meaning that they intend to overthrow the regime) and Separatists. Figure

3.1 shows how the political affiliations are distributed among the tweeps in respect to their number of followers (Hajizadegan & Jalaeipour, 2018).

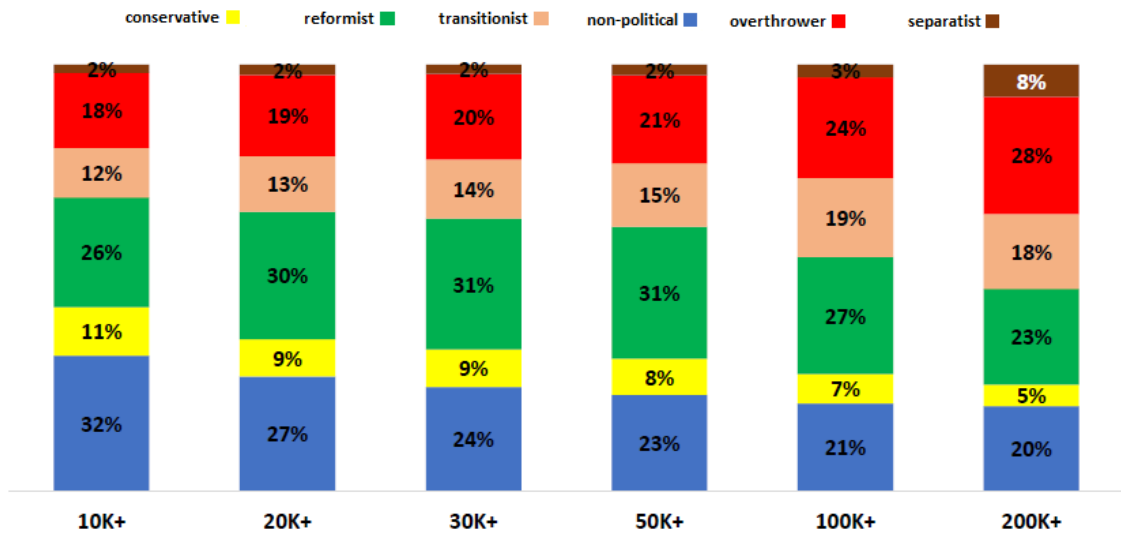


Figure 3.1 Proportion of Affiliations Among the Tweeps

- Organization/Individuals:** There are both tweeps who are affiliated with an organization (e.g. news channels, NGOs, etc.) and tweeps who are individuals sharing their personal views. The proportions of each affiliation's population is depicted in 3.2. As you see, the more we move toward more popular pages, the proportion of organizational pages increases (Hajizadegan & Jalaeipour, 2018):

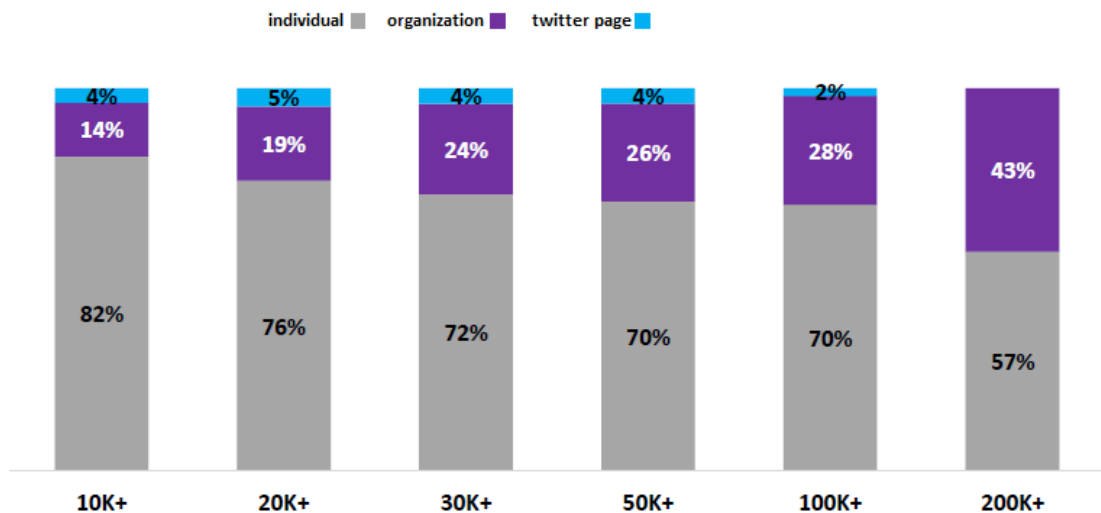


Figure 3.2 Proportion of Twitter page types

- Number of Followers:** Well... That's what it was.

We initially dropped separatists from our list due to their data-sparsity. We also omitted organizational tweeps since we are interested in studying the public opinion which is not supposed to be reflected in a tweet from a news channel or other organizations. At the end of the day, we were left with 1091 tweeps.

3.2 Tweets

For this project we worked on a dataset of 3,477,585 tweets which consists of all the tweets by the 1091 influencers from 08/20/2014 to 08/20/2019. The dataset is assembled in a csv file format obtained by Dr. Babak RezaeeDaryakenari as part of his project on Political Actors Social Media Accounts (PASMA). It consist of 48 columns which the more important ones for our task are briefly explained here:

- **'authorscreen_name'**: The user ID of the tweep as a string
- **'tweetfull_text'**: Tweet's full text as a string
- **'tweetcreated_at'**: Tweet's date and time of creation as a string in format of "yyyy-mm-dd hh:mm:ss".
- **'tweetin_reply_to_screen_name'**: A string showing the tweep ID of the person the tweet has been replied to. NaN if the tweet is not a reply.
- **'tweetlang'**: The language of the tweet in a two-character-string format (e.g. 'fa' for Farsi, 'en' for English)
- **'tweetretweet_count'**: The total number of retweets (integer)
- **'tweetfavorite_count'**: The total number of likes (integer)
- **'retweet_status'**: A boolean value which is 'True' if the tweet is a retweet and 'False' if the tweet is an original one.
- **'retweeted_statusauthorscreen_nam'**: A string showing the tweep ID of the person the tweet has been retweeted from. NaN if the tweet is not a retweet (i.e. as if 'Retweet_status'=False)
- **'hashtags'**: A string containing all the hashtags in the tweet within; separated by comma.

3.3 Text Preprocessing

We initially removed the data of pre-Trump's presidency (2016-11-08 00:00:00) to focus further on one of the most tense eras of the U.S.-Iran relationships. This has left us with 3,021,129 tweets. We also dropped the reply-tweets from our dataset being left with 1,972,404 tweets at the end of the day. Furthermore, as a prior step to our sentiment analysis, we defined a function which omitted all the punctuation-marks, numbers, carriage-returns and newline commands and links from our tweets' texts and converted all capital letters to lower-case letters and half-spaces to full-spaces. This type of generalization might have caused losing a very minimal set of information but contributed to our efforts in handling the data-sparsity problem.

3.4 Stratified Sampling on the Events

In the beginning, we took a random sample from all the tweets in the dataset for the labelling step. However, we faced the problem of extreme sparsity of +1 class tweets. It seems quite normal that in regular conditions, people are very unlikely to make a positive comment about their government.

To tackle this problem, we stratified our data by separating the three days around every seven foreign conflicts in our analysis. Then we selected a random sample in each stratification based on the table 3.1. This increased the proportion of +1 class to a more appropriate state. Later, it was followed by the upsampling method explained in the Methodology section and saved the day for our imbalanced data classification.

Events	from	till	Sample Size
Trump's pulling out of the JCPOA	5/8/2018 0:00	5/11/2018 0:00	800
The White House designation of IRGC as an FTO	4/8/2019 0:00	4/11/2019 0:00	800
Trump tweeting about the end of Iran	5/19/2019 0:00	5/22/2019 0:00	800
Iranian shoot-down of American drone	6/20/2019 0:00	6/23/2019 0:00	800
UK's seizure of the Iranian-flagged ship in Gibraltar	7/4/2019 0:00	7/7/2019 0:00	800
UK's seizure of Stena Impero by Iran	7/19/2019 0:00	7/22/2019 0:00	800
US imposition of sanctions on Zarif	7/31/2019 0:00	8/3/2019 0:00	800
General Sample			1000

Table 3.1 Stratified-Sampling's scheme

4. METHODOLOGY

In this chapter we will take a general look into our task of sentiment analysis and the tools and libraries we have used for this purpose next to the challenges we faced. At the end we will report on our accuracy for our different classification techniques in our two different approaches.

4.1 Sentiment Analysis

In order to prepare our dataset of tweets for a supervised-learning sentiment-analysis, we manually labelled around 7000 tweets in our dataset as -1, 0, +1 which corresponds to negative, neutral and positive sentiment toward the state. Then we trained Classical Natural Language Processing algorithms on our labeled section of data, and predicted the sentiment of all 2 million unlabeled tweets in our dataset. Consequently, we followed the path elaborated in the following subsection to fulfill our classification task.

4.1.1 Hazm Library

Hazm is a Python library for digesting Persian text. It is compatible with NLTK library and similarly could perform the following tasks on Persian texts (Sobhe, 2018):

- Text-Cleaning
- Text-Normalizing

- Sentence and Word Tokenizing
- Word-Stemming
- Word-Lemmatizing
- POS-Tagging
- Shallow-Parsing
- Dependency-Parsing
- Interfaces for Persian Corpora

Since we are looking into a dataset of tweets consisting of both Persian and English texts, we utilized the Hazm library as an alternative/complement of NLTK in Python by taking advantage of its word-lemmatizing and word-tokenizing ability.

However, one of the main issues with the current libraries for Persian texts is their lack of support for the informal language and different writing habits. For instance, in the following picture, Hazm is doing a fine job in lemmatizing the following word which is a formal and correct way of writing ‘I am going’:

```
▶ hazm.Lemmatizer().lemmatize('می‌روم')
```

☞ 'رفت#رو'

However, it is quite prevalent for tweeps on twitter to use informal Persian while writing. As you could see in the following picture, Hazm is unable to lemmatize the following word which is simply the informal version of the previous word.

```
▶ hazm.Lemmatizer().lemmatize('میرم')
```

☞ 'رم\u200cمی'

Moreover, some people do not use half-space in unofficial texts for comfort. They either ignore the half-space or replace it with a full-space. As you see, still Hazm is unable to support their lemmatization:



```
hazm.Lemmatizer().lemmatize('میرم')
```

```
↳ 'مرد#میر'
```



```
hazm.Lemmatizer().lemmatize('می رم')
```

```
↳ 'می رم'
```

Existence of these types of inconsistencies in writing which are the case for the majority of Persian verbs, and the inability of online tools such as Hazm in detecting them would cause undesirable sparsity in our Bag-of-Words matrix and consequently a less informative TF-IDF matrix. Future developments of such online tools like the Hazm library could cause a significant boost in Natural Language Processing tasks on Persian texts.

4.1.2 TF-IDF

Initially, after removing Persian and English Stop-Words, we extracted the text's TF-IDF features. TF-IDF is a common technique in the fields of information retrieval and text mining to evaluate the relationship for each word in the collection of documents.

The TF part in TF-IDF stands for term-frequency which refers to the occurrence of specific words in documents. Apparently, the more frequent a term in a document the more important it could be. On the other hand, the DF stands for document-frequency which implies the number of times a specific word appears in the collection of documents. Terms which have higher DF value are less likely to be important since they commonly appear in all documents (e.g. stop-words and common verbs like: did, to, be, etc.). Thus, IDF that is an inverse of DF, measures the importance of words in all documents. High IDF posits that the word is rare in all documents, therefore we might want to pay more attention to it (Kim & Gil, 2019).

For a term t in a document d , the weight W_{td} of term t in document d is given by

(Góralewicz, 2018):

$$W_{td} = TF_{td} \times \log\left(\frac{N}{DF_t}\right) \quad (4.1)$$

4.1.3 Classification

The labelled data was split into 80% train-set and 20% test-set. A 5fold cross-validation was run on the train-set to tune the best regularization-parameters for the classifiers based on the highest F1-Score (weighted). Among classical machine-learning algorithms, Logistic-Regression outperformed the rest and was chosen as our main classifier (see table 4.1).

We also tried another approach which we named Affiliation-Based-Approach; meaning that we classified tweets from each political affiliation separately rather than classifying them all at once. This approach made a slight improvement in our F-Score for tweets of Conservatives and Overthrowers(the radical anti-regime group). This could be somewhat intuitive considering that radical pro and anti-regime groups usually have a more distinct discourse and choice of words than other affiliations. Therefore, classifying them separately would make the classification easier for a classifier which is solely based on Bag-of-Words.

We finally classified Conservatives and Overthrowers based on the second approach and the rest of the affiliations based one the first approach.

4.1.4 The Challenge of Simultaneous Oversampling & Cross-Validation

One of the most common ways of handling imbalanced datasets is oversampling. In our dataset tweets labelled as +1 are pretty rare which is quite intuitive since people are usually less likely to state a positive sentiment about their government in usual cases.

In our initial confusion matrix for test-set, we were not able to detect any True-Positive or even False-Positive +1 labels with our best classifiers. This pushed us to try increasing the percentage of +1 labels firstly by stratified sampling from the weeks which are potentially suspicious for a 'Rally Round the Flag' case and sec-

only by performing oversampling in our labelled train-data. This provided us with reasonable accuracy for +1 class in a proportionate way to other classes' accuracy.

However, the latter practice has a downside. When we do oversampling in our train-set, a consequent cross-validation overestimates the test-set's accuracy. The clear reason is that oversampling creates repetition in a dataset on its discriminated classes. Thus, when running a k-fold cross-validation, the classifier is already being trained on a copy of some tweets in the k-1 training folds which makes it able to easily predict the other copy of those tweets in the 1 remaining test fold. This will make the cross-validation on the train-set unrepresentative of the test-set's accuracy.

In order to tackle this problem, we needed to adjust our steps' sequence as follows:

1. Splitting the train-set into k folds
2. Setting a regularization-parameter
3. Setting the first fold as the validation-set
4. Oversampling the data in the remaining k-1 training folds
5. Training the classifier on the oversampled k-1 folds with a regularization-parameter
6. Calculating the F-Score of our classifier on the validation-set
7. Setting the next fold as the validation-set
8. Returning to step 4 (If we aren't in the Kth fold)
9. Updating the regularization-parameter
10. Returning to step 3

For the coding part, we used utilized pipeline from *imblearn* library for this purpose akin to (Martin, 2019).

4.1.5 Grid Search

Since we used a 5-fold cross-validation and checked 10 regularization-parameters for each classifier, we did repeat oversampling and training every 50 times for each classifier. Figures 4.1 and 4.2 depict the plots for our grid search in tuning the regularization parameter for each classifier. For each plot, we have the classifier's

regularization parameter on the x-axis versus its corresponding F-Score(weighted) on the y-axis.

For Logistic-Regression we tuned the parameter named C which is the inverse of regularization parameter - A control variable that retains strength modification of regularization by being inversely positioned to the Lambda regulator.

Given how Scikit cites it as being:

$$C = \frac{1}{\lambda} \tag{4.2}$$

The relationship would be that lowering C - would strengthen the Lambda regulator which is the penalty factor for the linear optimization of the classifier (Rusin, 2019).

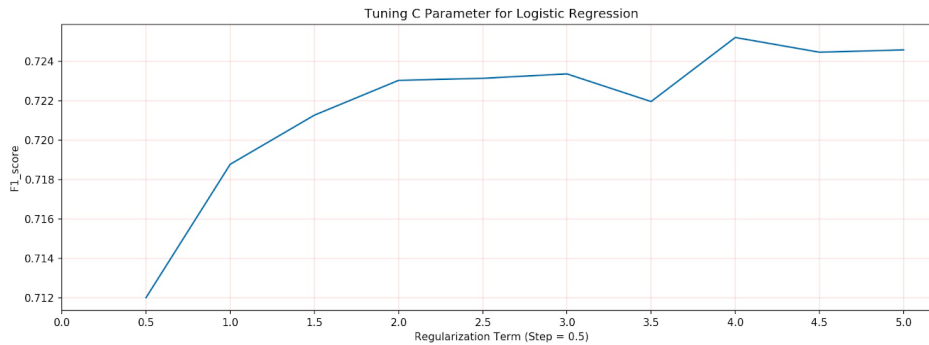


Figure 4.1 Grid-Search for tuning the best regularization-parameter on Logit

The alpha parameter in Naive Bayes is the smoothing parameter which is there to initially deal with the sparsity of data which causes zero parameters.

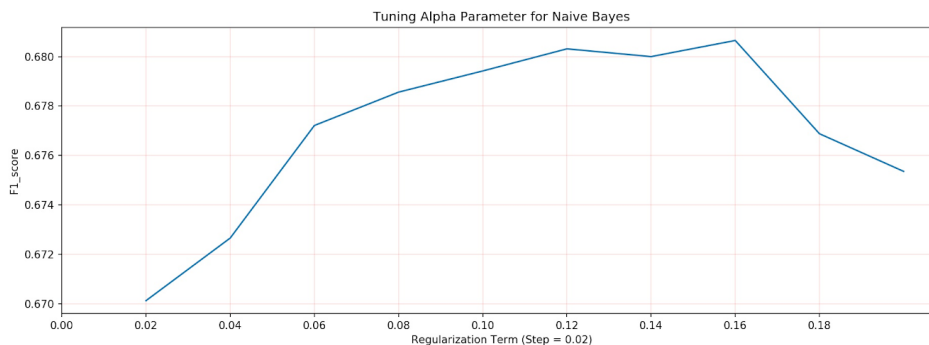


Figure 4.2 Grid-Search for tuning the best regularization-parameter on Naive-Bayes

4.1.6 Accuracy on the Test-Set

For performance-assessment, we initially looked into the F-Score (weighted) provided by each of the classifiers which convey the balance between the precision and the recall. For our first classification-approach (i.e. classifying tweets of all affiliations altogether) the accuracy measures of each classifier on the test-set are mentioned in table 4.1 (i.e. all metrics are set as *weighted*):

Classifier \ Accuracy Measure	Precision	Recall	F1_Score
Logit (C=4)	0.725	0.724	0.723
Multinomial-Naive-Bayes (alpha=0.16)	0.696	0.631	0.644
Random-Forest-Classifer	0.717	0.727	0.707

Table 4.1 Accuracy-Measures of each classifier (1st Approach)

Our results show compatibility with some of the previous literature claiming advantage for Logistic Regression over Naive Bayes Classifier when dealing with TF-IDF data (Pranckevicius & Marcinkevičius, 2017).

Since we are dealing with an imbalanced multi-class dataset, we are supposed to show the confusion matrices for a more realistic assessment of our classification power. Table 4.2 shows our confusion-matrix for our best classic classifier which is logistic regression with C=4. The rest of the confusion matrices could be found in Appendix B.

		Predicted		
		Sentiment	-1	0
Actual	-1	223	108	41
	0	79	628	58
	+1	30	47	99

Table 4.2 Confusion-Matrix for Logistic Regression on the whole test-data (1st approach)

In order to be able to compare our first classification approach with the affiliation-based one, we are also reporting its accuracy metrics and confusion matrices for each affiliation separately. Table 4.3 shows the precision, recall, and f-score for different affiliations in the test-set using the best-performed classifier of the first approach which is Logistic Regression with C=4.

As mentioned earlier, we also tried a second approach which we named as the 'affiliation-based' approach. Based on the intuition that people of the same affiliation have a more-or-less similar discourse and subsequently are more likely to use

Political Affiliation \ Accuracy Measure	Precision	Recall	F1_Score
Conservatives	0.621	0.587	0.589
Reformists	0.644	0.642	0.638
Transitionists	0.714	0.703	0.707
Non-Politicals	0.904	0.874	0.887
Overthrowers	0.793	0.695	0.724
Unknown	0.856	0.844	0.843

Table 4.3 Accuracy of the 1st approach on each affiliation Separately

similar words, we grouped the tweets by their political affiliation and tokenized their tweets separate from other groups. Then, as in the first approach, we tuned the best regularization parameter for the Logistic-Regression classifier on them and classified them separately. This means that we trained 6 more Logistic-Regression classifiers. Table 4.4 shows a report of the second approach’s accuracy on each affiliation in a comparable manner with the first approach’s report in table 4.3.

Political Affiliation \ Accuracy Measure	C	Precision	Recall	F1_Score
Conservatives	2	0.628	0.625	0.623
Reformists	0.4	0.620	0.607	0.611
Transitionists	0.64	0.644	0.645	0.644
Non-Politicals	0.20	0.849	0.871	0.859
Overthrowers	5	0.728	0.735	0.730
Unknown	1	0.664	0.697	0.667

Table 4.4 Accuracy of the 2nd approach on each affiliation separately

The rows highlighted as green in the 2nd approach are the ones who seem to outperform the 1st approach in terms of their F-Score. This outperformance holds for the Conservatives and the Overthrowers which are both known as the radical pro-regime and anti-regime camps. This is kind of an intuitive phenomenon since the more radical the political affiliation, the more unique their discourse might be; which makes it a better idea to classify them separately as done in the 2nd approach. Thus, we classified the Reformists, Transitionists, Non-Politicals, and Unknown ones by the first, and the Conservative and Overthrowers by the second approach.

At last, the F1-Score for every affiliation on the test-set was higher than 60% and the overall F1-Measure was higher than 70%. This accuracy was obtained on a super-messy dataset consisting of both Persian and English texts with plenty of misspelled words without space in Persian tweets. Next to that, the prevalence of both formal and informal Persian in the tweets and the lack of an appropriate library for catching those variations, made our sentiment-analysis task even more challenging.

After labeling all the tweets in the dataset as -1, 0, and +1, in each week of our analysis, we took the average of each tweep’s sentiments in that week (i.e. a tweep with three tweets of +1 sentiment and one tweet of -1 sentiment was calculated as +0.5 in that specific week). We kept track of the average of tweeps’ sentiments and tried to observe their changes in our interested consecutive weeks.

4.2 Hypothesis

What we are anticipating to observe in the times of foreign interventions is a significant increase in the average of tweeps’ sentiments.

Let $s_{i,t}$ denote the sentiment of a tweet from tweep i in time t and $n_{i,w}$ denote the total number of tweets from tweep i in week w . Thus we could calculate $S_{i,w}$ which is mean sentiment of tweep i in week w :

$$S_{i,w} = \left(\sum_{t \in w} s_{i,t} \right) / n_{i,w} \quad (4.3)$$

$$\Delta S_{i,w} = S_{i,w} - S_{i,w-1} \quad (4.4)$$

Hypothesis 1. In the weeks of foreign interventions the mean of the distribution of the parameter $\Delta S_{i,w}$, which denotes the changes made in tweep i ’s average sentiment in week w in comparison to week $w - 1$, shall be greater than zero in a statistically significant manner for either all the tweeps or some political affiliations.

$$\mathcal{H}_1 : \Delta S_{i,w} > 0 \quad (4.5)$$

The regression-analysis results in the conclusion section illustrate that for which tweep’s political affiliation and in what sort of events our hypothesis holds.

5. RESULTS

In this chapter, we will initially visualize our results of tweeps' sentiment toward the state and their possible shifts in cases of foreign intervention by some bubble charts. Then we measure the significance of the shifts in each event and our null-hypothesis by running a regression-analysis on the shifts of each affiliation in each foreign conflict event. Furthermore, we also try running an analysis of "likes" and investigate its possible implications for our research question.

5.1 Bubble-Charts

Before testing the statistical significance of our hypothesis, we came up with a visualizing idea of portraying the variations of our tweeps' sentiment towards the state.

We proposed a bubble chart for each week of our analysis, having the average of each tweep's sentiment towards the state in the x-axis, the total likes he/she have received on the y-axis, the total number of tweets as the size of the bubbles and the political affiliation on the colors.

Visually speaking, at the first glance, this chart could provide the audience with several implications regarding the 'rally round the flag'-status in the corresponding week while comparing it with its previous week's bubble-chart:

- If a specific color (affiliation) generally shifts to the right side - in comparison to its previous week's position - this week is a likely case of 'rally round the flag' for that political affiliation (i.e. that political affiliation is supporting the state on the issues happening this week).
- If a specific color generally mounts higher in the plot, we could imply a more

public favor towards the corresponding affiliation in that week which can give us a notion of 'rally round the flag' in the public tweeps rather than the mere influencers ones. For instance, if conservatives gain more 'likes' in a week, we could deduct that their position was welcomed and strengthened in those turn of events and probable sort of 'rally round the flag' in public may be inferred.

- The more the chart is populated with bigger bubbles of a color, the more the corresponding affiliation has been proactive in that week. This could be having something to do with their influencers feeling a more formidable position to talk.

Thus, in general, we would expect that in a week prone to a 'rally round the flag', we see a general shift to the right side - especially from moderate affiliations, Reformists and Transitionists, represented by blue and purple -, a higher altitude and size for red bubbles (Conservatives), and a lower altitude and size for black bubbles (Overthrowers).

The bubble-charts portrayed in figures 5.1, 5.2, 5.3, 5.4, 5.5, 5.6, 5.7, 5.8 and 5.9 depict some of the foreign-conflict weeks which had a closer behaviour to meeting our expectation. You could find all the remaining charts in the appendix.

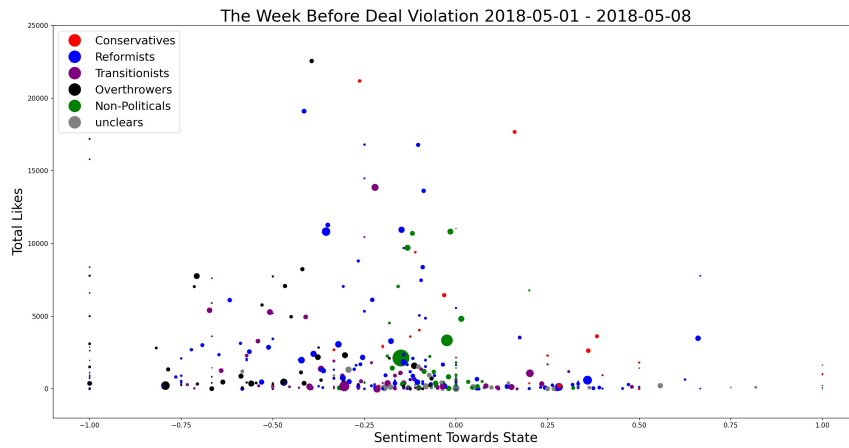


Figure 5.1 The week before U.S. pulled out of the JCPOA

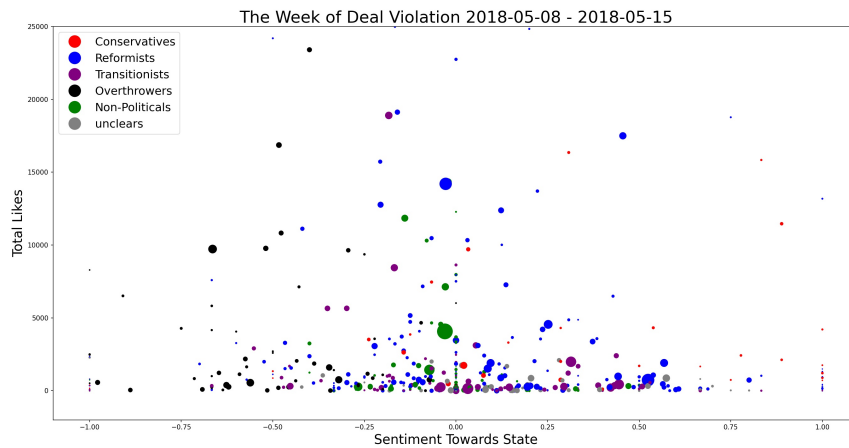


Figure 5.2 The week U.S. pulled out of the JCPOA (we anticipate the rally here)

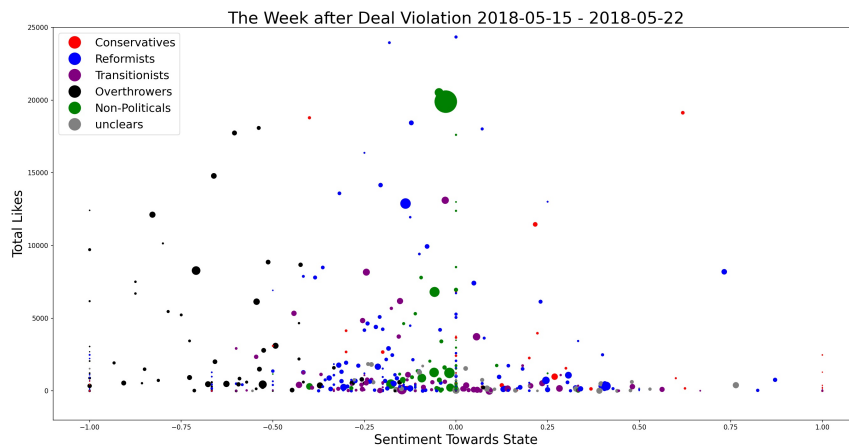


Figure 5.3 The week after U.S. pulled Out of the JCPOA

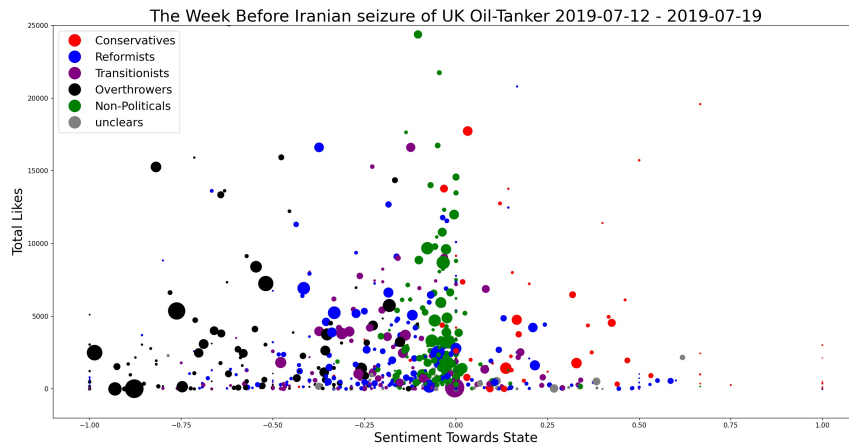


Figure 5.4 The week before Iranian seizure of the British oil-tanker

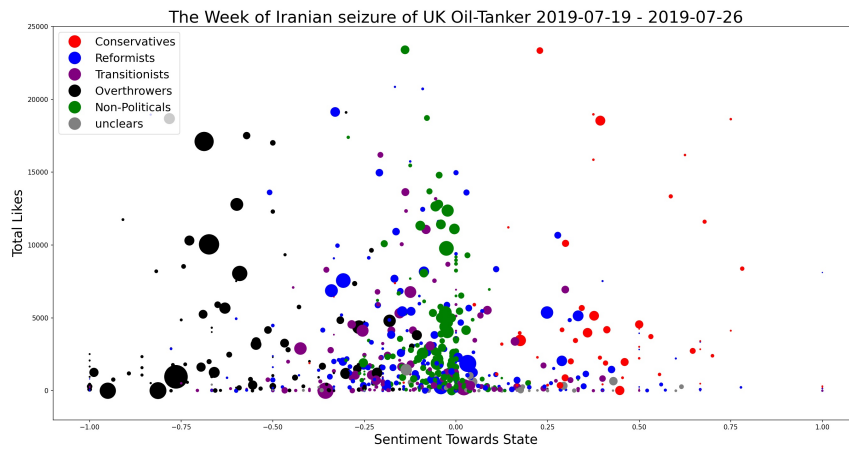


Figure 5.5 The week of Iranian seizure of the British oil-tanker

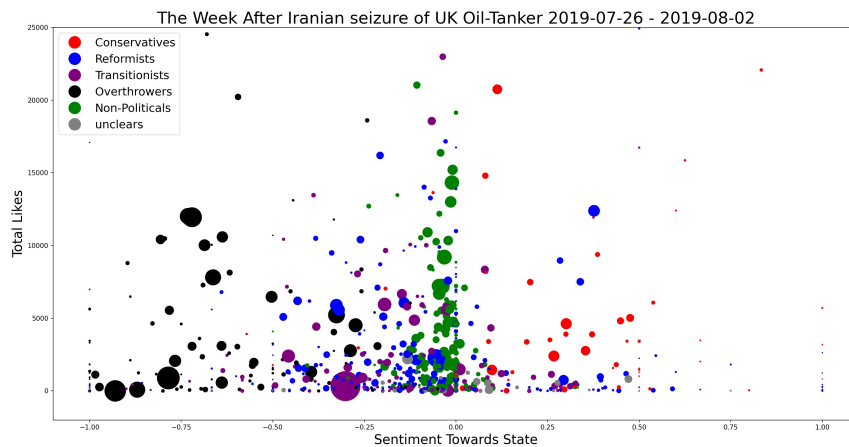


Figure 5.6 The week after Iranian seizure of the British oil-tanker

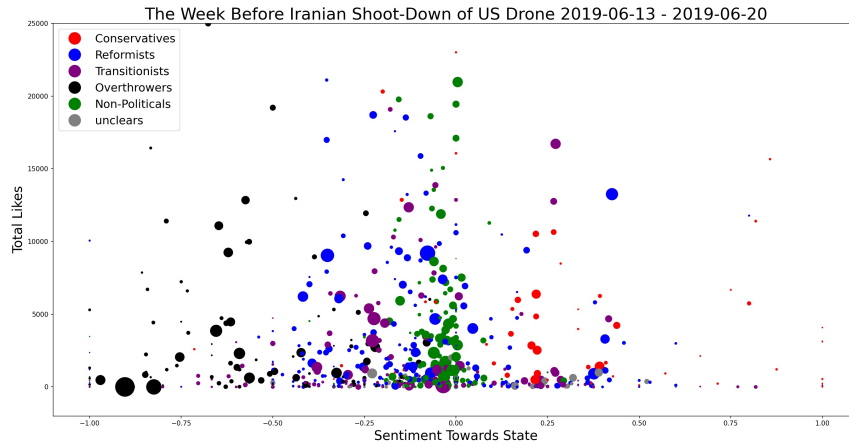


Figure 5.7 The week before Iranian shoot-down of the U.S. drone

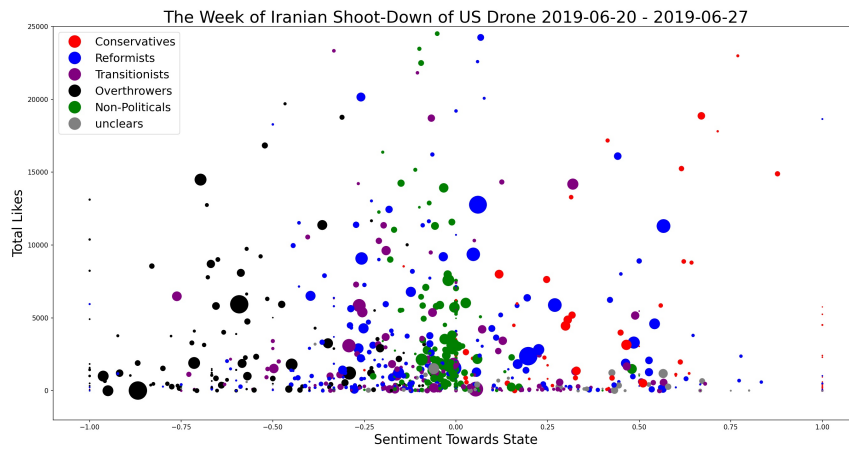


Figure 5.8 The week of Iranian shoot-down of the U.S. drone

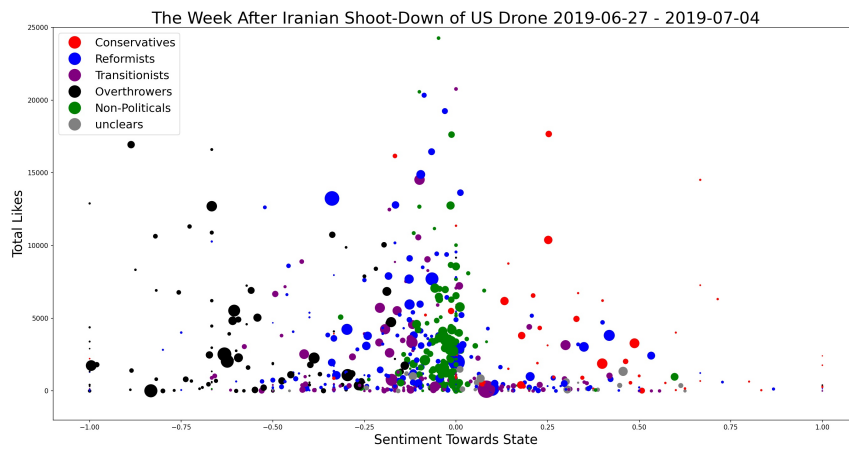


Figure 5.9 The week after Iranian shoot-down of the U.S. drone

5.1.1 Aggregated Bubble-Charts

To tackle the problem of the messiness of the bubble-charts and get a more collective view of the affiliations' behaviour rather than individual tweeps, we also imposed the aggregated version of bubble-charts as an extension to the regular ones.

In this set of charts, each affiliation is represented just by one collective dot/bubble. Thus, the x-axis value of each bubble shows the average sentiment of all tweets generated by all the tweeps from the corresponding affiliation of the bubble. Similarly, the y-axis, shows the average number of 'likes' that all tweets from that affiliation have won, and finally, the size of the bubbles represents the total number of tweets from that affiliation. Colors represent political affiliation just as before.

Figures 5.10, 5.11 and 5.12 show one case of our aggregated bubble-charts as a sample. The rest could be found in Appendix C.

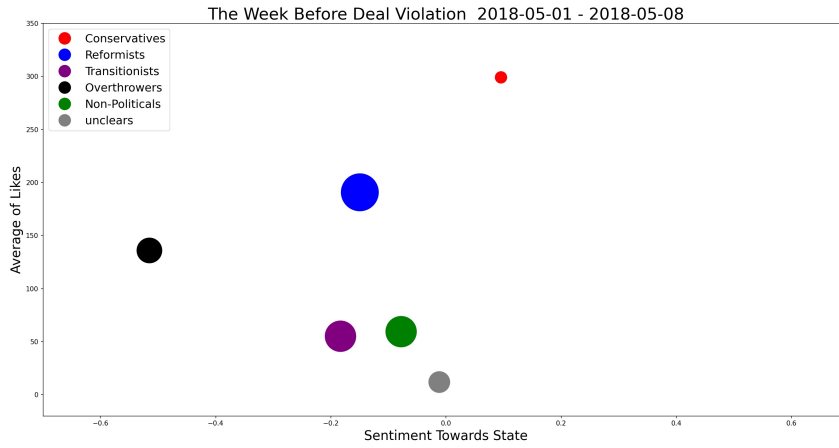


Figure 5.10 The week before U.S. pulled out of the JCPOA




Figure 5.11 The week U.S. pulled out of the JCPOA (we anticipate the rally here)




Figure 5.12 The week after U.S. pulled out of the JCPOA

A simple comparison between figure 5.10 and figure 5.11 shows that the Conservatives plus all moderate oppositions have rallied around the flag. However, we see a slight decrease in the altitude (average likes) of the Reformist camp positing a possible falter in their popular support. The intuition could be that the Nuclear Deal between Iran and 5+1 countries was signed by the Reformist government of President Hasan Rouhani. Rationally, the one-sided violation of the deal by the U.S. would be regarded as a sign of incompetency in the Reformist administration and discourse.

5.2 Average of Likes

The Bubble-Charts mainly focus on the political behaviour of the Influencers in cases of foreign interventions. Further than that, we are interested in investigating the public behaviour in twitter in such cases.

The 'likes', or formerly called 'favs', are one of the key elements for quantifying the public endorsement for a concept on Twitter, or more broadly, on the Social Media. Looking at the *total* 'likes' of each affiliation may give us some intuition about the support for their position at a specific time. However, the total 'likes' is correlated with the number of tweets generated by Influencers which is a parameter solely related to influencers' behaviour. Since we are interested in looking at the political support of camp while controlling the attitude of the Influencers, we consider the *average* of 'likes' as our measurement for popular support.

Figures 5.13, 5.14, 5.15, 5.16, 5.17, 5.18 and 5.19 depict the variations of average of 'likes' for different political camps in day around foreign conflicts. Our hypothesis is that we are likely to see an increase in Conservatives' mean of 'likes' for most cases.

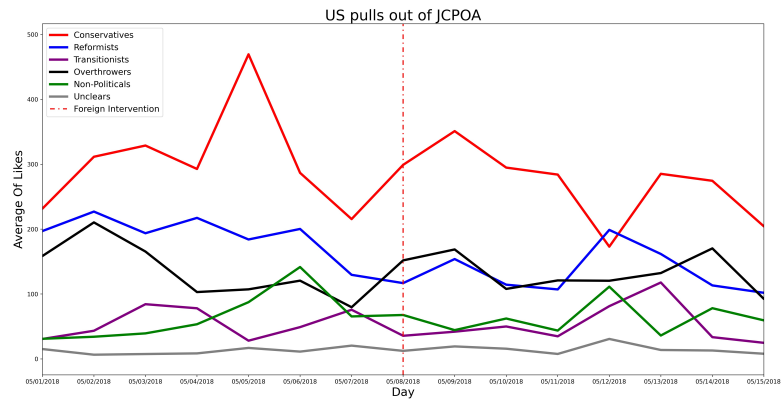


Figure 5.13 Average of likes - U.S. pulling out of JCPOA

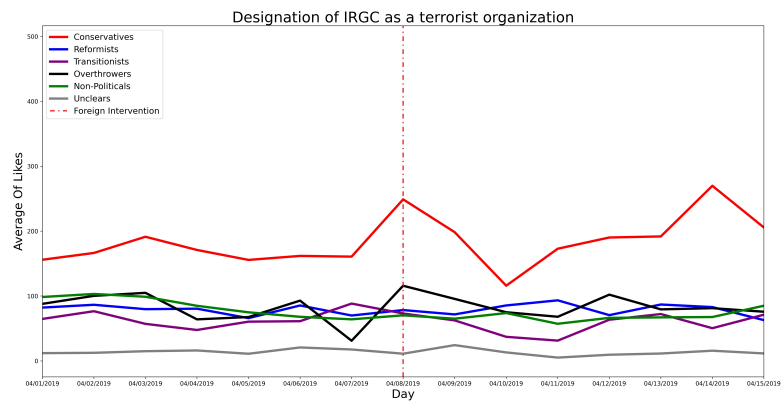


Figure 5.14 Average of likes - U.S. designation of IRGC as an FTO

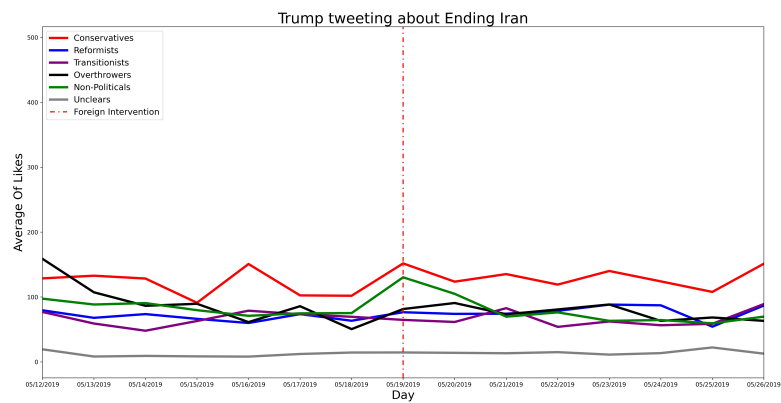


Figure 5.15 Average of likes - Trump's tweet about the 'End of Iran'

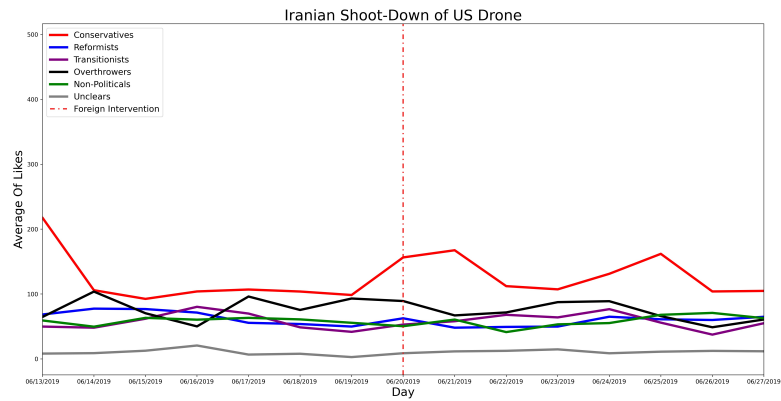


Figure 5.16 Average of likes - Iranian shoot-down of American drone

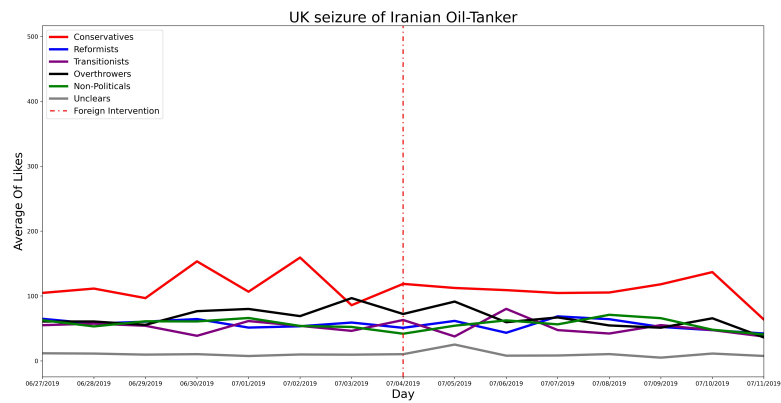


Figure 5.17 Average of likes - UK's seizure of Iranian-flagged ship in Gibraltar




Figure 5.18 Average of likes - Iranian's seizure of the British oil-tanker

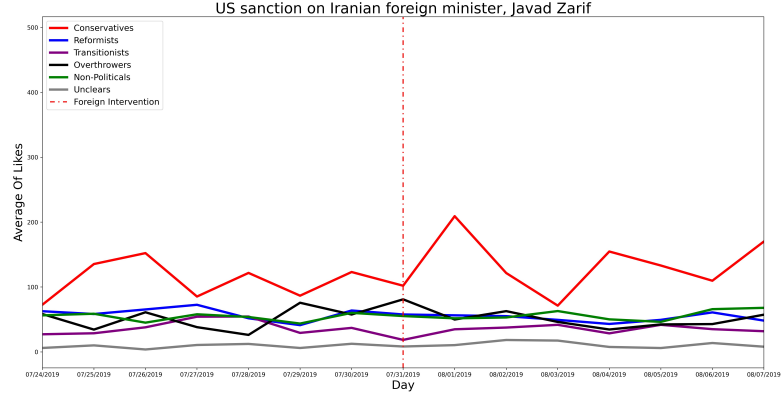


Figure 5.19 Average of likes - U.S. imposition of sanction on Javad Zarif

5.3 Regression Analysis

The previous charts provide the reader with an initial glance of 'Rally Round the Flag' status in times of foreign-interventions. To statistically measure the significance of the changes according to our hypothesis mentioned in equation 4.5, we ran a regression analysis on the tweets' sentiments as our descriptive measurement. Considering the linear equation 5.1, we inserted the time (week) as our independent-variable and the average of sentiment (S) for each tweet in week w as our dependent-variable. Thus, the coefficient of the linear regression illustrates the average changes in the tweets' sentiments as in equation 4.4.

$$S_w = aw + b \quad (5.1)$$

By looking at the coefficients' confidence-intervals, we could quantify the significance of the changes and examine our hypothesis on different events.

Figures 5.20, 5.21, 5.22, 5.23, 5.24, 5.25 and 5.26 show the 95% confidence-intervals for changes in sentiment towards the state across all the seven events of our analysis. The cases with a confidence-interval above zero are naturally considered as likely cases of foreign-intervention.

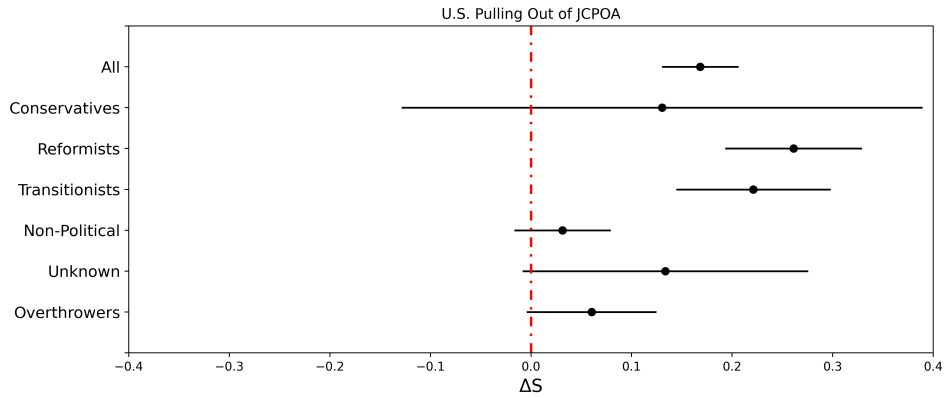


Figure 5.20 Regression Analysis - U.S. pulling out of JCPOA

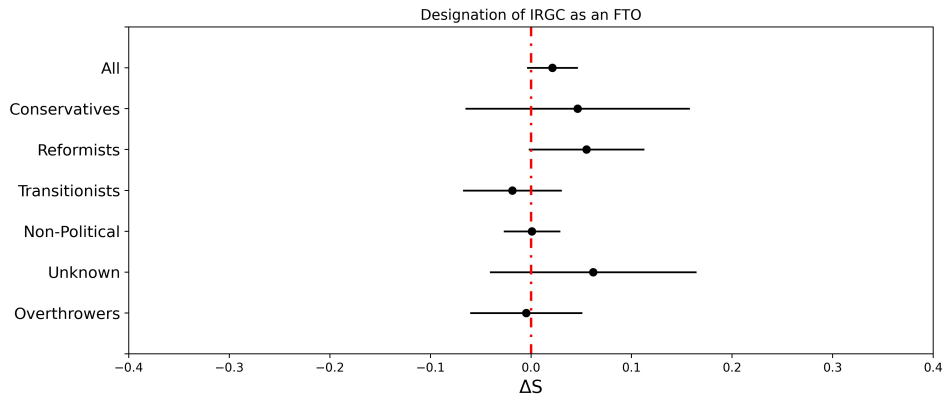


Figure 5.21 Regression Analysis - U.S. designation of IRGC as an FTO

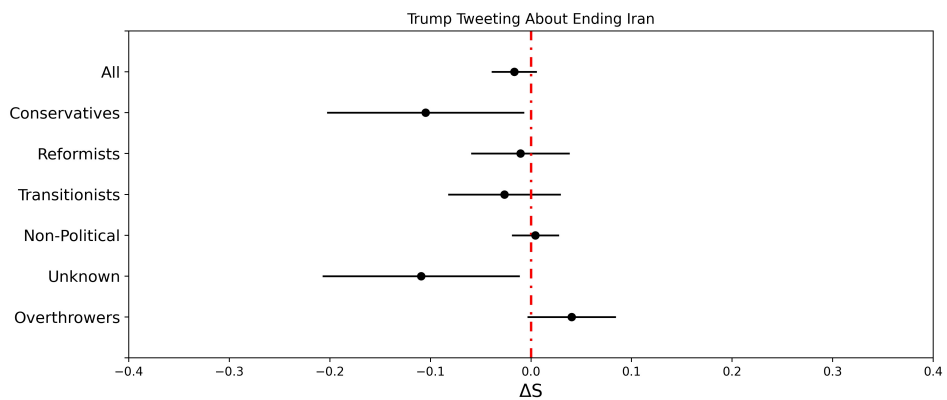


Figure 5.22 Regression Analysis - Trump's tweet about the 'End of Iran'

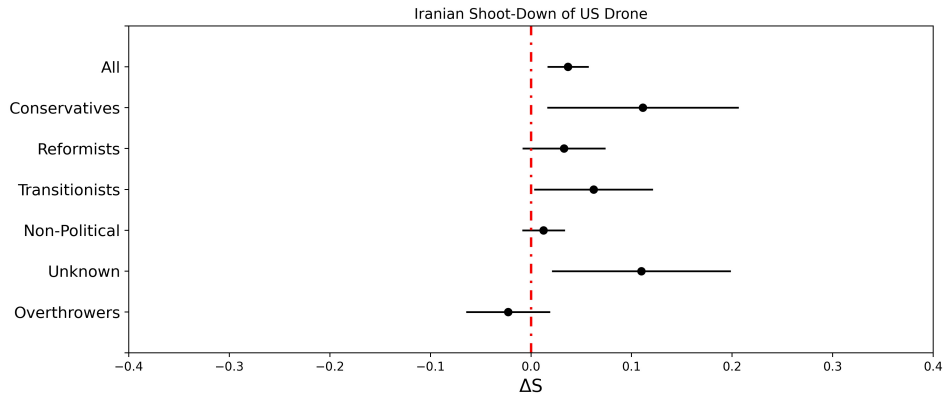


Figure 5.23 Regression Analysis - Iranian shoot-down of American drone

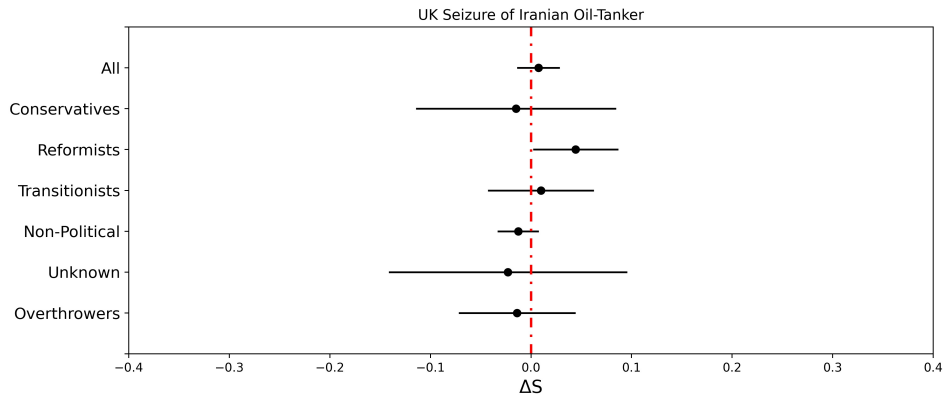


Figure 5.24 Regression Analysis - UK's seizure of Iranian-flagged ship in Gibraltar




Figure 5.25 Regression Analysis - Iranian's seizure of the British oil-tanker

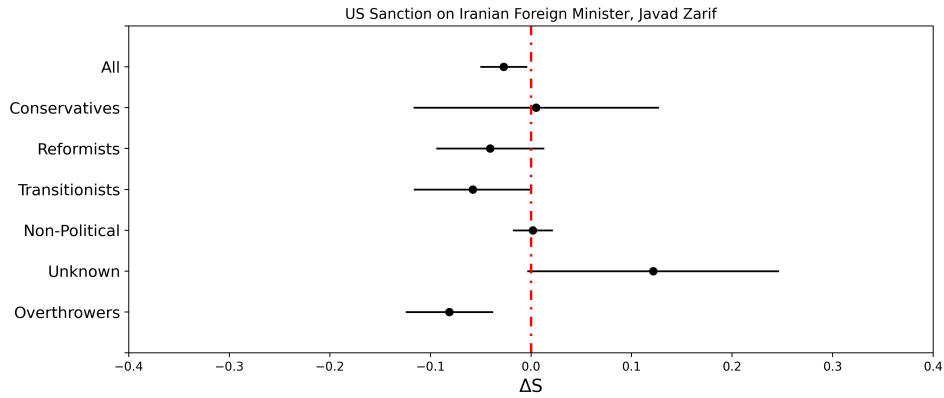


Figure 5.26 Regression Analysis - U.S. imposition of sanction on Javad Zarif

Finally, our hypothesis-testing shows that the positive changes in the sentiment towards the state in the events of 'JCPOA Violation', 'U.S. Designation of IRGC as an FTO', 'Iranian Shoot-Down of U.S. Drone' and 'Iranian Seizure of the British Oil-Tanker' are statistically significant. Although in some events there is no significant change in average of the users, we could still detect positive changes among some affiliations. In the conclusion section, we will look into the intuitions and implications which could be derived from our results.

6. CONCLUSION & FUTURE WORK

In this study, we proposed a new content-based approach for quantifying the 'Rally Round the Flag' phenomenon in Iran by utilizing social media. Furthermore, by comparison with public surveys, our analysis could also be used for investigating how representative of public opinion the Twitter influencers might be. For now, looking at the charts and regression-analysis we were able to infer the following conclusions:

- Based on our regression-analysis, the positive changes in the sentiment towards the state in the events of 'JCPOA Violation', 'U.S. Designation of IRGC as an FTO', 'Iranian Shoot-Down of U.S. Drone' and 'Iranian Seizure of the British Oil-Tanker' are statistically significant. In other words, one could claim that there is a 95% chance that in these three events, Iranian Influencers have rallied around the flag.
- Looking at the influencers, both Conservatives and moderate opposition groups (i.e. Reformists and Transitionists) are prone to 'Rally Round the Flag' effect in cases of foreign-interventions. However, the radical opposition group (i.e. Overthrowers) is more likely to hold their former positions.
- Regarding the case of 'the U.S. Pulling out of JCPOA', one intuition derived from the strange wide range of confidence-interval (i.e. which is due to the high variance of the values) for the Conservatives, could be that U.S. withdrawal from the JCPOA might have caused more various positive and negative sentiments among the Conservatives. Since the administration who signed the JCPOA with P5+1 was a Reformist one, the Conservatives are more likely to criticize the 'incompetency' of the government's diplomacy, next to their probable support for the state. This could cause a higher variance and longer confidence-interval.
- The comparison between the case of 'British Seizure of Iranian Oil-Tanker' and 'Iranian Seizure of British Oil-Tanker' shows that a case of military retaliation by Iran is much more popular (esp. among Conservatives) than the foreign

intervention cases which merely puts the Iranian government in the position of the oppressed victim. The massive 'Rally Round the Flag' effect for the case of 'Iranian Shoot Down of the U.S. Drone' further suggests this hypothesis which is also compatible with a key finding of the IranPoll institution's survey positing that "more than two-thirds of Iranians think their country should militarily respond if the United States, Saudi Arabia, or Israel were to attack an Iranian nuclear facility" (Smeltz, 2020).

- According to a recent poll, "many Iranians seem to have lost interest in the nuclear agreement reached between Iran and the P5+1 countries in 2015. Just 42 percent of Iranians approve of the nuclear agreement in the December 2019 poll, down from 76 percent approval in August 2015" (Smeltz, 2020). This could open an argument that the recent cases of foreign conflicts, is radicalizing the public opinion in Iran's political atmosphere strengthening either the Conservatives or Overthrowers camp. This could be observed in the boost for the average of likes in most cases of foreign-interventions for the Conservatives; which somewhat shows a temporary increase in their popular favor over moderates. However, this variation needs a future analysis of statistical significance before jumping into a certain conclusion.
- In cases of "Trump's Tweet" and "Sanction on Zarif", the reactions to internal tensions (e.g. the hunger strike of a political prisoner) superseded the reactions to a foreign conflict, generating more -1 class tweets and confusing both our descriptive and predictive-analytics in detecting the 'Rally Round the Flag' on Twitter.
- The comparison between the accuracy of our two classification-approaches, suggests that for such text-mining tasks, it is sometimes a good idea to group members of the same discourse and classify them separately rather than classifying all the discourses at once. This is most suggested for more polarized, radicalized, and unique discourses that have certain keywords in their manner of writing (i.e. in our task, Conservatives, and Overthrowers).

All of these findings are not only subject to a certain classification-error, but also subject to a high generalization and simplification of the classes. There are certainly much more complexities in the tweets in terms of sentiment towards the state than a simple mapping into -1, 0, and +1. Moreover, the Iranian administration is a complex of multiple governing actors. In some contexts, showing support for the Supreme Leader could mean slamming the current Reformist government and vice versa. However, all these cases are scaled while labelling.

The future work could also contain a comparison of our content-based approaches with a Twitter-Network-Analysis. This could be done either by looking into the network of followers or the network of retweets among the tweeps. Moreover, the content-based approach could be accompanied by an analysis of the hashtags next to the texts. An improvement in Persian text-mining libraries (e.g. Hazm library which we used) in the sense of enabling them to transform variations of Persian writing to each-other (e.g. transforming Persian informal texts to formal, dealing with the half-space variations in Persian texts) could offer a great boost to the classification accuracy in a content-based task.

Furthermore, the new case of the assassination of Gen. Qassem Soleimani followed by a massive rally in the streets and Iran's retaliation, provides another perfect case-study for the 'Rally Round the Flag' phenomenon. Unfortunately, this was missed in our dataset due to the early time-frame of our research. Future works could add this as an extension to our literature as well.

BIBLIOGRAPHY

- Back, M. D., Stopfer, J. M., Vazire, S., Gaddis, S., Schmukle, S. C., Egloff, B., & Gosling, S. D. (2010). Facebook profiles reflect actual personality, not self-idealization. *Psychological science*, *21*(3), 372–374.
- Baker, W. D. & Oneal, J. R. (2001). Patriotism or opinion leadership?: The nature and origins of the “rally ’round the flag” effect. *Journal of Conflict Resolution*, *45*(5), 661–687.
- Barberá, P., Jost, J. T., Nagler, J., Tucker, J. A., & Bonneau, R. (2015). Tweeting from left to right: Is online political communication more than an echo chamber? *Psychological Science*, *26*(10), 1531–1542. PMID: 26297377.
- Conover, M. D., Ratkiewicz, J., Francisco, M., Gonçalves, B., Menczer, F., & Flammini, A. (2011). Political polarization on twitter. In *Fifth international AAAI conference on weblogs and social media*.
- Goldstein, J. S., Pevehouse, J. C., & Sernau, S. (2008). *Principles of international relations*. Pearson Longman.
- Góralewicz, B. (2018). The tf*idf algorithm explained. <https://www.onely.com/blog/what-is-tf-idf/>. Accessed: 2018-03-06.
- Hajizadegan, A. (2019). Where the iranian twitter is heading towards? <https://www.magiran.com/article/3933137>.
- Hajizadegan, A. & Jalaepour, H. (2018). Identifying iranian twitter influencers. https://www.researchgate.net/publication/329182880_shnasayy_w_arzyaby_alghway_tathyrpdhyry_karbran_twytr_ayrany_az_ngrshhay_syasy_aynflwynsrha.
- Hetherington, M. J. & Nelson, M. (2003). Anatomy of a rally effect: George w. bush and the war on terrorism. *PS: Political Science amp; Politics*, *36*(1), 37–42.
- Jost, J. T. (2006). The end of the end of ideology. *American psychologist*, *61*(7), 651.
- Kim, S.-W. & Gil, J.-M. (2019). Research paper classification systems based on tf-idf and lda schemes. *Human-centric Computing and Information Sciences*, *9*.
- Kosinski, M., Stillwell, D., & Graepel, T. (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the national academy of sciences*, *110*(15), 5802–5805.
- Larson, J. M., Nagler, J., Ronen, J., & Tucker, J. A. (2019). Social networks and protest participation: Evidence from 130 million twitter users. *American Journal of Political Science*, *63*(3), 690–705.
- Martin, D. (2019). How to do cross-validation when upsampling data. <https://kiwidamien.github.io/how-to-do-cross-validation-when-upsampling-data.html>. Accessed: 2019-05-20.
- Mueller, J. E. (1970). Presidential popularity from truman to johnson. *American Political Science Review*, *64*(1), 18–34.
- Peksen, D. (2011). Does foreign military intervention help human rights? *PSN: Interstate Conflict (Topic)*, *65*.
- Pranckevicius, T. & Marcinkevičius, V. (2017). Comparison of naive bayes, random forest, decision tree, support vector machines, and logistic regression classifiers

- for text reviews classification. *Baltic Journal of Modern Computing*, 5.
- Rusin, M. (2019). What is the c parameter in logistic regression? <https://www.quora.com/What-is-the-C-parameter-in-logistic-regression#>.
- Smeltz, D. (2020). Majority of iranians oppose development of nuclear weapons. <https://www.thechicagocouncil.org/publication/majority-iranians-oppose-development-nuclear-weapons>.
- Sobhe (2018). Python library for digesting persian text. <http://www.sobhe.ir/hazm/docs/>.
- Weber, I., Garimella, V. R. K., & Batayneh, A. (2013). Secular vs. islamist polarization in egypt on twitter. In *Proceedings of the 2013 IEEE/ACM international conference on advances in social networks analysis and mining*, (pp. 290–297).
- Wood, R. M. (2008). "a hand upon the throat of the nation": Economic sanctions and state repression, 1976-2001. *International Studies Quarterly*, 52(3), 489–513.





Appendices

A. GRID-SEARCH PLOTS

We also tried tuning the regularization parameter for the logistic regression used on our affiliation-based classification. Meaning that we separated the tweets of each political affiliation's tweeps and used a separate classifier for each of them. The figures below show the tuning process for all of the political affiliations in our analysis:

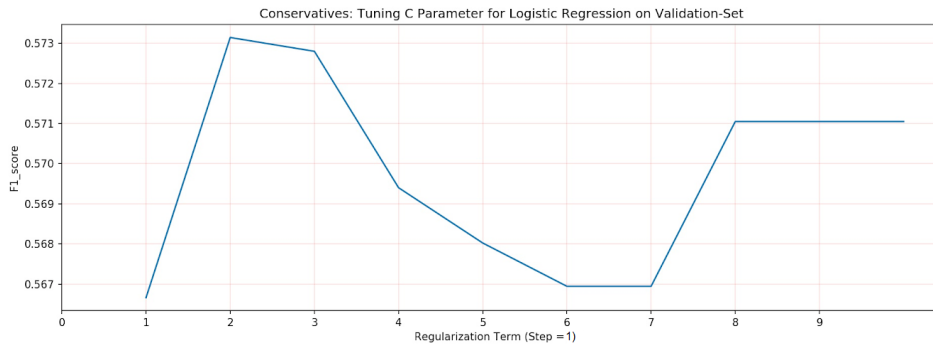


Figure A.1 Grid-Search for tuning the best regularization-parameter for Conservatives

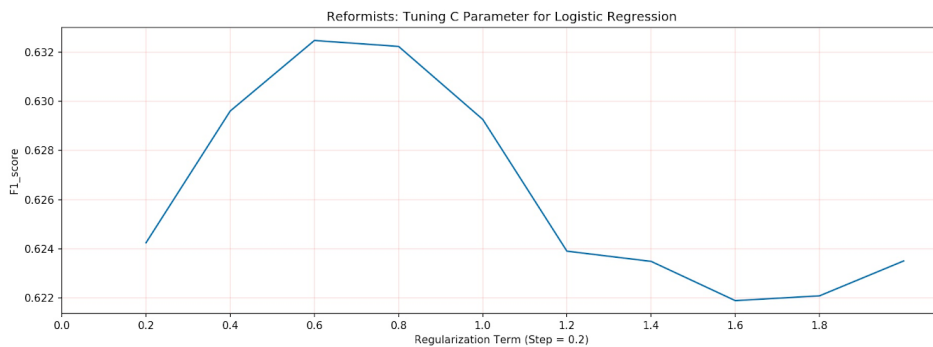


Figure A.2 Grid-Search for tuning the best regularization-parameter for Reformists

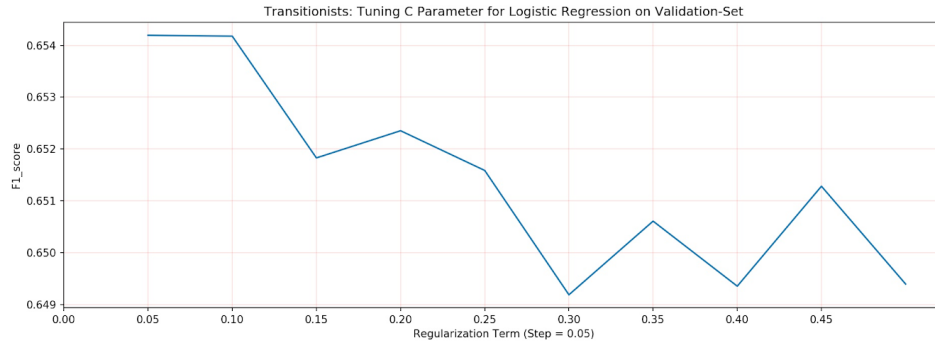


Figure A.3 Grid-Search for tuning the best regularization-parameter for Transitionists

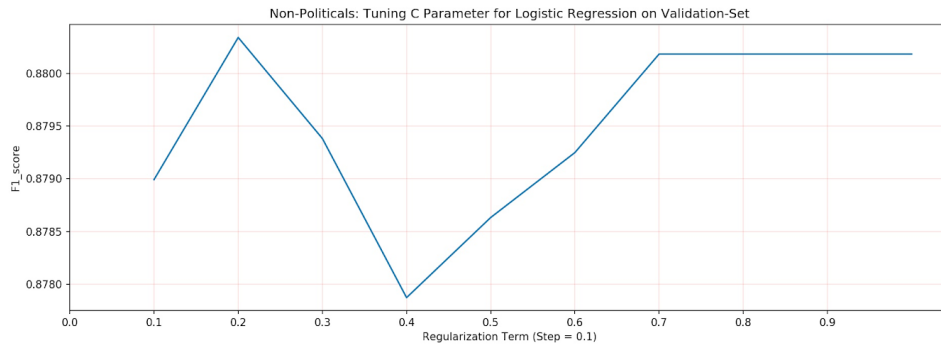


Figure A.4 Grid-Search for tuning the best regularization-parameter for Non-Politicals

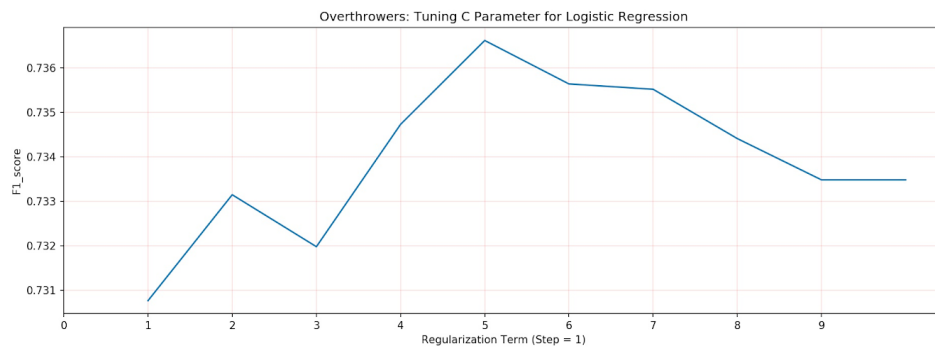


Figure A.5 Grid-Search for tuning the best regularization-parameter for Overthrowers

B. CONFUSION MATRICES

		Predicted		
		Sentiment	-1	0
Actual	-1	271	56	45
	0	197	457	111
	+1	50	25	101

Table B.1 Naive-Bayes Classifier on the whole test-set (1st approach)

		Predicted		
		Sentiment	-1	0
Actual	-1	191	161	20
	0	45	702	18
	+1	29	86	61

Table B.2 Random Forest Classifier on the whole test-set (1st approach)

		Predicted		
		Sentiment	-1	0
Actual	-1	9	6	2
	0	7	22	2
	+1	6	8	13

Table B.3 Logit on the Conservatives (1st approach)

		Predicted		
		Sentiment	-1	0
Actual	-1	42	32	7
	0	26	126	12
	+1	13	28	44

Table B.4 Logit on the reformists (1st approach)

		Predicted		
		Sentiment	-1	0
Actual	-1	61	22	4
	0	23	109	17
	+1	5	8	17

Table B.5 Logit on the Transitionists (1st approach)

		Predicted		
		Sentiment	-1	0
Actual	-1	5	9	1
	0	14	239	9
	+1	2	1	6

Table B.6 Logit on the Non-Politicals (1st approach)

		Predicted		
		Sentiment	-1	0
Actual	-1	5	1	3
	0	1	37	3
	+1	0	2	12

Table B.7 Logit on the Unclears (1st approach)

		Predicted		
		Sentiment	-1	0
Actual	-1	101	38	24
	0	8	95	15
	+1	4	0	7

Table B.8 Logit on the Overthrowers (1st approach)

		Predicted		
		Sentiment	-1	0
Actual	-1	8	9	1
	0	2	14	7
	+1	4	4	23

Table B.9 Logit(C=2) on the Conservatives (2nd approach)

		Predicted		
		Sentiment	-1	0
Actual	-1	42	16	13
	0	37	111	27
	+1	10	30	53

Table B.10 Logit(C=0.4) on the Reformists (2nd approach)

		Predicted		
		Sentiment	-1	0
Actual	-1	71	44	11
	0	39	171	19
	+1	14	17	20

Table B.11 Logit(C=0.64) on the Transitionists (2nd approach)

		Predicted		
		Sentiment	-1	0
Actual	-1	3	19	1
	0	11	258	3
	+1	0	5	3

Table B.12 Logit(C=0.2) on the Non-Politicals (2nd approach)

		Predicted		
		Sentiment	-1	0
Actual	-1	2	6	2
	0	2	40	3
	+1	0	7	4

Table B.13 Logit(C=1) on the Unclears (2nd approach)

		Predicted		
		Sentiment	-1	0
Actual	-1	119	30	2
	0	32	74	0
	+1	5	1	1

Table B.14 Logit(C=5) on the Overthrowers (2nd approach)

C. BUBBLE CHARTS

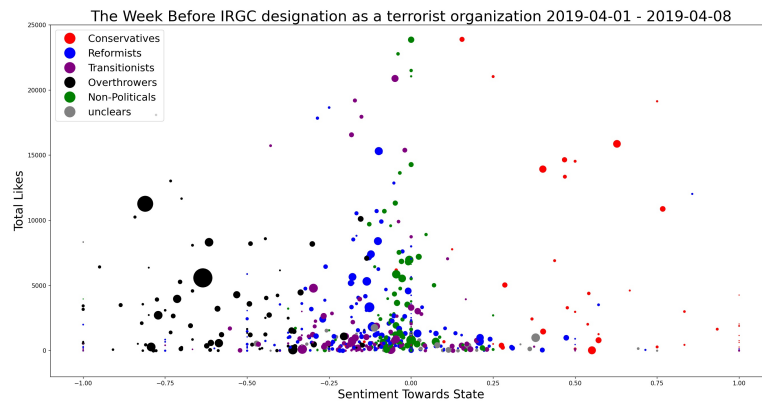


Figure C.1 The Week Before U.S. designated JCPOA as an FTO




Figure C.2 The week U.S. designated JCPOA as an FTO

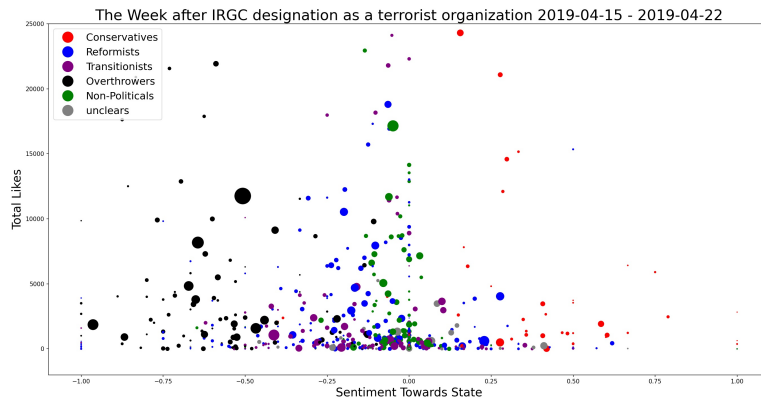


Figure C.3 The week after U.S. designated IRGC as an FTO



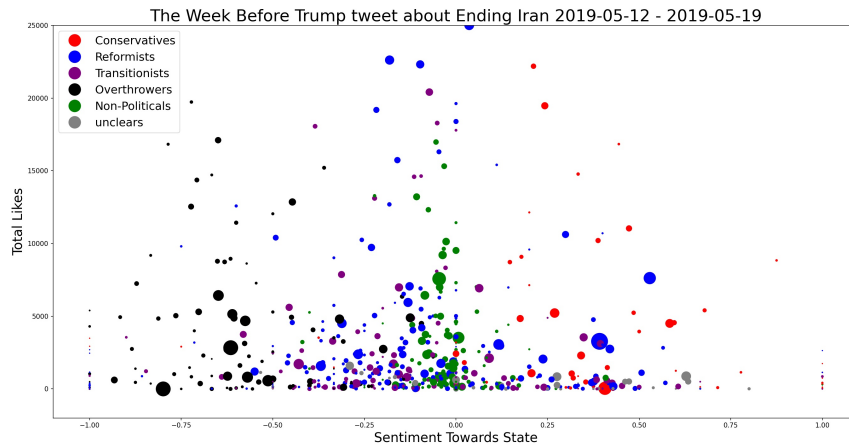


Figure C.4 The week before Trump's tweet about 'the End of Iran'

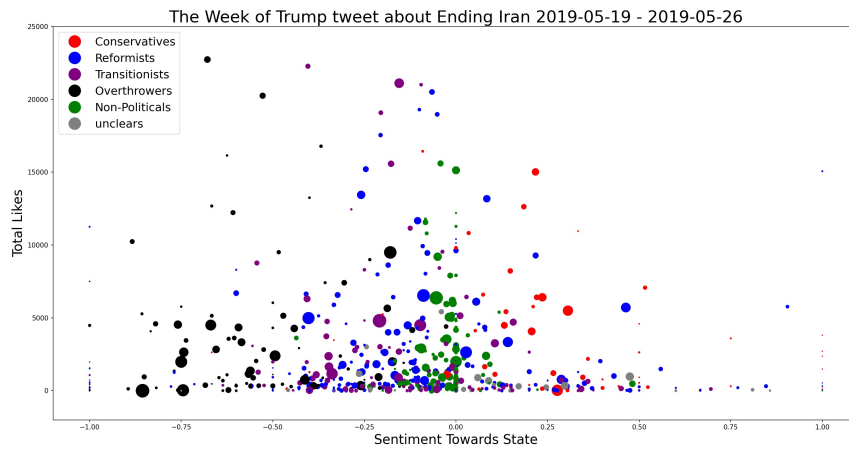


Figure C.5 The week of Trump's tweet about 'the End of Iran'

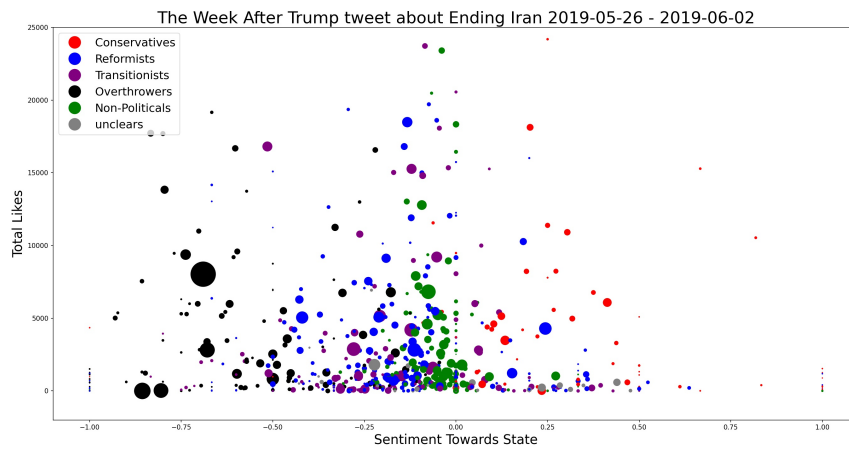


Figure C.6 The week after Trump's tweet about 'the End of Iran'




Figure C.7 The week before UK seizure of the Iranian oil-tanker




Figure C.8 The week of UK seizure of the Iranian oil-tanker




Figure C.9 The week after UK seizure of the Iranian oil-tanker




Figure C.10 The week before U.S. imposition of sanctions on Javad Zarif




Figure C.11 The week of U.S. imposition of sanctions on Javad Zarif




Figure C.12 The week after U.S. imposition of sanctions on Javad Zarif

D. AGGREGATED BUBBLE CHARTS




Figure D.1 The Week Before U.S. Pulled Out of the JCPOA

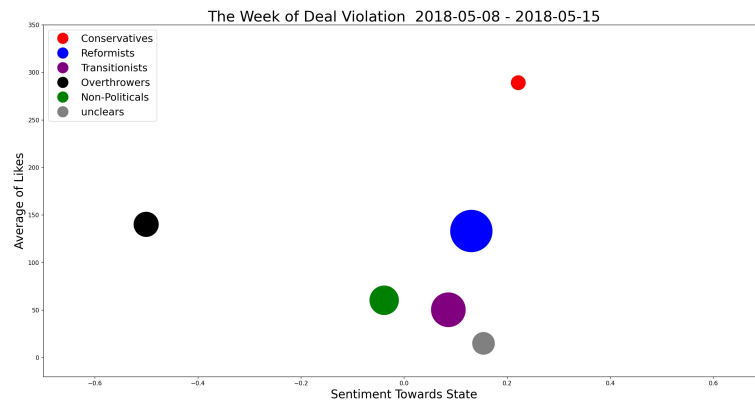


Figure D.2 The week U.S. pulled out of the JCPOA (we anticipate the rally here)




Figure D.3 The week after U.S. pulled out of the JCPOA






Figure D.4 The week before U.S. designated IRGC as an FTO

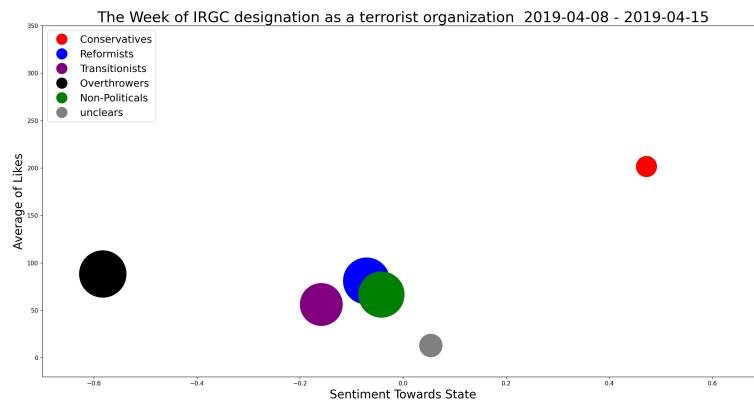


Figure D.5 The week U.S. designated IRGC as an FTO

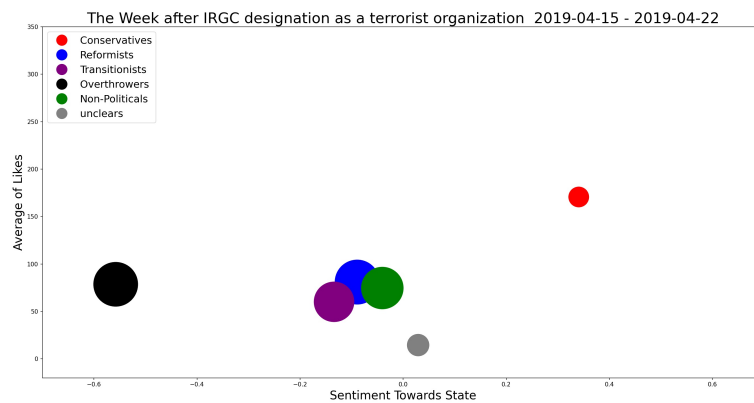


Figure D.6 The week after U.S. designated IRGC as an FTO

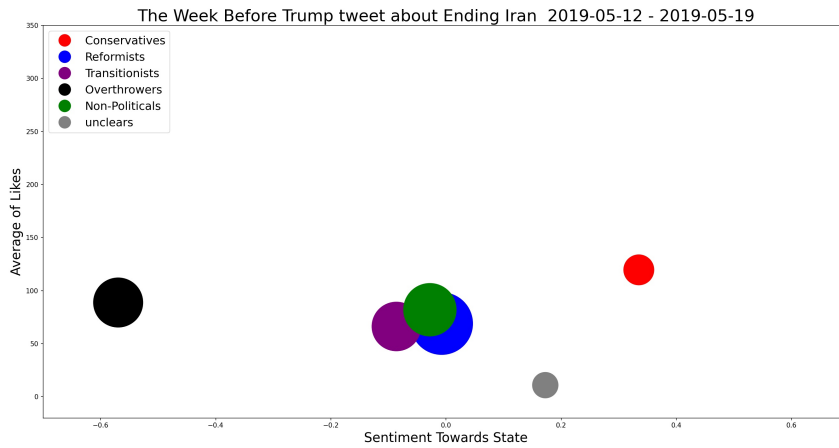


Figure D.7 The week before Trump's tweet about 'the End of Iran'

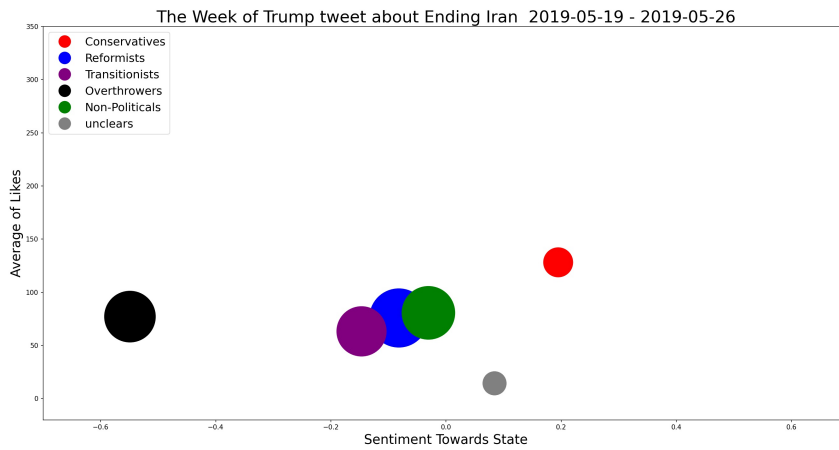


Figure D.8 The week of Trump's tweet about 'the End of Iran'

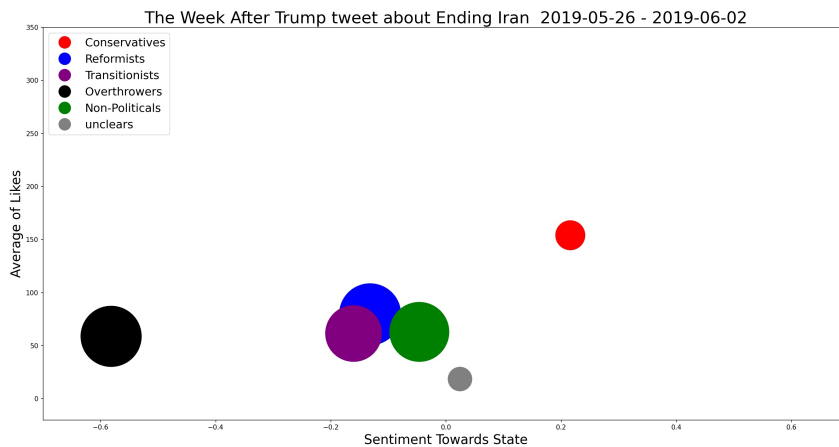


Figure D.9 The week after Trump's tweet about 'the End of Iran'




Figure D.10 The week before Iranian shoot-down of the U.S. drone

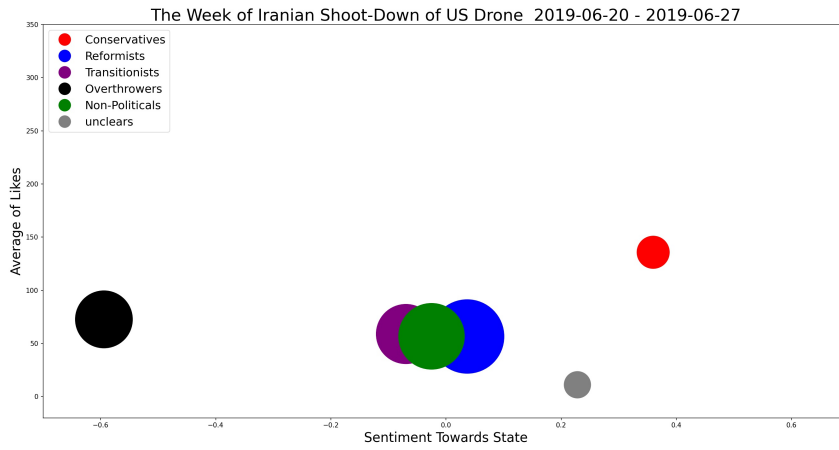


Figure D.11 The week of Iranian shoot-down of the U.S. drone

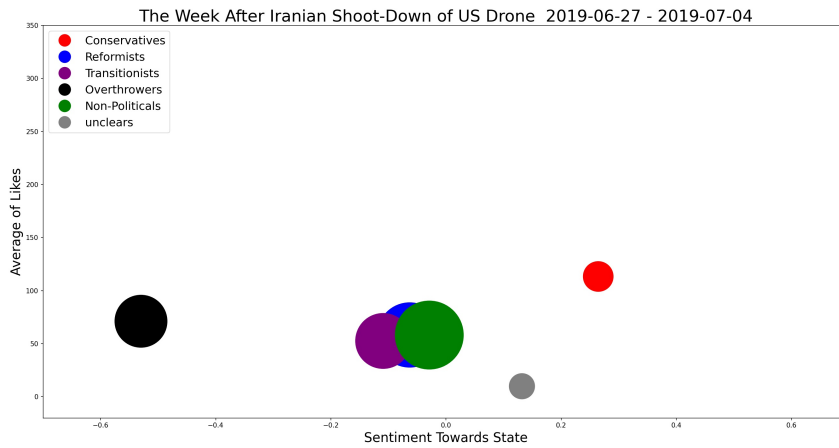


Figure D.12 The week after Iranian shoot-down of the U.S. drone

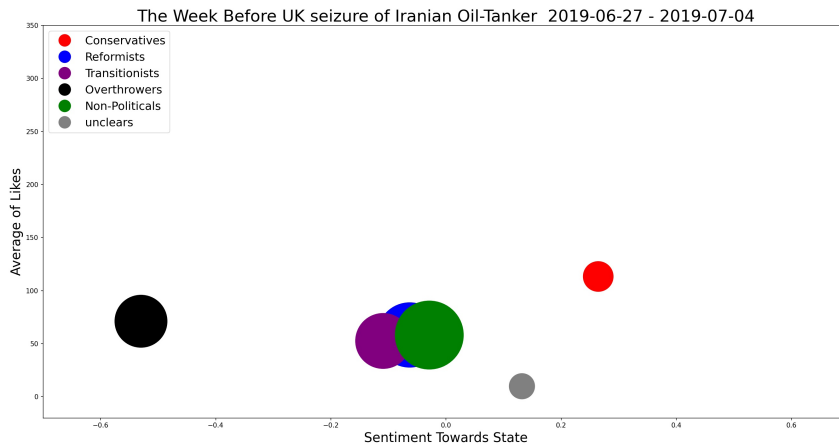


Figure D.13 The week before UK seizure of the Iranian oil-tanker

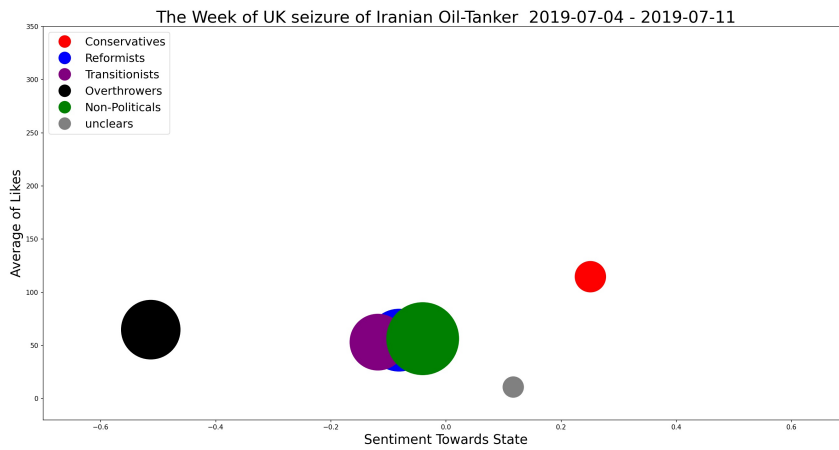


Figure D.14 The week of UK seizure of the Iranian oil-tanker

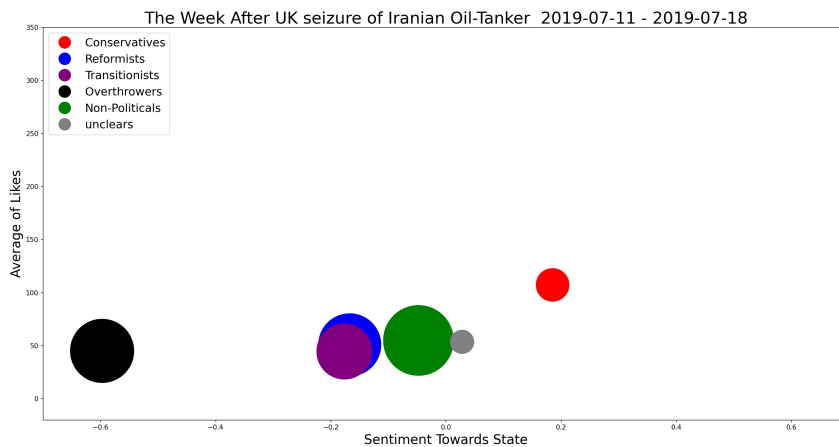


Figure D.15 The week after UK seizure of the Iranian oil-tanker

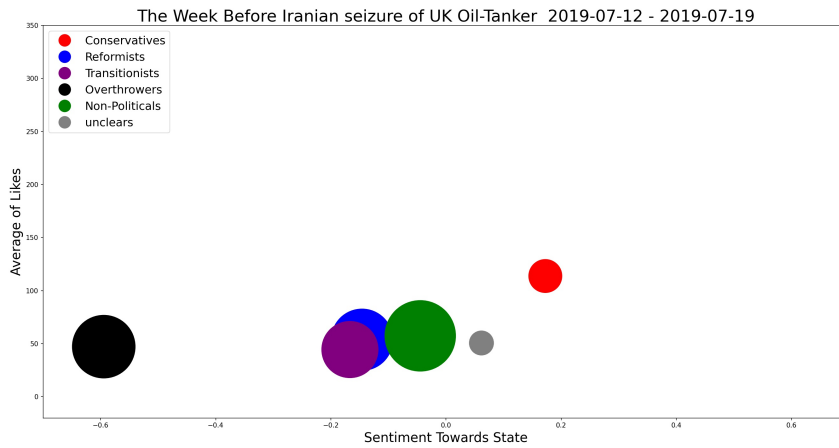


Figure D.16 The week before Iranian seizure of the British oil-tanker




Figure D.17 The week of Iranian seizure of the British oil-tanker

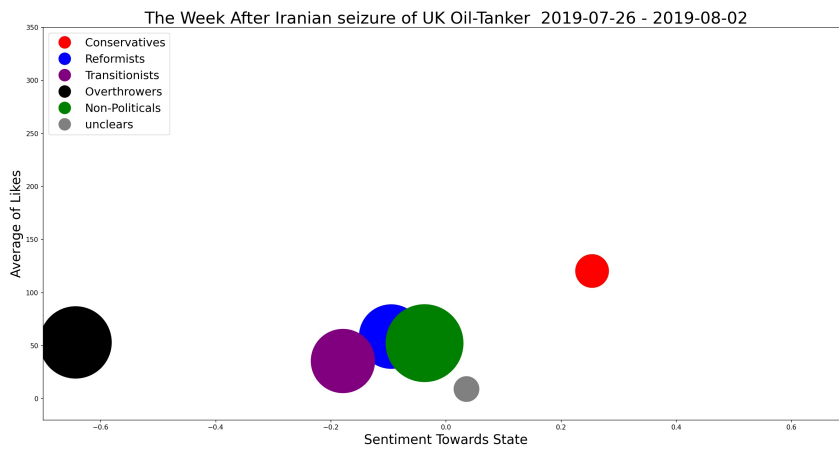


Figure D.18 The week after Iranian seizure of the British oil-tanker




Figure D.19 The week before U.S. imposition of sanctions on Javad Zarif




Figure D.20 The week of U.S. imposition of sanctions on Javad Zarif

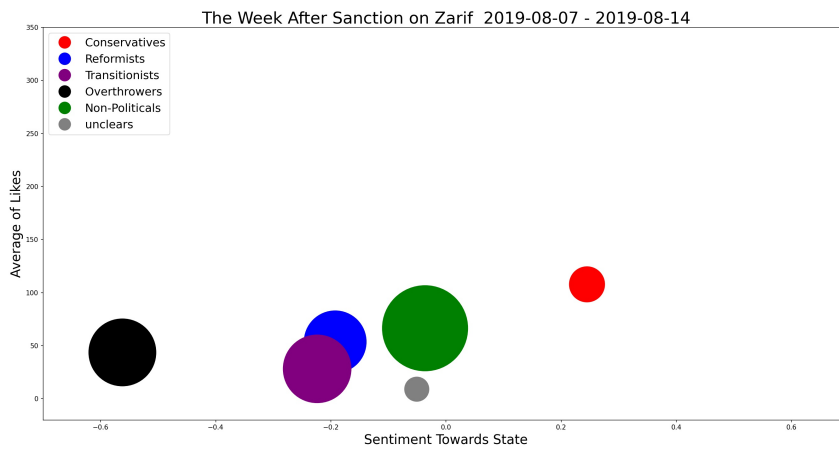


Figure D.21 The week after U.S. imposition of sanctions on Javad Zarif