Supplementary Information

This supplementary information file presents discussions about Twitter use in Venezuela, the role of traditional media during the event, an extended explanation of the data collection methodology, an extended discussion of the difference-in-differences research design, and a series of empirical extensions to the main analyses.

Twitter use in Venezuela

Twitter gained popularity in Venezuela after being adopted by Hugo Chavez in 2010 as "a tool for government", encouraging citizens to tweet concerns directly to him and employing 200 people to respond to citizens' messages.[1,2] In the early 2010s, Venezuela ranked among the top 20 countries in number of Twitter users[3] and ranked among the top 5 in terms of Twitter penetration[4] (with an estimate for 2012 of 21 percent). A Pew survey conducted in 2013 in Venezuela revealed that 73% of those aged 18-29, 52% of those aged 30-49, and 15% of those 50 or older, used social networking sites. Of these users, 49% reported using the sites to share views about politics, and 74% reported they had learned about others' political beliefs from something they posted on a social networking site (the highest rate among the 22 developing countries in the study).[5] Twitter has been used extensively by both government and opposition leaders, and has played an important role in the country's political developments.[6] Other sources regarding

[1] https://www.theguardian.com/world/2010/aug/10/hugo-chavez-twitter-venezuela
[2] His approach to social media made Chavez the second most popular head of state on Twitter in 2012, only behind Barack Obama (see http://www.digitaldaya.com/admin/modulos/galeria/pdfs/69/156_biqz7730.pdf).
[3] https://semiocast.com/publications/2012_07_30_Twitter_reaches_half_a_billion_accounts_140m_in_the_US
[4] https://www.comscore.com/Insights/Press-Releases/2011/4/The-Netherlands-Ranks-number-one-Worldwide-in-Penetration-for-Twitter-and-LinkedIn
[5] https://www.pewglobal.org/2014/02/13/emerging-nations-embrace-internet-mobile-technology/
[6] https://www.theatlantic.com/international/archive/2014/02/why-venezuelas-revolution-will-be-tweeted/283904/

Twitter and social media use in Venezuela can be found in Forelle et al (2015) and Munger et al (2018).

Traditional media

Though the analysis is centered on Twitter, traditional media also played a role in diffusing information related to the cancellation of the retweeting accounts. As previously highlighted, Maduro made a public announcement of the "attack" on national TV.[7] Using the *wayback machine*, one can observe too that the event was on the front page of the *TeleSUR* website, one of the main TV networks sponsored by the Venezuelan government, on November 2nd (see Figure A6).[8] In the leading article, Maduro proposes the creation of a new social network to combat Twitter's bias against his government. On the other hand, coverage of the event on private media (which is more sympathetic to the opposition), including newspaper *El Universal* and television station *Globovisión*, focused on Twitter's policies and possible reasons for the forced closure of the accounts, including the use of "phantom" users and behaviour consistent with "spam".[9] This coverage is consistent with the position of the opposition, and likely reinforced the effect of the event on citizens' perceptions of Maduro's popularity.

Though a comprehensive analysis of how the event was portrayed in older media forms is outside of the scope of this study, this snapshot reveals patterns consistent with those observed on Twitter, and is congruent with the context of a hybrid media system (Chadwick, 2017). It also suggests that traditional media was likely to have contributed to users' updating their beliefs

[7] A clip of which is available here: https://www.youtube.com/watch?v=ioTJQzTuPaU

[8] The wayback machine is a digital archive of the Internet which stores web content at different points in time. The first date at which the TeleSUR website was archived after the event was November 3rd: https://web.archive.org/web/20131103001549/https://www.telesurtv.net/

[9] See https://web.archive.org/web/20131104004753/http://globovision.com/articulo/sabe-usted-por-que-twitter-elimina-cuentas and http://www.eluniversal.com/nacional-y-politica/131102/estiman-que-twitter-elimino-seguidores-fantasmas-de-maduro

about the popularity of Maduro, and was one of the mechanisms supporting the unraveling of the existing 'spiral of silence'.

Data collection methodology

This section presents additional details on the collection methodology used to gather the dataset. I code a crawling tool using the python programming language and the 'requests' library (http://docs.python-requests.org). The tool makes calls to Twitter's advanced search engine, which are generally formatted in the URL as this:

```
https://twitter.com/search?f=tweets&vertical=default&q=twitter%20%40NicolasMaduro%20since%3A2013-11-01%20until%3A2013-11-02&src=typd&lang=en
```

This page returns the last 20 tweets posted on November 1st, which contain the @NicolasMaduro username and the keyword "twitter".[10] The tool makes calls for each date and each selected keyword and stores the text of the site (using the functionality of requests), from which one can then retrieve the tweet-id and the number of replies for each result. Depending on whether one wants tweets containing a username, or tweets written by a specific user, the format varies slightly (these can be verified by examining the results from Twitter's advanced search interface). The tweets' unique IDs are then used to make calls to the Twitter API and "hydrate" the tweets dataset. The "**f=tweets**" argument tells Twitter to show the "latest" tweets, as opposed to those which the platform deems most important, since this could lead to sampling bias; whereas the last twenty tweets can be expected to be more representative of all tweets. Doing this also allows me to proxy tweet volume using the time of publication of these last twenty tweets, as discussed in the text.

[10] Twitter search is not case or accent sensitive.

Network analysis of selected accounts

Figure A7 shows the Twitter network structure of the selected accounts of the political leaders, which can be used to validate the account classification.[11] Each node represents one of the accounts and an edge is drawn between two nodes if either of them follows the other (though the underlying network is directed, I represent it as undirected for simplicity). Government accounts, shown as red squares, are clustered strongly together. Opposition accounts, shown as blue circles, are also clustered together, though more weakly so. Finally, the control-group sports-related accounts, are shown as green triangles. The bigger node shown as a brown square represents @NicolasMaduro, who is not only central in the government cluster, but also in the network overall.[12]

Keyword analysis of political leaders

In addition to the analysis on engagement in the main text, Table A3 reports keyword counts during the event study window for this dataset. It is worth highlighting that, for opposition tweets, the frequency with which the keyword *maduro* appears increases significantly (as is the case for the keyword *venezuela*), while the frequency of the username of Henrique Capriles (@hcapriles; the main leader of the opposition), and the opposition promoted hashtag #quenadatedetenga, decreases. For tweets by government leaders, the use of both the *maduro* keyword and Maduro's username @NicolasMaduro, increase. These patterns are consistent with

[11] Twitter network structure has been shown to be a predictor of political ideology (Colleoni et al, 2014; Barberá, 2015; Halberstam and Knight, 2016).

[12] There is one government outlier account that appears in the opposition cluster, which is that of Luisa Ortega Díaz (@lortegadiaz). She was the Prosecutor General appointed by Chavez and was loyal to Nicolas Maduro until 2017. By 2018 (when the network data was collected) she was a Maduro critic and, as the network graph suggests, had made "new friends" among the opposition. Since the tweets used are for 2013, I keep her as a government account. A qualitative review of her tweets during the period of study reveal mostly neutral positions. In addition, the results do not change in any meaningful way when removing her tweets from the analysis.

the patterns documented for the other datasets and further show that the event was widely discussed by Venezuelan political leaders.

Validating the selection of pro-Maduro and anti-Maduro keywords

Together with a Research Assistant, we manually coded a subsample of tweets into "Pro-Maduro", "Neutral" and "Anti-Maduro" categories to validate the keyword choices for the tweets which mention @NicolasMaduro. Thirty tweets were selected (using a random number generator) for each of the 10 selected keywords. We then read each of these tweets and manually classified them across the categories. For the selected pro-Maduro keywords (*revolucion, camarada, comandante, chavez* and *victoria*), 65 percent of these are indeed supportive of Maduro, 29 percent are neutral, and 7 percent are critical of the president. On the other hand, for the selected anti-Maduro keywords (*regimen, ilegitimo, escasez, ladron,* and *maldito*), 92 percent are critical of Maduro, 7 percent are neutral, and 1 percent are supportive of the president. The complete results of this exercise, by keyword, are shown in Figure A8. Though the keywords are not perfectly predictive of the political stance of the tweets, they are strongly correlated with it.[13]

This measurement error in the political stance of tweets can lead to bias in the estimates of the effect of the accounts' closures on anti-Maduro and pro-Maduro tweet volume.[14] The estimates suggest that the volume in tweets using the pro-Maduro keywords increase by about 21 percent (from Table 2, column 6, multiplied by the estimated publication time to volume factor of 0.485) and those using the anti-Maduro keywords increase by about 41 percent (Table 2, column 5); a differential increase of 20 percent. If only a share of this increase actually captures the intended

[13] The correlation with the intended political stance is 0.83 coding the categories as -1 (anti-maduro), 0 (neutral), and 1 (pro-maduro).
[14] Note importantly that this is not "classical measurement error", the measure is instead an upper bound of the true increase in tweets of the intended sentiment.

sentiment, as suggested by the manual validation exercise, then the true increase in volume may be lower. A "back of the envelope" calculation using the results from the validation exercise would suggest that the increase in pro-Maduro tweets from these keywords is actually of 14 percent (21×0.65) and for anti-Maduro tweets of 38 percent (41×0.92); suggesting a differential increase of 24 percent more anti-Maduro tweets. Note importantly that since the anti-Maduro keywords capture the intended sentiment more precisely than the pro-Maduro keywords, this measurement error will tend to bias the estimates of the differential increase for anti-Maduro tweets downwards, suggesting that the estimates of the increase of anti-Maduro tweets relative to pro-Maduro tweets is a lower bound, or a conservative estimate for the actual change. [15]

Discussion of the difference-in-differences empirical methodology

This section presents additional details regarding the main identification strategy, as well as a placebo experiment, that help to illustrate the intuition behind the analysis.

The basic idea behind the research design can be illustrated using 2x2 tables. Table A4 shows the means of log(likes) for different periods, separately for the groups of interest. In panel A the periods coincide with the quasi-experiment of interest, comparing outcomes 15 days before the cancellation of the accounts, relative to 15 days after. In both of these time windows, opposition leaders (column 1) received on average more likes than the government (column 2, and the difference is presented in column 3). In addition, for both groups, log (likes) increased after the cancellation of the accounts (last row of panel A). However, the increase was larger for the

[15] For this exercise, an even more conservative calculation can use the normal approximation of the binomial distribution for the proportion of correctly assigned tweets (note the anti-Maduro sample of 150 tweets with 0.08 failure rate is just large enough to allow this). In that case, one could use the upper bound of the 95% confidence interval for the pro-Maduro sample (0.7), and the lower bound for the anti-Maduro sample (0.89), this would suggest a 15 percent increase for pro-Maduro tweets, 36 percent increase in anti-Maduro tweets, and a 21 percent differential increase.

opposition. Analogously, the gap between opposition and government was larger after the cancellation of the accounts. The difference in the differences of these means therefore captures the differential increase in likes that the opposition received, following the cancellation of the accounts. In particular, the estimates presented here suggest that opposition leaders received 0.242 more log (likes) relative to government leaders (column 3, last row of panel A, marked in bold).

One potential concern with these estimates could be that opposition leaders may have been on a faster growth path, even before the cancellation of the accounts. If this was the case, similar difference-in-differences estimates may be observed, but only because of these different pre-trends, not because of the cancellation of the accounts. In other words, the identification assumption of "parallel trends" would be violated. Figure 5 suggests that this was not the case, but a placebo experiment can provide additional evidence. Panel B repeats the exercise above but looking at estimates moving one 15-day period further back (October 2 to October 16). Note first that the gap between opposition and government persists for this period, but it is of a similar magnitude than the gap during the October 17 to October 31 period. Note too that there are no significant differences in engagement between this period and the next 15 days for either group (last row). Therefore, and reassuringly, the difference-in-differences estimate is statistically insignificant for this placebo experiment (-0.05).

Another potential bias affecting these difference-in-differences estimates could potentially arise if a very popular opposition leader (who on average gets lots of engagement) decides to tweet much more after the event, relative to less popular leaders. The preferred specification of the difference-in-differences estimator which presented in the main text includes user fixed-effects, which would account for this possible "selection into tweeting" mechanism. Therefore, the coefficients presented are within-user differences in these changes, that is, the difference is relative to engagement for the same user before the event. The estimates from the preferred specification which includes user fixed-effects, day fixed-effects and a series of controls, are smaller (0.214, Table 5, column 5), but similar in magnitude to the raw estimates presented here.

Lastly, Column 4 of Table A4 presents the means of log likes for the control group used in some specifications, the sports-related accounts. There are no statistically significant changes in the measure of engagement for these accounts, in neither the period of interest, nor the placebo experiment of the preceding time window.

Analysis of tweet volume in alternative keyword dataset

I analyze the volume to tweets which mention @NicolasMaduro to evaluate whether users' willingness to mention and criticize the president increased after the cancellation of the retweeting accounts. Directly addressing the president using his username @NicolasMaduro is an important act, as it can be viewed by him and others who follow him. Alternatively, however, one could simply look for tweets which contain the "maduro" keyword, regardless of whether they "tag" the president or not. Here I replicate the exercise of tweet volume using this alternative dataset. I collected tweets with the "maduro" keyword, as well as tweets with both "maduro" and each of the selected keywords (*venezuela, pueblo, twitter, revolucion, camarada, comandante, chavez, victoria, regimen, ilegitimo, escasez, ladron,* and *maldito*). The same restrictions as with the other datasets apply. The total number of tweets in this alternative dataset is 44,999.

I replicate tables 2 and 3 from the paper in this new dataset. These are presented in tables A5 and A6. The publication times suggest greater volumes of tweets overall for this new dataset, but the estimated changes in logs are generally similar in magnitude to those observed in the @NicolasMaduro dataset. The results from this alternative dataset suggest, relative to tweets including the @NicolasMaduro username, an even greater increase in the volume of tweets which mention the keyword *twitter* (Table A5, column 3, β=-2.119, p<0.001), which can be estimated to correspond to a 102 percent increase (-2.119×-0.485), and an even greater increase in tweets with anti-Maduro keywords (column 5, β=-0.918, p<0.001), or an estimated 45 percent (-0.918×-0.485), whereas there is a smaller and statistically insignificant increase in the volume of pro-Maduro tweets (column 6, β=-0.918, p<0.001). The results in table A6 suggest that the increase in anti-Maduro tweets relative to pro-Maduro tweets was statistically significant (column 6, β=-0.720, p<0.001), and represents an estimated differential increase of 35 percent.

The evidence presented in this section again confirms hypothesis H1a (increased criticism of the president). In fact, the support for the hypothesis is even stronger in this alternative dataset, which suggests even larger increases in the volume of anti-Maduro tweets, both overall and relative to pro-Maduro tweets.

Proxy measure of negative engagement

I use a proxy measure of negative engagement as an alternative outcome for the analysis. I construct the measure using log(replies + 1) – log(likes + 1). This measure considers the idea that tweets with a stronger negative reaction tend to receive more replies, but because many replies can also be positive, I use the number of likes to compensate for the strength of the positive feedback. This is illustrated with an example from current events using two recent tweets by Donald Trump in Figure A2. The tweets have similar number of likes and retweets but differ substantially in the number of replies. The more controversial tweet, which is subject to a strong negative reaction, has substantially more replies. The ratio of replies to likes is thus informative about negative reactions to a tweet.[16]

I replicate the main analyses with this measure as an outcome, which can help assess the statistical significance of the negative reaction. The results are presented in Table A7. For tweets by @NicolasMaduro, the proxy measure for negative engagement increases significantly after the cancelation of the retweeting accounts (column 2, β=0.163, p=0.047). For tweets by political leaders, there is an increase of about 11 percent in negative engagement for government leaders relative to sports-related accounts (column 3, β=0.106, p=0.059), and 16 percent greater negative engagement for government leaders relative to the opposition leaders (column 5, β=-0.156, p<0.001). The negative measure of engagement is not statistically different for the opposition relative to the sports-related accounts (column 4, β=-0.065, p=0.222). These findings are consistent suggest increased criticism of the government (congruent with H1a).

Heterogeneity across users' propensity to mention Nicolas Maduro

An additional empirical exercise studies whether the estimated effects are heterogeneous across users depending on their tendency to mention Maduro in their Twitter feeds. If users are more willing to express their relative support for the opposition because the president now appears less popular, then users who frequently mention Maduro may differentially benefit from the closure of the accounts (H3). I study this using both a difference-in-differences specification with a continuous treatment variable, as well as a triple-interaction framework as follows:[17]

[16] See also the definition of *#ratioed* (https://www.merriam-webster.com/words-at-play/words-were-watching-ratio-ratioed-ratioing).
[17] For an example of a triple-interaction design see for instance the analogous specification in Ferraz and Finan (2008), which examines whether government audits differentially affected municipalities with more local AM radio stations.

$$Y_{iut} = \beta_1 \cdot Post_t \cdot Opposed_u + \beta_2 \cdot Post_t \cdot MentionsMaduro_u$$

$$+ \beta_3 \cdot Post_t \cdot Opposed_u \cdot MentionsMaduro_u + X_{iut} \cdot \theta + \alpha_u + \alpha_t + \varepsilon_{iut}$$

where $MentionsMaduro_u$ is a variable that measures the frequency with which user u brings up the president in his or her tweets. The coefficient of interest, β_3, captures the differential engagement for users who mention Maduro more frequently, after the cancelation event, when coming from an opposition user. I present results both with and without fixed effects.

The analysis of heterogeneity across users' propensity to mention Maduro reveals that opposition leaders who mention Maduro more frequently experienced a larger relative increase in their number of likes (Table A8) relative to their peers (β=1.755, p=0.011), but the same is not true for government leaders. The results also reveal that the heterogeneity is differentially significant for opposition leaders relative to government leaders (β=1.502, p=0.034, for the preferred specification in column 6).

The coefficients suggest that an opposition political leader who mentions Maduro in ten percent of his tweets experienced a differential increase in tweet likes of around 17 percent after the accounts' cancelation, relative to an opposition leader who never mentions Maduro (1.755×0.1), and a differential increase of 29 percent relative to a government leader who also mentions Maduro in ten percent of his tweets (0.14 + 1.502×0.1). Figure A9 shows these marginal effects. The left panel shows the estimated change in log (likes) after the cancellation event based on the preferred fixed effects specification (Table A8, column 6). In addition, I show the results from a random effects model (right panel) that allows the estimation of the predicted log of likes, both before and after the event.[18] Before the cancellation of the accounts, higher

[18] Note that the *post* indicator is collinear with day-fixed effects in the preferred specification shown in Table A8, column 6, for this reason, I show the alternative model as well.

propensity to mention Maduro was negatively associated with engagement for members of the opposition, but this relationship flips in the days after the event. On the other hand, there is no significant relationship between engagement and users' tendency to mention the president for government leaders, neither before or after the event.

That opposition users who more frequently mentioned Maduro on the platform benefitted differentially from the cancelation of the accounts suggests that the change in the perception of Maduro's popularity was stronger for followers of these users, and therefore they experienced larger gains in political support, as measured by their follower engagement. This finding is consistent with hypothesis H3 and suggests that opposition followers who were differentially aware of Maduro's popularity on the platform reacted more strongly after the accounts were closed by Twitter.

Long-run analysis

As discussed in the main text, using a short window of time allows me to get closer to being able to infer a "causal relationship" between the cancellation of the accounts and the empirical facts documented. Many of the patterns in the medium and long-run can be viewed in the figures, but I formally replicate the main empirical exercises here using a 6-month long window (instead of a 30-day window), such that I compare outcomes in the 3-months after the cancellation of the accounts, relative to the 3-months before. Given that many other events can confound the analysis in the long-run, the exercise presented here should be viewed as descriptive. I replicate only the main results, and when possible, prioritize the difference-in-difference specifications (the pre/post analysis is more problematic due to seasonal trends, including the presence of Christmas and New Year's during the long-run window).

The results are presented in Table A9. In contrast with the short-run evidence, there is no observed increased differential volume for anti-Maduro tweets in the @NicolasMaduro dataset (column 3), or increased differential support for opposition users who more frequently mention @NicolasMaduro (column 7, this estimate is very imprecise but could be explained by the unraveling dynamics, such that the effects may initially be concentrated in these users but spread more broadly in the longer run). On the other hand, the patterns of increased replies for @NicolasMaduro (column 2), increased differential support for the opposition (column 5), and increased volume of anti-Maduro tweets in the alternative dataset (column 4, using “Maduro” keyword), are also present in the long-run.

Supplementary Information Tables and Figures

Table A1: @NicolasMaduro Followers

Date	Number of followers
March 2013	553,064
April 2013	983,123
May 2013	1,144,893
June 2013	1,207,997
July 2013	1,257,782
August 2013	1,320,013
September 2013	1,376,534
October 2013	1,418,953
November 2013	1,513,680
December 2013	1,575,707
January 2014	1,651,921

Source: twven.com

Table A2: Selected accounts for political leaders' dataset

Government	Opposition	Control
NicolasMaduro	hcapriles	MeridianoTV
dcabellor	leopoldolopez	SaschaFitness
jaarreaza	MariaCorinaYA	CarolinaPadron
JauaMiranda	liliantintori	leones_cbbc
TareckPSUV	ComandoSB	Pastormaldo
luislopezPSUV	alcaldeledezma	caroguillenESPN
jorgerpsuv	JJRENDON	greivisvasquez
taniapsuv	CarlosOcariz	LVBP_Oficial
JacquelinePSUV	hramosallup	Magallanes_bbc
gestionperfecta	HenriFalconLara	SalvadorPerez15
lortegadiaz	ramonmuchacho	6cichero6
HugoCabezas78	JulioBorges	la_grulla5
DanteRivasQ	luisvicenteleon	OzzieGuillen
ConCiliaFlores	VoluntadPopular	GarbiMuguruza
AmeliachPSUV	unidadvenezuela	BobKellyAbreu
DrodriguezVen	JuanRequesens	VizquelOmar13
NestorReverol	MiguelContigo	EUDeporte
garcesfrancisco	Diego_Arria	caraquistas
jdavidcabello	Pr1meroJusticia	porlagoma
Adan_Coromoto	rociosanmiguel	emanuelatleta
PartidoPSUV	GerardoBlyde	Magallanes_News
maperezpirela	ismaelprogreso	LuisAlvarez_1
psuvaristobulo	delsasolorzano	salorondon23
JuventudPSUV	carlosvecchio	aguilasdelzulia
blancaePSUV	PadreJosePalmar	Caracas_FC
ErikaPSUV	jcsosazpurua	Tomapapa
PanchoArias2012	PabloPerezOf	Lavinotintocom
irisvarela	FreddyGuevaraC	Alex_candal
FreddyBernal	VicenteDz	rubenoszki
HectoRodriguez	GenPenaloza	MaxCordaro
IzarraDeVerdad	DavidUzcategui	FerAlvarez
WalterDossier	Simonovis	CardenalesDice
vladimirpadrino	alfredoromero	RichardGol_espn
RobertSerraPSUV	manuelrosalesg	hturinese
jesusfariaPSUV	dsmolansky	OficialTigres
jchacon2021	antonioriverog	DIRECTVSportsVE
PedroCarreno_e	AndresVelasqz	cuantoacuanto
IsisPSUV	ENZOSCARANO	ElvisandrusSS1
jorgeamorin	BMarmoldeLeon	SeleVinotinto
AndreinaTarazon	alFranceschi	Arango_18
anat5	EvelingTrejo	milenagimon
danicabello11	juanjosemolina	JoseAltuve27
durancandanga	HimiobSantome	DvoTachira
tongorocho	plomoparejo	adriana_donghia
nicmaduroguerra	humbertotweets	JuanPaGalavis
PatriciaDorta40	orlandourdaneta	fpetrocelli
RALDAHIR	Gral_Vivas_P	GatoradeVzla
Cosole_Roja	mferreiratorres	DeportivoLara
MQuevedoF	MarceloNolla	dtcesarfarias
Marlenycdc	manocompa	MecheCelta

Table A3: Most frequently used keywords in tweets of political leaders

Opposition (period / num tweets)				Government (period / num tweets)			
Pre / 5,531		*Post / 5,949*		*Pre / 3613*		*Post / 4053*	
frequency	**keyword**	**frequency**	**keyword**	**frequency**	**keyword**	**frequency**	**keyword**
406	quenadatedetenga	476	venezuela	307	chavez	419	pueblo
382	hcapriles	455	maduro	286	forocandanga	368	chavez
310	venezuela	347	quenadatedetenga	283	nicolasmaduro	344	nicolasmaduro
310	maduro	319	hcapriles	238	venezuela	297	forocandanga
309	gobierno	314	gobierno	229	pueblo	237	hoy
299	pueblo	279	pueblo	212	psuv	211	venezuela
239	unidad	252	pais	152	maduro	207	psuv
238	pais	226	unidad	140	patria	196	maduro
237	caracas	164	venezolanos	131	revolucion	189	presidente
178	cambio	159	vecinos	130	gobierno	179	gobierno
161	vecinos	155	gracias	127	apoyoanicolasmaduro	175	contra
150	gracias	151	cambio	118	presidente	165	patria
136	progreso	133	contra	115	zulia	147	tropa
125	seguridad	132	plan	112	vivachavezcarajo	143	revolucion
125	baruta	127	regimen	112	bolivar	136	durancandanga
122	candidatos	125	ahora	111	simulacro	134	ubch
121	regimen	124	gente	105	contra	131	victoria
117	dias	121	dia	103	fotos	121	apoyoanicolasmaduro
117	contra	121	caracas	103	estado	119	fotos
110	bastaya	116	diputado	98	victoria	115	vivachavezcarajo
12	twitter	62	twitter	11	twitter	34	twitter

Notes: The table shows keyword counts for tweets by political leaders before and after the accounts' cancellations. Sample includes tweets by political leaders (government and opposition) in a 30-day window around October 31, 2013.

Table A4: Illustration of the difference-in-differences research design

	Mean of log (likes)			
	Opposition	Government	Difference (O-G)	Sports
	(1)	(2)	(3)	(4)
Panel A: Quasi-experiment of interest				
November 1 to November 15	1.233***	0.719***	0.514***	0.648***
	(0.134)	(0.104)	(0.168)	(0.075)
October 17 to October 31	0.903***	0.631***	0.273*	0.606***
	(0.127)	(0.091)	(0.156)	(0.085)
Difference (post - pre)	0.330***	0.088**	**0.242*** **	0.042
	(0.052)	(0.041)	**(0.066)**	(0.035)
Panel B: Placebo experiment				
October 17 to October 31	0.903***	0.631***	0.273*	0.606***
	(0.127)	(0.091)	(0.156)	(0.085)
October 2 to October 16	0.960***	0.638***	0.322**	0.643***
	(0.121)	(0.086)	(0.147)	(0.103)
Difference (post - pre)	-0.057	-0.007	**-0.05**	-0.038
	(0.039)	(0.032)	**(0.050)**	(0.039)

Notes: Sample includes tweets by political leaders (and sports-accounts) in the specified date ranges. Outcome measured is log(likes). Standard errors clustered at the user level in parentheses. Significant at (*) 90 percent, (**) 95 percent, (***) 99 percent confidence levels.

Table A5: Relationship between accounts' cancellation and publication times for tweets containing "Maduro"

Keyword sample:	None (last 20 overall)	*venezuela*	*pueblo*	*twitter*	Anti-Maduro (pooled)	Pro-Maduro (pooled)
	(1)	(2)	(3)	(4)	(5)	(6)
Post	-0.718***	-0.575**	-0.522	-2.119***	-0.929***	-0.217
	(0.196)	(0.255)	(0.425)	(0.366)	(0.151)	(0.135)
N	600	600	600	600	2,746	2,784
N-clusters	30	30	30	30	150	150
Keyword fixed-effects	No	No	No	No	Yes	Yes

Notes: Outcome measured is time of publication, measured in log seconds to the end of the day, as a proxy for tweet volume. Less time to the end of the day indicates greater tweet volume. All of the specifications include day of the week fixed effects. Standard errors clustered at the day level for columns 1-4, and clustered at the keyword-day level for columns 5-6. Significant at (*) 90 percent, (**) 95 percent, (***) 99 percent confidence levels.

Table A6: Differential change in publication times for tweets containing "Maduro"

Keyword sample:	Pro-Maduro vs. Neutral		Anti-Maduro vs. Neutral		Pro-Maduro vs. Anti-Maduro	
	(1)	(2)	(3)	(4)	(5)	(6)
Post	-0.617**		-0.610**		-0.265	
	(0.260)		(0.257)		(0.234)	
Pro-Maduro	1.882***					
	(0.234)					
Post * Pro-Maduro	0.358	0.416*				
	(0.348)	(0.240)				
Anti-Maduro			3.308***		1.392***	
			(0.196)		(0.177)	
Post * Anti-Maduro			-0.301	-0.294	-0.657**	-0.720***
			(0.298)	(0.216)	(0.280)	(0.171)
N	3,839	3,839	3,889	3,889	5,484	5,484
N-clusters	210	210	210	210	300	300
Day of the week fixed-effect	Yes	No	Yes	No	Yes	No
Keyword fixed-effects	No	Yes	No	Yes	No	Yes
Day fixed-effects	No	Yes	No	Yes	No	Yes

Notes: Outcome measured is time of publication, measured in standardized time until the end of the day, as a proxy for tweet volume. Less time to the end of the day indicates greater tweet volume. Standard errors clustered at the keyword-day level. Significant at (*) 90 percent, (**) 95 percent, (***) 99 percent confidence levels.

Table A7: Effects on proxy for negative engagement (log((replies+1)/(likes+1)))

Sample:	Tweets by @NicolasMaduro		Government and Sports	Opposition and Sports	Opposition and Government
	(1)	(2)	(3)	(4)	(5)
Post	0.195**	0.163**			
	(0.086)	(0.082)			
Post * Government			0.106*		
			(0.056)		
Post * Opposition				-0.065	-0.156***
				(0.052)	(0.041)
N	183	183	20,830	24,644	19,146
N-clusters			90	94	92
Controls	No	Yes	Yes	Yes	Yes
Day fixed-effects	No	No	Yes	Yes	Yes
User fixed-effects	No	No	Yes	Yes	Yes

Notes: All specifications use the proxy for negative engagement as the dependent variable. Sample includes tweets by @NicolasMaduro (columns 1 and 2) and tweets by prominent leaders (government, opposition and sports accounts; columns 3-5) in a 30-day window around October 31, 2013. Columns 1-2 use the pre/post regression specification and columns 3-5 use a difference-in-differences specification. Robust standard errors (columns 1-2) and standard errors clustered at the user level (columns 3-5) in parenthesis. Significant at (*) 90 percent, (**) 95 percent, (***) 99 percent confidence levels.

Table A8: Differential change in tweet likes by propensity of political leaders to mention Maduro

Sample:	Government		Opposition		Opposition vs. Government	
Model:	Difference-in-differences				Triple-difference	
	(1)	(2)	(3)	(4)	(5)	(6)
Post	0.024		0.240***		0.024	
	(0.054)		(0.056)		(0.054)	
MentionsMaduro	-0.699		3.281**		-0.699	
	(0.641)		(1.400)		(0.637)	
Post * MentionsMaduro	0.436	0.291	1.603*	1.755**	0.436	0.256
	(0.361)	(0.189)	(0.894)	(0.661)	(0.358)	(0.185)
Opposition					-0.037	
					(0.179)	
Post * Opposition					0.216***	0.140**
					(0.077)	(0.063)
Opposition * MentionsMaduro					3.980**	
					(1.532)	
Post * Opposition * MentionsMaduro					1.167	1.502**
					(0.959)	(0.699)
N	7,666	7,666	11,480	11,480	19,146	19,146
N-clusters	44	44	48	48	92	92
Controls (not shown)	No	Yes	No	Yes	No	Yes
Day fixed-effects	No	Yes	No	Yes	No	Yes
User fixed-effects	No	Yes	No	Yes	No	Yes

Notes: Dependent variable is natural log of number of likes for all columns. Sample includes tweets by political leaders in a 30-day window around October 31, 2013. Standard errors clustered at the user level in parenthesis. Significant at (*) 90 percent, (**) 95 percent, (***) 99 percent confidence levels.

Table A9: Descriptive long-run analysis (6-month window)

Sample:	Tweets by @NicolasMaduro		Tweets that mention @NicolasMaduro	Tweets that contain "Maduro" keyword	Tweets by opposition and government leaders		
Dependent variable:	Log retweets	Log replies	Log seconds to end of day		Log likes	Log replies	Log likes
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Post	-1.736***	0.599***					
	(0.021)	(0.046)					
Post*Anti-Maduro			-0.055	-0.302***			
			(0.058)	(0.071)			
Post*Opposition					0.319***	0.215***	0.338***
					(0.057)	(0.066)	(0.073)
Post*Opposition*MentionsMaduro							-0.795
							(1.092)
N	932	932	30,813	32,243	107,795	107,795	107,795
N-clusters			1,839	1,840	94	94	94
Original specification	Table 1, Column 2	Table 1, Column 6	Table 3, Column 6	Table A6, Column 6	Table 5, Column 5	Table 5, Column 6	Table A8, Column 6

Notes: The table shows the replication exercises for the main results in the 6-month long window. Refer to original tables for details on specification, controls, and clustering of standard errors. Results significant at (*) 90 percent, (**) 95 percent, (***) 99 percent confidence levels.

Figure A1: Political leaders' reactions to the closure of the accounts

Figure A2: Example of reply counts for controversial (right) vs non-controversial (left) tweet

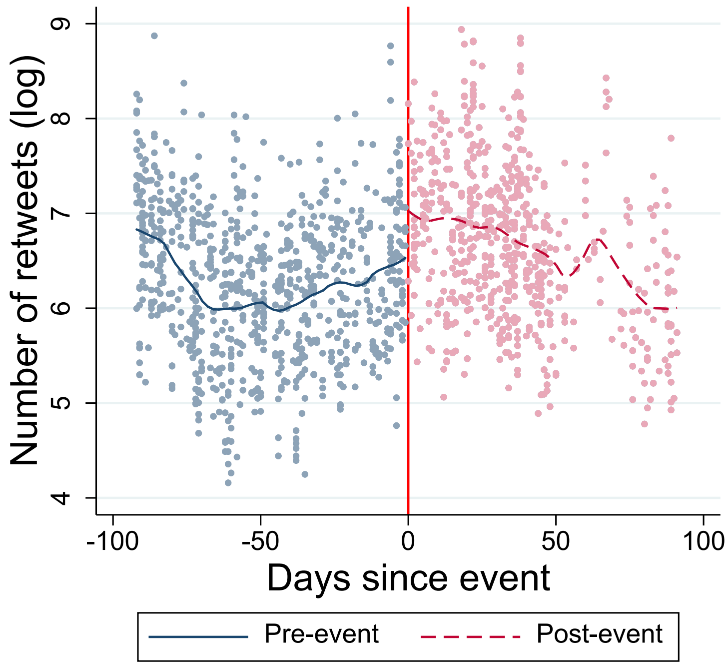

Figure A3: Retweets for the leader of the opposition Henrique Capriles (@hcapriles)

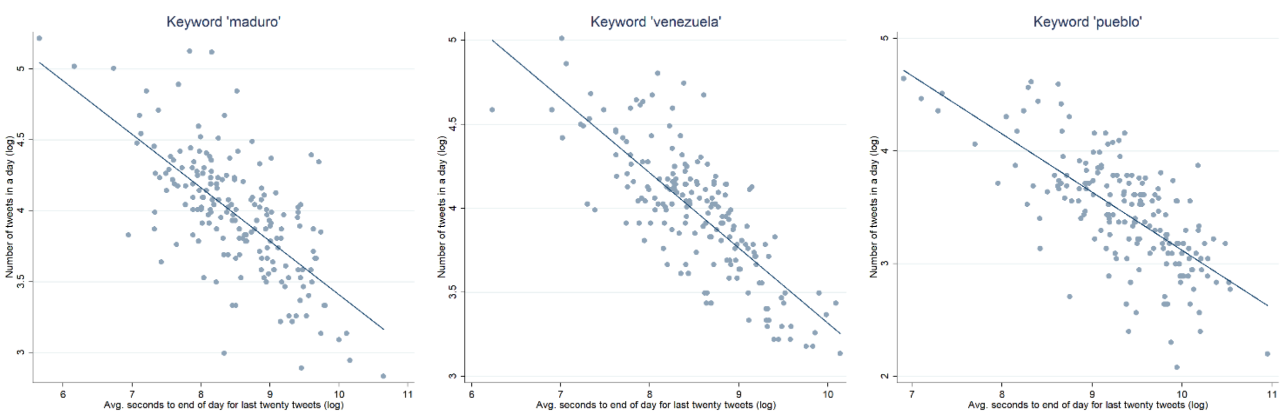

Figure A4: Relationship between number of tweets and time of publication of last twenty tweets at the daily level (each scatter point is one day)

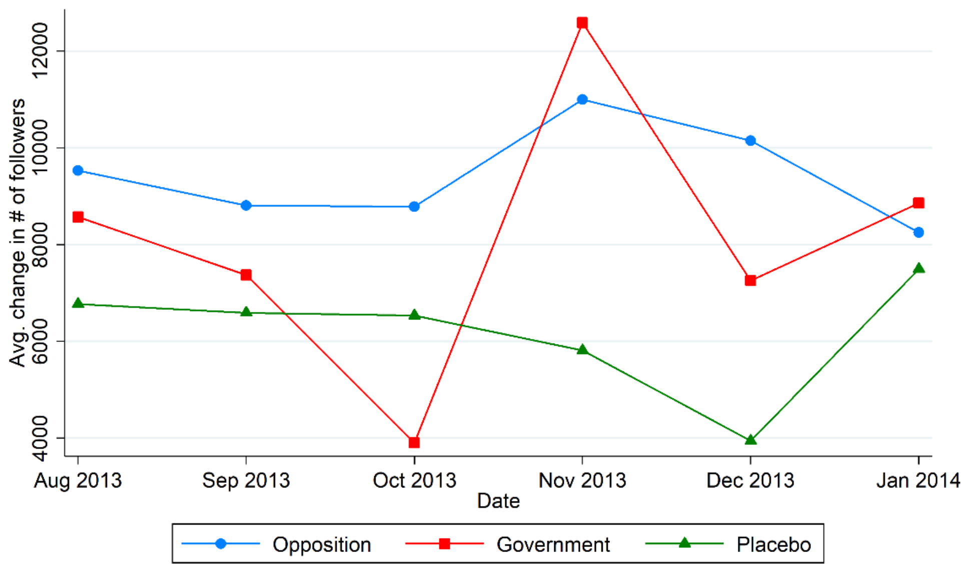

Figure A5: Average change in number of followers for selected accounts

Figure A6: Front page of the TeleSUR website on November 2nd/3rd (the website is dated November 2nd though the archive date is November 3rd)

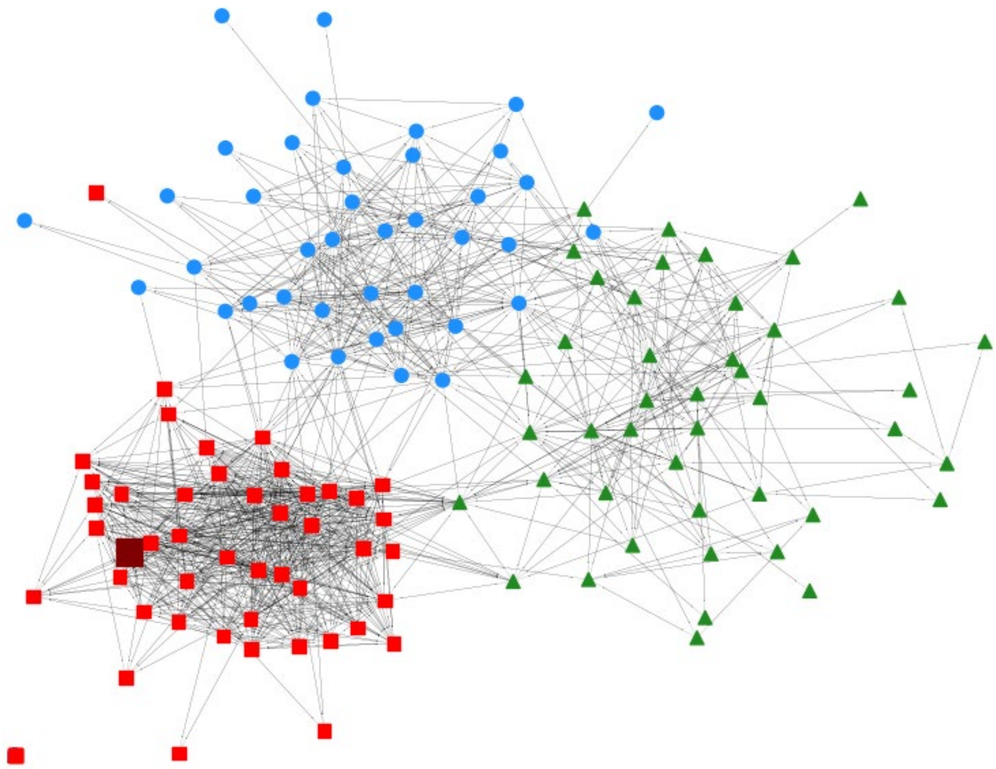

Figure A7: Twitter network of selected accounts

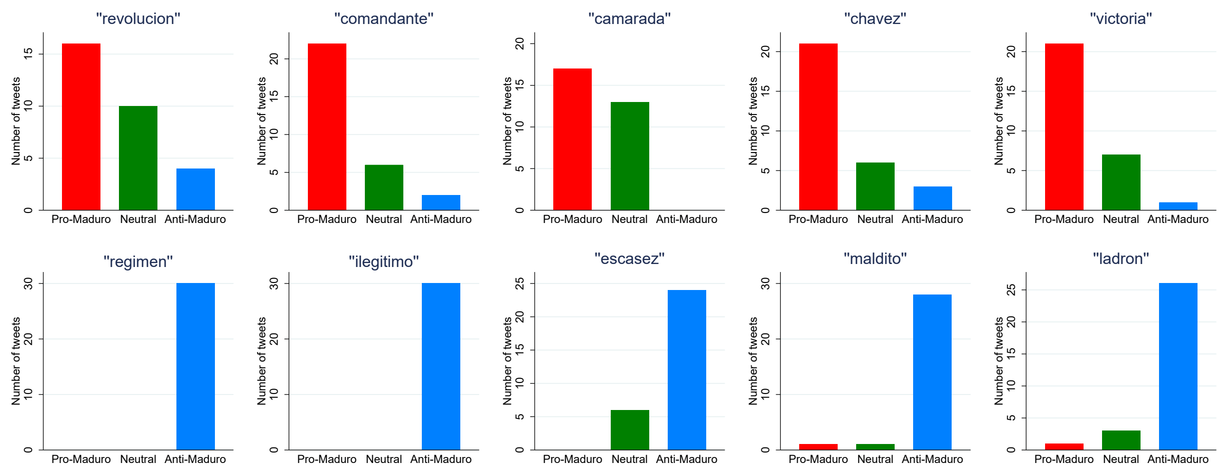

Figure A8: Results from manually coding a subsample of 300 randomly selected tweets which mention @NicolasMaduro, by keyword

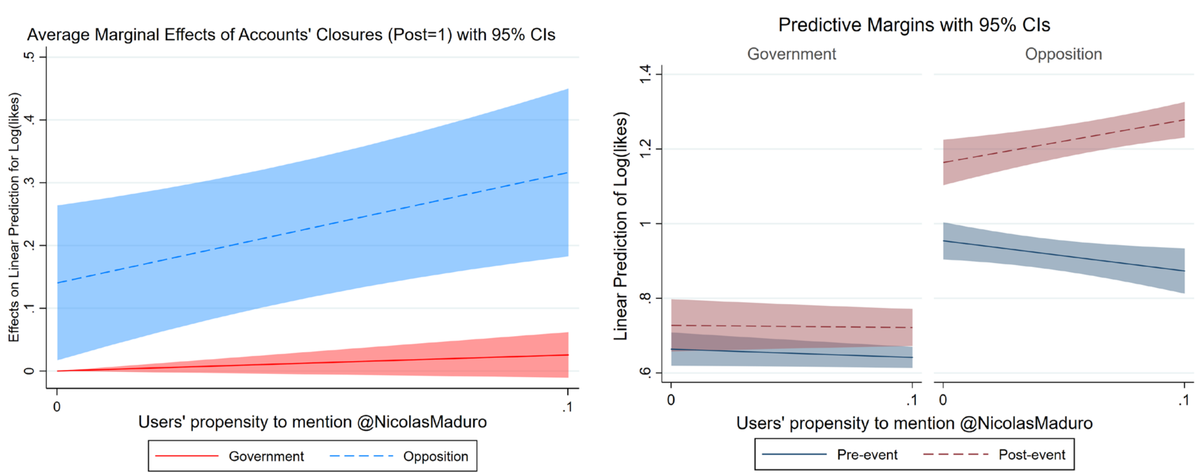

Figure A9: Heterogeneity by users' propensity to mention @NicolasMaduro

Model specification: Fixed effects (left) and Random effects (right)