

Fake It Til You Make It: A Natural Experiment to Identify European Politicians' Benefit from Twitter Bots*

Bruno Castanho Silva
Cologne Center for Comparative Politics
University of Cologne
bcsilva@wiso.uni-koeln.de

Sven-Oliver Proksch
Cologne Center for Comparative Politics
University of Cologne
so.proksch@uni-koeln.de

This paper has been published FirstView in the *American Political Science Review*, doi:10.1017/S0003055420000817.

Abstract

Social media giants stand accused of facilitating illegitimate interference with democratic political processes around the world. Part of this problem is malicious bots: automated fake accounts passing as humans. However, we lack a systematic understanding of which politicians benefit most from them. We tackle this question by leveraging a Twitter purge of malicious bots in July 2018 and a new dataset on Twitter activity by all members of national parliaments (MPs) in the EU in 2018. Since users had no influence on how and when Twitter purged millions of bots, it serves as an exogenous intervention to investigate whether some parties or politicians lost more followers. We find drops in follower counts concentrated among radical right politicians, in particular those with strong anti-EU discourse. This is the first set of empirical, causally-identified evidence supporting the idea that the radical right benefits more from malicious bots than other party families.

*We would like to thank Danielle Pullan, Pit Rieger, Rebecca Kittel, Leonie Diffené, Felix Reich, Barbara Zucchi Nobre Silva, Jasmin Spekkers, Noam Himmelrath, Lennart Schürmann and Lea Kaftan for excellent research assistance, and Leo Baccini, Jens Wäckerle, Bastian Becker, Ahmet Suerdem, participants of the PolText Conference at Waseda University in Tokyo, Sep. 13-15 2019, and the three anonymous reviewers for helpful comments and suggestions. All remaining errors are our own. This research has been funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – EXC 2126/1– 390838866. The replication material for this paper is available in the APSR Dataverse at <https://doi.org/10.7910/DVN/PAMABU>

Tech giants such as Google, Facebook, and Twitter have been under scrutiny for failing to stop illegitimate interferences with democratic processes on their platforms. These involve different attempts at manipulating public opinion, including automated fake accounts passing as humans (malicious bots). To better understand the potential democratic threat, researchers have studied politically active bots (e.g. Keller and Klinger, 2019; Shao et al., 2018), and their impact on public debate (Bovet and Makse, 2019; Ferrara, 2017). Our current understanding of social media bots in politics, however, relies primarily on single-country case studies – mostly the U.S. (exceptions include Keller and Klinger, 2019; Morales, 2020), and on bot-detection algorithms which may be prone to low accuracy (Rauchfleisch and Kaiser, 2020).

Based on a Twitter dataset including all members of national parliaments (MPs) in EU member states in 2018, we provide the first cross-national assessment of social media bots’ influence on European politics. Around July 11, 2018, Twitter purged millions of malicious bots without warning. We leverage this external intervention to investigate what kinds of parties lost the most followers. Findings indicate that radical right parties were significantly more affected, confirming untested hunches about how far right politicians rely on automated fake accounts to generate an appearance of popular support. Moreover, our analysis shows that the losses were larger among the most Euroskeptic of radical right politicians.

Twitter Bots

Individual popularity is the main currency on social media platforms, since high popularity cues trustworthiness and influence (Castillo, Mendoza and Poblete, 2011). Having more followers may lead to higher impact, since followers react to and spread politicians’ messages (Keller and von Königslöw, 2018; Popa et al., 2020). The relationship is reciprocal, and there is evidence that politicians are responsive to changes in issue attention of their Twitter followers (Barberá et al., 2019). The buzz also draws more attention from journalists who use Twitter to help build their agendas (Wihbey, Joseph and Lazer, 2019).

On top of its value as a communication tool, Twitter can also support parties' structural organization and resource collection during election campaigns (Jungherr, 2016). Therefore, political actors have an incentive to get as many followers as possible, including through illegitimate means. One such way relies on "social bots": a social media account run by a computer algorithm which tries to emulate human behavior on social media pretending to be a human user (Ferrara, 2017).¹ These bots engage in what Twitter calls "platform manipulation", defined as "attempts to disrupt the health of the public conversation via malicious automation and spam tactics" (Twitter, 2018). Such bots are the focus of this study, and we refer to them interchangeably as "social bots" or "malicious bots".

Who Benefits?

Social bots may follow and support a politician either because politicians themselves (or their social media teams) hire an individual or agency to set them up, or because someone else, for whatever reason or interest, sets up bots to act in favor of a politician or cause. Studies have shown that bots boosted Donald Trump's popularity (Onuchowska, Berndt and Samtani, 2019) and spread fake news favoring American conservatives (Bovet and Makse, 2019). The European Union took the matter seriously before the 2019 European Parliament elections and worked with social media platforms to try to ensure minimal external interference with the campaign process (European Commission, 2019). In the EU's words, "the evidence collected revealed a continued and sustained disinformation activity by Russian sources aiming to suppress turnout and influence voter preferences" (European Commission, 2019, p. 3).

It is generally impossible to know who is behind a Twitter bot. Nevertheless, we can still identify which politicians have more malicious bots among their followers, which leads them to being perceived as more popular than they are. Within the European context, our main hypothesis is that politicians from radical right parties are more likely to benefit from social bots to inflate their follower counts, and therefore experienced larger losses

¹Some bots do not disguise themselves as human. For example, news agencies have bots tweeting their headlines at regular intervals. We do not refer to these when talking about social or malicious bots.

of followers during the Twitter purge. The reasons are threefold. First, radical right parties rely on claims of representing the “silent majority” as the basis of their legitimacy (Mudde, 2019), meaning those parties have an added incentive to project an image of popularity, as the sheer number of supporters is at the core of their appeal. Second, they hold strong Euroskeptic positions (Vasilopoulou, 2018; Rooduijn, 2019), in line with the external agents active on European political social media to “attack the EU and its values” (European Commission, 2019, p. 3). Third, previous research has found that malicious Twitter bots favor (radical) right-wing political actors in countries including France (Ferrara, 2017) and the U.S. (Onuchowska, Berndt and Samtani, 2019).

The Purge

On July 11, 2018, news broke out that Twitter was conducting a massive deletion of bots engaging in platform manipulation. The New York Times reported that the “total combined follower count” on Twitter might fall by 6% by July 13, as the platform made efforts to curb the activities of automated profiles created to artificially inflate some users’ popularity (Confessore and Dance, 2018). In the words of Twitter’s Vice President for Trust and Safety, Del Harvey, “We don’t want to incentivize the purchase of followers and fake accounts to artificially inflate follower counts, because it’s not an accurate measure of someone’s influence”.² This purge followed the internal development of new methods to identify and suspend malicious automation and spamming (Roth and Harvey, 2018).

This event provides a unique opportunity to investigate the relationship between social media bots and politicians. Twitter’s actions were an exogenous shock for malicious bot makers. The deleted accounts had been suspended in the weeks before July 11, flagged as suspicious of being bots engaged in platform manipulation. Meaning, they were still included in follower counts up to deletion (July 11)³ but unable to create new posts since being suspended.⁴ Once an account was suspended, the user had the

²Cited in Confessore and Dance (2018).

³From that time onward, Twitter policy changed so that challenged accounts are removed from follower counts as soon as they are suspended – see Roth and Harvey (2018) and Confessore and Dance (2018).

⁴This is the reason why one cannot use this design to identify purge effects on retweets or likes – the accounts could not do so even before being officially deleted. In the Online Appendix we plot retweets per post for each party family over 2018. There is a drop for the radical right in May and June, tentatively

chance to verify that they were a real person, for example by phone message, and get their account restored (Roth and Harvey, 2018). Therefore, accounts deleted around July 11 had been suspended on a suspicion of being malicious bots, and the people who maintained them did not verify the accounts’ legitimacy. Thus, even if far-right Twitter users may behave in more “suspicious” ways compared to other users (Hjorth and Adler-Nissen, 2019) triggering more false positives, they could have avoided the July purge through verification. On the other hand, the purge was far-reaching and hit all types of celebrities,⁵ so we do not expect a “false negative” problem whereby malicious bots following mainstream politicians were less likely to be flagged. Nevertheless, because Twitter is not transparent about what exact parameters it used, we have to rely on their official communication (such as Roth and Harvey, 2018; Twitter, 2018), along with media accounts, that this was a) a general clean up of malicious bots not restricted to a single bot farm, and b) largely restricted to bots and not other kinds of terms of service violations.⁶

Most researchers studying Twitter bots have used detection algorithms (e.g. Keller and Klinger, 2019; Shao et al., 2018). While these might be accurate within countries, they are problematic across languages (Rauchfleisch and Kaiser, 2020). A bot-detection algorithm must not only tell apart bots from humans, but also neutral bots – say, news bots – from politically active ones passing as humans (Stukal et al., 2019). We rely on Twitter’s own detection algorithms to tell apart malicious bots from others and from humans, which we expect to be much more accurate than what any researcher could come up with using publicly available Twitter data. Moreover, we use the purge to get an immediate estimate of the number of bot followers for the population of all 77 million combined followers EU politicians had in July 2018. Due to Twitter’s rate limits, using a bot-detection algorithm on this list would take approximately 12 years (4,456 days) of

suggesting an effect of these accounts’ suspension, but the variance is very high over time so we refrain from making any claims.

⁵For example Barack Obama, Shaquille O’Neal, Oprah Winfrey, and Ellen DeGeneres all lost more than one million followers each (Jacobs, 2018).

⁶We acknowledge this limitation and welcome future studies using other methods or events to replicate the findings.

non-stop computing.⁷

Data

We have collected in real time all tweets posted by national members of parliament from all 28 EU member states between February and December 2018, using the streamR R package (Barberá, 2018) connected to the Twitter Streaming API. 4,964 out of more than 7,000 EU MPs had Twitter accounts, 2,550 of whom were active between June and August. The data contains the text of each tweet and the number of followers in the moment of posting, so that we rely on the follower counts for users who tweeted both before and after the purge to investigate its effects. Around 1,900 users tweeted at least once both right before and right after, meaning once between July 2 and July 9, and once between July 13 and July 20. From this sample, we take the last available follower count before July 9, and the first after July 13.⁸ Overall, there is a small reduction in the average number of politicians' followers: from 35,022 to 34,604, a loss of 1.2% within five days.

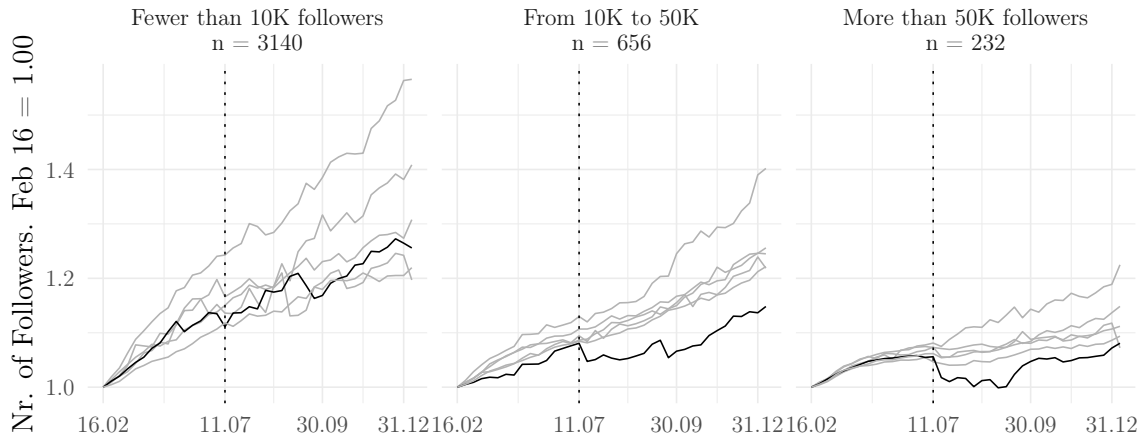
We use the Chapel Hill Expert Survey's (CHES, Bakker et al., 2015) classification of parties into party families to examine politicians' numbers of followers over the year.⁹ In the top panel of Figure 1 we split politicians into three groups, depending on their popularity, and set their follower count on February 16 to 1, thus allowing us to plot the weekly evolution of follower counts until December. The Figure shows a sharp decline among politicians of radical right parties exactly on July 11, particularly among users with more than ten thousand followers. The most popular radical right politicians on Twitter, with more than 50,000 followers, had on average 5% more followers in July than

⁷180 users every 15 minutes. In the Online Appendix we describe the usage of a bot detection algorithm on a random sample of followers from politicians in our data, to exemplify the differences to the purge.

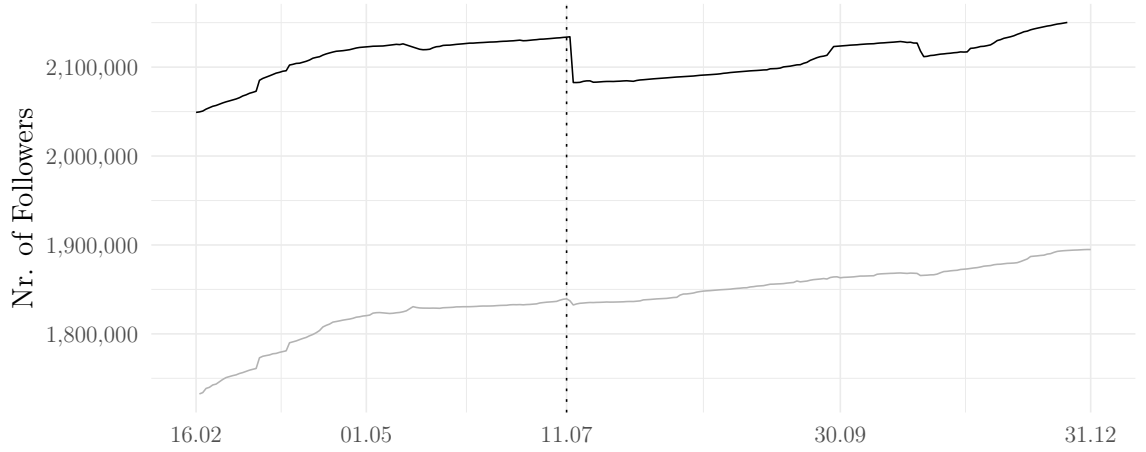
⁸News broke on July 11, mentioning that the deletion had already started. We move one more day further back to capture an estimate right before the purge. Moreover, Twitter never stated exactly when they stopped this one particular effort, so we allow for two days after news came out to look at follower counts.

⁹The CHES follows Hix and Lord (1997, p. 42–4) in classifying parties as radical right when their discourse features anti-immigration, nationalism, and right-wing stances on sociopolitical issues.

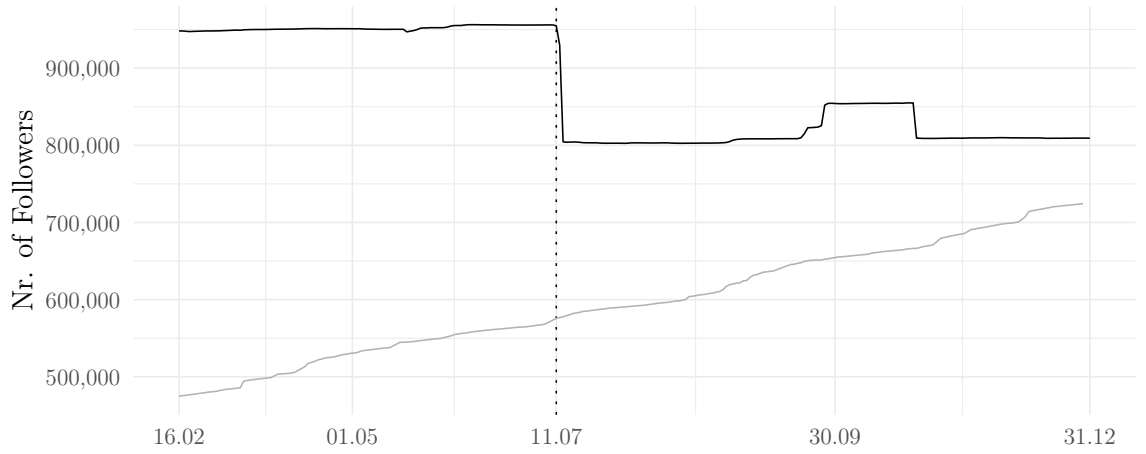
Figure 1: Trends in numbers of followers by party families and selected politicians



— Other families — Radical Right



— Jeremy Corbyn — Marine Le Pen



— Geert Wilders — Theresa May

in February, but the purge brought them immediately back to their February numbers.¹⁰ Moreover, the growth rate for radical right politicians after the purge is smaller than before.¹¹ A lower growth rate, departing from a smaller base, indicates that inflated follower counts were perhaps helping these politicians get more new followers than they would have organically.

The middle and bottom panels of Figure 1 focus on four well known politicians and their daily number of followers. Marine Le Pen and Geert Wilders, leaders of the radical right in France and the Netherlands, had a sharp drop in their follower counts right around July 11: Wilders lost 150,000 followers. In contrast Jeremy Corbyn and Theresa May, then leaders of Labour and Conservatives in Britain, show no signs of such discontinuity.¹²

Model

We test our hypothesis with multilevel models. The dependent variable is a user's percentage change in followers between July 9 and July 13.¹³ Percentage changes are used due to heavy skews of follower counts: most politicians have a number in the lower thousands, and a few are above 2 million. The main independent variable is belonging to the radical right party family, coded based on CHES data. We also add a dummy for the radical left as a control.

The second main independent variable is anti-EU discourse. Most radical right parties are Euroskeptic, but not all Euroskeptics are radical right. It is important to see whether the driving force is Euroskepticism or belonging to the radical right family. EU positions are estimated both at the individual MP level and at the party level. For individual MPs, we measure the sentiment of their tweets which are related to Europe before the purge, using the Lexicoder Sentiment Dictionary (Young and Soroka, 2012) translated into all

¹⁰In the Online Appendix we show the ten parties with the largest drop between July 9 and 13. Most are radical right.

¹¹Full numbers for all party families in the Online Appendix.

¹²Figure A.1 in the Online Appendix shows the mean and median change in number of followers by country. Median values are mostly around zero, while means are negative, indicating that the purge had zero effect for most politicians, but meant a large drop for a few.

¹³Using the last available number up to July 9, and the first after July 13. Analyses are restricted to tweets posted one week before/after news about the purge came out – meaning, July 2–9 and July 13–20.

EU languages by Proksch et al. (2019). Sentiment ranges from negative to positive based on the log of the ratio of positive to negative words in tweets mentioning Europe or the EU by each MP.¹⁴ We also use the standard EU position variable from the CHES dataset (Bakker et al., 2015), measuring party positions on EU integration from 1 (most negative) to 7 (most positive).

The models furthermore include standard legislative behavior control variables that are available cross-nationally. At the individual level they include an MP’s gender, number of terms served in parliament, and having any previous cabinet experience. At the party level, a party’s government-opposition status at the time of the purge and its parliamentary seat share. Due to the apparent larger effect of the purge on popular accounts seen in Figure 1, we also control for the logged number of followers. We also control for the sentiment of an MP in all their tweets, to control for users’ overall (negative) tone. Clearly, these do not represent all potential factors to influence MP’s online behavior, and we expect the party and country fixed effects to capture remaining heterogeneity.

Who Lost?

Table 1 shows the estimates from models on the percentage change in followers during the purge. Model 1 indicates that the purge is not associated with parties’ EU positions in general: neither EU sentiment on Twitter nor the CHES measure are predictors of percentage changes in followers. Model 2 replaces the CHES EU position variable with the two main Euroskeptic party family dummies: radical right and radical left. It shows that politicians belonging to radical right parties saw a significant drop in their following base ($\beta = -1.08$, or a loss of 1.08% in relation to other parties), which did not happen for MPs from radical left parties. Models 3 and 4 add an interaction between party family and EU sentiment on Twitter, visualized in Figure 2.¹⁵ It is significant for the radical

¹⁴First, we subset the corpus to include only tweets that mention a keyword that signals it refers to Europe, and then apply the sentiment dictionary to this EU-related corpus of tweets. The list of keywords is in the Online Appendix.

¹⁵Those observations on the most positive end, meaning radical right MPs who are positive about Europe, are mostly Polish MPs from the ruling PiS. According to CHES, this party is moderately Euroskeptic, but is more divided than average on the European integration issue (Bakker et al., 2015).

right (Model 3), but not the left (Model 4), indicating that losses were concentrated among radical right politicians with the most negative discourse about Europe on Twitter. This is in line with the EU Commission’s warning of actors working to spread anti-EU messages (European Commission, 2019).

The Online Appendix contains several robustness and sensitivity checks: first, we test models using a larger time-window before and after the purge (two and four weeks on each side), increasing the number of observations but reducing the validity of causal claims. Substantive results hold. Second, we run the models from Table 1 dropping outliers, defined as users who had percentage change in their follower count above 30% during the purge (41 MPs in total), and results remain the same. Third, we present a difference-in-differences model testing the effects on total follower counts instead of percentage changes: only the radical right has a significant loss of followers ($\beta = -3251.52$ fewer followers). Fourth, we use a placebo approach of “moving the purge” to all dates between June 1 and August 31, and re-run the same models, to see if we observe coefficients as large as those we see on the real dates. At no other time were the coefficients for either the main radical right variable nor the interaction with EU discourse nearly as large as the effects of the actual purge. Finally, we also test the effects of party family on “friends”, meaning how many accounts a user follows. Politicians from radical right parties also had a significant drop in the number of accounts they follow around the time of the purge.

Conclusions

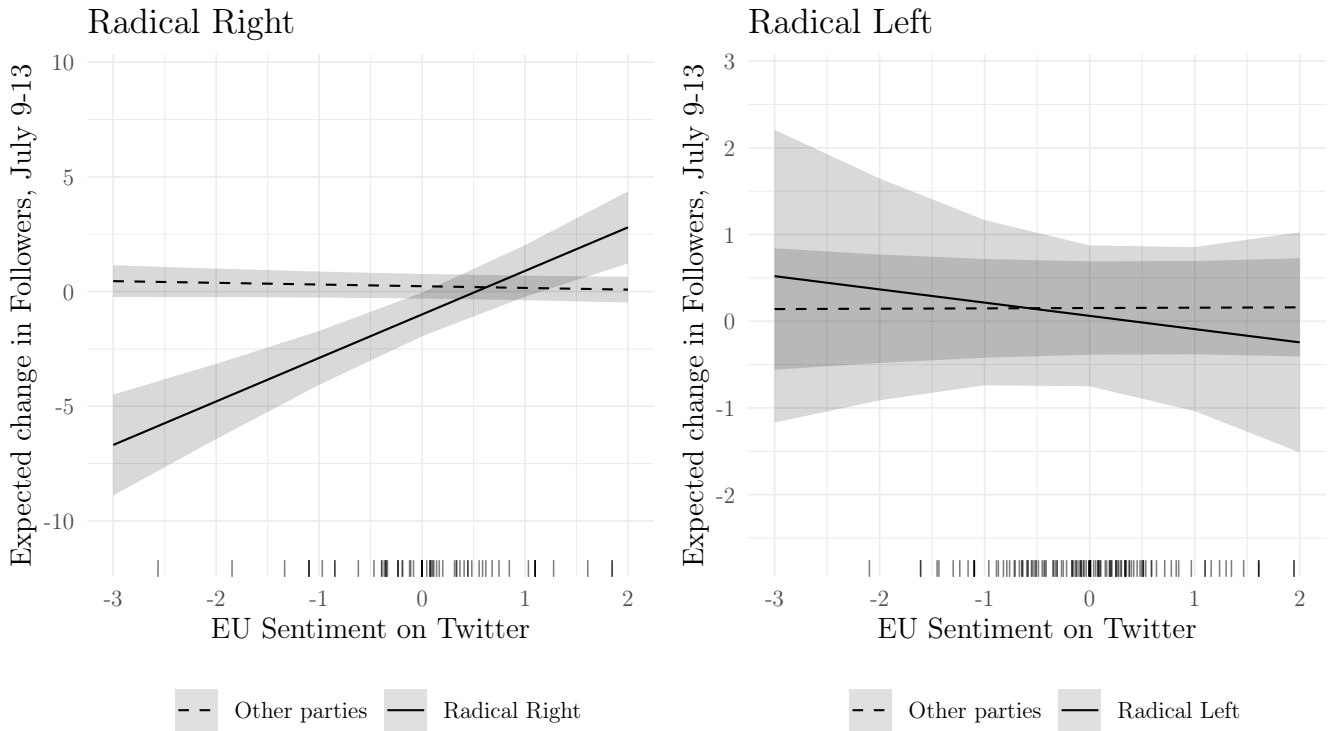
Winning the social media popularity contest allows political actors to claim they are widely successful among the public. In this paper we set out to test which parties in the EU benefit most from social bots to artificially inflate their popularity on Twitter. We identify significant and substantive following of radical right politicians in Europe by malicious bots. Using a large purge of such bots by Twitter in July 2018, we find that belonging to a radical right party is associated with an average loss of up to 5% of an MP’s followers within that week, particularly among the most popular and louder

Table 1: Individual and Party-Level Determinants of Percentage Changes in Followers – July 9-13

	Model 1	Model 2	Model 3	Model 4
Intercept	3.86*	4.50*	4.49*	4.51*
	[2.88; 4.81]	[3.33; 5.68]	[3.39; 5.65]	[3.38; 5.67]
Male	-.03	-.02	-.01	-.02
	[-.24; .19]	[-.22; .17]	[-.22; .18]	[-.24; .20]
Terms in office	-.02	-.02	-.02	-.02
	[-.10; .06]	[-.11; .06]	[-.10; .06]	[-.10; .06]
Cabinet experience	.10*	.10*	.10*	.10*
	[.01; .21]	[.00; .20]	[.00; .21]	[.00; .20]
Twitter Sentiment	-.50*	-.51*	-.51*	-.51*
	[-.89; -.03]	[-.96; -.08]	[-.94; -.08]	[-.93; -.07]
Twitter EU Sentiment	-.01	-.01	-.07	.00
	[-.13; .12]	[-.13; .12]	[-.21; .05]	[-.13; .13]
Nr. of Followers (log)	-.40*	-.40*	-.41*	-.40*
	[-.48; -.32]	[-.49; -.32]	[-.49; -.33]	[-.48; -.32]
Seat share	.00	.00	.00	.00
	[-.01; .02]	[-.01; .02]	[-.01; .01]	[-.01; .02]
In government	-.10	-.13	-.05	-.13
	[-.49; .27]	[-.50; .25]	[-.45; .34]	[-.52; .30]
EU Position	.01	-.10	-.09	-.10
	[-.09; .12]	[-.24; .05]	[-.23; .04]	[-.24; .04]
Radical right		-1.08*	-1.23*	-1.08*
		[-1.98; -.24]	[-2.01; -.38]	[-2.01; -.21]
Radical left		-.09	-.06	-.09
		[-.72; .56]	[-.66; .55]	[-.73; .57]
Twitter EU Sentiment * Radical right			1.97*	
			[1.31; 2.66]	
Twitter EU Sentiment * Radical left				-.16
				[-.66; .34]
AIC	8391.54	8389.47	8357.87	8391.97
BIC	8463.79	8472.83	8446.79	8480.89
Num. obs.	1915	1915	1915	1915
N. parties	119	119	119	119
N. countries	25	25	25	25

* 0 outside the bootstrapped 95% confidence interval.

Figure 2: Expected Change in Followers for Radical Right Parties their EU Sentiment (Models 3 and 4)



anti-EU radical right politicians.

These findings are the first evidence that radical right parties benefit more than other parties from social media bots to inflate their popularity, and thus artificially get more attention than they should. This has implications for how the EU continues to tackle online abuses and misinformation, and for journalists who use Twitter popularity as a source, which might lead to giving radical right parties more attention than they deserve. Considering how media attention can boost the radical right’s electoral performance (de Jonge, 2019), fake accounts may help those politicians be electorally successful. It seems possible that having enough fake followers helps attract real ones. Finally, we also introduce a novel dataset on Twitter use by politicians in EU countries. With complete coverage of tweets posted by members of 28 national parliaments over several months, the analysis performed here is one of many possibilities.¹⁶

We must highlight that our findings only show that radical right parties benefited from

¹⁶Within Twitter’s Terms of Use, we cannot publish tweets’ text and metadata, but only the tweets’ IDs. From the IDs, researchers can use the Twitter API to download all tweets which are still available.

fake followers to get an inflated follower count. We do not claim that parties themselves are setting up or buying these followers. This is something we cannot test. Nevertheless, knowing who benefits from these bots, and the consequences for politics, is of absolute relevance for keeping democratic processes clean and fair.

References

- Bakker, Ryan, Catherine de Vries, Erica Edwards, Liesbet Hooghe, Seth Jolly, Gary Marks, Jonathan Polk, Jan Rovny, Marco Steenbergen and Milada Anna Vachudova. 2015. “Measuring party positions in Europe: The Chapel Hill expert survey trend file, 1999-2010.” *Party Politics* 21(1):143–152.
- Barberá, Pablo. 2018. *streamR: Access to Twitter Streaming API via R*. R package version 0.4.3.
URL: <https://cran.r-project.org/web/packages/streamR>
- Barberá, Pablo, Andreu Casas, Jonathan Nagler, Patrick J. Egan, Richard Bonneau, John T. Jost and Joshua A. Tucker. 2019. “Who Leads? Who Follows? Measuring Issue Attention and Agenda Setting by Legislators and the Mass Public Using Social Media Data.” *American Political Science Review* 113(4):883–901.
- Bovet, Alexandre and Hernán A. Makse. 2019. “Influence of fake news in Twitter during the 2016 US presidential election.” *Nature Communications* 9(7):1–14.
- Castillo, Carlos, Marcelo Mendoza and Barbara Poblete. 2011. “Information Credibility on Twitter.” *Proceedings of the 20th International Conference on World Wide Web – WWW ’11*.
- Confessore, Nicholas and Gabriel J.X. Dance. 2018. “Battling Fake Accounts, Twitter to Slash Millions of Followers.” *The New York Times* July 11.
URL: <https://www.nytimes.com/2018/07/11/technology/twitter-fake-followers.html>

- de Jonge, Léonie. 2019. “The Populist Radical Right and the Media in the Benelux: Friend or Foe?” *The International Journal of Press/Politics* 24(2):189–209.
- European Commission. 2019. *Report on the implementation of the Action Plan Against Disinformation*. Brussels: JOIN(2019).
- Ferrara, Emilio. 2017. “Disinformation and Social Bot Operations in the Run Up to the 2017 French Presidential Election.” *First Monday* 22(8):1–33.
- Hix, Simon and Christopher Lord. 1997. *Political Parties in the European Union*. New York: Macmillan Education.
- Hjorth, Frederik and Rebecca Adler-Nissen. 2019. “Ideological Asymmetry in the Reach of Pro-Russian Digital Disinformation to United States Audiences.” *Journal of Communication* 69(2):168–192.
- Jacobs, Julia. 2018. “In Twitter Purge, Top Accounts Lose Millions of Followers.” *The New York Times* July 12.
URL: <https://www.nytimes.com/2018/07/12/technology/twitter-followers-nyt.html>
- Jungherr, Andreas. 2016. “Four Functions of Digital Tools in Election Campaigns: The German Case.” *The International Journal of Press/Politics* 21(3):358–377.
- Keller, Tobias R. and Katharina Kleinen von Königslöw. 2018. “Followers, Spread the Message! Predicting the Success of Swiss Politicians on Facebook and Twitter.” *Social Media + Society* 4(1):1–11.
- Keller, Tobias R. and Ulrike Klinger. 2019. “Social Bots in Election Campaigns: Theoretical, Empirical, and Methodological Implications.” *Political Communication* 36(1):171–189.
- Morales, Juan S. 2020. “Perceived Popularity and Online Political Dissent: Evidence from Twitter in Venezuela.” *The International Journal of Press/Politics* 25(1):5–27.
- Mudde, Cas. 2019. *The Far Right Today*. Cambridge: Polity.

- Onuchowska, Agnieszka, Donald J. Berndt and Sagar Samtani. 2019. Rocket Ship or Blimp? – Implications of Malicious Accounts Removal on Twitter. In *Proceedings of the 27th European Conference on Information Systems (ECIS), Stockholm & Uppsala, Sweden, June 8–14*.
- Popa, Sebastian Adrian, Zoltán Fazekas, Daniela Braun and Melanie-Marita Leidecker-Sandmann. 2020. “Informing the Public: How Party Communication Builds Opportunity Structures.” *Political Communication* 37(3):329–349.
- Proksch, Sven-Oliver, Will Lowe, Jens Wäckerle and Stuart Soroka. 2019. “Multilingual Sentiment Analysis: A New Approach to Measuring Conflict in Legislative Speeches.” *Legislative Studies Quarterly* 44(1):97–131.
- Rauchfleisch, Adrian and Jonas Kaiser. 2020. “The False Positive Problem of Automatic Bot Detection in Social Science Research.” *Berkman Klein Center Research Publication* March(2020-3).
- Rooduijn, Matthijs. 2019. “State of the field: How to study populism and adjacent topics? A plea for both more and less focus.” *European Journal of Political Research* 58(1):362–372.
- Roth, Yoel and Del Harvey. 2018. “How Twitter is fighting spam and malicious automation.”
URL: https://blog.twitter.com/official/en_us/topics/company/2018/how-twitter-is-fighting-spam-and-malicious-automation.html
- Shao, Chengcheng, Giovanni Luca Ciampaglia, Onur Varol, Kai-Cheng Yang, Alessandro Flammini and Filippo Menczer. 2018. “The spread of low-credibility content by social bots.” *Nature Communications* 9:1–9.
- Stukal, Denis, Sergey Sanovich, Joshua A. Tucker and Richard Bonneau. 2019. “For Whom the Bot Tolls: A Neural Networks Approach to Measuring Political Orientation of Twitter Bots in Russia.” *SAGE Open* 9(2):1–16.

Twitter. 2018. “Transparency Reports: Platform Manipulation.”.

URL: <https://transparency.twitter.com/en/platform-manipulation.html#platform-manipulation-jan-jun-2018>

Vasilopoulou, Sofia. 2018. The Radical Right and Euroskepticism. In *The Oxford Handbook of the Radical Right*, ed. Jens Rydgren. Oxford: Oxford University Press pp. 122–140.

Wihbey, John, Kenneth Joseph and David Lazer. 2019. “The social silos of journalism? Twitter, news media and partisan segregation.” *New Media & Society* 21(4):815–835.

Young, Lori and Stuart N. Soroka. 2012. “Affective news: The automated coding of sentiment in political texts.” *Political Communication* 29(2):205–231.