The Returns to Viral Media: The Case of US Campaign Contributions*

Johannes Böken, University of Warwick
Mirko Draca, University of Warwick
Nicola Mastrorocco, University of Bologna
Arianna Ornaghi, Hertie School

June 27, 2025

Abstract

Social media has created new, highly competitive markets for attention. But to what extent does attention on social media generate tangible economic returns and how are these returns characterized? Using a daily dataset of Twitter activity and campaign contributions for US Members of Congress (2019-2020), we show that attention on Twitter, as measured by likes, increases small donations. However, the effect is highly skewed: only a few members benefit substantially, consistent with a winner-takes-all market. These results are confirmed using a geography-based causal design tracking donation patterns across counties, showing that the increase in donations from attention on Twitter comes disproportionately from high Twitter usage areas.

JEL Codes: D72, P00
Keywords: Social Media, Twitter, Campaign Contributions

*Johannes Böken: johannes.boken@warwick.ac.uk; Mirko Draca: m.draca@warwick.ac.uk; Nicola Mastrorocco: nicola.mastrorocco@unibo.it; Arianna Ornaghi: ornaghi@hertie-school.org. We thank Nathan Canen, Adam di Lizia, Rafael Jiménez-Durán, Karsten Müller, Nicola Persico, Carlo Schwarz, Camille Urvoy, Niko Wolf, and seminar and conference participants at the 2023 NICEP Conference, the 12th CRC TRR 190 Retreat, London Business School, Navarra University, the IEB Political Economy Workshop, Brown University, New York University, Free University of Bolzano, Northwestern Kellogg, the 8th Economics of Media Bias Workshop, Benelux PECO, RFBerlin, University of Potsdam, POLECONUK Webinar, Charles University Faculty of Law, Prague University of Economics and Statistics, ifo Institute, the CESifo Political Economy Workshop, and the University of Warwick.

1 Introduction

Social media has created an age of attention: a time in which influence, value, and success across domains increasingly hinge on the ability to capture and sustain public focus. What makes social media unique is that, in contrast to earlier broadcast media technologies such as newspapers, radio, and TV, social media is massively multi-channel due to its low barriers to entry. Any user can set up an information feed and popularize it. As a result, massive amounts of information are produced and shared on social media on a daily basis. This has made attention an increasingly scarce, precious, and contested resource. In this new environment, whether in media, business or entertainment, those who can cut through the noise hold a distinctive advantage. Politics is no exception.[1]

In this paper, we characterize the market for attention on social media by studying political communication on the Twitter platform (now known as 'X'). More precisely, we examine whether attention on Twitter—measured through user engagement—translates into real-world economic returns, specifically in the form of campaign contributions to Members of Congress (MOCs) in the United States. We focus on representatives who served in the 116th Congress, which includes the 2-year period leading up to the 2020 election when a record $4 billion was donated to congressional candidates (Federal Election Commission, 2021*b*). Our main finding is that attention on Twitter generates tangible returns by increasing donations, but that the relationship is driven by viral days of extreme popularity. This highlights that the returns to attention on Twitter are highly skewed, in line with Twitter being a superstar market.

Technically, the modeling problem that we face to identify returns to Twitter attention is analogous to the 'lift' problem in advertising, that is, attributing causal links between persuasive messages transmitted to targeted audiences and their resulting financial decisions, in this case the propensity to donate (Aral, 2021). We address this challenge by using different levels of aggregation and types of variation to establish the pattern of linkages between attention on Twitter and political donations. Our analysis also shows how this link is related to the networked, endogenous feedback mechanisms

[1] As Ezra Klein (New York Times) observes: *"Attention, not money, is now the fuel of American politics."*

intrinsic to Twitter, specifically the skewed or viral nature of attention on the platform.

The first empirical approach that we adopt is based on a daily panel of sitting MOCs where we relate the amount of daily small donations below \$1000 received by an MOC to contemporaneous attention on Twitter, that we proxy using likes. We document a positive impact of attention on donations. The relationship is highly robust and has a modest magnitude when considered on a within-person and month basis. However, if we decompose this relationship according to a step function based on different thresholds of attention, we see that it is driven by activity in the very top of the likes distribution. Specifically, we estimate that appearing in the top 10% of the likes distribution (our practical definition of going viral) is associated with a 0.6 percentage point increase in the probability of receiving a donation and approximately a 5.3% increase in the amount of donations at the intensive margin. But this effect dissipates quickly—there is no significant association below the 80th percentile of likes.

This shows that Twitter is a technology that can be effective for raising donations, but not for everyone. The nature of the platform, with its emphasis on engagement and visibility, tends to amplify the popularity of certain posts over others and it is only these viral messages that seem to translate into concrete financial returns. Through a simple back-of-the-envelope calculation, we show that the top 5 MOCs in terms of Twitter attention receive between 10% and 25% higher donations through the electoral cycle than MOCs outside the top 50. We find limited evidence that returns to attention provide direct incentives to adopt ideologically polar positions on average, although viral winners themselves are more likely be extreme. Especially if other returns to Twitter are similarly concentrated, these tournament-style dynamics might still incentivize MOCs to take up more extreme positions.

Also in line with the market for Twitter attention being ‘winner-takes-all’ in nature, we document a crowding out effect of viral days on attention. When looking at two MOCs sharing 20% of their followers, the probability of one of them having a viral day decreases by about 1 percentage point (10% of the baseline mean) if the other MOC is currently going viral. This is not driven by similar ideology or geographic proximity. However, conditional on still going viral, returns to attention in terms of campaign donations are not impacted by other MOCs' virality.

We also exploit our MOC-level daily panel to test the robustness of our analysis. First, we study the timing of the relationship using a specification including leads and lags of viral days. This indicates that the donations-likes relationships is a temporally concentrated one, unfolding mainly over three days. We also show limited impacts of lead terms, which would be an indicator of systematic reverse causality from donations to likes, for example, cases where donations were used to fund offline campaign activity (e.g., rallies) or online advertising which then generated likes. Second, we look at other attention shocks that could be correlated with attention on Twitter and donations: coverage on traditional media and MOCs' in-real-life (IRL) activities. The effect of viral days on donations is barely affected by the inclusion of these controls. Interestingly, the magnitude of the effect of going viral on Twitter and of being mentioned on cable news or by the New York Times are remarkably similar. This suggests that the Twitter effect is not conflated with the news cycle and that media markets might be segmented, with different attention shocks reaching different audiences.

The second modeling strategy that we adopt is based on a county-by-MOC-by-period panel. This allows us to implement a geography-based design that tests whether the increase in donations that MOCs experience when they gain more attention on Twitter is indeed driven by countries with higher cross-sectional levels of Twitter usage. This has the additional advantage of allowing us to include a much richer set of fixed effects (namely, MOC-by-period fixed effects), that take into account many of the concerns highlighted thus far. For identification, following Müller and Schwarz (2023) and Fujiwara et al. (2024), we use the shock associated with the 2007 South by Southwest (SXSW) festival as an instrumental variable (IV) for Twitter usage across areas. The 2007 SXSW festival was a critical diffusion event in the history of Twitter that shaped the geographical pattern of the network's usage rates in long run. Both the OLS and the IV show that in periods in which MOCs receive more attention on Twitter, they experience more donations from counties with higher Twitter usage. This confirms that the increase in donations that MOCs experience when they receive attention on Twitter is indeed the result of a platform-specific channel.

This paper contributes to three main literatures. First and foremost, we contribute to the growing literature studying the effects of social media on political outcomes, including political participation (Bond et al., 2012), voting (Fujiwara et al., 2024; De Luca et al., 2021; Battiston et al., 2025), polarization (Allcott et al., 2020), mobilization (Enikolopov et al., 2020), and, more recently, politicians' behavior (Bessone et al., 2022; Schöll et al., 2023). While we only mention a handful of examples here, we refer to Zhuravskaya et al. (2020) and Lorenz-Spreen et al. (2023) for two recent systematic reviews of work in this area. In this paper, we focus on a different political behavior: campaign donations. Closest to our work is Petrova et al. (2021), who show that when politicians running for Congress in the United States in the 2009-2014 period open a Twitter account, they experience a significant increase in the campaign contributions they receive. The goal of this paper is different. We aim to characterize not only *if* but also *how* attention on Twitter affects campaign contributions at a more mature stage of the market: when all politicians have a Twitter account with a large number of followers and the overall user base has reached a plateau.[2] In such a market, as we document, it is not being on the platform per se that matters, but how much attention one is able to get. In line with this, we show that who benefits on the platform is different at different stages. While new entrants are particularly able to benefit in the earlier period (Petrova et al., 2021), as Twitter emerges as a novel communication channel, by the end of the 2010s only the people who go viral very regularly are able to harness sufficient attention to make a difference in terms of donations. In fact, Twitter displays a level of concentration in line with that cable news or the New York Times. However, because individuals who are successful on Twitter are different than those who are successful on traditional media, Twitter might still broaden the playing field. Related to our work is also Rotesi (2023), who looks at the effect of Twitter penetration on political participation and donations around US presidential elections in 2008, 2012, 2016. In addition to focusing on donations given by citizens rather than donations received by politicians, this paper also looks at an earlier period when adoption was still growing.

In addition, we contribute to the literature in economics and political science on the determinants of US campaign contributions and their effect on electoral outcomes (Gerber, 2004; Snyder Jr,

[2]The number of monthly active Twitter users started plateauing around 66 to 69 millions in 2016 (Statista, 2023).

1990; Fouirnaies and Hall, 2014; Fouirnaies and Fowler, 2022, see Dawood (2015) for an extensive review). While earlier work in this area focused on large donors (Heerwig, 2016; Rhodes et al., 2018; Teso, 2022) researchers have recently shifted their attention to small donors, who are becoming more and more relevant in terms of number and volume of donations, but might display different behavior and motives (Alvarez et al., 2020; Bouton, Cagé, Dewitte and Pons, 2024; Bouton, Castanheira and Drazen, 2024; Culberson et al., 2019)). Although the focus on our paper is on social media, our results provide interesting evidence as to why individuals engage in small donations. The fact that short-term shocks to attention on social media are sufficient to impact donations is highly suggestive of expressive (rather than strategic) motives being behind the decision to send a monetary contribution to a campaign, but also that pull factors matter for donation decisions, in line with what Bouton, Cagé, Dewitte and Pons (2024) also find.[3]

Finally, we contribute to the literature on superstar markets, as initiated by Rosen (1981) and explored in contributions such as Célérier and Vallée (2019), Gabaix and Landier (2008), Koenig (2023), Krueger (2005) and Terviö (2009). The prominence of network-based sharing in the design of the Twitter platform fits the characteristics of scale-related technical change (Koenig, 2023); and the skewed pattern of attention and resulting contributions that we observe in our data clearly reflects the structure of a typical superstar market setting. Adding even further, the winner-takes-all nature of the platform is also evidenced by our crowding out results.

The remainder of the paper proceed as follows. Section 2 describes the data. Section 3 provides information about the background, and descriptive characterization of attention on Twitter. Section 4 reports our empirical approach and our main results, while Section 5 describes the analysis that also exploits geographic variation. We conclude in Section 6.

[3] Bouton, Cagé, Dewitte and Pons (2024) first find that the number of small donors and their total contributions have been growing over the period 2005-2020. Second, they also show who are these small donors and when they are more likely to donate. In particular, they find that pull factors such as TV and Facebook ads affect their behavior.

2 Data

2.1 Data Sources

We combine data from various sources described below. For our first analysis we build a daily panel that combines MOCs' activity and popularity on Twitter with detailed data on political donations and additional measures of coverage by traditional media. For our second analysis, we further distinguish donations by geography, adding data on donors' location reported by the FEC and local Twitter penetration obtained from Fujiwara et al. (2024).

Twitter Data. Our main dataset is based on the social media messages (tweets) of sitting MOCs during the 116th Congress between January 2019 and October 2020. We collect the tweets of 533 MOCs across their personal, congressional, and campaign accounts. After excluding MOCs for whom we do not have the complete Twitter activity for the period of interest, our final database comprises over 1.1 million Congressional tweets from 501 different MOCs.[4] In addition to the text of the tweets, we collect the following standard metrics: likes (click-based expressions of approval), retweets (forwarding of messages to individual's followers), and replies (messages sent in response to a tweet). We also collect the Twitter IDs of each MOC's followers to compute the overlap of followers across different MOCs.[5]

FEC Contributions. Our data on campaign contributions comes from the public FEC database on donations by individuals (Federal Election Commission, 2021*a*). We include all direct and indirect donations going to a candidate's committee. Specifically, we focus on donations from individuals who contribute more than \$200 over the electoral cycle. Additional information includes the donation amount, date of receipt, the sender's location (ZIP code), and whether the donation went directly to a candidate's campaign or came through a conduit.[6] Following Petrova et al. (2021), we

[4] Specifically, since our data collection and analysis were conducted in early 2021, we lack information on a small number of MOCs who exited Congress after the November 2020 election and retired their accounts or were otherwise suspended from the platform. In addition, we drop 25 individuals who did not record individual contributions during the study period, 2 non-voting delegates, and 3 MOCs who did not serve a full term.

[5] Due to a technical glitch in the data collection, we do not have the followers of Representative Seth Moulton.

[6] The FEC's website defines conduits as follows: "Anyone who receives and forwards an earmarked contribution to

focus on small donations, which we define as donations smaller than $1000. This is motivated by large donors probably having different motives and being unlikely to be affected by social media, but we show robustness of our results to other definitions in Appendix Section C. Overall, small donations under $1000 account for 37% of contributions from individuals, and 94% of donors in our data.

Other Shocks to Attention. To measure attention on traditional media, we collect daily mentions of MOCs from cable news, the New York Times, and local newspapers. For cable news, we obtain measures of coverage using the transcripts of cable news programs of three major cable channels (CNN, FNC, MSNBC) available from the Internet Archive. For each MOC, we count the number of mentions of their full name ("Nancy Pelosi") or their title and last name ("Representative Pelosi") and define our measure of cable news coverage as the sum across the three channels. We use article-level metadata from the New York Times' API to measure MOCs mentions in the newspaper. We define an article to be about an MOC if it includes the name and surname of the MOC as one of its keywords, in the persons category. Finally, local newspaper mentions of MOCs come from newslibrary.com. To create a measure of MOC coverage, we search for the full name of the MOC in all headlines and lead first paragraphs of the local newspapers included in the library.

We also collect information on MOCs' IRL activities. First, we measure MOCs' activity in Congress using information on Congressional speeches from Bellodi (2023). Second, we collect information on MOCs' trips to their home state from the Clerk of the House of Representatives.[7] Finally, we retrieve all Presidential and Vice-Presidential visits to the MOCs' home state from Facebook and Wikipedia.[8]

Local Twitter Usage. We obtain data on Twitter usage and early Twitter adoption from Fujiwara et al. (2024). Twitter users are measured by the number of unique users tweeting from

a candidate or a candidate's committee is considered a conduit or intermediary." Conduits include online fundraising platforms such as ActBlue and WinRed.

[7]This information is available from https://disclosures-clerk.house.gov/GiftTravelFilings.

[8]The intuition for doing this is that not only a visit from the President and Vice President is likely to generate attention by itself, but also quite likely for the Representatives for the state (who, for example, is likely to attend it or, at a minimum, comment on it).

a given location in a 75% sample of all geo-tagged tweets from the US between 2014 and 2015 (Kinder-Kurlanda et al., 2017). To measure early adoption of Twitter, we use the number of local users following the SXSW Twitter account (@SXSW) who joined in March 2007, in the wake of the 2007 SXSW festival. Furthermore, we use data on SXSW followers who joined before 2007 as an additional control. For our main analysis, we aggregate the data at the county level. Additional details, validation, and discussion can be found in Müller and Schwarz (2023) and Fujiwara et al. (2024).

Other Data. We obtain information on MOCs' demographics and tenure from the @unitedstates github repository and the Library of Congress. To measure MOCs' ideology, we use the first dimension of the voteview DW-nominate score (Lewis et al., 2022), which captures the traditional liberal-conservative spectrum in economic matters. County-level characteristics are from Fujiwara et al. (2024). We also use data from the 2020 wave of the American National Election Studies (American National Election Studies, 2021), a representative survey of eligible voters.

2.2 Descriptive Statistics

Table 1 reports summary statistics for our main dataset. It covers 501 MOCs over 670 days from January 1st, 2019 to October 31st, 2020. MOCs receive on average just under 3000 likes per day across all of their accounts, although these numbers are very skewed: the median number of daily likes is 31 and only 9% of MOC-days yield more than 2000 likes. MOCs' accounts send on average almost 3 tweets per day and receive over 320 replies and 710 retweets.

Other attention shocks are similarly skewed as Twitter popularity. On an average day only 11% of the MOCs are mentioned on cable news, 0.1% in the New York Times, and 0.6% in local newspapers. There is some overlap between Twitter and other attention shocks, but it is by no means complete. For example, out of those MOCs mentioned by cable news, only 37% also receive more than 2000 likes on Twitter in the same day.

Contributions also follow a skewed pattern and are very sparse. MOCs in our sample receive over $1200 per day in small donations from 15 distinct donors but no donations at the median.

Table 1: Descriptive Statistics

	Mean	SD	Median
Twitter			
Likes	2990.811	24843.474	31.000
Likes > 2000	0.092	0.289	0.000
Replies	319.470	2265.853	6.000
Retweets	709.071	5017.033	9.000
Tweets	2.931	4.135	2.000
Other Shocks to Attention			
Cable news mentions > 0	0.109	0.312	0.000
Mentioned on cable news & Top 10% Twitter	0.043	0.203	0.000
Mentioned in the NYT	0.011	0.102	0.000
Mentioned in the NYT & Top 10% Twitter	0.008	0.089	0.000
Mentioned in local newspapers	0.064	0.245	0.000
Mentioned in local newspapers & Top 10% Twitter	0.023	0.151	0.000
Donations			
All donations	3273.546	22723.499	0.000
Small donations	1219.040	11788.448	0.000
Small donations, if donations > 0	2725.097	17508.605	525.000
All donors	15.946	186.223	0.000
Small donors	15.026	183.437	0.000
Small donors, if donors > 0	33.590	273.126	4.000

This table presents summary statistics on daily Twitter activity, coverage on cable news, in the New York Times, and in local newspapers, and campaign donations received for the 501 MOCs part of our sample. The dataset is an MOC-by-date panel.

Conditional on receiving at least one donation, the average amount of small donations received is \$2725 per day from 33 donors. Note that our definition of small donors excludes disproportionately large donations. Including large donors adds on average less than one additional donor (6%) but increases the average donation by over \$2000 (170%) relative to the small-donation sample.

3 Background

3.1 Overlap Between Twitter Usage and the Propensity to Donate

The potential influence of Twitter on political donations is naturally a function of the size and composition of its reach. Commentators such as Klein (2022) characterize Twitter during this period as the primary social network for elite professionals across knowledge-based industries.[9] Indeed, using data from the 2020 wave of ANES (a representative survey of eligible voters), we find that Twitter users have higher incomes, are more urban, are more educated, and are more likely to identify with the Democratic party than the population average (Appendix Table A1).[10] Most importantly in our setting, Twitter users are more likely to donate to political causes, in particular to individual candidates; occasional Twitter users are 13% more likely and daily Twitter users are 57% more likely to have donated to a candidate in 2020 than the population average. This effect is not driven by other demographics: a simple linear probability model predicts that occasional Twitter users are 32% more likely and daily Twitter users are 74% more likely to have donated to a candidate than the sample mean, holding income, education, location, and party affiliation fixed.[11] This data indicates that the Twitter network provides politicians a direct channel to a very active, politically engaged audience.

3.2 The Characteristics of Twitter Attention and Political Donations

Twitter is in theory very accessible and the first barrier to entry, the creation of an account, seems low. However, obtaining and maintaining a large number of active followers is challenging because Twitter, like many other (online) social networks, exhibits high degrees of concentration, positive degree assortativity, and ideological homophily (Mislove et al., 2007; Antonakaki et al., 2021; Zhuravskaya et al., 2020). As shown in Appendix Figure A1, the distribution of likes on Twitter between MOCs is very concentrated, and the degree of concentration is comparable to traditional cable news and of the New York Times: the respectively 25 (50) most popular MOCs receive 77% (89%) of likes on Twitter, 78% (88%) of mentions on cable news, and 75% (86%) of New

[9] As Klein (2022) notes: 'Twitter might have a smaller user base than Facebook, Instagram and even Snapchat, but it shapes the dominant narratives in key industries like politics, media, finance and technology more than any other platform. Attention–particularly that of elite leaders in these industries–is a valuable resource, one that Twitter manages and trades in.'

[10] This is in line with findings from a Pew Research Center survey (Wojcik and Hughes (2019)).

[11] We run an OLS regression of an indicator equal to one if an individual contributed to a candidate on indicator variables for occasional and daily Twitter use and controls for income, education, location, political leaning (details on the variables are in the notes of Appendix Table A1). The estimated coefficients of those indicators are 0.051 and 0.118 (both significant at $p<0.001$ with robust SEs); the omitted category is no reported Twitter use. The sample mean of the outcome variable is 0.1597.

York Times mentions. Mentions on local newspapers are much less concentrated, underlying the importance of local news in providing coverage of less popular MOCs. There is some overlap in Twitter and other media popularity, but it is not extreme: out of the top 50 MOCs by Twitter likes, 56% are also in the top 50 by cable news mentions, 53% in the top 50 by New York Times mentions, and 37% by local newspapers mentions. This incomplete overlap can also be seen in Appendix Table A2, that reports the top 10 MOCs by attention channel.

3.3 Who and What Goes Viral?

Twitter is a platform where information spreads rapidly and widely. Investigating what goes viral is therefore helpful for understanding the mechanisms and patterns of information dissemination. This can be examined in terms of characteristics of the sender as well as the content of tweets. In this section, we establish patterns regarding the drivers of Twitter attention. This is not only interesting per se, but will also assist us in studying the heterogeneity of our results across these characteristics.

We begin by considering *who* goes viral, by estimating the correlation between being in the top 10% of the likes distribution and different MOC characteristics. Appendix Figure A2 panel (a) reveals several interesting patterns. First, looking at ideology, we can see that MOCs from the Democratic party are more likely to go viral. This is unsurprising given the political leaning of the platform's user base. Regardless of party, MOCs at the extremes of the ideological distribution (defined as being in the bottom and top decile in the distribution of the DW nominate score) also get more attention on the platform. Second, female MOCs do not appear to receive differential attention of the platform, while younger MOCs do. Third, MOCs serving in the Senate receive more attention on Twitter but that having a longer tenure has no effect.[12] Finally, as we would mechanically expect, MOCs with more followers are significantly more likely to go viral. Appendix Figure A2 panel (b), (c) and (d) show that there is some heterogeneity in what characteristics predict attention on

[12]That MOCs serving in the Senate are more likely to go viral might be surprising. However, they also tend to have a higher number of followers, which might explain the pattern we see here. In line with this, running a horse race between the two characteristics cuts the effect of being in the Senate by three and leaves it not statistically significant at conventional levels.

different traditional media platforms, which suggests that different MOCs might be successful in different attention markets.[13]

We then focus on *what* goes viral by looking at the content of the tweets. Appendix Figure A3 shows the share of tweets that contain a specific type of content, separately by whether the tweet is in the top 10% of the likes distribution or not. First, we observe that more localized tweets are slightly less likely to be popular. Second, certain types of content resonate more with people, leading to more sharing and engagement. Using a dictionary-based method (Valence Aware Dictionary for Sentiment Reasoning or VADER) and Google's Perspective API, we show that tweets that go viral are 60% more likely to display negative sentiment and three times more likely to be toxic than non-viral tweets.[14] While these results are not surprising to those familiar with the Twitter platform we think they are notable in terms of a) the magnitude of the implied relationship; and b) the potential to create perverse incentives when factored into the capacity to attract donations. Finally, MOCs rarely use Twitter to directly solicit donations: less than 0.5% of tweets include a direct link to a donation platform. This suggests that it is unlikely that our effects will be explained by strategic soliciting behavior on part of MOCs.

4 Daily Panel

4.1 Modeling Framework

The first part of our analysis follows a panel structure with daily variation in our variables of interest. In particular, we estimate the following baseline specification:

$$y_{it} = \beta * likes_{it} + \alpha_{im(t)} + \tau_{p(i)t} + \epsilon_{it}, \tag{1}$$

[13]Specifically, MOCs who are Senate members receive disproportionate attention across all platforms, MOCs with extreme ideologies are only overrepresented on Twitter and on social media. Similarly, we only see a Democratic advantage on Twitter.

[14]We define a tweet be toxic if it has aa toxicity score in the top 5% of the distribution, or higher than 0.22. To put this in perspective, toxicity scores higher than 0.30 are generally considered to be suspect and higher than 0.80 are considered to be hate speech. Perhaps reassuringly, only 174 tweets (our of more than 1.1 millions) have a toxicity score higher than 0.80. As a comparison, 5.6% of tweets were found to meet the threshold in a random sample of English language tweets (Jiménez Durán, 2023).

where y_{it} is the log + 1 of the total amount of donations below $1000 received by MOC i on day t; $likes_{it}$ is the log + 1 of the total Twitter likes received by MOC i on day t; $\alpha_{im(t)}$ are MOC-by-month fixed effects; $\tau_{p(i)t}$ are date-by-party fixed effects. Standard errors are clustered at the MOC level.

The strength of this specification is that it offers a within-MOC analysis. That is, we control for heterogeneity in both the amount of small donations and the amount of Twitter attention that an MOC receives on average. The inclusion of more restrictive MOC-by-month fixed effects, $\alpha_{im(t)}$, ensures that the within-person comparison occurs in sub-periods where the level of donations is comparable. This is important in our setting as donation levels vary widely over time, and different MOCs might be exposed to different donation cycles. The date-by-party fixed effects, $\tau_{p(i)t}$, allow MOCs belonging to different parties to be on different non-parametric trends.

The first innovation that we introduce to this specification is to divide the distribution of likes according to different thresholds and estimate a step function. This allows us to measure instances of viral attention, which we define as cases where an MOC-day observation is at or above the 90th percentile in the overall likes distribution. We also use this discretized approach more generally to deal with the many zeroes issue, for example, by formulating extensive margin specifications based on binary variables that flag when donations or attention are positive valued or cross particular thresholds.

The second innovation allows us to investigate the dynamic relationship between virality on Twitter and donations. To do so, we estimate a dynamic version of the baseline equation including leads and lags of our main independent variable, which allows us to track the pattern of donations before and after a viral day. In particular, we estimate the following specification:

$$y_{it} = \sum_{k=1}^{10} \delta_{-k} * \text{Top 10\% Twitter}_{it-k} + \sum_{k=0}^{10} \delta_{+k} * \text{Top 10\% Twitter}_{it+k} + \alpha_{im(t)} + \tau_{p(i)t} + \epsilon_{it}, \quad (2)$$

where Top 10% Twitter$_{it}$ is an indicator for the number of likes that MOC i receives on day t being in the top 10% of the overall likes distribution and all other variables are defined as before. The δ_{-k} and δ_{+k} parameters measure any systematic movements in outcome y_{it} before, on, and after a viral day.

Table 2: Twitter Likes and Small Donations

	Small donations						
	Log + 1					Dummy	Log
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log(Likes + 1)	0.352***	0.039***	0.042***	0.011***			
	(0.024)	(0.006)	(0.006)	(0.003)			
Top 10% Twitter					0.094***	0.006*	0.053***
					(0.019)	(0.003)	(0.014)
Date FE	X	X					
MOC FE		X	X				
Date-by-party FE			X	X	X	X	X
MOC-by-month FE				X	X	X	X
Observations	335670	335670	335670	335670	335670	335670	149200
MOCs (clusters)	501	501	501	501	501	501	496
Mean dep. variable	2.799	2.799	2.799	2.799	2.799	0.447	6.247

This table shows the relationship between Twitter likes and small donations. In columns (1) to (4), we regress the log + 1 of the total amount of donations below $1000 that an MOC receives on a given day on the log + 1 of the total Twitter likes the same MOC receives on the same day. Column (1) estimates a pooled specification that includes date fixed effects only, column (2) adds MOC fixed effects, column (3) includes date-by-party fixed effects, and column (4) includes both date-by-party and MOC-by-month fixed effects (equation (1)). Column (5) estimates the same specification as column (4) but using as the independent variable an indicator for the number of likes that the MOC receives on the same day being in the top 10% of the likes distribution (1681 likes or more). While using the same specification and the same independent variable, column (6) defines the outcome as an indicator variable equal to one if an MOC receives positive donations and column (7) as the log of the total amount of donations below $1000 that an MOC receives on a given day. In column (7), the sample is conditional on receiving positive donations. Standard errors are clustered at the MOC level.

4.2 Daily Panel Results

Table 2 presents our results on the relationship between small donations and likes. Column (1) estimates a specification that only includes date fixed effects. There is a strong and sizable correlation between the amount of donations received by an MOC and the MOC's popularity on Twitter. This seems to be driven by variation across MOCs: including MOC fixed effects cuts the effect by a factor of ten (column (2)). Including date-by-party fixed effects does not substantially impact the point estimate (column (3)). We report our preferred specification including month-by-MOC fixed effects in column (4). This highly restrictive specification yields a coefficient of 0.01.

The use of a log + 1 transformation complicates the interpretation of the magnitude of our estimates, as the recovered coefficients cannot be interpreted as elasticities (see Chen and Roth, 2024;

Mullahy and Norton, 2024). To make progress on this, we begin by discretizing the independent variable, and estimate the effect of being in the top 10% of the likes distribution. Column (5) shows that being in the top 10% of the likes distribution increases donations by 0.094 log + 1 points. We then look at the extensive and intensive margin separately. We find that Twitter attention affects both margins. Column (6) and (7) show that being in the top 10% of the likes distribution increases the probability of receiving at least one donation by 0.6 percentage points (1.35% of the baseline mean) and, conditional on receiving donations, it increases the amount received by 5.3%.

Viral Impacts. Discretizing the dependent variable helps us understand magnitudes, but it can also be substantively relevant if returns are skewed. We assess the gradient of the basic donations-likes relationship using a step function approach and report the relationship between likes on Twitter and small donations for different percentiles of the likes distribution, our measure of virality, in Figure 1. There is no relationship between small donations and likes in the bottom 80% of the likes distribution. However, days with a number of likes in the top 20% of popularity are associated with increased donations and, within this group, the returns to popularity are high. Being in the 95-99th and the 99-99.9th percentile group results in coefficients in the 0.20-0.30 range. These results show how Twitter can be considered a winner-takes-all market. The nature of the platform, with its emphasis on engagement, visibility, and virality, tends to amplify the popularity of certain tweets over others and this in turn seems to translate into concrete financial returns.

Importantly, increasing returns for viral attention is not mechanically explained by viral tweets receiving more likes: we show in Appendix Figure A4 that the coefficient of small donations to the number of likes is almost double for tweets in the highest percentiles of attention. This is in line with the fractal structure of returns found in other empirical studies of superstar markets such as Koenig (2023).

Dynamics. Figure 2 illustrates the dynamic effect by estimating the lag structure for a viral day across 10 leads and lags. This can be interpreted as a test for systematic reverse causality with, for example, donations underwriting activity which would then generates likes. The forward terms are neither individually nor jointly significant in this model. Specifically, we can reject a joint test

Figure 1: Twitter Likes and Small Donations, by Virality

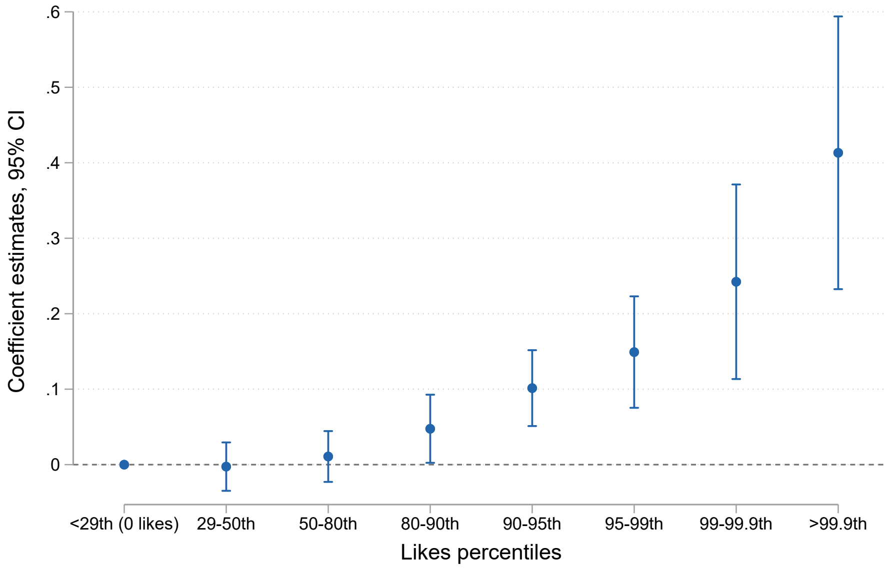

This figure shows the relationship between Twitter likes and small donations, for different levels of virality (i.e., for different percentiles of the likes distribution). In particular, we regress the log + 1 of the total amount of donations below $1000 that an MOC receives on a given day on a series of indicators for the number of likes that the MOC receives on the same day being in different percentiles of the likes distribution, date-by-party fixed effects, and MOC-by-month fixed effects (similar to equation (1)). The omitted category $< 29th$ includes days in the 29th percentile and below (corresponding to days in which the MOCs receive 0 likes); $29 - 50th$ indicates days in the 51st to 80th percentile (between 32 and 350 likes); $50 - 80th$ indicates days in the 51st to 80th percentile (between 32 and 347 likes); $80 - 90th$ indicates days in the 81st to 90th percentile (between 348 and 1657 likes); $90 - 95th$ indicates days in the 91st to 95th percentile (between 1658 and 7,696 likes); $95 - 99th$ indicates days in the 95th to 99th percentile (between 7,698 and 68,138 likes); $99 - 99.9th$ indicates days in the 100th percentile, excluding the top 0.1% of the distribution (between 68,139 and 320,217 likes); $> 99.9th$ indicates the top 0.1% of the distribution (between 321,134 and 3,808,126 likes). Standard errors are clustered at the MOC level.

that 3 days of forward terms are significant with an F-statistic of 0.79 (p-value $= 0.50$). We also find a significant difference between the forward and contemporaneous effect of virality (p-value $=$ 0.015).

Figure 2 also provides information on the timing of the effect. The impact of a viral day is relatively short-lived. After an initial spike, the effect tends to become smaller and fades after three days. Since the leads and lags are included simultaneously, their coefficients can be interpreted as partial effects at k days after an event and we can sum them together to get the total dynamic impact of a viral day shock. Based on Figure 2, there is an overall 0.24 log + 1 points increase in donations over the 3 days following a viral day, which is substantially larger than the contemporaneous effect. The extensive and intensive margin of donations follow similar dynamics (see Appendix Figure A5).

Figure 2: Twitter Likes and Small Donations, Leads and Lags

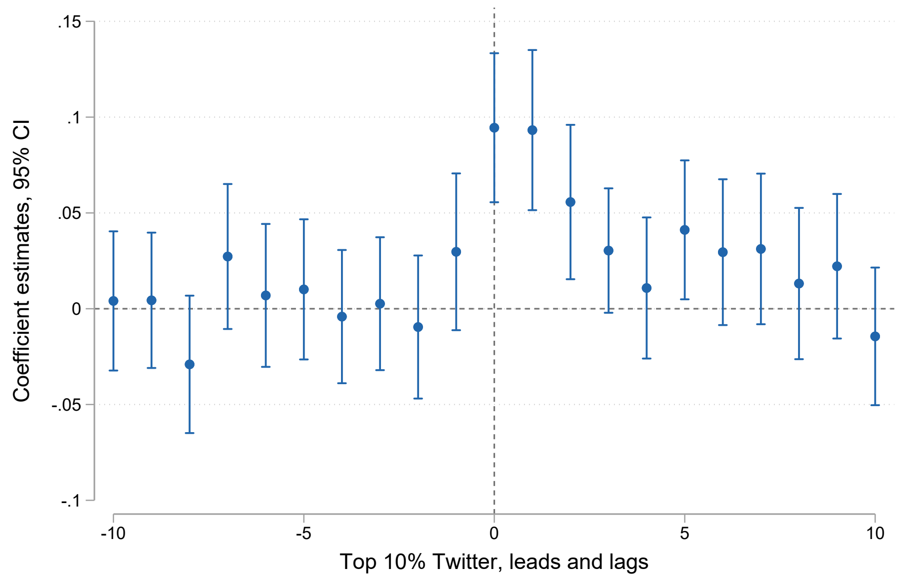

This figure shows the dynamic relationship between likes on Twitter and small donations. In particular, we regress the log + 1 of the total amount of donations below \$1000 that an MOC receives on a given day on ten leads and lags of an indicator for the number of likes that the MOC receives on the same day being in the top 10% of the likes distribution (1681 likes or more), date-by-party fixed effects, and MOC-by-month fixed effects (equation (2)). Standard errors are clustered at the MOC level.

Robustness Checks. We discuss the robustness of our main results to different transformations of Twitter attention and small donations and to different functional forms in Appendix C.

4.3 Other Shocks to Attention

A potential confounding factor in our estimation are events that drive attention towards specific MOCs independently of Twitter, such as coverage on traditional media and real-life events. This attention could direct more small political donations to their campaign, but might also drive Twitter likes, as people start paying more attention to MOCs on social media. To understand the relationship between these other attention shocks and campaign donations, we estimate a series of regressions similar to equation 1 but substitute $likes_{it}$ with the relevant attention measure. We also check that the donations-likes relationship is robust to including all of these controls by performing a 'horse race' including all variables simultaneously.

We present the results of this analysis in Figure 3. Panel (a), which reports estimates from the univariate regressions, shows that traditional media coverage does indeed affect donations. Being

Figure 3: Effect of Other Attention Shocks on Donations

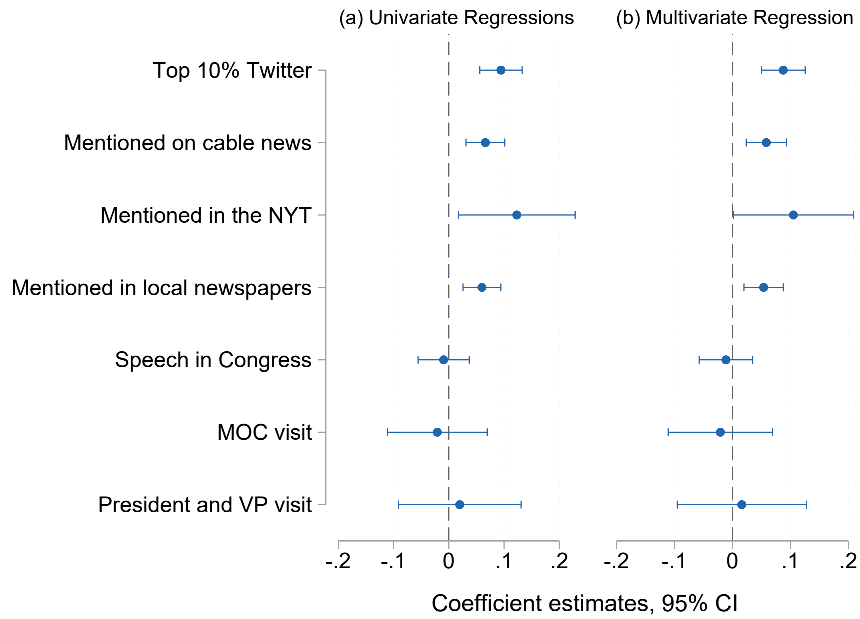

This figure shows the relationship between different attention shocks and small donations. Panel (a) reports estimates from a series of univariate regressions in which we regress the log + 1 of the total amount of campaign donations below $1000 that an MOC receives on a given day on a proxy for the attention shock, date-by-party fixed effects, and MOC-by-month fixed effects (similar to equation (1)). The proxies for the attention shocks that we use are: an indicator for the number of likes that the MOC receives on the same day being in the top 10% of the likes distribution (1681 likes or more); an indicator for MOC being mentioned on cable news; an indicator for MOC being mentioned in the New York Times; an indicator for the MOC being mentioned in local newspapers; an indicator for the MOC delivering a Congressional speech; an indicator for the MOC visiting his/her home state; and an indicator for the President or the Vice-President visiting the home state of the MOC. In panel (b), we perform a "horse race" by including all the attention shocks in the same regression. In all regressions, the sample excludes one MOC not included in the Congressional speeches dataset. Standard errors are clustered the MOC level.

mentioned on cable news, in the New York Times, or in local newspapers positively affects donations. Interestingly, the returns from going viral on Twitter and being mentioned on traditional media are comparable in magnitude. Instead, IRL activities such as giving a speech in Congress or visits to the MOC's home state, whether by the MOCs themselves or the President/Vice-President, do not have an effect. Panel (b) shows, however, that the effect of going viral on Twitter goes beyond the news cycle. When controlling for all these extra shocks to attention, our main estimate remains virtually unchanged. The effect of different types of traditional media coverage are also similar in the 'horse race' regression. This suggests that there might be segmentation in media markets, so that shocks to attention on different platforms are likely to affect distinct audiences.

4.4 Discussion

Harvesting. It is possible that what viral tweets do is move donations across time: in other words, it is possible that what we are estimating is harvesting across the electoral cycle, rather than an overall increase in donations. To test for intertemporal harvesting in the short run, we re-estimate our baseline specification but using as the outcome donations measured over different time periods: the day of the tweet, the week after the tweet (excluding the day of the tweet itself), and the month after the tweet (excluding the first week). Appendix Table A4 shows limited evidence of harvesting, at least in the short-run. Attention on Twitter has a positive and significant effect on contemporaneous donations (columns (1) and (4)). In line with our event study, attention on Twitter also has a positive effect within one week (columns (2) and (5)). However, Twitter attention has no offsetting negative effect on donations after that point (i.e., after one week but within one month, columns (3) and (6)).

Crowding-out of Attention. The concentration in attention can in part be explained by a crowding out effect of viral days. In Appendix Table A5, we illustrate how MOCs' likelihood of going viral depends on virality around them. To do so, we regress an MOC's probability of going viral on the number of other politicians going viral on the same day, for different ways of proxying for the MOCs' network.

Users' attention on the platform appears to be limited. Independently on whether we measure an MOCs' network using Twitter followers' overlap (column (1)), ideology (column (2)), or geography (column (3)), we find that the number of other MOCs going viral is positively associated with a lower probability of the MOC going viral himself/herself. The spillover effects along the Twitter network are larger in magnitude and not mechanically explained by our relative definition of viral days, as the absolute number of likes significantly decreases as well (column (5)).[15] Instead, the spillover effects based on ideology and geography are smaller in magnitude and less robust (see again column (5)).

[15] Specifically, an MOC's probability of having a viral day decreases by 1 percentage point if a different MOC who is followed by 20% of the MOC's followers goes viral, or 5.2 percentage points if 100% of their followers also follow the other MOC.

Finally, we check in column (6) whether viral days have further negative effects on donors' propensity to contribute on top of the negative impact on attention. Conditional on still going viral, donations do not further decrease. As a result, the negative spillover effects on donations are just proportional to the effects on attention.

Returns in a Concentrated Market. Is the magnitude of the returns we estimate large enough to be consequential for MOCs' fundraising efforts? We answer this question by performing a simple back-of-the-envelope calculation. First note the extent to which viral attention is concentrated: the five MOCs with the largest number of likes over the sample period go viral 605 days on average, those with the 6th to 50th highest number of likes 242 days on average, and everyone else goes viral 31 days on average. Then consider the size of the effect that we model. A viral day increases donations by 5.3% on the intensive margin if we only consider the contemporaneous effect of going viral and by 12% if we compute the short-term cumulative effect over three days (see Table 2 and Appendix Figure A5). At a daily mean of $2725 in small donations, these effects correspond to a marginal effect of respectively $144 (one day) and $327 (three days). In turn, this means that MOCs in positions 1 to 5 earn an additional $83261 to $201688 in small contributions raised over the full sample period from going viral on Twitter. This corresponds to around 10% to 25% higher average donations over the electoral cycle. Importantly, this suggests that Twitter is a technology that can be effective for raising donations, but not for everyone: only very few representatives are able to harness attention on the platform to substantially increase donations from small donors. But for the winners, returns can be large.

Perverse Incentives. In Section 3.2 we showed that there is substantial heterogeneity in *who* and *what* goes viral. To have a complete picture of whether Twitter favors MOCs with certain characteristics or distorts incentives to create specific types of messaging and content, we need to understand whether returns are also skewed. Interestingly, as Appendix Figure A6 and A7 show, we do not find evidence of heterogeneity in the size of effects according to different characteristics.[16]

[16] Note that we can use the analysis of heterogeneous effects by content reported in Appendix Table A7 to address two specific concerns regarding the interpretation of our findings. First, it does not appear to be the case that our findings can be explained by MOCs using Twitter to solicit donations. This is apparent from the fact that this type of content is

This suggests that the main channel through which these content or MOC characteristics impact donations is by changing the probability of going viral. But are these differences substantial enough to create perverse incentives for MOCs?

We elaborate on this point considering returns to Twitter for those with extreme ideology, which is important insofar that is allows us to consider whether Twitter can account for the bias of small donors' towards political extremists (Bouton, Cagé, Dewitte and Pons, 2024). The probability of going viral for MOCs from the extreme ideological deciles (1 or 10) is 80% higher than for MOCs from the central deciles (4-7), meaning they go viral on 46 more days on average.[17] But this higher rate of going viral is not enough to deliver a substantial advantage to the average ideologically extreme candidate. Adapting our earlier simple back-of-the-envelope calculation, this difference in virality only implies an additional $6673 to $16193 over the full sample period.

However, because returns to attention in this market are so concentrated, understanding perverse incentives requires moving beyond the *average*. As can also be seen in Appendix Table A2, Twitter winners are significantly more likely to be ideologically extreme: 3% of extreme MOCs are in the top 10 of Twitter attention, while this share is only 1.75% of non-extreme MOCs. Especially if other returns from Twitter attention are similarly concentrated, these tournament-style dynamics might incentivize MOCs to adopt more extreme positions.

Persuasion Rate. To interpret the magnitude of our estimates, we calculate persuasion rates (DellaVigna and Gentzkow, 2010). The treatment that we consider is exposure to an MOC's viral tweets, while the behavior is making a donation to the same MOC. Following Fujiwara et al. (2024), we note that for marginal changes in exposure the persuasion rate can be approximated as $f = \beta * \frac{y}{e(1-y)}$, where β is the estimate of the (conditional) semi-elasticity of viral days on the log number of donors that we estimate from equation (1), e is the share of the population

extremely rare, as we mentioned in Section 3.2. Second, it also does not appear to be the case that our findings are driven by external events such as rallies or MOCs' activities in their home district, an omitted variable that might drive both attention on Twitter and donations. Not only are local tweets slightly less likely to go viral, but they also have exactly the same returns in terms of campaign donations.

[17] The probability of going viral is 15.3% for the extreme deciles and 8.5% for the central deciles.

who is on Twitter, and y is the share of the population who donates.[18] Setting $\beta = 0.045$, $e = 0.32$, and $y = 0.16$, we estimate that the persuasion rate is approximately 2.6%. This persuasion rate is lower than the average persuasion rates reported in DellaVigna and Gentzkow (2010) but slightly larger than the persuasion rate of an MOC opening a Twitter account on donations estimated by Petrova et al. (2021) and the persuasion rate of political advertising on vote share from Spenkuch and Toniatti (2018).[19]

5 Geography-based design

5.1 Modeling Framework

Our MOC-by-day analysis establishes that the more attention a MOC receives on Twitter, the higher the campaign donations that they receive. While these results are compelling. there is still scope to argue that the donations-likes relationship may be driven by other unobservable factors. For example, the full set of attention shocks related to a specific MOC may not be well proxied by mentions on traditional media. Alternatively, it could be a matter of MOCs' campaign efforts generating donations such that Twitter attention is a secondary by-product of that effort rather than being the actual mechanism that catalyses donations (that is, reverse causality).

In response to these concerns we develop a geography-based design that allows cross-sectional variation in Twitter usage to play a role. This is facilitated by the fact that the FEC data contains the location of donors, which allows us to build a dyadic panel where we observe the donations that a MOC receives from 3,097 different counties located in contiguous US states (as well as the District of Columbia). The objective of this geography-based design is to test whether the effect of Twitter

[18]While the main outcome throughout the paper is the amount of small donations that a candidate receives in a day, persuasion rates are conceptually more appropriate to scale effects on binary behaviors (in this case, donating versus non-donating). This is why we focus on the effect on number of donations in this discussion. To estimate the marginal effect of likes on donors, we estimate equation (1) using the log. number of donors (conditional on receiving at least one donation) as outcome and an indicator for going viral as the independent variable. This gives us a coefficient of 0.043, significant at the 99% level, see Appendix Table A6. This number represents a more accurate estimate of the magnitude, but also a lower bound as it disregards the higher likelihood of receiving donations on high-popularity days.

[19]Note that this rate again represents a lower bound, as it does not account for additional intensive margin effects and no lagged effects on the following days (like in Figure 2).

attention on donations is indeed driven by counties with higher Twitter penetration, as we would expect under the assumption that exposure to a MOC's content on the social media platform is the mechanism that produces our results.

We use our MOC-by-county-by-time panel to estimate the following augmented version of our previous specification:

$$y_{ict} = \beta * (likes_{it} \times users_c) + \gamma * (likes_{it} \times X_c) + \tau_{it} + \theta_{ct} + \delta_{ic} + \varepsilon_{ict}, \quad (3)$$

where y_{ict} is an indicator variable equal to one if MOC i receives any small donations from county c in time period t; $likes_{it}$ is the log + 1 total likes received by MOC i in period t; $users_c$ is the log + 1 number of Twitter users in county c; X_c are county-level controls; τ_{it} are MOC-by-period fixed effects; θ_{ct} are county-by-period fixed effects; and δ_{ic} are MOC-by-county fixed effects. Standard errors are clustered by county. An estimate of $\beta > 0$ in this specification indicates that the change in donations when likes for a given MOC is larger in counties with higher levels of Twitter usage.

Implementing this geography-based design requires several modifications from our baseline approach, driven by the fact that introducing variation across counties exponentially increases the dimensionality of our dataset. The structure of 3,097 counties and 501 MOCs creates over 1.5 million distinct dyads. This is then multiplied by the number of time periods, which can alternatively be defined at the monthly, weekly, or daily level. We concentrate on results from a MOC-by-county-by-month panel as a tractable level of estimation, but present weekly and daily models as robustness checks in the Appendix. In addition, introducing the variation across counties also increases the sparsity of our outcome. [20] To take this into account, we focus on the extensive margin of donations.

Adding this geographical dimension allows us to include a comprehensive set of fixed effects that absorb a substantial number of potential confounding factors into our basic specification. First, we can include MOC-by-period fixed effects (τ_{it}) to absorb all observed and unobserved temporal shocks to MOCs' general newsworthiness or other sources of attention. Second, we can control for

[20]Even in the cross-section of 1.5 million MOC-county dyads, only 6.3% of the dyads presents non-zero donations when calculated over our full period.

area shocks using county-by-period fixed effects (θ_{ct}) to account for common shifts in donations from a county across all candidates. This also picks up any change in donations driven by Twitter content that is not MOC-specific. Finally, we can control for MOC-by-county fixed effects. This takes into account the fact that a given MOC may be more popular in some counties compared to others on an ongoing basis.

Technically, the setup in equation 3 is analogous to a shift-share design where Twitter usage provides county-level exposure to the aggregate shock of MOC-specific Twitter attention. The main concern with this approach is that Twitter usage (our local exposure measure) might be correlated with other cross-sectional factors that might drive the effect of MOC attention on donations. We address this concern in two main ways. We begin by including a large range of interactions between the MOC-level time series of Twitter likes and various county-level characteristics such education levels or demographic structure. In the interests of scientific continuity, we include the full set of 32 cross-sectional county controls used by Fujiwara et al. (2024), which are reported in Appendix Table A9, in addition to Census region and population decile fixed effects.

Including these interactions helps to isolate a Twitter-specific channel from other confounders, but cannot exclude that some unobservable factors might be at play. We therefore ensure that we are only exploiting exogenous variation in Twitter usage through an instrumental variables design in which Twitter usage is instrumented by attendance at the 2007 South-by-Southwest (SXSW) festival in Austin, Texas. In particular, we use the log + 1 number of SXSW followers in 2007 in county c ($SXSWFollowers_c^{2007}$) interacted with the log + 1 number of likes that MOC i receives on date t to instrument $likes_{it} \times users_c$ in equation (3) within a Two-Stage Least Squares (2SLS) framework.

The SXSW festival is an annual event that showcases various forms of media, including music, film, and interactive technology. As noted by Müller and Schwarz (2023) and Fujiwara et al. (2024) who introduced the identification strategy, SXSW 2007 is considered a tipping point for Twitter's popularity and an important early influence on the evolution of its network structure. Twitter's advertising campaign during the 2007 festival resulted in a large increase in the number of daily tweets during the festival and in the year following it. About 60% of early Twitter adopters were connected to SXSW and the platform's growth accelerated disproportionately in counties with SXSW followers who joined Twitter during the 2007 festival. To the extent that early adoption of social media platforms displays persistence, we can use the locations of Twitter's early adopters at SXSW as an instrument for Twitter usage in later periods.[21]

Continuing with the shift-share analogy, identification in our geography-based design comes from our instrument for local exposure being exogenous (Goldsmith-Pinkham et al., 2020). Specifically, conditional on the controls, the number of SXSW followers in 2007 should be uncorrelated with the error term. As noted in Borusyak et al. (2025), this is akin to a parallel trends assumption: counties more or less exposed to MOC attention on Twitter (as captured by the variation in Twitter usage operating via the number of SXSW followers in 2007) would have experienced similar changes in MOC-specific donations. Given the restrictive specification that we estimate—in particular, the fact that we include extensive county-level controls including Census region fixed effects interacted with MOC likes that ensure we are only exploiting within-region variation—this appears to be a plausible assumption. Still, one could be worried about counties that send several people to the SXSW festival in 2007 are different along unobservable characteristics such as openness to innovation. To address this final identification concern, we also control for the intrinsic propensity of individuals from county c to attend the festival by controlling for the number of followers of the SXSW account that joined Twitter before 2007. As a result, in our more restrictive and therefore preferred specification, our instrument should pick up idiosyncratic differences in attendance specific to the 2007 SXSW edition.

5.2 Results

We begin by exploring descriptively the relationship between small donations and Twitter attention separately for counties that have high and low Twitter usage. In particular, Figure 4 shows a binned scatterplot of the relationship between the probability that we observe at least one donation from a

[21] In their papers, Müller and Schwarz (2023) and Fujiwara et al. (2024) provide a thorough discussion and extensive evidence of both the relevance of SXSW followers for the growth of local Twitter activity and the exogeneity of this shock to Twitter's geographic diffusion.

Figure 4: Twitter Likes and Small Donations, High- versus Low-Twitter-Usage Counties

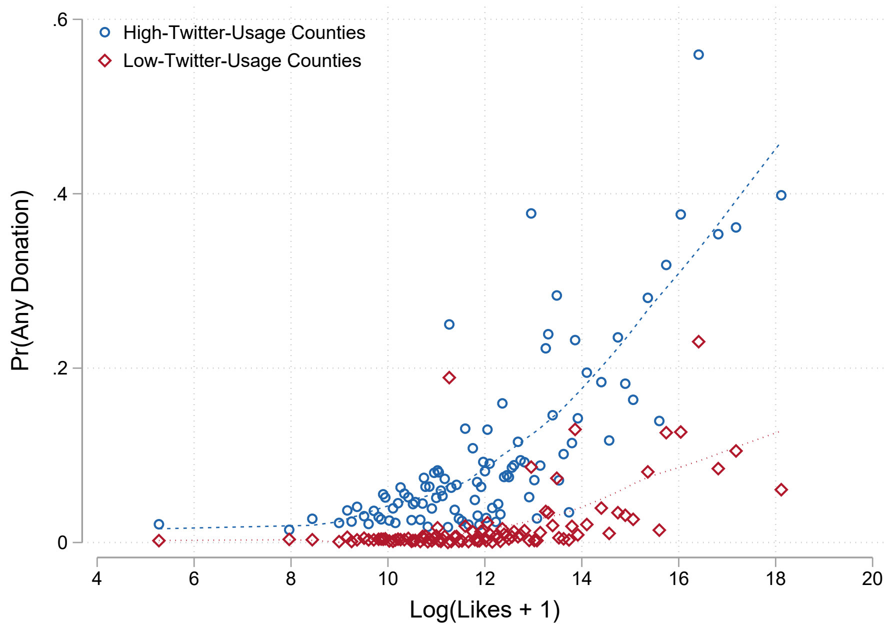

This figure shows a binned scatter plot of the relationship between the probability that an MOC receives a donation from a country and the log + 1 number of Twitter likes that the MOC receives over our sample period, separately for counties with Twitter users above and below the median. The relationship is estimated using a county-MOC level dataset.

given county to a given MOC and the overall Twitter likes that the MOC receives over the entire sample period, separately for counties with number of Twitter users above and below the median. In line with Twitter being the channel driving our effect, the relationship between small donations and Twitter likes is stronger for counties with a higher number of Twitter users.

Our MOC-by-county-by-month model is implemented in Table 3. Column (1) reports the baseline OLS estimates and shows that, in months when they have higher Twitter popularity, MOCs receive significantly more donations from counties with higher Twitter usage. Adding interactions between an MOC's monthly likes and the full set of 32 cross-sectional county characteristics used in Fujiwara et al. (2024) leaves our estimates basically unchanged (column (2)). Given that the average log number of Twitter users is equal to 5.287, the marginal effect of a 1 unit increase in log + 1 likes for the average county corresponds to a 0.26 percentage points increase in the probability of receiving a donation.

While the fact that adding the controls does not impact the estimate is reassuring that we are indeed picking up the effect of Twitter usage, it is possible that high-Twitter usage counties might differ along unobservable dimensions. To address this concern, we move on to our instrumental variable strategy. In columns (3) and (4) we study the reduced form relationship between the probability of a county donating to an MOC and the SXSW followers instrument. The interaction between attention on Twitter and SXSW followers is positively and significantly associated with the probability that an MOC receives positive donations (column (3)) and is not affected by the inclusion of the interaction of Twitter likes and pre-2007 followers of the SXSW account (column (4)). The corresponding first stage specifications are reported in columns (7) and (8). The first stage is strong as per the conventional benchmark.

Finally, column (5) and (6) report our 2SLS estimates. Again, we find that the effect of attention on Twitter on the probability than an MOC receives at least one small donation from a given county is larger for counties that have an (exogenously) higher number of Twitter users. The 2SLS estimates are higher than the OLS estimate, in line with the earlier results of Fujiwara et al. (2024) and Müller and Schwarz (2023). This pattern is suggestive of either a higher local average treatment effect (LATE) for counties with greater levels of SXSW 2007 sign-up or selection into Twitter usage leading to downward bias in the OLS.

In the Appendix, we investigate different levels of time aggregation. In particular, in Appendix Table A7 we estimate models at the weekly level (N=149 million) and in Appendix Table A8 we estimate models at the daily level, focusing on the last 180 days before the election (N=279 million). Very reassuringly, the results are consistent across all levels of aggregation. Overall, we find that when MOCs go viral on Twitter, they receive significantly more small donations from areas with higher Twitter use. Using a shock to the early diffusion of Twitter as an instrument, we can show that this appears to be a Twitter-specific channel and not driven by other county characteristics.

Table 3: Twitter and Small Donations, Geography-Based Design

	Pr(Any Small Donation)						Likes X Twitter users	
	OLS				2SLS		First Stage	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log(Likes + 1)*Log(Twitter Users + 1)	0.0006***	0.0005***			0.0015***	0.0018***		
	(0.0001)	(0.0001)			(0.0003)	(0.0005)		
Log(Likes + 1)*Log(Twitter Users SXSW07 + 1)			0.0008***	0.0009***			0.5596***	0.5177***
			(0.0002)	(0.0003)			(0.0597)	(0.0657)
Log(Likes + 1)*Log(Twitter Users pre06 + 1)				-0.0002		-0.0004		0.1104
				(0.0004)		(0.0005)		(0.1065)
MOC-by-Month Fixed Effects	X	X	X	X	X	X	X	X
County-by-Month Fixed Effects	X	X	X	X	X	X	X	X
MOC-by-County Fixed Effects	X	X	X	X	X	X	X	X
Basic Controls	X	X	X	X	X	X	X	X
Extended Controls		X	X	X	X	X	X	X
Observations	34,135,134	34,135,134	34,135,134	34,135,134	34,135,134	34,135,134	34,135,134	34,135,134
Counties (clusters)	3097	3097	3097	3097	3097	3097	3097	3097
MOCs	501	501	501	501	501	501	501	501
F-statistic							87.842	62.191

This table shows the relationship between Twitter likes and small donations, by county-level Twitter usage. In column (1), we regress an indicator variable for an MOC receiving at least one donation from a given county in a given month on the log + 1 of the number of Twitter likes the MOC receives in the same month interacted with the log + 1 number of Twitter users in the county, log + 1 likes interacted with county population decile fixed effects and Census region fixed effects, MOC-by-month fixed effects, county-by-month fixed effects, and MOC-by-county fixed effects (equation (3)). Column (2) additionally controls for log + 1 Twitter likes interacted with a set of 32 county-level characteristics. Columns (3) estimates a similar specification but including as the main explanatory variable the log + 1 Twitter likes interacted with the log + 1 SXSW followers who joined in March 2007 (column (3)) and also with the log + 1 SXSW followers who joined in 2006 (column (4)). Columns (5) reports estimates from a 2SLS specification where the interaction between log + 1 Twitter likes and log + 1 Twitter users is instrumented using the interaction between the log + 1 Twitter likes and the log + 1 SXSW followers who joined in March 2007. Columns (6) additionally controls for the log + 1 Twitter likes interacted with the log + 1 SXSW followers who joined in 2006. Columns (7) and (8) report the first stage for the 2SLS specifications corresponding to columns (5) and (6) respectively. All specifications are estimated on an MOC-by-county-by-month panel. Standard errors are clustered at county level.

6 Conclusion

Social media has transformed the structure of mass communication, introducing new mechanisms of fast endogenous feedback, sensitivity to behavioral biases, and massive multi-channel competition to many information-related industries. In particular, social media now allows politicians to communicate directly and instantaneously with potential voters and partisan supporters and compete for attention in unprecedented ways relative to earlier broadcast media.

The Twitter platform is a pre-eminent vehicle for this type of communication in the US. Our analysis shows that MOC-specific attention on Twitter is strongly associated with donations in a way that is both temporally concentrated around given increases in attention and skewed in its structure. Specifically, the relationship only becomes strongly evident at heightened levels of viral attention. This relationship is further validated by our geography-based analysis that shows that when Twitter attention increases, the flow of donations that follows come from ares with higher levels of Twitter usage.

While our evidence indicates that Twitter is a technology that can be effective for raising donations, this does not occur for everyone: only a very small number of representatives are able to harness attention on the platform to increase donations from small donors. This is reinforced by crowding out effects in attention—when a given MOC goes viral the probability of their peers gaining attention is reduced. Our results therefore point strongly to the operation of a superstar-style market amongst MOCs on the Twitter platform. Especially if other returns from Twitter attention are similarly concentrated, these tournament-style dynamics might be strong enough to incentivize MOCs to adopt more extreme positions.

Our findings also have some implications for our understanding of attention markets and our broader understanding of Twitter's role in electoral politics. Although the platform has changed in recent years there are general implications of our findings that are likely to be persistent. Our analysis comparing returns to attention on Twitter and on traditional media platforms strongly suggests that media market are highly segmented. In addition, while Twitter is similarly concentrated in attention as traditional media platforms (with the exception of local newspapers), individuals who are successful on Twitter might be different than those who are able to effectively communicate on other platforms. As a result, platforms such as Twitter might still open the playing field to a different set of players who are able to harness this novel communication technology.

References

Allcott, H., Braghieri, L., Eichmeyer, S. and Gentzkow, M. 2020. 'The Welfare Effects of Social Media', *American Economic Review* 110(3), 629–676.

Alvarez, R. M., Katz, J. N. and Kim, S.-y. S. 2020. 'Hidden Donors: The Censoring Problem in US Federal Campaign Finance Data', *Election Law Journal: Rules, Politics, and Policy* 19(1), 1–18.

American National Election Studies. 2021, 'ANES 2020 Time Series Study Full Release [dataset and documentation], February 10, 2022 version.'. www.electionstudies.org.

Antonakaki, D., Fragopoulou, P. and Ioannidis, S. 2021. 'A Survey of Twitter Research: Data Model, Graph Structure, Sentiment Analysis and Attacks', *Expert Systems with Applications* 164, 114006.

Aral, S. 2021, *The Hype Machine: How Social Media Disrupts our Elections, Our Economy and Our Health–and How We Must Adapt*, Harper Collins.

Battiston, G., Boffa, F., Levi, E. and Stillman, S. 2025. 'Strategic Use of Social Media by Political Parties: Evidence from Italy', *Working Paper* .

Bellodi, L. 2023, Bureaucratic Information in Congress. Working Paper.

Bessone, P., Campante, F. R., Ferraz, C. and Souza, P. 2022, Social Media and the Behavior of Politicians: Evidence from Facebook in Brazil. NBER Working Paper # 30306.

Bond, R. M., Fariss, C. J., Jones, J. J., Kramer, A. D., Marlow, C., Settle, J. E. and Fowler, J. H. 2012. 'A 61-Million-Person Experiment in Social Influence and Political Mobilization', *Nature* 489(7415), 295–298.

Borusyak, K., Hull, P. and Jaravel, X. 2025. 'A Practical Guide to Shift-Share Instruments', *Journal of Economic Perspectives* 39(1), 181–204.

Bouton, L., Cagé, J., Dewitte, E. and Pons, V. 2024, Small Campaign Donors. NBER Working Paper # 30050.

Bouton, L., Castanheira, M. and Drazen, A. 2024. 'A Theory of Small Campaign Contributions', *The Economic Journal* 134(662), 2351–2390.

Célérier, C. and Vallée, B. 2019. 'Returns to Talent and the Finance Wage Premium', *Review of Financial Studies* 32(10), 4005–4040.

Chen, J. and Roth, J. 2024. 'Logs with Zeros? Some Problems and Solutions', *The Quarterly Journal of Economics* 139(2), 891–936.

Culberson, T., McDonald, M. P. and Robbins, S. M. 2019. 'Small Donors in Congressional Elections', *American Politics Research* 47(5), 970–999.

Dawood, Y. 2015. 'Campaign Finance and American Democracy', *Annual Review of Political Science* 18, 329–348.

DellaVigna, S. and Gentzkow, M. 2010. 'Persuasion: Empirical Evidence', *Annual Review of Economics* 2(1), 643–669.

De Luca, G., Huning, T. and Santos Monteiro, P. 2021, Britain has had enough of experts? Social networks and the Brexit referendum. Working Papers.

Enikolopov, R., Makarin, A. and Petrova, M. 2020. 'Social Media and Protest Participation: Evidence from Russia', *Econometrica* 88(4), 1479–1514.

Federal Election Commission. 2021*a*, 'Contributions by Individuals, 2019-2020 [dataset]'.

Federal Election Commission. 2021*b*, 'Statistical Summary of 24-Month Campaign Activity of the 2019-2020 Election Cycle'.

Fouirnaies, A. and Fowler, A. 2022. 'Do Campaign Contributions Buy Favorable Policies? Evidence from the Insurance Industry', *Political Science Research and Methods* 10(1), 18–32.

Fouirnaies, A. and Hall, A. B. 2014. 'The Financial Incumbency Advantage: Causes and Consequences', *The Journal of Politics* 76(3), 711–724.

Fujiwara, T., Müller, K. and Schwarz, C. 2024. 'The Effect of Social Media on Elections: Evidence from the United States', *Journal of the European Economic Association* 22(3), 1495–1539.

Gabaix, X. and Landier, A. 2008. 'Why Has CEO Pay Increased So Much?', *Quarterly Journal of Economics* 123(1), 49–100.

Gerber, A. S. 2004. 'Does Campaign Spending Work? Field Experiments Provide eEidence and

Suggest New Theory', *American Behavioral Scientist* 47(5), 541–574.

Goldsmith-Pinkham, P., Sorkin, I. and Swift, H. 2020. 'Bartik Instruments: What, When, Why, and How', *American Economic Review* 110(8), 2586–2624.

Heerwig, J. A. 2016. 'Donations and Dependence: Individual Contributor Strategies in House Elections', *Social Science Research* 60, 181–198.

Jiménez Durán, R. 2023, The Economics of Content Moderation: Theory and Experimental Evidence from Hate Speech on Twitter. Working Paper.

Kinder-Kurlanda, K., Weller, K., Zenk-Möltgen, W., Pfeffer, J. and Morstatter, F. 2017. 'Archiving Information from Geotagged Tweets to Promote Reproducibility and Comparability in Social Media Research', *Big Data & Society* 4(2), 2053951717736336.

Klein, E. 2022, Elon Musk Might Break Twitter. Maybe That's a Good Thing. The Ezra Klein Show.

Koenig, F. 2023. 'Technical Change and Superstar Effects: Evidence from the Rollout of Television.', *American Economic Review: Insights* 5(2), 207–223.

Krueger, A. 2005. 'The Economics of Real Superstars: The Market for Rock Concerts in the Material World', *Journal of Labor Economics* 23(1), 1–30.

Lewis, J. B., Poole, K., Rosenthal, H., Boche, A., Rudkin, A. and Sonnet, L. 2022, 'Voteview: Congressional Roll-Call Votes Database'. https://voteview.com/.

Lorenz-Spreen, P., Oswald, L., Lewandowsky, S. and Hertwig, R. 2023. 'A Systematic Review of Worldwide Causal and Correlational Evidence on Digital Media and Democracy', *Nature Human Behaviour* 7(1), 74–101.

Mislove, A., Marcon, M., Gummadi, K. P., Druschel, P. and Bhattacharjee, B. 2007, Measurement and Analysis of Online Social Networks, *in* 'Proceedings of the 7th ACM SIGCOMM Conference on Internet Measurement', pp. 29–42.

Mullahy, J. and Norton, E. C. 2024. 'Why Transform Y? The Pitfalls of Transformed Regressions with a Mass at Zero', *Oxford Bulletin of Economics and Statistics* 86(2), 417–447.

Müller, K. and Schwarz, C. 2023. 'From Hashtag to Hate Crime: Twitter and Antiminority

Sentiment', *American Economic Journal: Applied Economics* 15(3), 270–312.

Petrova, M., Sen, A. and Yildirim, P. 2021. 'Social Media and Political Contributions: The Impact of New Technology on Political Competition', *Management Science* 67(5), 2997–3021.

Rhodes, J. H., Schaffner, B. F. and La Raja, R. J. 2018. 'Detecting and Understanding Donor Strategies in Midterm Elections', *Political Research Quarterly* 71(3), 503–516.

Rosen, S. 1981. 'The Economics of Superstars.', *American Economic Review* 71(5), 845–858.

Rotesi, T. 2023, The Impact of Twitter on Political Participation. Working Paper.

Schöll, N., Gallego, A. and Le Mens, G. 2023. 'How Politicians Learn from Citizens' Feedback: The Case of Gender on Twitter', *American Journal of Political Science* .

Snyder Jr, J. M. 1990. 'Campaign Contributions as Investments: The US House of Representatives, 1980-1986', *Journal of Political Economy* 98(6), 1195–1227.

Spenkuch, J. L. and Toniatti, D. 2018. 'Political Advertising and Election Results', *The Quarterly Journal of Economics* 133(4), 1981–2036.

Terviö, M. 2009. 'Superstars and Mediocrities: Market Failure in the Discovery of Talent', *Review of Economic Studies* 76(2), 829–850.

Teso, E. 2022. 'Influence-Seeking in US Corporate Elites' Campaign Contribution Behavior', *The Review of Economics and Statistics* pp. 1–34.

Wojcik, S. and Hughes, A. 2019, Sizing Up Twitter Users. Pew Research Center Report.

Zhuravskaya, E., Petrova, M. and Enikolopov, R. 2020. 'Political Effects of the Internet and Social Media', *Annual Review of Economics* 12, 415–438.

A Appendix Tables

Table A1: ANES statistics

	No Twitter	Some Twitter	Daily Twitter	Total
Donation (candidate)	0.138	0.182	0.251	0.160
Donation (party)	0.102	0.094	0.132	0.104
Income $100k+	0.372	0.514	0.511	0.418
College	0.299	0.471	0.470	0.354
City	0.301	0.313	0.354	0.309
Democrat	0.363	0.447	0.581	0.404
Share	0.680	0.208	0.112	0.518

This table shows characteristics of Twitter and non-Twitter users using survey data from the 2020 wave of ANES. We report the mean of different variables by respondents' self-reported Twitter use, where No Twitter indicates zero reported Twitter use, Some Twitter indicates occasional Twitter use (between sporadic visits and multiple times per week), and Daily Twitter indicates individuals who use the website at least once per day. Share describes the shares of the different groups of Twitter usage in the data. The variables are coded as follows: Donation (candidate) is an indicator equal to one if the respondent reported contributing money to an individual candidate running for public office during the election year. Donation (party) is an indicator equal to one if the respondent reported contributing money to political party during the election year. Income $> \$100k$ is an indicator equal to one if the respondent reports a family income of $100,000 or higher. College is an indicator equal to one if the respondent reports having a bachelor's degree or higher. City is an indicator equal to one if the respondent reports living in a city. Democrat is an indicator equal to one if the respondent feels connected to the Democratic Party (compared to not feeling connected to a party or feeling connected to the Republican Party).

Table A2: Top 10 MOCs by Attention Channel

Rank	Twitter	Cable news	NYTimes	Local newspapers
1	Alexandria Ocasio-Cortez	Bernard Sanders	Elizabeth Warren	Bernard Sanders
2	Bernard Sanders	Nancy Pelosi	Nancy Pelosi	Elizabeth Warren
3	Ilhan Omar	Elizabeth Warren	Mitch McConnell	John Lewis
4	Adam B. Schiff	Mitch McConnell	Amy Klobuchar	Nancy Pelosi
5	Ted Lieu	Adam B. Schiff	Cory A. Booker	Ilhan Omar
6	Nancy Pelosi	Lindsey Graham	Alexandria Ocasio-Cortez	Amy Klobuchar
7	Eric Swalwell	Charles E. Schumer	Adam B. Schiff	Cory Gardner
8	Charles E. Schumer	Amy Klobuchar	Ilhan Omar	Mitch McConnell
10	Brian Schatz	Mitt Romney	Lindsey Graham	Alexandria Ocasio-Cortez

This table lists the top 10 MOCs by overall number of likes on Twitter, mentions on cable news, mentions in the New York Times, and in local newspapers over the 2019-2020 period.

Table A3: Other shocks to attention

Econometric Model	Small donations	
	Univariate	Multivariate
	(1)	(2)
Top 10% Twitter	0.094***	0.088***
	(0.019)	(0.019)
Mentioned on cable news	0.066***	0.058***
	(0.018)	(0.018)
Mentioned in the NYT	0.123**	0.105*
	(0.054)	(0.053)
Mentioned in local newspaper	0.060***	0.053**
	(0.018)	(0.017)
Speech in Congress	-0.010	-0.011
	(0.024)	(0.024)
MOC visit	-0.021	-0.021
	(0.046)	(0.046)
President and VP visit	0.020	0.016
	(0.057)	(0.057)
Date-by-party FE	X	X
MOC-by-month FE	X	X
Observations	335,000	335,000
MOC (clusters)	500	500
Mean dep. variable	2.796	2.796

This table shows the relationship between different attention shocks and small donations. Column (1) reports estimates from a series of univariate regressions in which we regress the log + 1 of the total amount of campaign donations below $1000 that an MOC receives on a given day on a proxy for the attention shock, date-by-party fixed effects, and MOC-by-month fixed effects (similar to equation (1)). The proxies for the attention shocks that we use are: an indicator for the number of likes that the MOC receives on the same day being in the top 10% of the likes distribution (1681 likes or more); an indicator for MOC being mentioned on cable news; an indicator for MOC being mentioned in the New York Times; an indicator for the MOC being mentioned in local newspapers; an indicator for the MOC delivering a Congressional speech; an indicator for the MOC visiting his/her home state; and an indicator for the President or the Vice-President visiting the home state of the MOC. In column (2), we perform a "horse race" by including all the attention shocks in the same regression. In all regressions, the sample excludes one MOC not included in the Congressional speeches dataset. Standard errors are clustered the MOC level.

Table A4: Viral Shocks and Donation Harvesting

	Small donations					
	Day of	Week after	Month after	Day of	Week after	Month after
	(1)	(2)	(3)	(4)	(5)	(6)
Log(Likes + 1)	0.011***	0.004***	0.000			
	(0.003)	(0.002)	(0.001)			
Top 10% Twitter				0.098***	0.027**	-0.010
				(0.020)	(0.011)	(0.008)
MOC-by-month FE	X	X	X	X	X	X
Date-by-party FE	X	X	X	X	X	X
Observations	320139	320139	320139	320139	320139	320139
MOCs (clusters)	501	501	501	501	501	501
Mean dep. variable	2.723	6.312	7.498	2.723	6.312	7.498

Notes: This table shows the relationship between Twitter likes and small donations measured for different time periods. In particular, in columns (1) to (3) we regress the log + 1 total amount of donations below $1000 that an MOC receives over the specified time horizon on the log + 1 of the total Twitter likes the same MOC receives on a given day, day-by-party fixed effects, and MOC-by-month fixed effects (similar to equation (1)). Columns (4) to (6) estimate an analogous specification but using as the independent variable an indicator for the number of likes that the MOC receives on the same day being in the top 10% of the likes distribution (1681 likes or more). Standard errors are clustered at the MOC level. The outcome is measured as follows. Columns (1) and (4) use contemporaneous same day donations as the outcome, while columns (2) and (5) calculate the outcome as the cumulative sum of donations received over seven days following a given day t (day t excluded). Analogously, columns (3) and (6) calculate the outcome in the period between seven and thirty days after a given day t.

Table A5: Virality Spillovers

	Top 10% Twitter				Likes	Donations
	(1)	(2)	(3)	(4)	(5)	(6)
Number of viral tweets, weighted	-0.052***			-0.052***	-0.163***	-0.008
	(0.004)			(0.004)	(0.037)	(0.018)
Number of viral tweets, same ideology		-0.001**		-0.001**	0.005	0.004
		(0.001)		(0.001)	(0.006)	(0.005)
Number of viral tweets, same state			-0.002**	0.001	-0.003	0.004
			(0.001)	(0.001)	(0.010)	(0.007)
Top 10% Twitter						0.090***
						(0.020)
Date-by-party FE	X	X	X	X	X	X
MOC-by-month FE	X	X	X	X	X	X
Observations	335000	335000	335000	335000	335000	335000
MOCs (clusters)	500	500	500	500	500	500
Mean dep. variable	0.100	0.100	0.100	0.100	3.479	2.802

This table shows the crowding out of Twitter attention by estimating the relationship between MOCs' own and others' virality. In column (1), we regress an indicator for the number of likes that an MOC receives on a given day being in the top 10% of the likes distribution (1681 likes or more) on the number of other MOCs also in the top 10% of the likes distribution weighted by the share of Twitter follower overlap that the original MOC has with each politician, date-by-party fixed effects, and MOC-by-month fixed effects. In columns (2) and (3), we use as the independent variable the number of other MOCs in the same DW-nominate decile (column (2)) or from the same state (column (3)) who are also in the top 10% of the likes distribution on the same day. Column (4) reports estimates that includes the three independent variables from columns (1) to (3) in a single specification. Note that due to a technical glitch in the data collection, the followers' overlap could not be constructed for one MOC (Rep. Scott Moulton), who was therefore omitted from the sample for this analysis. Standard errors are clustered at the MOC level.

Table A6: Twitter Likes and Small Donors

	Small Donors			
	Log + 1		Dummy	Log
	(1)	(2)	(3)	(4)
Log(Likes + 1)	0.004***			
	(0.001)			
Top 10% Twitter		0.051***	0.006*	0.045***
		(0.008)	(0.003)	(0.009)
Date-by-party FE	X	X	X	X
MOC-by-month FE	X	X	X	X
Observations	335670	335670	335670	149200
MOCs (clusters)	501	501	501	496
Mean dep. variable	0.874	0.874	0.447	1.652

This table shows the relationship between Twitter likes and the number of small donors. In column (1), we regress the log + 1 of the total number of donors who make a contribution below $1000 to an MOC on a given day on the log + 1 of the number of Twitter likes the MOC receives on the same day, date-by-party fixed effects, and MOC-by-month fixed effects (similar to equation (1)). Column (2) estimates the same specification as column (1) but using as the independent variable an indicator for the number of likes that the MOC receives on the same day being in the top 10% of the likes distribution (1681 likes or more). While using the same specification and the same independent variable, column (3) defines the outcome as an indicator variable equal to one if an MOC receives a donation from at least one small donor and column (4) as the log of the total number of small donors that donate to an MOC in a day. In column (4), the sample is conditional on receiving at least one donation. Standard errors are clustered at the MOC level.

Table A7: Twitter and Small Donations, MOC-County-Weekly Design

	Pr(Any Small Donation)						Likes X Twitter users	
	OLS				2SLS		First Stage	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log(Likes + 1)*Log(Twitter Users + 1)	0.0004***	0.0003***			0.0013***	0.0014***		
	(0.0000)	(0.0000)			(0.0001)	(0.0002)		
Log(Likes + 1)*Log(Twitter Users SXSW07 + 1)			0.0008***	0.0007***			0.5596***	0.5177***
			(0.0001)	(0.0001)			(0.0597)	(0.0657)
Log(Likes + 1)*Log(Twitter Users pre06 + 1)				0.0000		-0.0001		0.1104
				(0.0002)		(0.0002)		(0.1065)
MOC-by-Week Fixed Effects	X	X	X	X	X	X	X	X
County-by-Week Fixed Effects	X	X	X	X	X	X	X	X
MOC-by-Week Fixed Effects	X	X	X	X	X	X	X	X
Basic Controls	X	X	X	X	X	X	X	X
Extended Controls		X	X	X	X	X	X	X
Observations	148,953,312	148,953,312	148,953,312	148,953,312	148,953,312	148,953,312	148,953,312	148,953,312
Counties (clusters)	3097	3097	3097	3097	3097	3097	3097	3097
MOCs	501	501	501	501	501	501	501	501
F-statistic							87.843	62.191

This table shows the relationship between Twitter likes and small donations, by county-level Twitter usage. In column (1), we regress an indicator variable for an MOC receiving at least one donation from a given county in a given week on the log + 1 of the number of Twitter likes the MOC receives in the same week interacted with the log + 1 number of Twitter users in the county, log + 1 likes interacted with county population decile fixed effects and Census region fixed effects, MOC-by-week fixed effects, county-by-week fixed effects, and MOC-by-county fixed effects (equation (3)). Column (2) additionally controls for log + 1 Twitter likes interacted with a set of 32 county-level characteristics. Columns (3) estimates a similar specification but including as the main explanatory variable the log + 1 Twitter likes interacted with the log + 1 SXSW followers who joined in March 2007 (column (3)) and also with the log + 1 SXSW followers who joined in 2006 (column (4)). Columns (5) reports estimates from a 2SLS specification where the interaction between log + 1 Twitter likes and log + 1 Twitter users is instrumented using the interaction between the log + 1 Twitter likes and the log + 1 SXSW followers who joined in March 2007. Columns (6) additionally controls for the log + 1 Twitter likes interacted with the log + 1 SXSW followers who joined in 2006. Columns (7) and (8) report the first stage for the 2SLS specifications corresponding to columns (5) and (6) respectively. The sample covers all weeks from January 2019 to October 31st 2020 (96 weeks). Standard errors are clustered at county level.

Table A8: Twitter and Small Donations, MOC-County-Daily Design

	Pr(Any Small Donation)						Likes X Twitter users	
	OLS				2SLS		First Stage	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Log(Likes + 1)*Log(Twitter Users + 1)	0.0001***	0.0001***			0.0005***	0.0005***		
	(0.0000)	(0.0000)			(0.0001)	(0.0001)		
Log(Likes + 1)*Log(Twitter Users SXSW07 + 1)			0.0003***	0.0003***			0.5597***	0.5178***
			(0.0000)	(0.0000)			(0.0597)	(0.0657)
Log(Likes + 1)*Log(Twitter Users pre06 + 1)				0.0001		0.0000		0.1104
				(0.0001)		(0.0001)		(0.1065)
MOC-by-Day Fixed Effects	X	X	X	X	X	X	X	X
County-by-Day Fixed Effects	X	X	X	X	X	X	X	X
MOC-by-Day Fixed Effects	X	X	X	X	X	X	X	X
Basic Controls	X	X	X	X	X	X	X	X
Extended Controls		X	X	X	X	X	X	X
Observations	279,287,460	279,287,460	279,287,460	279,287,460	279,287,460	279,287,460	279,287,460	279,287,460
Counties (clusters)	3097	3097	3097	3097	3097	3097	3097	3097
MOCs	501	501	501	501	501	501	501	501
F-statistic							87.864	62.211

This table shows the relationship between Twitter likes and small donations, by county-level Twitter usage. In column (1), we regress an indicator variable for an MOC receiving at least one donation from a given county on a given day on the log + 1 of the number of Twitter likes the MOC receives on the same day interacted with the log + 1 number of Twitter users in the county, log + 1 likes interacted with county population decile fixed effects and Census region fixed effects, MOC-by-day fixed effects, county-by-day fixed effects, and MOC-by-county fixed effects (equation (3)). Column (2) additionally controls for log + 1 Twitter likes interacted with a set of 32 county-level characteristics. Columns (3) estimates a similar specification but including as the main explanatory variable the log + 1 Twitter likes interacted with the log + 1 SXSW followers who joined in March 2007 (column (3)) and also with the log + 1 SXSW followers who joined in 2006 (column (4)). Columns (5) reports estimates from a 2SLS specification where the interaction between log + 1 Twitter likes and log + 1 Twitter users is instrumented using the interaction between the log + 1 Twitter likes and the log + 1 SXSW followers who joined in March 2007. Columns (6) additionally controls for the log + 1 Twitter likes interacted with the log + 1 SXSW followers who joined in 2006. Columns (7) and (8) report the first stage for the 2SLS specifications corresponding to columns (5) and (6) respectively. The sample covers the 180 days up to October 31st 2020. Standard errors are clustered at county level.

Table A9: List of Control Variables for Additional Interactions.

Variable
Population density
Log(County area)
Distance from Austin, TX (in miles)
Distance from NYC (in miles)
Distance from San Francisco (in miles)
Distance from Washington, DC (in miles)
% aged 20-24
% aged 25-29
% aged 30-34
% aged 35-39
% aged 40-44
% aged 45-49
% aged 50+
Population growth, 2000-2016
% white
% black
% native American
% Asian
% Hispanic
% unemployed
% below poverty level
% employed in IT
% employed in construction/real estate
% employed in manufacturing
% with high school degree
% with college education
% watching CNN
% watching Fox News
% watching prime time TV
Exposure to sport competition
Share of noncitizens
Average charitable index

This table reports the baseline county characteristics that we interact with the log + 1 number of likes in our geography-based design.

B Appendix Figures

Figure A1: Concentration of Twitter Likes and Other Attention Channels

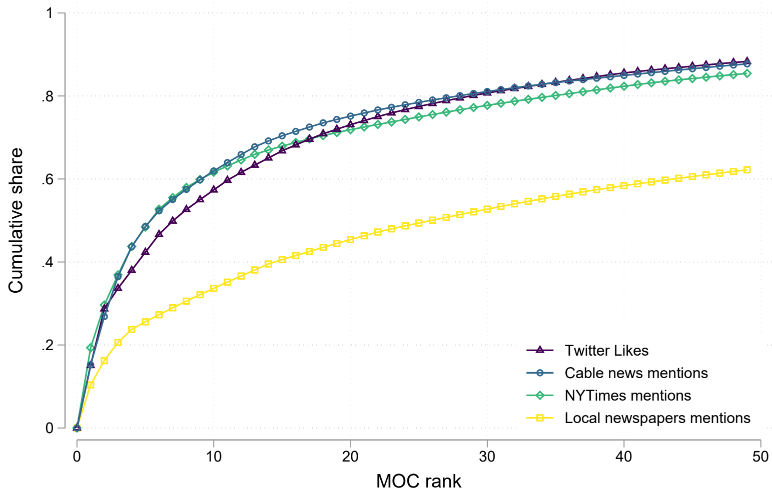

This figure plots the cumulative distribution of MOCs' total Twitter likes and cable news, New York Times, and local newspapers mentions by their relative rank, going from the most mentioned (liked) to least. The sample is restricted to the top 50 MOCs for each attention channel.

Figure A2: Correlates of Attention by Attention Channel

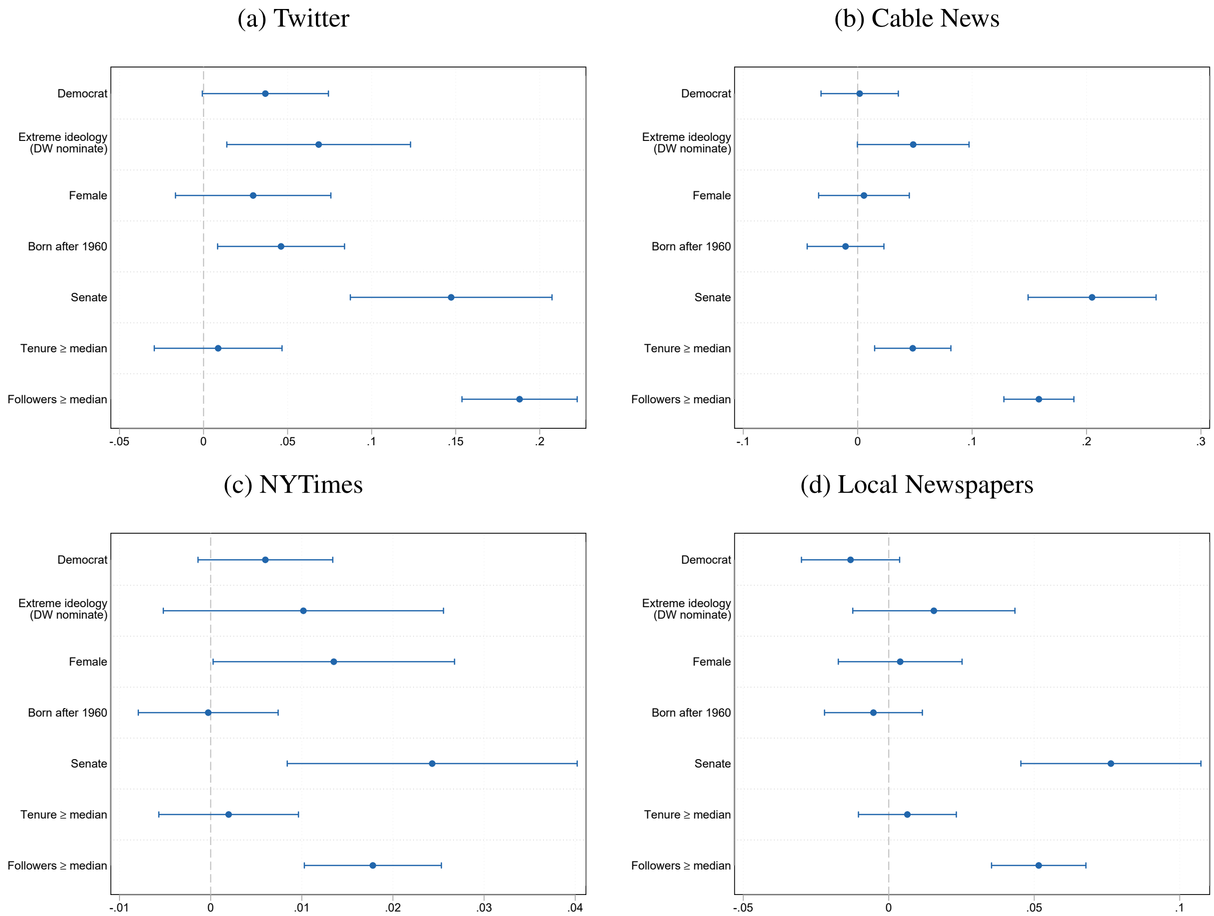

This figure shows the correlation between the attention received by MOCs across different attention channels and their characteristics. Using an MOC-by-date panel, we regress the outcomes described below on a specific MOC characteristic and date fixed effects. The outcome is an indicator for the number of likes that the MOC receives on the same day being in the top 10% of the likes distribution in panel (a); an indicator variable for the MOC being mentioned on cable news in panel (b); an indicator variable for the MOC being mentioned in the New York Times in panel (c); and an indicator variable for the MOC being mentioned in a local newspaper in panel (d). The effect of each characteristic is estimated from a separate regression. Standard errors are clustered at the MOC level. We define MOCs as having an extreme ideology if their DW nominate score is in the bottom or top decile of the distribution; the definition of all other characteristics should be self-explanatory.

Figure A3: Likes by Tweet Characteristics

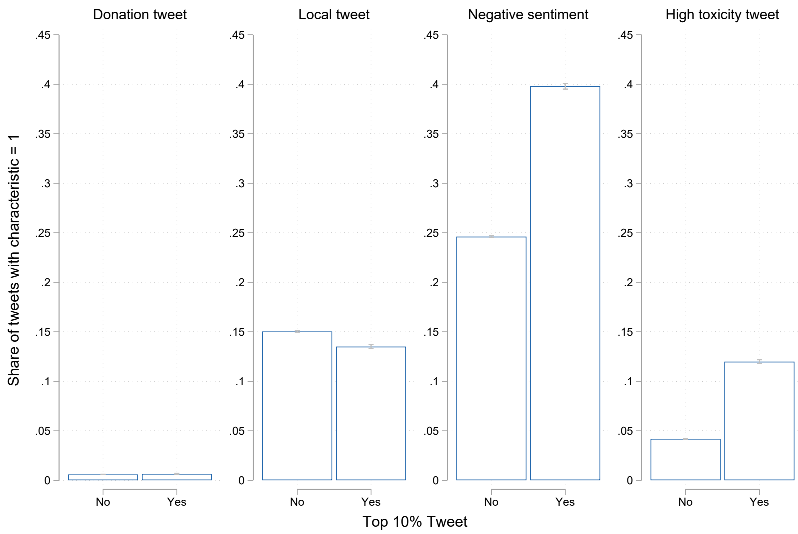

This figure shows the share of tweets with a given content by whether the tweet is in the top 10% of the likes distribution or not. A donation tweet is a tweet that links to a donation portal (e.g., WinRed or ActBlue); a local tweet is tweet in which the MOC mentions the name of at least one municipality located in their congressional district; a negative sentiment tweet is a tweet with negative polarity score according to VADER; a high toxicity tweet is a tweet with a toxicity score in the top 5% of the toxicity score distribution according to Google's Perspectives API.

Figure A4: Twitter Likes and Small Donations, Returns by Virality

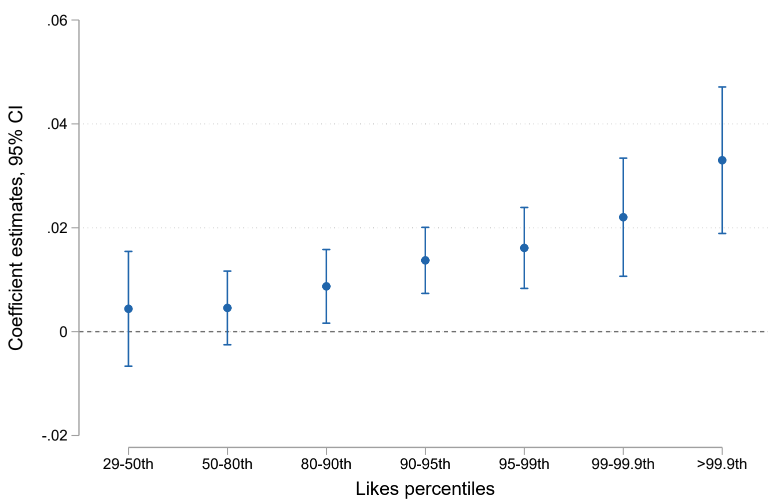

This figure shows how the relationship between Twitter likes and small donations varies depending on the level of virality. In particular, we regress the log + 1 of the total amount of donations below $1000 that an MOC receives on a given day on a series of indicators for the number of likes that the MOC receives on the same day being in different percentiles of the likes distribution interacted with the log + 1 of the total Twitter likes received, date-by-party fixed effects, and MOC-by-month fixed effects (similar to equation (1)). The category $29 - 50th$ includes days with positive likes in the 50th percentile and below, or having between 1 and 31 likes. $50 - 80th$ indicates days in the 51st to 80th percentile (between 32 and 350 likes). $80 - 90th$ indicates days in the 81st to 90th percentile (between 351 and 1680 likes). $90 - 95th$ indicates days in the 91st to 95th percentile (between 1681 and 7,776 likes). $95 - 99th$ indicates days in the 95th to 99th percentile (between 7,777 and 67,989 likes). $99 - 99.9th$ indicates days in the 100th percentile, excluding the top .1% of the distribution (between 67,996 and 319,364 likes). $> 99.9th$ indicates the top .1% of the distribution (between 319,601 and 3,808,126 likes). Standard errors are clustered at the MOC level.

Figure A5: Extensive and Intensive Margin of Donations, Leads and Lags

(a) Donations Dummy

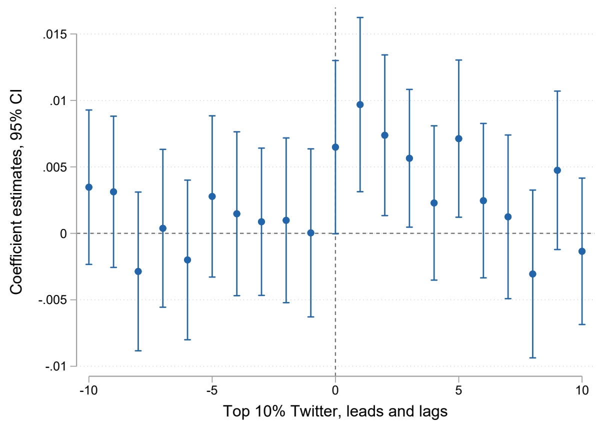

(b) Log Donations

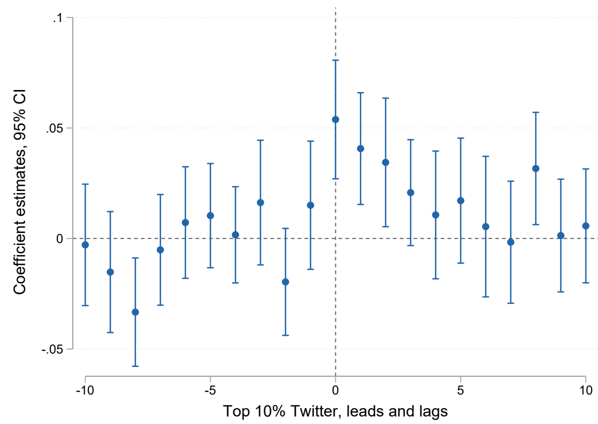

This figure shows the dynamic relationship between likes on Twitter and small donations, separately estimated along the extensive and intensive margin. In panel (a), we regress an indicator variable equal to 1 if an MOC receives positive donations on a given day on ten leads and lags of an indicator for the number of likes that the MOC receives on the same day being in the top 10% of the likes distribution (1681 likes or more), date-by-party fixed effects, and MOC-by-month fixed effects (equation (2)). Panel (b) estimates the same specification but using the log + 1 of the total amount of donations below $1000 that an MOC receives in a day as the outcome, while restricting the sample to days in which the MOC received positive donations. Standard errors are clustered at the MOC level.

Figure A6: Effect of Twitter Likes on Small Donations, Heterogeneity by MOC Characteristics

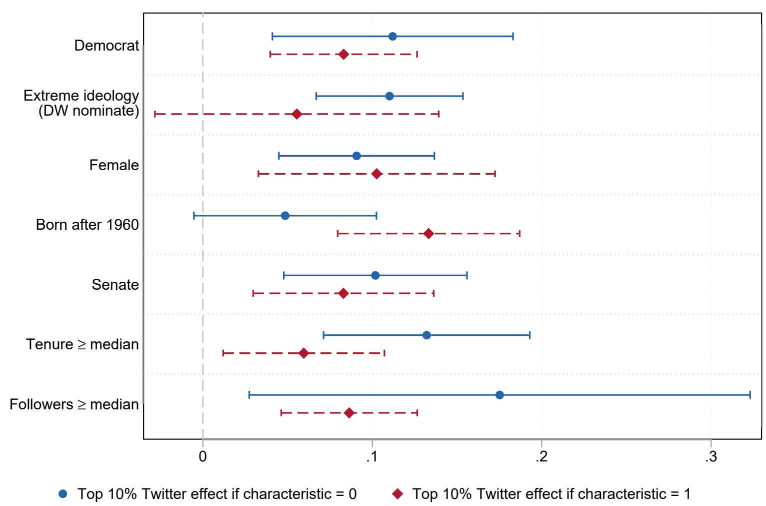

This figure shows heterogeneity of the relationship between Twitter likes and small donations by MOC characteristics. We regress the log + 1 of the total amount of donations below $1000 that an MOC receives on a given day on an indicator for the number of likes that the MOC receives on the same day being in the top 10% of the likes distribution (1681 likes or more), the interaction between the same indicator and a specific MOC characteristic, date-by-party fixed effects, and MOC-by-month fixed effects (similar to equation (1)). We estimate a separate regression for each characteristic and we report in the figure the effect for MOCs who have or do not have the specific characteristic. Standard errors are clustered at the MOC level. We define MOCs as having an extreme ideology if their DW nominate score is in the bottom or top decile of the distribution; the definition of all other characteristics should be self-explanatory.

Figure A7: Effect of Twitter Likes on Small Donations, Heterogeneity by Tweet Content

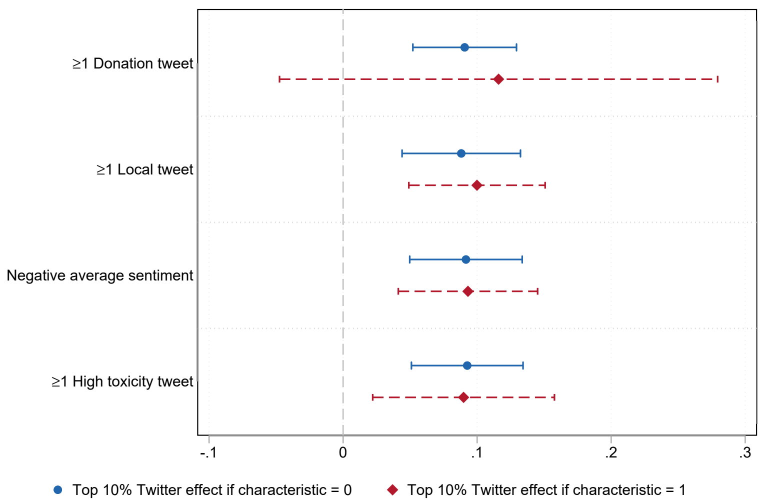

This figure shows heterogeneity of the relationship between Twitter likes and small donations by tweet content. We regress the log + 1 of the total amount of donations below \$1000 that an MOC receives on a given day on an indicator for the number of likes that the MOC receives on the same day being in the top 10% of the likes distribution (1681 likes or more), the interaction between the same indicator and an indicator variable for the characteristic of the tweet content, date-by-party fixed effects, and MOC-by-month fixed effects (similar to equation (1)). We estimate a separate regression for each characteristic and we report in the figure the effect for MOC-days that do or do not display the specific characteristic. Standard errors are clustered at the MOC level. The characteristics of tweet content are defined as follows: ≥ 1 donation tweet is an indicator variable equal to 1 if the MOC posted the link to a donation portal (e.g., WinRed or ActBlue); ≥ 1 local tweet is an indicator variable equal to 1 if the MOC mentioned the name of a municipality located in their congressional districts at least once; negative average sentiment is an indicator variable equal to 1 if the average sentiment of the tweets is negative (where the sentiment of the tweets is calculated using VADER); ≥ 1 high toxicity tweet is an indicator variable equal to 1 if the MOC had at least one tweet with a toxicity score in the top 5% of the toxicity score distribution according to Google's Perspectives API.

C Robustness Checks

Robustness to Variables' Definitions. We check the robustness of our results to the use of different measures of Twitter attention and small donations in Appendix Table A10. We find a positive relationship between small donations and tweets' other popularity metrics (retweets, replies, quotes). Interestingly, more intensive interactions such as quotes have higher returns. We also find similar effects for various measures of donations, such all individual donations (i.e., including contributions $\geq$ \$1000) and donations through conduits.

Functional Form. In Appendix Table A11, we show further robustness to functional form assumptions. Reassuringly, we find a significant and positive link between Twitter attention and likes across multiple alternative specifications, including linear probability models (LPM) with indicator variables for high-likes and high-donation days (columns (2) and (3)), using logs on a subsample of frequent non-zero donations (column (4)), and the untransformed levels-levels linear relationship (column (5)).

Event Studies. In Appendix Figure A8 we estimate variations of this basic dynamic model using viral days defined according to different thresholds as well as the continuous measure of likes. These all yield similar leads and lag structures to that seen in Figure 2.

Table A10: Twitter Likes and Small Donations, Robustness to Different Measures of Attention and Donations

	Donations				
	Small			All	Conduit
	(1)	(2)	(3)	(4)	(5)
Replies	0.018***				
	(0.004)				
Retweets		0.016***			
		(0.003)			
Quotes			0.024***		
			(0.005)		
Log(Likes + 1)				0.013***	0.010***
				(0.003)	(0.003)
Date-by-party FE	X	X	X	X	X
MOC-by-month FE	X	X	X	X	X
Observations	335670	335670	335670	335670	335670
MOCs (clusters)	501	501	501	501	501
Mean dep. variable	2.799	2.799	2.799	3.447	2.102

This table shows the robustness of the relationship between Twitter likes and small donations to using different measures of Twitter popularity and different definitions of the outcome variable. In column (1) we regress the log + 1 of the total amount of donations below $1000 that an MOC receives on a given day on the log + 1 of Twitter replies the MOC receives on the same day, date-by-party fixed effects, and MOC-by-month fixed effects (similar to equation (1)). In column (2) and (3), we use as the main independent variable the log + 1 of retweets and the log + 1 of quotes that the MOC receives on the same day respectively. In column (4) we do not restrict donations to be below $1000, while in column (5) we only consider donations received through the conduits Winred and Actblue. Standard errors are clustered at the MOC level.

Table A11: Twitter Likes and Small Donations, Robustness to Functional Form Issue

	Small donations				
	Log + 1	$\mathbb{1}(> \$1k)$	$\mathbb{1}(> \$5k)$	Log	Linear
	(1)	(2)	(3)	(4)	(5)
Log(Likes + 1)	0.011***				
	(0.003)				
Top 10% Twitter		0.017***	0.010***		
		(0.004)	(0.002)		
Log likes				0.013***	
				(0.003)	
Likes					0.003**
					(0.001)
Date-by-party FE	X	X	X	X	X
MOC-by-month FE	X	X	X	X	X
Observations	335670	335670	335670	119817	335670
MOCs (clusters)	501	501	501	486	501
Mean dep. variable	2.799	0.154	0.039	6.320	1219.040

This table shows the robustness of the relationship between Twitter likes and small donations to different functional forms of the main variables. Column (1) replicates our baseline specification (see Table 2 column (4)). In particular, we regress the log + 1 of the total amount of donations below $1000 that an MOC receives on a given day on the log + 1 of the total Twitter likes the same MOC receives on the same day, date-by-party fixed effects, and MOC-by-month fixed effects (equation (1)). Columns (2) and (3) use indicator variables for high-contribution and high-popularity observations. The dependent variables are, respectively, an indicator variable equal to 1 if the total amount of donations below $1000 than an MOC receives in a given day is above $1000 or above $5000. The independent variable is an indicator for the number of likes that the MOC receives on the same day being in the top 10% of the likes distribution (1681 likes or more). Column (4) uses the log transformation on a subsample where zero donation days are rare: the three months leading up to the 2020 election and MOCs who have less than 5% zero-donation and zero-like days. Column (5) presents results from the linear specification using the untransformed variables. Standard errors are clustered at the MOC level.

Figure A8: Twitter Likes and Small Donations, Leads and Lags Robustness

(a) Log Likes + 1

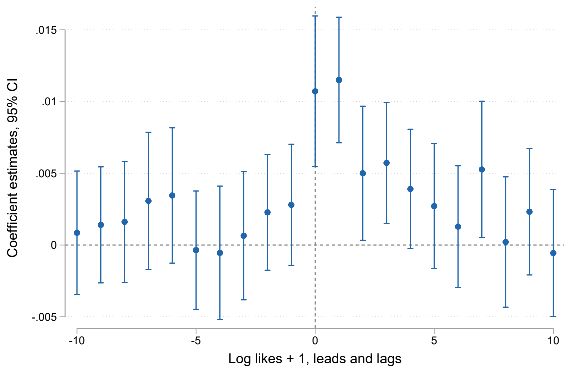

(b) Top 1% Likes

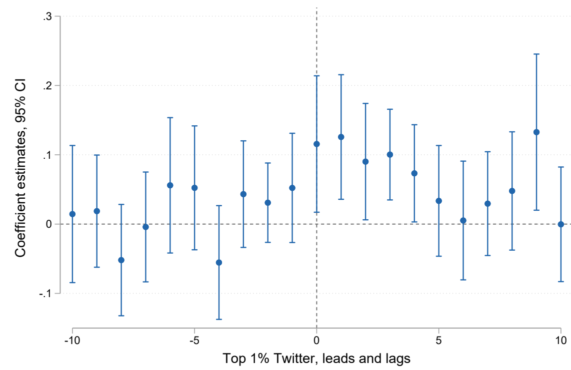

(c) Only 2019

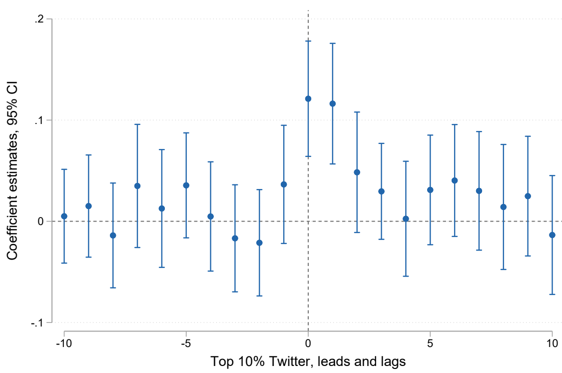

(d) By party

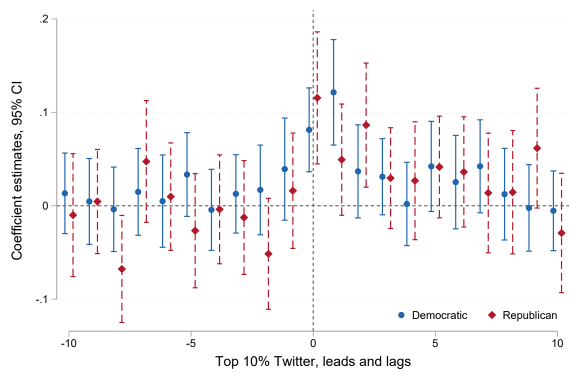

This figure shows the robustness of the dynamic relationship between likes on Twitter and small donations. In panel (a), we regress the log + 1 of the total amount of donations below $1000 that an MOC receives in a day on ten leads and lags of the log + 1 of the total Twitter likes the same MOC receives on the same day, date-by-party fixed effects, and MOC-by-month fixed effects. In panel (b), the independent variables are ten leads and lags of an indicator for the number of likes that the MOC receives on the same day being in the top 10% of the likes distribution (67996 likes or more). Panel (c) estimates our main specification (equation (2)) restricting the sample to 2019 (thus, the pre-COVID and pre-election period). Finally, panel (d) estimates our main specification (equation (2)) splitting the sample by whether the MOC is part of the Democratic (blue) or Republican (red) party. Standard errors are clustered at the MOC level.