THE EFFECT OF SOCIAL MEDIA ON ELECTIONS: EVIDENCE FROM THE UNITED STATES

Thomas Fujiwara
Princeton University

Karsten Müller
National University of Singapore, Business School

Carlo Schwarz
Bocconi University, Department of Economics

This appendix presents further details on data construction, the SXSW festival, additional robustness exercises, further results, and the LATE extrapolation:

- Appendix A provides additional details on the data.
- Appendix B discusses the SXSW instrument.
- Appendix C shows additional robustness checks.
- Appendix D provides further results.
- Appendix E discusses extrapolation for the average treatment effect.

Appendix A: Additional Details on Data

E-mail: fujiwara@princeton.edu (Fujiwara); kmueller@nus.edu.sg (Müller); carlo.schwarz@unibocconi.it (Schwarz)

Journal of the European Economic Association
Preprint prepared on 4 October 2023 using jeea.cls *v1.0.*

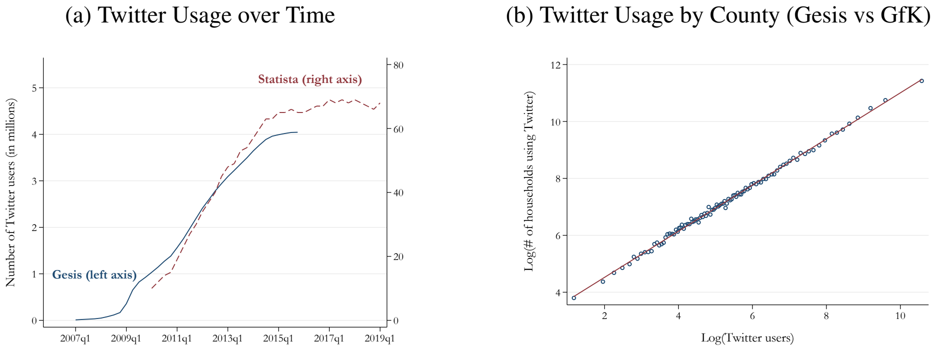

FIGURE A.1. Validation of Twitter usage measure. This graph shows two validation exercises for the Twitter usage measure in the Gesis data (Kinder-Kurlanda et al. 2017). Panel (a) plots the number of Twitter users in the Gesis data and the number of active monthly users reported by Statista based on Twitter's own reporting. Panel (b) plots the percentiles of the number of Twitter users in the Gesis data at the county-level against the number of users based on the GfK Media Survey.

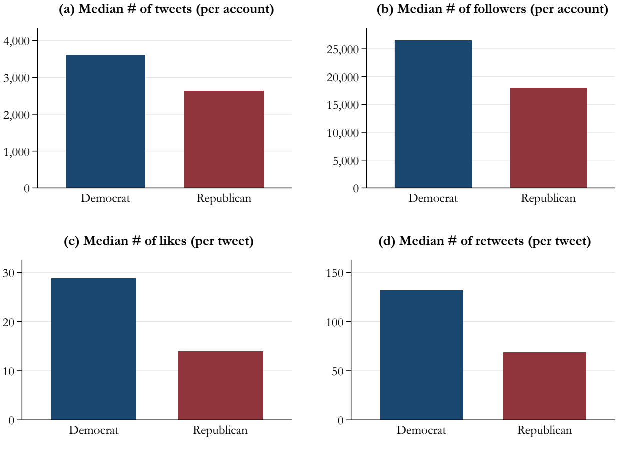

FIGURE A.2. Twitter reach by party (median). This figure plots data on the Twitter reach of Congress members. The sample includes all 901 senators and House representatives who were in office between 2007 and 2019 for whom we could identify a Twitter account. For each account, we plot the median number of tweets and followers, and the median number of "likes" and retweets of their tweets. The data were collected from Twitter in November 2019.

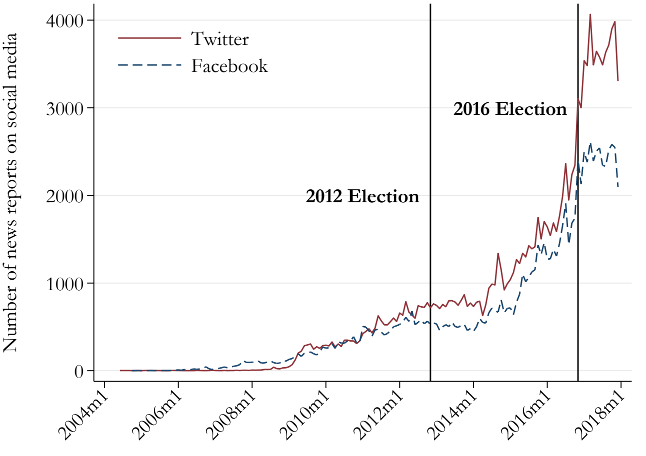

FIGURE A.3. News reports about social media. This graph plots the number of times the terms "Twitter" and "Facebook" are mentioned in USA Today, The Washington Post, The New York Post, and The New York Times based on data from Nexis.

TABLE A.1. Summary statistics (county-level).

	Mean	**Std. Dev.**	**Min.**	**Median**	**Max.**	**N**
Vote outcomes and Twitter data						
Republican two-party vote share (2016)	0.46	0.17	0.08	0.45	0.95	3,065
Change in Republican two-party vote share, 2000-2016	-0.02	0.10	-0.33	-0.03	0.45	3,065
Republican two-party vote share (2020)	0.47	0.17	0.09	0.45	0.96	3,065
Change in Republican two-party vote share, 2000-2020	-0.01	0.10	-0.34	-0.02	0.48	3,065
Log(Twitter users)	8.22	1.99	0.00	8.45	12.35	3,065
Log(SXSW followers, March 2007)	0.69	1.13	0.00	0.00	4.98	3,065
Log(SXSW followers, Pre)	0.33	0.73	0.00	0.00	3.61	3,065
Geographical controls						
Population density	1925.15	6342.94	0.10	508.30	69468.40	3,065
Log(County area)	6.72	0.92	3.26	6.64	9.91	3,065
Distance from Austin, TX (in miles)	1731.21	653.61	5.04	1750.86	3098.88	3,065
Distance from Chicago (in miles)	1246.18	813.51	7.16	1103.75	3040.38	3,065
Distance from NYC (in miles)	1600.47	1255.14	6.48	1285.98	4191.67	3,065
Distance from San Francisco (in miles)	2841.16	1231.98	41.11	3157.16	4565.01	3,065
Distance from Washington, DC (in miles)	1448.51	1175.55	7.88	1047.13	3983.08	3,065
Demographic controls						
% aged 20-24	0.07	0.02	0.01	0.06	0.27	3,065
% aged 25-29	0.07	0.01	0.03	0.07	0.15	3,065
% aged 30-34	0.07	0.01	0.03	0.06	0.12	3,065
% aged 35-39	0.06	0.01	0.03	0.06	0.10	3,065
% aged 40-44	0.06	0.01	0.02	0.06	0.10	3,065
% aged 45-49	0.06	0.01	0.02	0.06	0.09	3,065
% aged 50+	0.36	0.06	0.11	0.35	0.75	3,065
Population growth, 2000-2016	0.14	0.19	-0.43	0.10	1.32	3,065
% white	0.65	0.21	0.03	0.68	0.98	3,065
% black	0.12	0.12	0.00	0.08	0.85	3,065
% native American	0.01	0.03	0.00	0.00	0.90	3,065
% Asian	0.05	0.06	0.00	0.03	0.37	3,065
% Hispanic	0.15	0.15	0.01	0.09	0.96	3,065
% unemployed	5.31	1.42	1.80	5.10	24.10	3,065
Socioeconomic controls						
% below poverty level	15.11	5.34	1.40	15.10	53.30	3,065
% employed in IT	0.02	0.02	0.00	0.02	0.21	3,065
% employed in construction/real estate	0.07	0.03	0.00	0.07	1.00	3,065
% employed in manufacturing	0.11	0.08	0.00	0.08	0.72	3,065
% adults with high school degree	28.13	7.41	8.30	27.50	54.80	3,065
% adults with college degree	20.98	3.66	8.70	20.90	35.60	3,065
% watching Fox News	0.26	0.01	0.23	0.26	0.30	3,064
% watching prime time TV	0.43	0.01	0.40	0.43	0.47	3,064
China shock controls						
Exposure to Chinese import competition	2.63	2.02	-0.63	2.10	43.08	3,065
Share of routine occupations	31.87	2.36	22.23	32.14	36.66	3,065
Average offshorability index	-0.02	0.50	-1.64	0.09	1.24	3,065
1996 election control						
Republican two-party vote share (1996)	0.41	0.11	0.10	0.41	0.79	3,065

Notes: This table presents descriptive statistics for the main estimation sample, weighted by the turnout in the 2000 presidential elections.

Table A.2. Summary statistics (2016 CCES individual-level).

	Mean	**Std. Dev.**	**Min.**	**Median**	**Max.**	**N**
Vote outcome						
Voted for Trump	0.49	0.50	0.00	0.00	1.00	146,579
Twitter data						
Log(Twitter users)	8.30	1.91	0.69	8.45	12.35	146,579
Log(SXSW followers, March 2007)	0.69	1.12	0.00	0.00	4.98	146,579
Log(SXSW followers, Pre)	0.32	0.71	0.00	0.00	3.61	146,579
Individual control variables						
Log(Age)	3.89	0.37	2.89	3.99	4.61	146,579
Family income dummies	7.12	3.65	1.00	7.00	13.00	146,579
Female dummy	1.52	0.50	1.00	2.00	2.00	146,579
Education dummies	3.54	1.54	1.00	3.00	6.00	146,579
Marital status dummies	2.36	1.71	1.00	1.00	5.00	146,579
Interest in news dummies	1.61	0.98	1.00	1.00	7.00	146,579

Notes: This table presents descriptive statistics for the CCES estimation sample, weighted by survey weights.

Table A.3. Summary statistics (Gallup individual-level).

	Mean	**Std. Dev.**	**Min.**	**Median**	**Max.**	**N**
Candidate approval outcomes						
Approve of Trump?	0.34	0.48	0.00	0.00	1.00	64,764
Approve of Kasich?	0.60	0.49	0.00	1.00	1.00	8,735
Approve of Rubio?	0.50	0.50	0.00	1.00	1.00	6,201
Approve of Cruz?	0.41	0.49	0.00	0.00	1.00	11,504
Approve of Sanders?	0.57	0.50	0.00	1.00	1.00	27,137
Approve of Clinton?	0.43	0.50	0.00	0.00	1.00	36,367
Twitter data						
Log(Twitter users)	8.29	1.97	0.00	8.48	12.35	64,764
Log(SXSW followers, March 2007)	0.72	1.15	0.00	0.00	4.98	64,764
Log(SXSW followers, Pre)	0.34	0.73	0.00	0.00	3.61	64,764
Individual control variables						
Income dummies	6.99	2.38	1.00	7.00	10.00	64,764
Female dummy	1.50	0.50	1.00	2.00	2.00	64,764
Education dummies	3.58	1.60	1.00	4.00	6.00	64,764
Marital status dummies	1.98	0.94	1.00	2.00	5.00	64,764
Age deciles	4.45	2.68	1.00	4.00	10.00	64,764

Notes: This table presents descriptive statistics for the Gallup estimation sample, weighted by survey weights.

TABLE A.4. Instrument balancedness.

	March 2007 *and* Pre (1)	March 2007 *only* (2)	Pre *only* (3)	Difference in means (2) - (3)	p-value	Šidàk p-value
Population density	5192.27	1021.39	1998.35	-976.96	0.07*	0.91
Log(County area)	6.30	6.63	6.54	0.09	0.73	1.00
Distance from Austin, TX (in miles)	1775.99	1749.38	1626.64	122.74	0.48	1.00
Distance from Chicago (in miles)	1439.45	1329.47	1214.42	115.05	0.53	1.00
Distance from NYC (in miles)	1685.31	1594.99	1510.05	84.94	0.78	1.00
Distance from San Francisco (in miles)	2751.83	2900.11	2833.01	67.10	0.83	1.00
Distance from Washington, DC (in miles)	1558.55	1450.23	1397.05	53.18	0.85	1.00
% aged 20-24	0.07	0.08	0.08	0.00	0.92	1.00
% aged 25-29	0.09	0.07	0.07	-0.00	0.51	1.00
% aged 30-34	0.08	0.07	0.07	-0.00	0.58	1.00
% aged 35-39	0.07	0.06	0.06	-0.00	0.82	1.00
% aged 40-44	0.06	0.06	0.06	0.00	0.82	1.00
% aged 45-49	0.07	0.06	0.06	0.00	0.89	1.00
% aged 50+	0.32	0.35	0.35	-0.00	0.97	1.00
Population growth, 2000-2016	0.18	0.18	0.15	0.03	0.56	1.00
% white	0.50	0.65	0.67	-0.02	0.62	1.00
% black	0.18	0.12	0.08	0.04	0.20	1.00
% native American	0.01	0.01	0.02	-0.02	0.02**	0.45
% Asian	0.10	0.05	0.05	-0.01	0.55	1.00
% Hispanic	0.20	0.16	0.15	0.01	0.80	1.00
% unemployed	4.86	5.05	4.51	0.54	0.07*	0.91
% below poverty level	15.71	15.82	13.69	2.14	0.17	1.00
% employed in IT	0.04	0.02	0.02	-0.00	0.98	1.00
% employed in construction/real estate	0.06	0.07	0.07	0.01	0.39	1.00
% employed in manufacturing	0.07	0.09	0.07	0.02	0.16	0.99
% adults with high school degree	21.76	25.99	25.77	0.22	0.88	1.00
% adults with college degree	18.89	21.16	21.20	-0.04	0.97	1.00
% watching Fox News	0.25	0.26	0.26	-0.00	0.91	1.00
% watching prime time TV	0.42	0.43	0.43	0.00	0.91	1.00
Exposure to Chinese import competition	2.55	2.46	2.79	-0.32	0.54	1.00
Share of routine occupations	32.47	31.38	31.25	0.13	0.82	1.00
Average offshorability index	0.37	-0.07	-0.05	-0.02	0.84	1.00
Republican two-party vote share (1996)	0.36	0.42	0.42	-0.00	0.90	1.00

Notes: This table presents the averages for the main control variables separately for the three types of counties relevant for our identification strategy: 1) the 47 counties with SXSW followers that joined Twitter both in March 2007 and the "pre-period"; 2) the 108 counties with SXSW followers that joined in March 2007 (but none in the "pre-period"); and 3) the 20 counties with SXSW that joined in the "pre-period" (but none in March 2007). The demographic and socioeconomic controls are measured in 2016. We report p-values from a two-sided t-test for the equality of means between the counties with the key identifying variation, as well as Šidàk-corrected values to account for multiple hypothesis testing. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A.1.. Additional Details on the Logistic Regression Classifier

We train a separate machine learning classifier for each of the three election years in our data using the Python sci-kit package (Pedregosa et al. 2011). These classifiers help us to determine whether tweets are more likely to be sent by Democratic-leaning or Republican-leaning users. The classification process starts with the preparation of the underlying Twitter data. The inputs are the text of each of the 4,300,579 tweets from U.S. Congress members. To focus on election-related tweets, we restrict the sample to tweets that were sent either in the election year or in the year leading up to the election and mention at least one of the presidential candidates.

The target variable y for the classifier is an indicator variable equal to one for tweets sent by Republicans and zero otherwise. The feature matrix X for the classifier are created by count-vectorizing the texts of the tweets. In other words, we transform the text of the tweets into $n \times v$ matrix, where n is the number of tweets and v is the number of unique 1,2-grams that occur in the tweets. In preparation for this step, we removed common words (stopwords), links, and special characters from the tweets. Additionally, we reduced the words in the tweets to their morphological roots using a lemmatizer, which improves the performance of the classifier. As an example, the lemmatizer changes words like "walking" and "walked" to "walk". Lastly, we reweight the n-grams v of tweet i using term frequency–inverse document frequency (tf-idf):

$$tfidf(f_{i,v}) = (1 + \ln(f_{i,v}) \cdot (\ln(\frac{1+T}{1+d_v}) + 1) \tag{A.1}$$

where d_v is the number of tweets n-gram v appears in. This reweighting reduces the importance of words that appear frequently in many tweets, which help little to discriminate between tweets. The tf-idf vectorizer also normalizes the feature matrix by its L2-norm. The vector y and the matrix X then serve as the input for a $L2$ regularized logistic regression classifier. The classifier minimizes the following cost function[1]:

$$\min_{w,c} C \sum_{i=1}^{n} \log(\exp(-y_i (X_i^T w + c) + 1) + \frac{1}{2} w^T w \tag{A.2}$$

where w are the weights (coefficients) of the logistic regression, c is a constant (intercept), and C is the inverse of the regularization strength. Larger values of C imply weaker regularization. For $C \longrightarrow \infty$, the classifier converges towards a normal logistic regression. As is standard in most machine learning applications, we choose the optimal regularization strength C using 10-fold cross-validation. This involves randomly splitting the training data into ten equal slices. Nine of the ten slices are then used to train the classifier, while the out-of-sample performance is evaluated against the remaining slice using F1-scores.

The final classifiers achieves an out-of-sample F1-score of 0.916 in 2012, of 0.843 in 2016, and of 0.904 in 2020. The classifiers, therefore, accurately predict the party affiliation of Congress members. We then take these classifiers and apply it to the

1. Note that this formulation of the cost function assumes that y_i takes values $-1; 1$. We use this formulation in line with the sci-kit documentation.

(a) 2012 Election

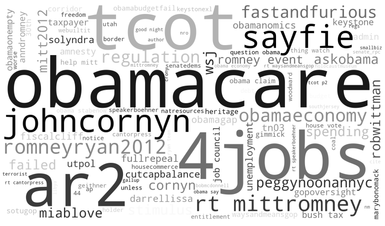

(b) 2016 Election

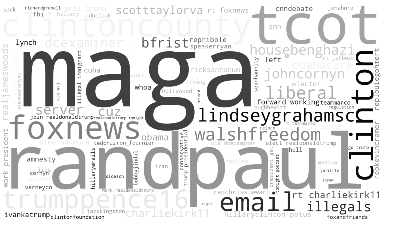

(c) 2020 Election

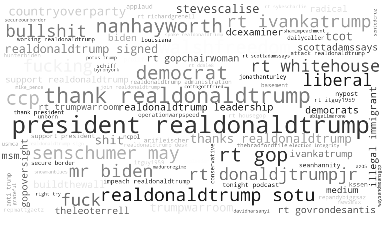

FIGURE A.4. Most predictive terms of "Republican" Tweets by election. This word cloud plots the n-grams most predictive of tweets sounding like those of Republican Congress members, as identified by the logistic regression classifier for each election cycle. The size of the word represents the magnitude of the coefficients.

universe of tweets sent during the 2012, 2016, and 2020 presidential elections. For each tweet in the election data, the classifiers provide us with a predicted class (either Democrat or Republican) and a probability for this class label. To avoid that our results are driven by tweets for which the classifier is "uncertain," we code tweets with a predicted class probability below 60% as neutral. This adjustment has no bearing on our findings. In spirit, this approach is similar to the work of Gentzkow and Shapiro (2010). While they identify expressions that are more frequently used by Democrats and Republicans by hand, we use a machine learning classifier to identify n-grams in the tweets of Congress members that help us to differentiate between the two parties.

We visualize the most predictive n-grams identified by the classifiers for each election cycle in Figure A.4. Overall, the classifiers identify words, hashtags, and Twitter handles that are intuitively associated with a Republican slant for each election year. Among the most predictive term are the hashtags "tcot" (Top Conservatives on Twitter) and "maga" (Make America great again) and particularly in 2020 many references to Donald Trump's Twitter account ("realdonaldtrump").

Appendix B: Additional Details on the SXSW Festival

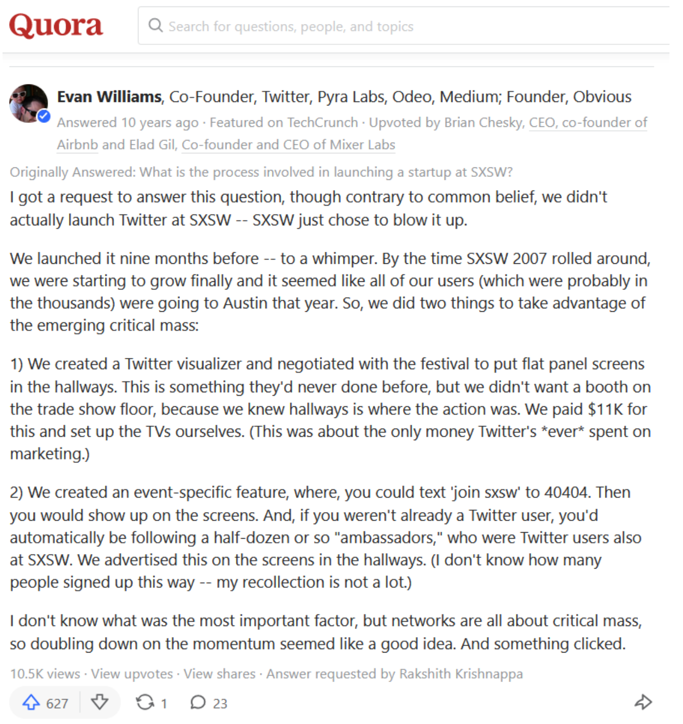

Quora

Search for questions, people, and topics

Evan Williams, Co-Founder, Twitter, Pyra Labs, Odeo, Medium; Founder, Obvious

Answered 10 years ago · Featured on TechCrunch · Upvoted by Brian Chesky, CEO, co-founder of Airbnb and Elad Gil, Co-founder and CEO of Mixer Labs

Originally Answered: What is the process involved in launching a startup at SXSW?

I got a request to answer this question, though contrary to common belief, we didn't actually launch Twitter at SXSW -- SXSW just chose to blow it up.

We launched it nine months before -- to a whimper. By the time SXSW 2007 rolled around, we were starting to grow finally and it seemed like all of our users (which were probably in the thousands) were going to Austin that year. So, we did two things to take advantage of the emerging critical mass:

1) We created a Twitter visualizer and negotiated with the festival to put flat panel screens in the hallways. This is something they'd never done before, but we didn't want a booth on the trade show floor, because we knew hallways is where the action was. We paid $11K for this and set up the TVs ourselves. (This was about the only money Twitter's *ever* spent on marketing.)

2) We created an event-specific feature, where, you could text 'join sxsw' to 40404. Then you would show up on the screens. And, if you weren't already a Twitter user, you'd automatically be following a half-dozen or so "ambassadors," who were Twitter users also at SXSW. We advertised this on the screens in the hallways. (I don't know how many people signed up this way -- my recollection is not a lot.)

I don't know what was the most important factor, but networks are all about critical mass, so doubling down on the momentum seemed like a good idea. And something clicked.

10.5K views · View upvotes · View shares · Answer requested by Rakshith Krishnappa

627 1 23

FIGURE B.1. Screenshot quote from Twitter founder. This screenshot shows the full post of Twitter co-founder Evan Williams posted on Quora on January 4, 2011 describing the role of the SXSW festival in the platform's rise to popularity (Quora 2011).

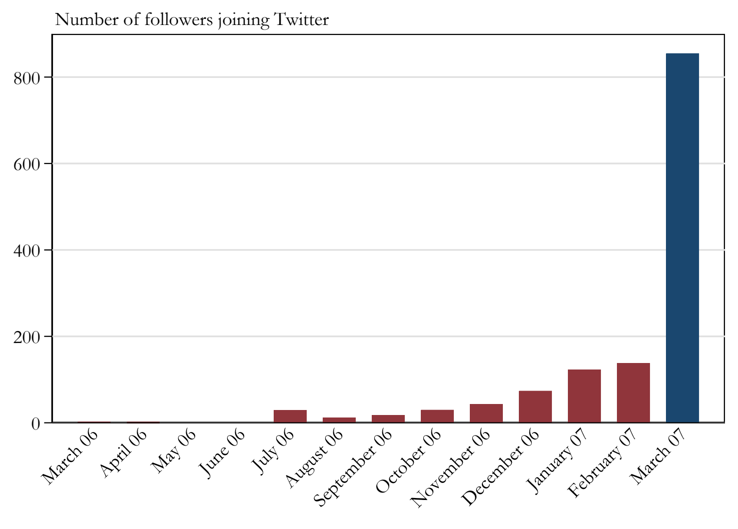

FIGURE B.2. Number of SXSW follower joining each month. The figures shows the number of SXSW Follower joining in each month.

(a) Long-term Effects of the 2007 SXSW on Twitter Adoption

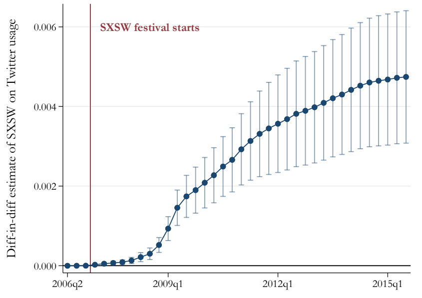

(b) Connections to the SXSW festival

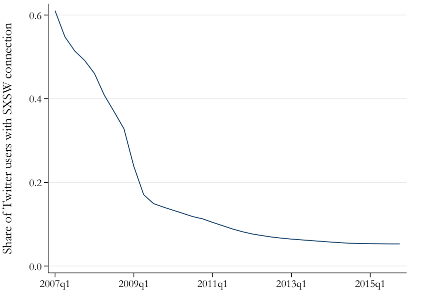

FIGURE B.3. Additional evidence for the impact of the SXSW festival. The figures provide evidence on the long-term impact of the SXSW festival on Twitter usage across the United States. Panel (a) plots the β_τ from the panel event study regression $users_{ct} = \sum_\tau \beta_\tau SXSW_c^{March2007} \times 1(t = \tau) + \sum_\tau \delta_\tau SXSW_c^{Pre} \times 1(t = \tau) + \theta_c + \gamma_t + \varepsilon_{ct}$ where $users_{ct}$ is the number of Twitter users per capita in county c on quarter t, $SXSW_c^{March2007}$ is the logarithm of (one plus) the number of SXSW followers in county c who joined Twitter in March 2007 and $SXSW_c^{Pre}$ is a similarly defined variable for followers who joined Twitter before March 2007. We standardize the variables to have a mean of zero and standard deviation of one. The whiskers represent 95% confidence intervals based on standard errors clustered by state, where 2006q4 serves as excluded period. While the confidence intervals for 2006q2 and 2006q3 cannot be seen, they include zero. Panel (b) plots the share of Twitter who either follow SXSW or follow a user that follows SXSW.

TABLE B.1. Balancedness of SXSW counties' user characteristics.

First names (Corr. = 0.63)		**Terms used in bio (Corr. = 0.89)**	
Pre-period	March 2007	Pre-period	March 2007
michael	michael	http	http
paul	john	founder	com
mike	chris	com	digital
chris	jeff	tech	founder
eric	matt	product	medium
justin	brian	co	director
ryan	david	digital	tech
kevin	alex	director	music
jeff	jason	design	social
david	kevin	social	marketing

Notes: This table presents the ranking of the most common first names and terms used in a Twitter user's "bio" among users who follow "South by Southwest" on Twitter, depending on whether they joined during March 2007 or in the pre-period.

Table B.2. Are Twitter users in counties with SXSW followers different?

User first names (Corr. = 0.97)		**Terms used in user bio (Corr. = 0.94)**	
Other counties	SXSW counties	Other counties	SXSW counties
michael	michael	love	co
chris	david	life	love
john	chris	co	life
david	john	http	http
sarah	alex	http co	http co
mike	mike	god	music
emily	matt	ig	lover
ryan	sarah	music	ig
matt	ryan	university	de
alex	andrew	like	like

Notes: This table compares the individual characteristics of Twitter users from counties with "South by Southwest" followers who joined in March 2007 ("SXSW counties") to Twitter users from all other U.S. counties ("Other counties"). We plot the ranking of the most common first names and terms used in a Twitter user's "bio".

Appendix C: Additional Robustness Checks

TABLE C.1. LASSO variable selection for 2016/2020 election.

	Dep. var.: Republican vote share in...			
	Census Region FE		State FE	
	Controls (1)	Controls^2 (2)	Controls (3)	Controls^2 (4)
Panel A: 2SLS 2016 Election				
Log(Twitter users)	-0.030**	-0.039***	-0.031***	-0.046***
	(0.012)	(0.013)	(0.012)	(0.013)
Observations	3,064	3,064	3,064	3,064
Nr. Controls	48	609	93	654
Nr. selected controls	30	70	50	87
Panel B: 2SLS 2020 Election				
Log(Twitter users)	-0.026*	-0.035***	-0.029**	-0.039***
	(0.014)	(0.013)	(0.013)	(0.013)
Observations	3,064	3,064	3,064	3,064
Nr. Controls	48	609	93	654
Nr. selected controls	31	73	51	91

Notes: This table presents county-level regressions where the dependent variable is the Republican vote share in the 2016 or 2020 presidential election. *Log(Twitter users)* is instrumented using the number of users who started following SXSW in March 2007 (in logs with 1 added inside). Columns 1 includes census region fixed effects and allows all potential control variables discussed in the text (48 controls) to be selected by the LASSO procedure. Columns 2 includes census region fixed effects and allows for interactions of all control variables with each other (609 potential controls). Column 3 includes state fixed effects and allows all control variables to be selected (93 potential controls). Column 4 includes state fixed effects and allows all interactions of control variables with each other to be selected (654 potential controls). In all regressions the included controls are selected using the partialing-out LASSO procedure from Chernozhukov et al. (2015a,b). Standard errors in parentheses are clustered at the state level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

TABLE C.2. Twitter and the Republican vote share – Robustness.

	No regression weights (1)	Election year weights (2)	Pre-period control polynomial (3)	Pre-period control deciles (4)	No zero SXSW user counties (5)	Spatial standard errors (6)	Per Capita Twitter Usage (7)
Panel A: Republican vote share in 2016							
Log(Twitter users)	-0.037***	-0.021**	-0.019***	-0.020***	-0.030**	-0.037***	
	(0.011)	(0.008)	(0.007)	(0.007)	(0.012)	(0.011)	
Log(Twitter users p.c.)							-0.091**
							(0.044)
Observations	3,064	3,064	3,064	3,064	165	3,064	3,064
Mean of DV	0.64	0.46	0.46	0.46	0.34	0.64	0.46
Robust F-stat.	72.94	94.50	125.40	114.91	23.46	54.13	24.50
Panel B: Republican vote share in 2020							
Log(Twitter users)	-0.036**	-0.020**	-0.018**	-0.019**	-0.033**	-0.036***	
	(0.014)	(0.010)	(0.008)	(0.009)	(0.016)	(0.012)	
Log(Twitter users p.c.)							-0.086*
							(0.047)
Observations	3,064	3,064	3,064	3,064	165	3,064	3,064
Mean of DV	0.65	0.47	0.47	0.47	0.35	0.65	0.47
Robust F-stat.	72.94	88.47	125.40	114.91	23.46	54.13	24.50

Notes: This table presents 2SLS results estimated using equation (3). The dependent variable is the vote share of the Republican party in the 2016 and 2020 presidential elections in panel A and B, respectively. *Log(Twitter users)* is instrumented using the number of users who started following SXSW in March 2007 (in logs with 1 added inside). All regressions except columns 1, 2, and 6 are weighted by turnout in the 2000 presidential election. Column 1 omits regression weights. Column 2 weights by the turnout in the election year (2016 or 2020) instead of 2000. Columns 3 and 4 control for a fifth-order polynomial and deciles of SXSW followers who joined Twitter before the SXSW 2007 event, respectively. Column 5 drops all counties that had no SXSW followers joining Twitter in March 2007 or in the period before. Column 6 uses spatial standard errors based on the method proposed in Colella et al. (2019), implemented in Stata as $acreg$, using a 200 miles cutoff. Column 7 uses the number of Twitter users per capita (in logs with 1 added inside). All regressions include the controls from columns 5 and 10 in Table 2. In columns 1 to 5, 7, and 8, standard errors in parentheses are clustered by state. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

TABLE C.3. Twitter and changes in Republican vote share, 2004–2020.

	Dep. var.: ΔRepublican vote share between...				
	2000-04 (1)	2000-08 (2)	2000-12 (3)	2000-16 (4)	2000-20 (5)
Panel A: Reduced form					
Log(SXSW followers, March 2007)	-0.002	-0.003	-0.001	-0.009**	-0.008**
	(0.002)	(0.003)	(0.003)	(0.004)	(0.004)
Log(SXSW followers, Pre)	-0.000	-0.001	-0.002	0.000	-0.002
	(0.003)	(0.005)	(0.003)	(0.004)	(0.005)
Panel B: 2SLS					
Log(Twitter users)	-0.004	-0.006	-0.002	-0.017**	-0.015**
	(0.004)	(0.006)	(0.006)	(0.007)	(0.007)
Log(SXSW followers, Pre)	0.000	-0.001	-0.002	0.001	-0.001
	(0.003)	(0.005)	(0.003)	(0.005)	(0.005)
Observations	3,064	3,064	3,064	3,064	3,064
Mean of DV	0.03	-0.02	-0.01	-0.02	-0.01
Robust F-stat.	121.18	121.18	121.18	121.18	121.18

Notes: This table presents county-level regressions where the dependent variable is the change in the vote share of the Republican party between 2000 and the indicated year. *Log(SXSW followers, March 2007)* is the number of Twitter users (in logs, with 1 added inside) who joined in March 2007 and follow South by Southwest (SXSW). *SXSW followers, Pre* is the number of SXSW followers who registered at some point in 2006, defined similarly. *Twitter users* are the number of users in 2014-2015. All regressions control for population deciles, Census region fixed effects, and the full set of controls (as in columns 5 and 10 of Table 2). The first-stage regressions for 2SLS results (Panel B) are presented in Table 1, with the F-stat for the excluded instrument in the bottom row. On Panel (a), the coefficient on *Log(SXSW followers, March 2007)* for 2004, 2008, and 2012 are jointly statistically insignificant (p-value=0.398). Further, the average effect in 2016 and 2020 is statistically distinct from the average effect in 2004, 2008, and 2012 (p-value=0.001). Observations are weighted by turnout in the 2000 presidential election. Standard errors in parentheses are clustered by state. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

FIGURE C.1. Twitter and the Republican vote share, 1924–1996 (Reduced form).

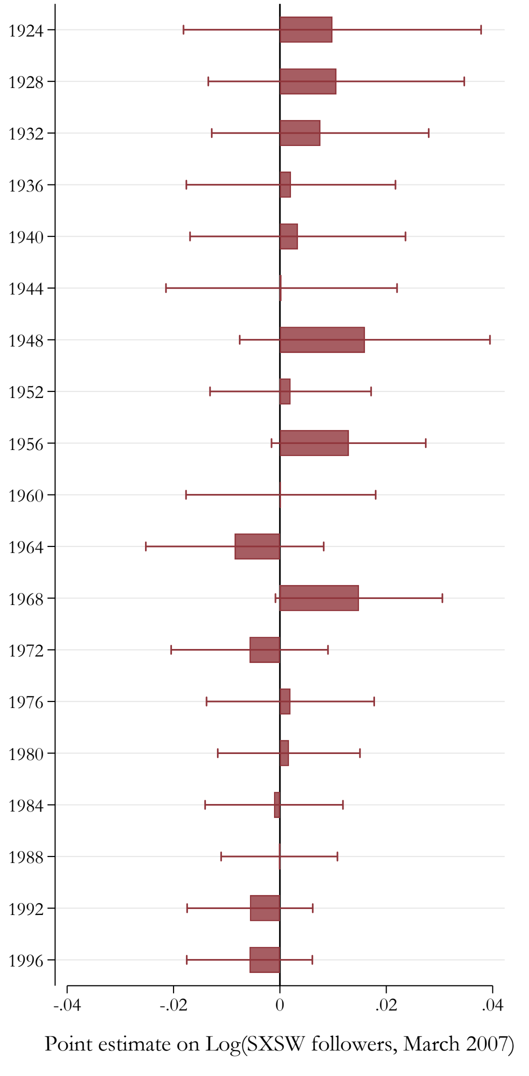

Notes: This figure plots reduced form estimates $\hat{\beta}'$ from county-level regressions as in equation (2). These estimates reflect the correlation of *Log(1 + SXSW followers, March 2007)* with the Republican vote share in presidential elections while controlling for *Log(1 + SXSW followers, Pre)*. All regressions control for population deciles and Census region fixed effects, and the full set of controls except 1996 Election controls (same as columns 4 and 9 of Table 2). Regressions are weighted by turnout in the 2000 presidential election. Whiskers represent 95% confidence intervals based on standard errors clustered by state.

TABLE C.4. Time-varying Twitter usage and changes in vote shares.

	Dep. var.: ΔRepublican vote share between...				
	2000-04 (1)	2000-08 (2)	2000-12 (3)	2000-16 (4)	2000-20 (5)
Panel A: First stage					
Log(SXSW followers, March 2007)	0.644***	0.644***	0.546***	0.523***	0.523***
	(0.055)	(0.055)	(0.047)	(0.048)	(0.048)
Panel B: Reduced form					
Log(SXSW followers, March 2007)	-0.002	-0.003	-0.001	-0.009**	-0.008**
	(0.002)	(0.003)	(0.003)	(0.004)	(0.004)
Log(SXSW followers, Pre)	-0.000	-0.001	-0.002	0.000	-0.002
	(0.003)	(0.005)	(0.003)	(0.004)	(0.005)
Panel C: 2SLS					
Log(Twitter users)	-0.003	-0.005	-0.002	-0.017**	-0.015**
	(0.003)	(0.005)	(0.005)	(0.007)	(0.007)
Log(SXSW followers, Pre)	0.000	-0.001	-0.002	0.001	-0.001
	(0.003)	(0.005)	(0.003)	(0.005)	(0.005)
Observations	3064	3064	3064	3064	3064
Mean of DV		-0.024	-0.008	-0.018	-0.009
Robust F-stat.	135.88	135.88	134.88	121.18	121.18
Twitter usage measured in	*2008*	*2008*	*2012*	*2016*	*2016*

Notes: This table presents county-level regressions where the dependent variable is the change in the vote share of the Republican party between 2000 and the indicated year (except for Panel B, where the dependent variable is *Twitter users*. *Log(SXSW followers, March 2007)* is the number of Twitter users (in logs, with 1 added inside) who joined in March 2007 and follow South by Southwest (SXSW). *SXSW followers, Pre* is the number of SXSW followers who registered at some point in 2006, defined similarly. Differently from other tables, *Twitter users* varies over time, as opposed to being fixed to 2014-2015. All regressions control for population deciles, Census region fixed effects, and the full set of controls (as in columns 5 and 10 of Table 2). Observations are weighted by turnout in the 2000 presidential election. Standard errors in parentheses are clustered by state. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

TABLE C.5. Twitter and the Ross Perot vote.

	Dep. var.: Vote share Ross Perot in...			
	1992		1996	
	(1)	(2)	(3)	(4)
Panel A: Reduced form				
Log(SXSW followers, March 2007)	0.003	0.003	0.000	-0.000
	(0.003)	(0.003)	(0.002)	(0.002)
Panel B: 2SLS				
Log(Twitter users)	0.006	0.007	0.000	-0.001
	(0.006)	(0.006)	(0.003)	(0.003)
Population deciles	Yes	Yes	Yes	Yes
Census region FE	Yes	Yes	Yes	Yes
Geographical controls	Yes	Yes	Yes	Yes
Demographic controls	Yes	Yes	Yes	Yes
Socioeconomic controls	Yes	Yes	Yes	Yes
China shock controls	Yes	Yes	Yes	Yes
1996 election control		Yes		Yes
Observations	3,064	3,064	3,064	3,064
Mean of DV	0.20	0.20	0.10	0.10
Robust F-stat.	118.21	121.18	118.21	121.18

Notes: This table presents county-level regressions where the dependent variable is the third party vote share in the 1992 or 1996 presidential election. *Log(SXSW followers, March 2007)* is the number of Twitter users (in logs, with 1 added inside) who joined in March 2007 and follow South by Southwest (SXSW). *SXSW followers, Pre* is the number of SXSW followers who registered at some point in 2006, defined similarly. *Twitter users* are the number of users in 2014-2015. The first-stage regressions for 2SLS results (Panel B) are presented in Table 1, with the F-stat for the excluded instrument in the bottom row. Observations are weighted by turnout in the 2000 presidential election. Standard errors in parentheses are clustered by state. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

TABLE C.6. Twitter and vote shares in Democratic primaries.

	Dep. var.: Vote share in Democratic Primary of...							
	Clinton 2016	Sanders 2016	Warren 2020	Biden 2020	Sanders 2020	Buttigieg 2020	Bloomberg 2020	Klobuchar 2020
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Reduced form								
Log(SXSW followers, March 2007)	0.005	-0.004	0.002	-0.009	0.017***	-0.003	0.001	-0.003
	(0.010)	(0.010)	(0.006)	(0.014)	(0.006)	(0.002)	(0.004)	(0.002)
Panel B: 2SLS								
Log(Twitter users)	0.008	-0.006	0.003	-0.014	0.025**	-0.004	0.001	-0.004
	(0.014)	(0.014)	(0.009)	(0.021)	(0.010)	(0.003)	(0.006)	(0.002)
Observations	2,656	2,656	2,769	2,769	2,769	2,769	2,769	2,769
Mean of DV	0.55	0.43	0.06	0.56	0.24	0.02	0.06	0.01
Robust F-stat.	67.94	67.94	73.68	73.68	73.68	73.68	73.68	73.68

Notes: This table presents county-level regressions where the dependent variable is the vote share of the indicated candidate in the Democratic party primaries in 2016 or 2020. *Log(SXSW followers, March 2007)* is the number of Twitter users (in logs, with 1 added inside) who joined in March 2007 and follow South by Southwest (SXSW). *SXSW followers, Pre* is the number of SXSW followers who registered at some point in 2006, defined similarly. *Twitter users* are the number of users in 2014-2015. All regressions control for population deciles, Census region fixed effects, and the full set of controls (as in columns 5 and 10 of Table 2). The first-stage regressions for 2SLS results (Panel B) are analogous to the one presented in Table 1, except for the different sample of counties for which primary results are available. The F-stat for the excluded instrument is provided in the bottom row. Observations are weighted by turnout in the 2000 presidential election. Standard errors in parentheses are clustered by state. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

TABLE C.7. Twitter and vote decisions in the 2016 CCES – Robustness.

	Dep. var.: Voted for Trump in 2016				
	(1)	(2)	(3)	(4)	(5)
	Baseline	Verified vote	Vote intention	Intended Trump vote	Intended other vote
Log(Twitter users)	-0.135***	-0.154***	-0.133**	-0.231***	-0.064*
	(0.045)	(0.054)	(0.052)	(0.082)	(0.034)
Observations	146,579	56,375	46,418	14,723	24,354
Mean of DV	0.492	0.497	0.455	0.991	0.137
Marginal effect	[-0.049]	[-0.055]	[-0.048]	[-0.005]	[-0.013]

Notes: This table presents results estimated using IV probit models, as in equation (4). The dependent variable is a dummy for individuals in the CCES who voted for Trump in 2016. *Log(Twitter users)* is instrumented using the number of SXSW followers that joined Twitter in March 2007. All regressions control for the (log) number of SXSW followers that joined Twitter at some point in 2006, family income, gender, education levels, marital status, news interest, and age, as well as county-level population deciles and Census region fixed effects. Regressions are weighted by survey weights. Standard errors in parentheses are clustered by state. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Appendix D: Further Results

TABLE D.1. Additional outcomes.

	Switching prob.	*ΔCampaign don., 2000-16*		*Trump approval, 2017*	
	Obama to Trump (1)	Democrats (2)	Republicans (3)	Democrats (4)	Republicans (5)
Log(Twitter users)	-0.138***	0.866***	0.168	-0.011**	-0.037***
	(0.048)	(0.185)	(0.235)	(0.005)	(0.009)
Observations	3,065	2,250	2,446	2,727	2,920
Mean of DV	0.105	1.943	1.096	0.066	0.850
Robust F-stat.	74.65	62.57	64.34	60.20	66.44

Notes: This table presents results from county-level regressions of equation (3). Column 1 shows results from an IV probit regression where the dependent variable is a dummy equal to 1 for the 217 counties for which both Obama and Trump gained the majority of votes in 2008 and 2016, respectively. In columns 2 and 3, the dependent variable is the difference in the natural logarithm of campaign donations to the Democratic and Republican party, respectively, between 2000 and 2016. In columns 4 and 5, the dependent variable is the share of respondents in the Gallup Daily Poll approving of Trump in 2017. *Log(Twitter users)* is instrumented using the number of users who started following SXSW in March 2007. All regressions control for population deciles and Census region fixed effects and geographical controls. Regressions are weighted by turnout in the 2000 presidential election. Standard errors in parentheses are clustered by state. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

TABLE D.2. Twitter and changes in voter turnout, 2004–2020.

	ΔVotes cast/voting age pop.				
	2000-04 (1)	2000-08 (2)	2000-12 (3)	2000-16 (4)	2000-20 (5)
Panel A: Reduced form					
Log(SXSW followers, March 2007)	-0.000	-0.000	-0.001	0.001	0.007**
	(0.003)	(0.003)	(0.004)	(0.004)	(0.003)
Log(SXSW followers, Pre)	-0.000	0.000	-0.002	-0.001	-0.005
	(0.006)	(0.004)	(0.005)	(0.005)	(0.005)
Panel B: 2SLS					
Log(Twitter users)	-0.000	-0.000	-0.001	0.002	0.014**
	(0.006)	(0.006)	(0.008)	(0.008)	(0.006)
Log(SXSW followers, Pre)	-0.000	0.000	-0.002	-0.001	-0.006
	(0.006)	(0.005)	(0.005)	(0.005)	(0.005)
Observations	3,063	3,063	3,063	3,063	3,063
Mean of DV	0.088	0.079	0.053	0.057	0.126
Robust F-stat.	121.23	121.23	121.23	121.23	121.23

Notes: This table presents county-level regressions where the dependent variable is the change in the voter turnout (as a share of voting age population) between 2000 and the indicated year. *Log(SXSW followers, March 2007)* is the number of Twitter users (in logs, with 1 added inside) who joined in March 2007 and follow South by Southwest (SXSW). *SXSW followers, Pre* is the number of SXSW followers who registered at some point in 2006, defined similarly. *Twitter users* are the number of users in 2014-2015. All regressions control for population deciles, Census region fixed effects, and the full set of controls (as in columns 5 and 10 of Table 2). The first-stage regressions for 2SLS results (Panel B) are presented in Table 1, with the F-stat for the excluded instrument in the bottom row. Observations are weighted by turnout in the 2000 presidential election. Standard errors in parentheses are clustered by state. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

TABLE D.3. Twitter and congressional elections – Reduced form estimates.

Panel A: House elections

	Δ*Republican vote share in House election between...*									
	2000-02 (1)	2000-04 (2)	2000-06 (3)	2000-08 (4)	2000-10 (5)	2000-12 (6)	2000-14 (7)	2000-16 (8)	2000-18 (9)	2000-20 (10)
Log(SXSW followers, March 2007)	0.010	0.002	0.017*	0.018*	0.015	0.017	0.024*	0.009	0.019*	0.017
	(0.011)	(0.010)	(0.010)	(0.009)	(0.011)	(0.011)	(0.012)	(0.012)	(0.011)	(0.011)
Observations	3,046	3,007	3,035	2,993	3,027	3,047	3,042	3,044	3,047	3,046
Mean of DV	0.02	0.01	-0.04	-0.06	0.03	-0.01	0.02	0.00	-0.04	-0.01

Panel B: Senate elections

	Δ*Republican vote share in Senate election between...*									
	1996-02 (1)	1998-04 (2)	2000-06 (3)	1996-08 (4)	1998-10 (5)	2000-12 (6)	1996-14 (7)	1998-16 (8)	2000-18 (9)	1996-20 (10)
Log(SXSW followers, March 2007)	0.009	0.007	0.003	0.005	-0.008	0.010	-0.013	-0.009	-0.004	-0.008
	(0.011)	(0.014)	(0.010)	(0.010)	(0.012)	(0.014)	(0.014)	(0.016)	(0.010)	(0.009)
Observations	2,247	2,049	1,832	2,247	2,049	1,832	2,247	2,049	1,832	2,247
Mean of DV	0.01	-0.02	-0.06	-0.05	0.02	-0.06	0.02	-0.07	-0.10	-0.00

Notes: This table presents reduced form results, as in equation (2). For House elections in Panel A, the dependent variable is the change in the Republican vote share since 2000. For Senate elections in Panel B, the dependent variable is the change in the Republican vote share from six, twelve, or eighteen years ago (to accommodate senators' 6-year terms). The main independent variable is the number of users who started following SXSW in March 2007 (in logs with 1 added inside). All regressions control for the (log) number of SXSW followers that joined Twitter at some point in 2006, population deciles and Census region fixed effects and the full set of controls (as in columns 5 and 10 of Table 2). Regressions are weighted by turnout in the 2000 presidential election. Standard errors in parentheses are clustered by state. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

TABLE D.4. Twitter and the Republican vote share in swing and safe counties.

	Swing counties (1)	Republican counties (2)	Democratic counties (3)	Safe counties (4)
Panel A: ΔRepublican vote share 2000-2016				
Log(Twitter users)	-0.073***	-0.006	-0.008	-0.005
	(0.024)	(0.008)	(0.008)	(0.006)
Log(SXSW followers, Pre)	0.013	-0.001	0.000	-0.002
	(0.015)	(0.014)	(0.008)	(0.005)
Observations	716	1,990	358	2,348
Mean of DV	-0.033	0.021	-0.040	-0.012
Robust F-stat.	14.70	11.97	105.57	99.87
Panel B: ΔRepublican vote share 2000-2020				
Log(Twitter users)	-0.066***	-0.017**	-0.013	-0.007
	(0.019)	(0.007)	(0.009)	(0.006)
Log(SXSW followers, Pre)	0.006	-0.005	-0.001	-0.003
	(0.012)	(0.013)	(0.009)	(0.006)
Observations	716	1,990	358	2,348
Mean of DV	-0.027	0.026	-0.026	-0.002
Robust F-stat.	14.70	11.97	105.57	99.87

Notes: This table presents results estimated using 2SLS, as in equation (3). The dependent variable is the change in the vote share of the Republican party between the 2000 and 2016/2020 presidential elections in Panels A and B, respectively. *Swing counties* are those that were not consistently won by either Republicans or Democrats between 2000 and 2012; *Republican* and *Democratic* counties are those who voted consistently. Safe counties are the counties from columns (2) and (3) combined. *Log(Twitter users)* is instrumented using the number of users who started following SXSW in March 2007. *SXSW followers, Pre* is the number of SXSW followers who registered at some point in 2006. All regressions control for population deciles and Census region fixed effects and the full set of controls (as in columns 5 and 10 of Table 2). Regressions are weighted by turnout in the 2000 presidential election. Standard errors in parentheses are clustered by state. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

(a) Tweets about Republican Presidential Candidates

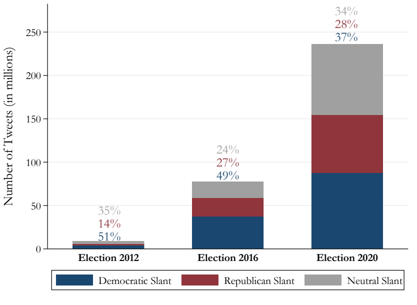

(b) Tweets about Democratic Presidential Candidates

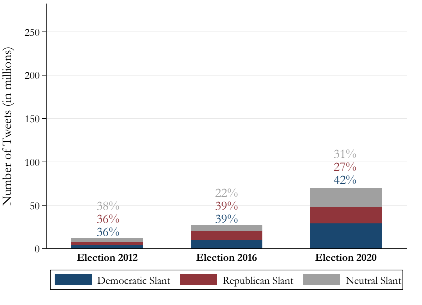

FIGURE D.1. Twitter's partisan slant (Tweet measure). These figures present the number of tweets (as opposed to the number of "likes" of such tweets in Figure 8) that contain the last name of the candidates in the 2012, 2016 and 2020 presidential elections, depending on whether the tweet was classified as having a Republican (instead of Democratic) slant. We classify the slant of a tweet based on the Twitter network of the user who sent the tweet. If the user follows more Democratic than Republican Congress members, they will be classified as a Democrat, and vice versa. Users who follow an equal number of Democrats and Republican or no Congress members are classified as neutral.

(a) Likes for Tweets about Republican Presidential Candidates

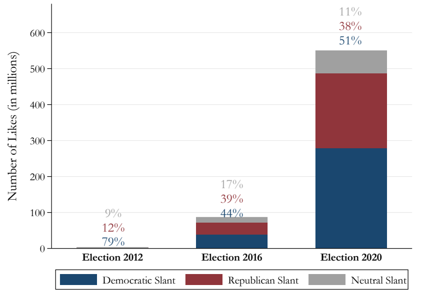

(b) Likes for Tweets about Democratic Presidential Candidates

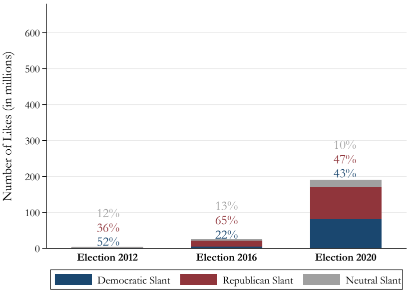

(c) Tweets about Republican Presidential Candidates

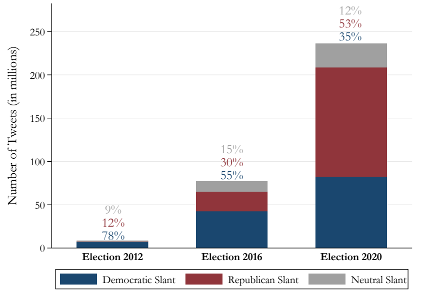

(d) Tweets about Democratic Presidential Candidates

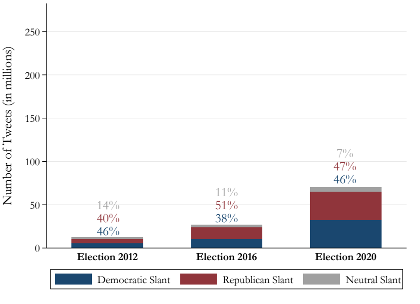

FIGURE D.2. Twitter's partisan slant (Text-Based Classifier). These figures present the number of "likes" received by tweets, or the number of tweets, that contain the last name of the candidates in the 2012, 2016 and 2020 presidential elections, depending on whether the tweet was classified as having a Republican (instead of Democratic) slant. We classify the slant of a tweet based on similarity in the language to that of a congressional Republican or Democrat, using a L2 regularized logistic regression classifier using the tweets sent by Congress members. Optimal normalization strength is chosen using 10-fold cross-validation. Tweets with a predicted class probability below 60% are coded as neutral. See Appendix A.1. for details.

(a) 2016 Presidential Election Results

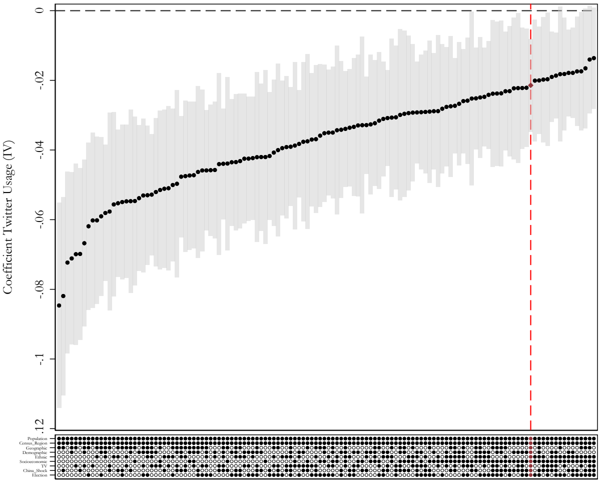

(b) 2020 Presidential Election Results

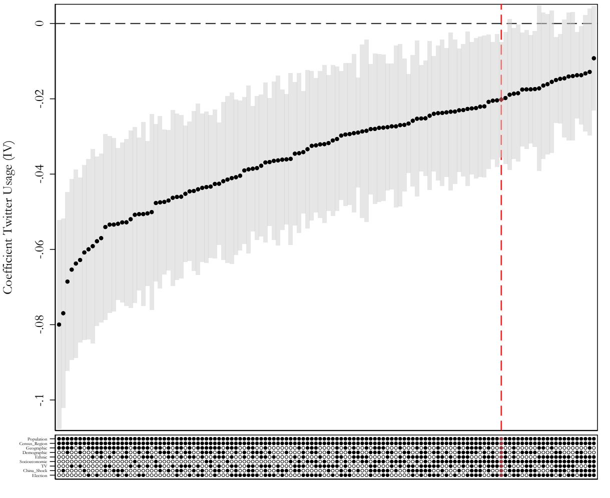

FIGURE D.3. Specification curve. These figures plot the 2SLS estimates and 95% confidence intervals from a regression of the Republican vote share in 2016 on $Log(\textit{Twitter users})$, instrumented with $\textit{SXSW}^{\textit{March 2007}}$. All regressions include population deciles, census region fixed effects, and $\textit{SXSW}^{\textit{Pre}}$. The combination of the other included control variables is shown at the bottom; filled circles mean a set of controls was included. The baseline specification with all controls is marked by the vertical line.

Appendix E: Additional Details on the Extrapolation for the Average Treatment Effect

Andrews and Oster (2019) show how selection into participating in an experiment can be used to make extrapolations regarding the external validity of an experiment. When a set of covariates **X** is observed for both the “experimental sample" and “population," Andrews and Oster (2019) provide a procedure that uses effect heterogeneity based on **X** estimated within the experimental sample to extrapolate to the average treatment effect for the “population."

We build on their procedure and argue that we can similarly use heterogeneity in the treatment effect within the counties that “identify" our results to extrapolate the treatment effect to all other counties in the US. Column (5) of Table C.2 show that we obtain similar estimates to our baseline when we only compare counties with SXSW followers that joined Twitter in March 2007 to counties with followers that joined in the pre-period, while excluding those counties in neither group. We can use this subsample of counties as the “experimental sample," and extrapolate effects to the “population" of all other counties.

Since Andrews and Oster (2019) approach is designed for a binary treatment, we adjust our regression framework by defining a treatment indicator variable equal to 1 for counties with SXSW followers who joined in March 2007 and 0 for the counties with followers that joined in the pre-period. We estimate the treatment effect for the subsample of counties that do not have zero SXSW followers in both periods using the regression specification $y_c = \alpha + \beta \cdot \mathbb{1}[SXSW_c^{March\ 2007} > 0] + \varepsilon_c$. The resulting treatment effect estimate is -0.075, which is similar Table 2 Panel B column (1). We then perform a linear prediction of this treatment effect based all observable variables in Table A.4 within this subsample. The resulting predicted treatment effect is -0.085. Last, we extrapolate the treatment effect for the rest of US counties. Based on the variation in observable characteristics we would predict an ATE of -0.184 for the US overall.

Note that this extrapolation is based on adjusting our reduced-form estimates to use a binary indicator variable for treatment thus the coefficients are not directly comparable to our baseline estimates. The approach further assumes quasi-random treatment assignment with in the counties with SXSW variation. Taken this into account, the extrapolation should therefore be viewed as suggestive, but confirming the notion that the effect for all US counties would be larger.

References

Andrews, Isaiah and Emily Oster (2019). “A Simple Approximation for Evaluating External Validity Bias.” *Economics Letters*, 178(C), 58–62.

Chernozhukov, Victor, Christian Hansen, and Martin Spindler (2015a). “Post-selection and post-regularization inference in linear models with many controls and instruments.” *American Economic Review*, 105(5), 486–90.

Chernozhukov, Victor, Christian Hansen, and Martin Spindler (2015b). “Valid post-selection and post-regularization inference: An elementary, general approach.” *Annu. Rev. Econ.*, 7(1), 649–688.

Colella, Fabrizio, Rafael Lalive, Seyhun Orcan Sakalli, and Mathias Thoenig (2019). “Inference with Arbitrary Clustering.” IZA Discussion Papers 12584, Institute of Labor Economics (IZA), URL https://EconPapers.repec.org/RePEc:iza:izadps:dp12584.

Gentzkow, Matthew and Jesse M Shapiro (2010). “What Drives Media Slant? Evidence from US Daily Newspapers.” *Econometrica*, 78(1), 35–71.

Kinder-Kurlanda, Katharina, Katrin Weller, Wolfgang Zenk-Möltgen, Jürgen Pfeffer, and Fred Morstatter (2017). “Archiving Information from Geotagged Tweets to Promote Reproducibility and Comparability in Social Media Research.” *Big Data & Society*, 4(2), 2053951717736336.

Pedregosa, F., G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay (2011). “Scikit-learn: Machine Learning in Python.” *Journal of Machine Learning Research*, 12, 2825–2830.

Quora (2011). “What is the process involved in launching a start-up at SXSW?”