Does Fake News Affect Voting Behavior? Evidence from Big College Football Games

Alden Cheng

March 11, 2025

Abstract

The issue of fake news has been hotly debated in recent years, with some claiming that it played a role in 2016 U.S. presidential elections. Despite these claims, there has been limited evidence to date linking fake news directly to voting behavior. In this project, I seek to provide credible evidence on this question by using big college football games as an instrument for fake news consumption, during the lead-up to the 2016 presidential election. During this period, I find that search volumes for pro-Trump fake news terms were lower in counties close to college football teams that played a big game when a piece of pro-Trump fake news went viral, and also that these counties were less likely to vote for Trump. The magnitude of these estimates suggests that a one-standard deviation increase in search volume for pro-Trump fake news terms increased Trump's vote share in 2016 by about 5.4–7.4 percentage points. In addition, I find that the relationship between fake news exposure and voting behavior is approximately linear, by developing and estimating a novel nonparametric IV estimation method.

I am very grateful to Alberto Abadie, David Autor, Amy Finkelstein, Thomas Wollmann and participants at MIT's Labor Lunch for their valuable comments.

1 Introduction

Fake news has become an increasingly hotly debated topic in recent years, with claims in the popular media that it may have influenced the results of the U.S. presidential election and the Brexit vote in 2016 (Jesdanun 2019). This debate over the spread of fake news has led to numerous changes, with Facebook and other social media platforms altering the way they present news information (e.g., removing newsfeed sidebars, and the labelling of certain stories on Facebook as "disputed"). These changes notwithstanding, some policymakers think that these do not go far enough, e.g., with Senator Elizabeth Warren calling for the breakup of Facebook (Kang and Kaplan 2019).

However, despite the perception that fake news have the potential to affect election results, there is limited evidence directly linking fake news to voting behavior. Moreover, recent polarization in politics (Gentzkow 2016) casts doubt on whether such a link exists. In particular, partisan voters may vote for their preferred candidate regardless and may be more susceptible to believe fake new stories favoring their candidate, whereas one might conjecture that swing voters may be more circumspect and are less likely to be fooled by fake news.

There are at least two reasons for the relative paucity of empirical evidence on this topic. First, any attempt to determine whether fake news has a causal effect on voting outcomes using observational estimates faces the threat of selection bias. If partisan voters are more likely to be exposed to fake news favoring their candidate (e.g., by consciously seeking out news favoring their candidate, or recommendations made by social media algorithms), then OLS estimates of the effect of fake news on voting behavior will be biased upwards. Second, most data on fake news consumption suffer from substantial measurement error, which will tend to attenuate OLS estimates towards zero. For instance, Allcott and Gentzkow (2017) find a high false recall rate for fake news stores in a survey they conducted: 15 percent of respondents recalled seeking a fake news story that surfaced around the election period, although 14 percent also recalled seeing a fake news story that was invented by the authors and never circulated.

In this paper, I study the effect of fake news on voting behavior in the 2016 U.S. presidential election using a novel empirical strategy. To deal with both selection bias from omitted variables and attenuation bias from measurement error, I employ an IV strategy that leverages exogenous variation in fake news consumption originating from big college football games.

In particular, the presence of a big college football game may reduce the exposure of voters living nearby to a piece of fake news going viral at the same time through two channels. First, the social media feed of voters in these areas may be dominated by sports-related news, crowding out concurrent fake news stories (Eisensee and Stromberg 2007). Second, these voters may substitute away from online activity towards in-person activities such as attending the game or watching it with friends (Dahl and DellaVigna 2009). At the same time, the exclusion restriction also seems plausible, given that scheduling of college football games occurs prior to the start of the season, at which time it would have been difficult to predict when fake news stories will break.

First stage IV results show that indeed, search volumes for pro-Trump fake news terms were lower in counties close to college football teams that played a big game shortly before the election, and while the reduced form results indicate that these counties were also less likely to vote for Trump. The magnitude of these estimates suggest that a one-standard deviation increase in search volume for pro-Trump fake news terms increased Trump's vote share by about 5.4–7.4 percentage points (percentage points).

This effect could have been the result of fake news changing voters' minds, or differentially influencing turnout for potential Trump and Clinton voters. While I can rule out that the effect of fake news on voting behavior is completely explained by third-party voters switching to Trump, estimate of the effect on turnout are too noisy to be informative. In addition, I do not find evidence that pro-Trump fake news had down-ballot effects.

This paper relates to several studies on whether fake news resulted in Trump winning the 2016 election. Boxell, Gentzkow, and Shapiro (2018) find suggestive evidence that fake news is unlikely to have played a major role in electing Trump. In particular, they find that Trump did disproportionately well among voters who are least likely to use the Internet, compared to previous Republican candidates. While this sheds light on voting behavior of specific subpopulations, the link between voting behavior and fake news is somewhat indirect. The present paper seeks to establish a more direct link by using direct measures of fake news. Gunther, Beck, and Nisbet (2018) reach a different conclusion, concluding that fake news increased support for Trump based on correlational evidence from a survey of Obama voters. However, one may be concerned that this study suffers from omitted variables bias, which my paper addresses using exogenous variation in exposure to fake news.

My paper is also related to a literature on behavioral economics and sports. This paper is similar in spirit to Eisensee and Stromberg (2007) and Tabakovic and Wollman (2019), in the sense that inattention to certain news due to sports events can have material consequences. In particular, Eisensee and Stromberg find that natural disasters that occurred during the Olympics or World series tended to receive less foreign aid, while Tabakovic and Wollman use unexpected NCAA football outcomes as an instrument for research support received by universities to study the effect of research expenditures on scientific productivity and downstream technology. Another paper that studies the behavioral consequences of football games is Card and Dahl (2011), who find that upset losses lead to an increase in violence by men against their wives and girlfriends.

This paper proceeds as follows. Section 2 gives a brief background on fake news and the 2016 election, and describes my data sources. I discuss my empirical strategy in section 3, before presenting results in section 4. Section 5 concludes.

2 Setting and Data

2.1 Fake News and the 2016 U.S. Presidential Election

The 2016 U.S. presidential election was a hotly contested one, and several pieces of fake news in the period leading up to the election were shared widely on Facebook (Silverman 2016). There are several key takeaways from Figure 1, which plots the number of Facebook shares of pro-Trump and pro-Clinton fake news stories over time using data from Allcott and Gentzkow (2017). First, we observe that the number of Facebook shares number in the millions for several pieces of fake news stories, suggesting that a non-trivial proportion of voters were exposed to fake news. Second, the vast majority of Facebook shares were for pro-Trump fake news stories, so I focus on the effect of pro-Trump fake news stories in my analysis

Third, the greatest number of high-profile fake news stories appeared from October onwards. Moreover, given that older pieces of fake news may fade from voters' mind and become less relevant by the time of the election, I focus on the period just before the election from October onwards. However, I exclude dates after October 28, so as to not conflate the effect of fake news with the effect of Comey's announcement on October 28 that the FBI was reopening its investigation of Clinton's emails.

Figure 1: Facebook Shares of Fake News Stories in the Lead-up to the 2016 U.S. Presidential Election

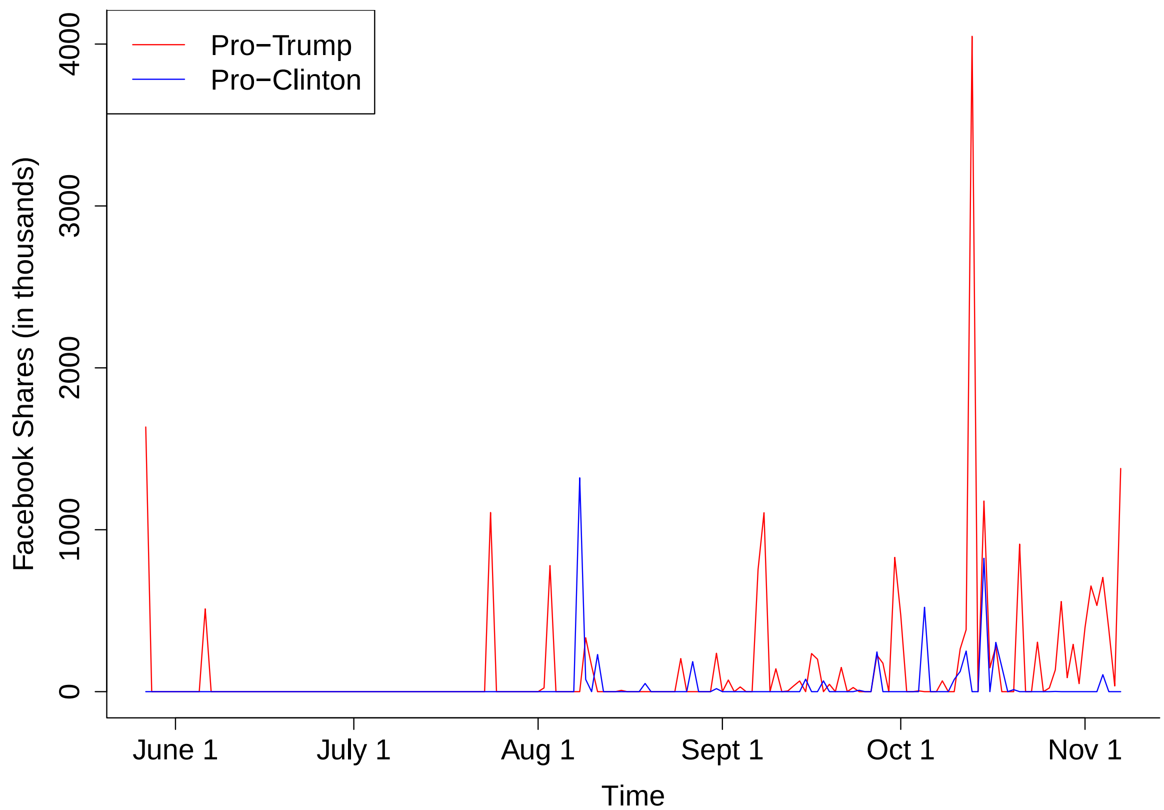

Notes: This figure shows the number of Facebook shares for various fake news stories favoring Trump and Clinton in the lead-up to the 2016 US presidential election.

2.2 Data

First, I use data on voting returns for the 2016 US presidential election, which is available at the county level. For this reason, I also conduct most of my analysis at the county level, weighting counties by their population.

Second, I collect data on college football games by hand. Specifically, I identify 25 big college football games according to an article by Sports Illustrated (Dellenger 2019). before searching for the county that each football team is located in, and merging this to the voting data. This procedure returns 18 unique counties that were home to a big college football game during my sample period, so only a relatively small number of counties are “treated” by this definition of the instrument. Hence, I also consider a definition of the instrument which includes neighboring counties. Figure 2a shows that the treated counties are fairly geographically spread out across the US.

Figure 2: Home Counties to Big Football Games in the Lead-up to the 2016 Election

(a) Home Counties to Big Football Games in the Lead-up to the 2016 Election

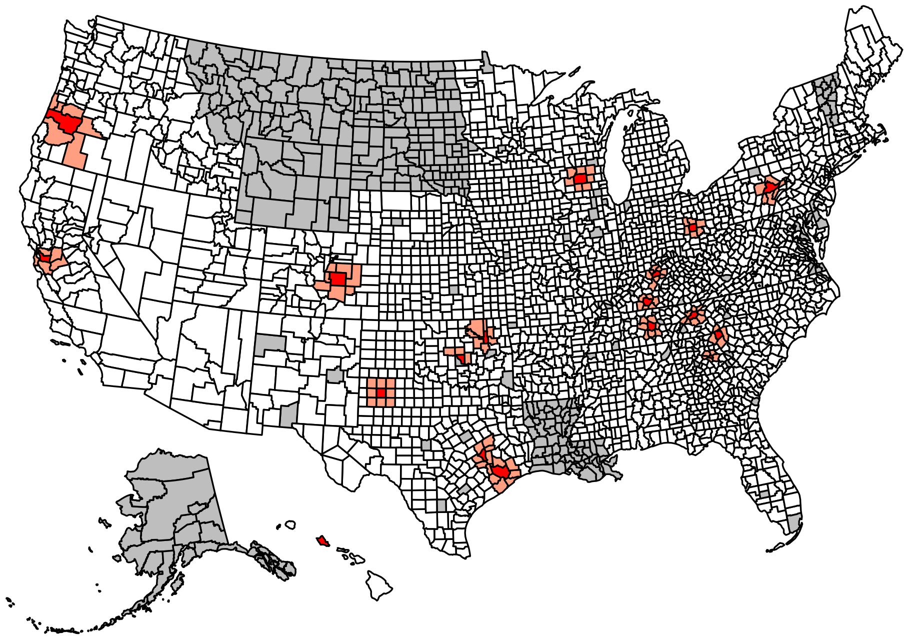

(b) Search Intensity for Fake News Terms in the Lead-up to the 2016 Election

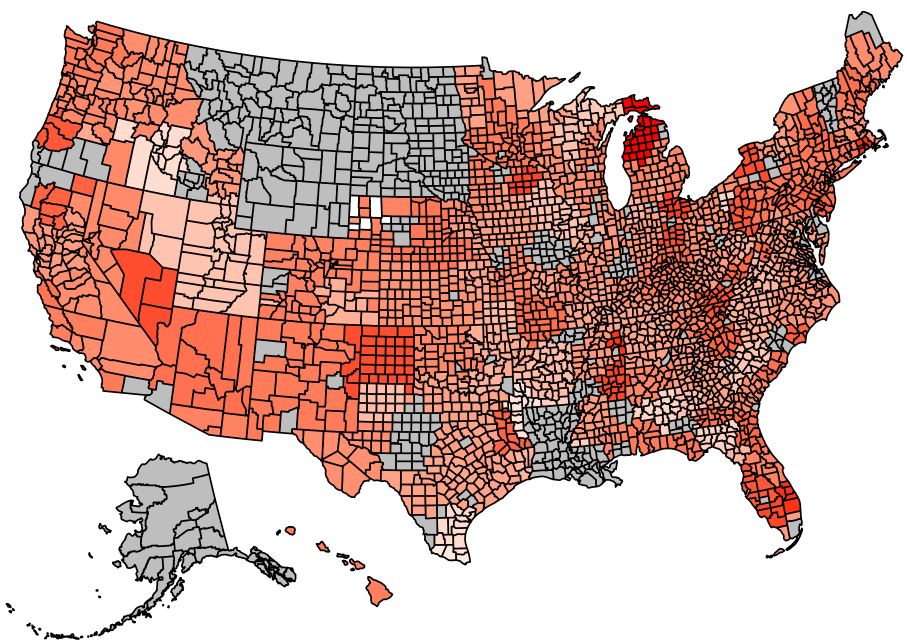

Notes: Panel (a) shows counties that were not included in the Google Trends results in grey, counties that hosted a big college football game in the lead-up to the 2016 presidential election at the time a major pro-Trump fake news story broke in red, adjacent counties in orange, and the remaining counties in white. Panel (b) shows a heat map of searches for pro-Trump fake news stories in the lead-up to the 2016 presidential election at the time a major pro-Trump fake news story broke, with darker shades of red indicating more searches, and grey indicating counties that were not included in the Google Trends results.

Third, I measure fake news consumption using Google Trends data, using relative search volumes for various fake news terms. I then merge the DMA-level search data with the county-level vote data. However, there are several difficulties with using the Google Trends data. One issue is that the search volume data is based on a sample of all searches and this sample changes, so I may get different results when I conduct the same search on different days. I deal with this by conducting multiple searches over different days.

Another issue is that there is a privacy threshold for the Google Trends data, so if there are not enough searches for the term in some DMA, Google Trends returns a missing value. In practice, this happens quite often for various fake news terms, since there are relatively few searches for these terms (even during the sample period). In order to mitigate the missing values problem, I use a procedure from Stephens-Davidowitz (2014) and Oster (2018):

1. Obtain the (relative) number of searches for a common term (call this term 1), e.g., chair.

2. Get the number of searches for the term of interest (call this term 2) or the common term.

3. Take the difference between the two:

$$[Searches\ for\ terms\ 1\ or\ 2] - [Searches\ for\ term\ 1].$$

I follow these procedures using several different common search terms, before averaging and standardizing the results to obtain the variable for fake news consumption. Figure

2b shows geographic variation in the search intensity for various fake news terms.

2.3 Summary Statistics

Table 1 presents summary statistics for counties, for the entire sample in column 1, and then for the subsamples of untreated and treated counties in columns 2 and 3 respectively (where I consider counties with big games and adjacent counties as treated). Finally, column 4 shows results from a t-test of equality of means for the two subsamples defined by the instrument.

We observe that unemployment rates in the two subsamples are very similar, but big-game counties tend to have larger populations. So, in my subsequent analyses, I weight counties by population.[1] However, the differences in vote shares at big-game counties and counties without big games nearby are statistically insignificant, which supports the notion that the instrument is as good as randomly assigned.

[1] Differences in vote shares in past elections tend to be positive, since Republican counties have smaller populations on average.

Table 1: Summary Statistics

	All	No Big Games	Big Games	Diff: (2) - (3)
	(1)	(2)	(3)	(4)
Unemployment in 2015 (percent)	5.471	5.490	5.022	0.5
	(1.937)	(1.961)	(1.199)	(0.2)
Unemployment in 2016 (percent)	5.173	5.193	4.722	0.5
	(1.823)	(1.846)	(1.123)	(0.2)
Population (thousands)	103.5	97.62	236.2	-138.6
	(334.9)	(324.0)	(509.0)	(30.5)
Difference in Vote Share in 2012 (percent)	21.50	21.39	23.93	-2.5
	(29.32)	(29.35)	(28.71)	(2.7)
Difference in Vote Share in 2008 (percent)	15.48	15.31	19.42	-4.1
	(27.43)	(27.43)	(27.40)	(2.5)
Difference in Vote Share in 2004 (percent)	21.84	21.69	25.06	-3.4
	(24.95)	(24.97)	(24.32)	(2.3)
Difference in Vote Share in 2000 (percent)	17.45	17.37	19.08	-1.7
	(23.60)	(23.67)	(22.14)	(2.2)
Number of Counties	2,955	2,830	125	-

Notes: Column 1 contains summary statistics for the full sample, whereas columns 2 (and 3) contains summary statistics for counties without (respectively, with) a big college football game nearby (i.e., within the county, or adjacent to the county) at the time a major fake news story broke. Column 4 shows the difference in means between columns 2 and 3, and the number in parenthesis corresponds to the t-test for equality of means. The unit of observation is a county. Difference in vote shares refers to the vote share of the main Republican presidential candidate minus the vote share of the main Democratic presidential candidate.

3 Empirical Strategy

Consider a voter i living in county j, and suppose that whether they decide to vote for Trump follows a linear probability model:

$$Y_i = \beta_0 + \beta_1 X_i + W_i' \gamma + \epsilon_i. \tag{1}$$

In this equation, the variable Y_i is an indicator for whether i votes for Trump, X_i is a measure of her exposure to pro-Trump fake news, W_i is a vector of the voter's characteristics, and ϵ_i is an idiosyncratic shock. The coefficient β_1 represents the causal effect of fake news on voter behavior.

Since I only have county-level vote data, I aggregate equation (1) to the county level:

$$\bar{Y}_j = \beta_0 + \beta_1 \bar{X}_j + \bar{W}_j' \gamma + \bar{\epsilon}_j. \tag{2}$$

As I mentioned earlier, it is likely to $Cov(X_i, \epsilon_i) > 0$, so the omitted variables bias if we estimate equation (2) using OLS will bias the estimate of β_1 upwards. On the other hand, $\bar{X}_j$ is likely to be measured with substantial error, so the attenuation bias will attenuate the estimate towards zero.

To address both of these issues, I instrument $\bar{X}_j$ with a dummy variable Z_j which is equal to one if there is a big college football game played in county j around the time a pro-Trump fake news breaks. The idea behind this instrument is that a big football game may affect voters' online activity, thereby influencing their exposure to fake news, as captured by the coefficient α_1 in the first stage for this IV in equation (3):

$$\bar{X}_j = \alpha_0 + \alpha_1 Z_j + \bar{W}_j' \delta + \eta_j. \tag{3}$$

We may expect a big football game to reduce nearby voters' exposure to concurrent fake news ($\alpha_1 < 0$) for two reasons. First, increased news content related to the football game may crowd out fake news in these voters' social media feeds (Eisensee and Stromberg 2007). Second, voters may substitute online activity for in-person activities, such as attending the game or watching the game with friends (Dahl and DellaVigna 2009).

The exclusion restriction is that shocks in voting behavior $\bar{\epsilon}_j$ are uncorrelated with the instrument Z_j. Equivalently, the instrument Z_j can only influence voting behavior $\bar{Y}_j$ through its effect on fake news consumption. One possible violation of this assumption is if the instrument Z_j is not randomly assigned, and counties near big football games have a greater tendency to vote for a particular party independent of exposure to fake news. Indeed, the summary statistics in Table 1 suggests that such counties tended to vote Republican in past elections, even though these differences are statistically insignificant. Hence, to address this, I include controls for voting behavior in past presidential elections in my IV specifications. Moreover, even if controls for past voting behavior fail to eradicate the link between big-game counties and underlying tendency to vote Republican, the direction of the 2SLS estimate will be biased downwards, so that we would underestimate the effect of fake news on voting behavior.

Another possibility is that the presence of big games affects behavior other than fake news consumption. For example, perhaps individuals in these counties substitute away from TV news shows towards sports channels on such nights. Nonetheless, it seems relatively unlikely that watching the news to a lesser degree for a day or two would substantially affect voting behavior, unless there were some other events in the news important enough to sway voting that were occurring contemporaneously with the breakout of fake news stories.

Finally, relatively few counties had a big college football game around the time a pro-Trump fake news story broke, so we consider another definition of the instrument Z_j which includes adjacent counties. This is also related to potential concerns about measurement error in the instrument. Nonetheless, if we model the measurement errors in Z_j and fake news searches as classical and assume that they are uncorrelated, the IV estimates will not be affected by attenuation bias.[2]

4 Results

4.1 Estimates on the Effect of Fake News on Voting Behavior

Table 2 shows results for my IV strategy, with the three panels corresponding to the first stage, reduced form, and 2SLS results respectively. In all regressions, I control

[2]An easy way to see this is to note that the IV estimate is given by the Wald ratio, and the attenuation factor for the coefficient on the instrument in the reduced form and first stage equations are the same, and will thus cancel out.

for Republican vote shares in the past four elections, and in some specifications I also control for the county unemployment rate in 2015 to increase precision (since voters tend to punish incumbent parties for perceived poor economic performance).

The first stage estimates in panel A show that searches for fake news terms are lower in big-game counties, consistent with the crowdout or substitution effects (Eisensee and Stromberg 2007; Dahl and DellaVigna 2009) described earlier, while the reduced form estimates in panel B indicate that voters in big-game counties were also less likely to vote for Trump, with the broader definition of the instrument resulting in more precise estimates. Finally, the 2SLS estimates in panel C suggest that a one-standard deviation increase in search volume for pro-Trump fake news terms increased Trump's vote share by at least between 5.4 and 7.4 percentage points.

Table 2: IV Estimates of the Effect of Fake News on Trump Vote Share

	Standardized Searches for Fake News			
Panel A: First Stage	(1)	(2)	(3)	(4)
Big-Game County	-0.356**	-0.348**		
	(0.161)	(0.161)		
Big-Game or Adjacent County			-0.345**	-0.341**
			(0.149)	(0.149)
Vote Shares in Past 4 Elections	X	X	X	X
Control for 2015 Unemployment Rate		X		X
Number of Observations	2,747	2,747	2,747	2,747

Notes: Observations are weighted by the total population of the county. Robust standard errors clustered at the DMA level are shown in parentheses. *** p<0.01, ** p<0.05, * p<0.1

	Trump Vote Share			
Panel B: Reduced Form	(1)	(2)	(3)	(4)
Big-Game County	-0.026*	-0.019		
	(0.014)	(0.016)		
Big-Game or Adjacent County			-0.025**	-0.020*
			(0.010)	(0.011)
Vote Shares in Past 4 Elections	X	X	X	X
Control for 2015 Unemployment Rate		X		X
Number of Observations	2,747	2,747	2,747	2,747

Notes: Observations are weighted by the total population of the county. Robust standard errors clustered at the DMA level are shown in parentheses. *** p<0.01, ** p<0.05, * p<0.1

	Trump Vote Share			
Panel C: Two-Stage Least Squares	(1)	(2)	(3)	(4)
Standardized Searches for Fake News Terms	0.074***	0.054*	0.073***	0.058**
	(0.028)	(0.032)	(0.026)	(0.025)
Anderson-Rubin CI	[0.023, 0.15]	[-0.046, 0.11]	[0.043, 0.16]	[0.012, 0.13]
Instrument Definition: Big-Game Counties Only	X	X		
Instrument Definition: Includes Adjacent Counties			X	X
Controls for Vote Shares in Past Elections	X	X	X	X
Control for 2015 Unemployment Rate		X		X
F-Statistic	4.86	4.66	5.39	5.24
Number of Observations	2,747	2,747	2,747	2,747

Notes: Observations are weighted by the total population of the county. Robust standard errors clustered at the DMA level are shown in parentheses. *** p<0.01, ** p<0.05, * p<0.1

While the first stage estimates are statistically significant at the 5 percent significance level, the we observe that the F-statistics are less than 10, suggesting that weak instruments may be an issue. Hence, in panel C, I also show Anderson-Rubin confidence intervals, which are robust to weak instruments (Anderson and Rubin 1949; Andrews, Stock and Sun 2019). In three of the four specifications, the 90 percent confidence intervals only contain positive values, and in all cases, they are centered on a positive value.[3]

We can compare the IV estimates in Table 2 with the OLS estimates in Appendix Table B.1. While the OLS estimates are positive and statistically significant, the points estimates indicate that a one-standard deviation increase in search for fake news terms is associated with a 1.1 percentage point increase in Trump vote share, which is much smaller than the IV estimates of a 5.4–7.4 percentage point effect. This difference between the OLS and IV estimates can be explained by attenuation bias in the OLS estimates due to measurement error in searches for fake news terms.

4.2 Mechanisms and Down-Ballot Effects

Next, we consider potential mechanisms through which fake news may have affected vote shares. First, it could have been due to voters changing their vote from Clinton or a third-party candidate to Trump, or it could have resulted in higher (respectively, lower) turnout among voters who would more likely have voted for Trump (Clinton).

To gain some insight into this, I estimate the same IV specifications as above, but replacing Trump's vote share with Clinton's vote share and the difference between the Trump and Clinton vote shares in Appendix Tables B.2 and B.3 respectively. These results suggest that higher fake news searches also resulted in a lower Clinton vote share, and the sum of the effects on Trump and Clinton vote shares is roughly equal to the estimated effect on the difference in vote shares. This indicates that the increase

[3] Note that in general the Anderson-Rubin confidence intervals are not centered on the 2SLS estimates.

in Trump's vote share due to fake news exposure is unlikely to have been solely due to potential third-party voters switching to Trump.

To distinguish between the effect of fake news on turnout and switching from Clinton to Trump, I estimate the effect of fake news on turnout using the same empirical strategy, but replacing vote shares with turnout. Unfortunately, the estimates in Appendix Table B.4 show that the effects on turnout are too noisily estimated to be informative.

Finally, I consider the possibility of down-ballot effects. In particular, given our finding the voters who consumed more pro-Trump fake news are more likely to vote for Trump, it seems plausible that these voters may also be more likely to vote for other Republican candidates in non-presidential races. To explore this possibility, I consider House elections and estimate the effect of fake news on vote shares for Republican House candidates using the same IV specifications. The IV and OLS estimates on the effect of pro-Trump fake news on Republican house vote share in Appendix Table B.5 are not statistically significant, and the standard errors for the IV estimates are rather large.

4.3 Nonlinear Treatment Effects

Previously, in equation (2), we assumed that the effect of fake news on Trump's vote share is linear. However, if an outsized share of voters in a county are exposed to pro-Trump fake news, it could plausibly lead to groupthink, in which case the relationship between fake news and vote share would be convex.

Hence, to study nonlinear treatment effects, I consider a more general model, given by:

$$\bar{Y}_j = g(\bar{X}_j) + \bar{W}_j'\gamma + \bar{\epsilon}_j,$$

where the relationship between searches and voting behavior is given by the potentially nonlinear function $g(\cdot)$, we assume $\bar{\epsilon}_j \perp Z_j$, and the treatment effect is given by the derivative $g'(\bar{X}_j)$. This setup may seem somewhat reminiscent of the literature on marginal treatment effects (Heckman and Vytlacil 2005), but instead of a dummy treatment variable and a continuous instrument, in this setting, we have a continuous treatment variable and a dummy instrument.

To elaborate on the identification of $g'(\cdot)$, for each value x in the support of $\bar{X}_j$, define a new dummy variable $D_j(x) \equiv \mathbb{I}[\bar{X}_j \geq x]$, and suppose that there exists coefficients $\alpha(x)$ and $\beta(x)$ such that:

$$Y = \alpha(x) + \beta(x)D(x) + \bar{W}_j'\gamma + u(x), \mathbb{E}[u(x)|Z] = 0,$$

or equivalently, that satisfy:

$$\mathbb{E}[g(X) - \alpha(x) - \beta(x)D(x) - \bar{W}_j'\gamma | Z = 1] - \mathbb{E}[g(X) - \alpha(x) - \beta(x)D(x) - \bar{W}_j'\gamma | Z = 0] = 0.$$

In this case, for any given x, we can estimate $\alpha(x)$ and $\beta(x)$ using 2SLS. We can then show that the relationship between $\beta(x)$ and $g'(x)$ is given by:

$$g'(x) = \left(f_{X|Z=1}(x) + f_{X|Z=0}(x)\right)\beta(x),$$

where $f_{X|Z}$ is the density function of $\bar{X}_j$ conditional on Z (weighted by population). A formal statement and proof of this result is given in Appendix Section A.

Figure 3 shows estimates of the marginal effect $g'(\bar{X}_j)$ following this approach, as well as pointwise 95 percent confidence intervals. The results in this figure support the linear specification in equation (2): indeed, we observe that the marginal effect of fake news on voting behavior ($g'(x)$) is approximately constant throughout most of the distribution of counties' fake news search intensity (x).

Figure 3: Marginal Effect of Fake News on Voting Behavior

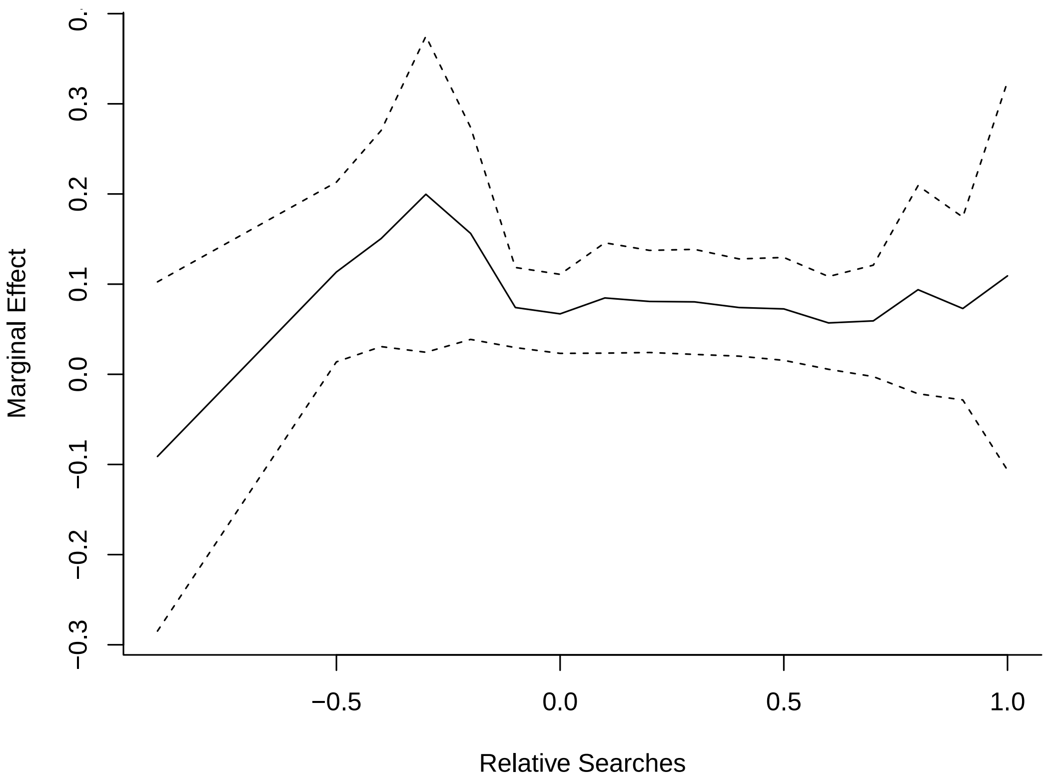

Notes: This figure shows the marginal effects of fake news on republican vote share with a solid black line, and 95 percent confidence intervals shown as dashed black lines. The instrument used is counties with big football games and adjacent counties, and the controls include Republican vote shares in 2000, 2004, 2008, and 2012, as well as unemployment rates in 2015. Standard errors from the IV regressions are clustered at the DMA level, and the variance for the kernel density estimates is obtained using nonparametric bootstrap.

5 Conclusion

This paper studies the effect of fake news consumption on voting behavior in the context of the 2016 U.S. presidential election. Leveraging quasi-experimental variation in fake news consumption induced by big college football games, I find that a one standard deviation increase in searches for pro-Trump fake news stories led to a 5.4–7.4 percentage points increase in Trump's vote share. More broadly, a similar empirical strategy can be used in future studies that seek to quantify how individuals' awareness of specific news events may affect their behavior, and this paper also provides a way to increase estimate potentially nonlinear treatment effects in an IV setting with a continuous treatment variable and dummy instrumental variable.

References

[1] Anderson, T.W., and Herman Rubin. "Estimation of the Parameters of a Single Equation in a Complete System of Stochastic Equations." *The Annals of Mathematical Statistics*, 20(1): 46–63.

[2] Andrews, Isaiah, James H. Stock, and Liyang Sun, 2019. "Weak instruments in instrumental variables regression: Theory and practice." *Annual Review of Economics*, 11:727-753.

[3] Allcott, Hunt, and Matthew Gentzkow, 2017. "Social Media and Fake News in the 2016 Election." *Journal of Economic Perspectives*, 31(2): 211–236.

[4] Boxell, Levi, Matthew Gentzkow, and Jesse M. Shapiro, 2018. "A note on internet use and the 2016 U.S. presidential election outcome." *PLoS ONE* 13(7): e0199571.

[5] Card, David, and Gordon B. Dahl, 2011. "Family violence and football: The effect of unexpected emotional cues on violent behavior." *The Quarterly Journal of economics*, 126(1): 103–143.

[6] Dahl, Gordon, and Stefano DellaVigna, 2009. "Does movie violence increase violent crime?." *The Quarterly Journal of Economics*, 124(2): 677-734.

[7] Dellenger, Ross, 2019. "The Greatest Rivalries in College Football History." *Sports Illustrated*, August 12, 2019. from https://www.si.com/college/2019/08/12/best-rivalries-college-football-history.

[8] Eisensee, Thomas, and David Strömberg, 2007. "News droughts, news floods, and US disaster relief." *The Quarterly Journal of Economics*, 122(2): 693–728.

[9] James J. Heckman and Edward J. Vytlacil, 2005. "Structural equations, treatment effects, and econometric policy evaluation". *Econometrica*, 73(3): 669–738.

[10] Jesdanun, Anick. 2019. "Did Facebook Data Help Trump? 'Great Hack' Explores Scandal." *Associated Press*, July 24, 2019. https://apnews.com/article/5537b5caa9644cd1bb3b041affc44682.

[11] Kang, Cecilia, and Thomas Kaplan, 2019. "Warren Dares Facebook With Intentionally False Political Ad." *New York Times*, October 12, 2019. https://www.nytimes.com/2019/10/12/technology/elizabeth-warren-facebook-ad.html.

[12] Oster, Emily, 2018. "Does disease cause vaccination? Disease outbreaks and vaccination response." *Journal of Health Economics*, 57: 90–101.

[13] Silverman, Craig. 2016. "This Analysis Shows How Viral Fake Election News Stories Outperformed Real News on Facebook." BuzzFeed News, November 16, 2016. https://www.buzzfeednews.com/article/craigsilverman/viral-fake-election-news-outperformed-real-news-on-facebook.

[14] Stephens-Davidowitz, S., 2014. "The cost of racial animus on a black candidate: evidence using Google search data." *Journal of Public* Economics, 118: 26–40.

[15] Tabakovic, Haris, and Thomas G. Wollmann, 2019. "The impact of money on science: Evidence from unexpected NCAA football outcomes." *Journal of Public Economics*, 178: 104066.

Appendix

A Nonparametric IV

Consider the following nonparametric IV model:

$$Y = g(X) + U,\ \mathbb{E}[U|Z] = 0,$$

where X is a univariate continuous random variable with (potentially infinite) support $(\underline{x},\bar{x})$, and Z is a dummy variable. The function $g(x)$ is the CEF, and we are assuming that the treatment effect is constant for individuals with each value of X. Now, for each value x in the support of X, define a new dummy variable $D(x) \equiv \mathbb{I}[X \geq x]$, and suppose that there exists coefficients $\alpha(x)$ and $\beta(x)$ such that:

$$Y = \alpha(x) + \beta(x)D(x) + \epsilon(x), \mathbb{E}[\epsilon(x)|Z] = 0,$$

or equivalently, that satisfy:

$$\mathbb{E}[g(X) - \alpha(x) - \beta(x)D(x)|Z = 1] - \mathbb{E}[g(X) - \alpha(x) - \beta(x)D(x)|Z = 0] = 0.$$

Then, we can estimate $\beta(x)$ using 2SLS.

Consider the assumptions:

1. First stage assumption: $Cov(Z, D(x)) > 0$ for all x.

2. Exclusion restriction: $Cov(Z, \epsilon(x)) = 0$ for all x.

3. Y, X, and Z have finite first two moments.

4. $g(x)$ is continuously differentiable for all x and X has density f_X.

5. $f_X(x)$ is strictly positive and continuous on its support.

Proposition 1. *Under Assumptions 1–5, $\alpha(x)$, $\beta(x)$, and $g'(x)$ are identified, with:*

$$g'(x) = \big(f_{X|Z=1}(x) + f_{X|Z=0}(x)\big)\,\beta(x). \tag{4}$$

Proof. Under Assumptions 1–3, $\alpha(x)$ and $\beta(x)$ are identified as estimates from a two-stage least squares regression with Y as the outcome, and Z as the instrument for $D(x)$. So, we have:

$$\beta(x) = \frac{\mathbb{E}[Y|Z=1] - \mathbb{E}[Y|Z=0]}{\mathbb{E}[D(x)|Z=1] - \mathbb{E}[D(x)|Z=0]} = \frac{\mathbb{E}[g(X)|Z=1] - \mathbb{E}[g(X)|Z=0]}{\Delta(x)},$$

where $\Delta(x) \equiv \mathbb{E}[D(x)|Z=1] - \mathbb{E}[D(x)|Z=0] = Pr(X \geq x|Z=1) - Pr(X \geq x|Z=0)$.

Next, to relate $\beta(x)$ to $g'(x)$, we consider local perturbations in $\mathbb{E}[g(X)|Z=z]$ and $\Delta(x)$ around x. In particular, we have:

$$\begin{aligned}
\beta(x) &= \lim_{\delta \to 0^+} \frac{\left[g(x) + \frac{\delta}{2} g'(x)\right] - \left[g(x) - \frac{\delta}{2} g'(x)\right]}{\left[\delta f_{X|Z=1}(x)\right] - \left[-\delta f_{X|Z=0}(x)\right]} \\
&= \frac{g'(x)}{f_{X|Z=1}(x) + f_{X|Z=0}(x)}.
\end{aligned}$$

Rearranging terms, we obtain the expression for $g'(x)$ in equation (4), as desired. □

Note that this implies that $g(x)$ is identified as well, since we can simply integrate $g'(x)$ and use $\mathbb{E}[g(X)] = \mathbb{E}[Y]$ to obtain the constant of integration.

Online Appendix

B Appendix Figures and Tables

Table B.1: OLS Estimates of the Effect of Fake News on Trump Vote Share

	Trump Vote Share		Clinton Vote Share		Difference in Vote Share	
	(1)	(2)	(3)	(4)	(5)	(6)
Standardized Searches for Fake News Terms	0.011**	0.011**	-0.006*	-0.006*	0.018**	0.018**
	(0.006)	(0.006)	(0.003)	(0.003)	(0.009)	(0.009)
Controls for Vote Shares in Past 4 Elections	X	X	X	X	X	X
Control for 2015 Unemployment Rate		X		X		X
Number of Observations	2,747	2,747	2,747	2,747	2,747	2,747

Notes: Observations are weighted by the total population of the county. Robust standard errors clustered at the DMA level are shown in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table B.2: IV Estimates of the Effect of Fake News on Clinton Vote Share

	Standardized Searches for Fake News			
Panel A: First Stage	(1)	(2)	(3)	(4)
Big-Game County	-0.360**	-0.351**		
	(0.162)	(0.163)		
Big-Game or Adjacent County			-0.351**	-0.347**
			(0.150)	(0.152)
Controls for Dem. Vote Shares in Past 4 Elections	X	X	X	X
Control for 2015 Unemployment Rate		X		X
Number of Observations	2,747	2,747	2,747	2,747

Notes: Observations are weighted by the total population of the county. Robust standard errors clustered at the DMA level are shown in parentheses. *** p<0.01, ** p<0.05, * p<0.1

	Clinton Vote Share			
Panel B: Reduced Form	(1)	(2)	(3)	(4)
Big-Game County	0.017	0.012		
	(0.017)	(0.018)		
Big-Game or Adjacent County			0.017	0.013
			(0.012)	(0.013)
Controls for Dem. Vote Shares in Past 4 Elections	X	X	X	X
Control for 2015 Unemployment Rate		X		X
Number of Observations	2,747	2,747	2,747	2,747

Notes: Observations are weighted by the total population of the county. Robust standard errors clustered at the DMA level are shown in parentheses. *** p<0.01, ** p<0.05, * p<0.1

	Clinton Vote Share			
Panel C: Two-Stage Least Squares	(1)	(2)	(3)	(4)
Standardized Searches for Fake News Terms	-0.048	-0.034	-0.049*	-0.039
	(0.035)	(0.040)	(0.028)	(0.030)
Anderson-Rubin CI	[-0.1, 0.06]	[-0.081, 0.13]	[-0.109, 0.01]	[-0.093, 0.04]
Instrument Definition: Big-Game Counties Only	X	X		
Instrument Definition: Includes Adjacent Counties			X	X
Controls for Dem. Vote Shares in Past 4 Elections	X	X	X	X
Control for 2015 Unemployment Rate		X		X
F-Statistic	4.95	4.66	5.44	5.21
Number of Observations	2,747	2,747	2,747	2,747

Notes: Observations are weighted by the total population of the county. Robust standard errors clustered at the DMA level are shown in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table B.3: IV Estimates of the Effect of Fake News on Difference Between Trump and Clinton Vote Shares

Panel A: First Stage	Standardized Searches for Fake News (1)	(2)	(3)	(4)
Big-Game County	-0.361**	-0.355**		
	(0.162)	(0.163)		
Big-Game or Adjacent County			-0.352**	-0.349**
			(0.150)	(0.151)
Controls for Diff. Vote Shares in Past 4 Elections	X	X	X	X
Control for 2015 Unemployment Rate		X		X
Number of Observations	2,747	2,747	2,747	2,747

Notes: Observations are weighted by the total population of the county. Robust standard errors clustered at the DMA level are shown in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Panel B: Reduced Form	Difference Between Trump and Clinton Vote Shares (1)	(2)	(3)	(4)
Big-Game County	-0.040	-0.029		
	(0.029)	(0.031)		
Big-Game or Adjacent County			-0.040*	-0.032
			(0.022)	(0.022)
Controls for Diff. Vote Shares in Past 4 Elections	X	X	X	X
Control for 2015 Unemployment Rate		X		X
Number of Observations	2,747	2,747	2,747	2,747

Notes: Observations are weighted by the total population of the county. Robust standard errors clustered at the DMA level are shown in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Panel C: Two-Stage Least Squares	Difference Between Trump and Clinton Vote Shares (1)	(2)	(3)	(4)
Standardized Searches for Fake News Terms Fake News Terms	0.112*	0.083	0.114**	0.092*
	(0.058)	(0.066)	(0.048)	(0.049)
Anderson-Rubin CI	[-0.031, 0.24]	[-0.148, 0.18]	[0.025, 0.23]	[-0.015, 0.2]
Instrument Definition: Big-Game Counties Only	X	X		
Instrument Definition: Includes Adjacent Counties			X	X
Controls for Diff. Vote Shares in Past 4 Elections	X	X	X	X
Control for 2015 Unemployment Rate		X		X
F-Statistic	4.93	4.74	5.51	5.35
Number of Observations	2,747	2,747	2,747	2,747

Notes: Observations are weighted by the total population of the county. Robust standard errors clustered at the DMA level are shown in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table B.4: Effect of Fake News on Turnout

	Voter Turnout (Percent)					
	IV	IV	IV	IV	OLS	OLS
	(1)	(2)	(3)	(4)	(5)	(6)
Standardized Searches for Fake News Terms	-34.549	-65.571	-2.342	-1.177	0.147	0.142
	(270.574)	(765.855)	(7.814)	(6.487)	(0.182)	(0.180)
Instrument Definition: Big-Game Counties Only	X	X				
Instrument Definition: Includes Adjacent Counties			X	X		
Control for Turnout in 2012 Election	X	X	X	X	X	X
Control for 2015 Unemployment Rate		X		X		X
Number of Observations	2,612	2,612	2,612	2,612	2,612	2,612

Notes: Observations are weighted by the total population of the county. Robust standard errors clustered at the DMA level are shown in parentheses.
*** p<0.01, ** p<0.05, * p<0.1

Table B.5: Down-Ballot Effects of Fake News

	Republican House Vote Share					
	IV	IV	IV	IV	OLS	OLS
	(1)	(2)	(3)	(4)	(5)	(6)
Standardized Searches for Fake News Terms	-0.010	-0.022	-0.064	-0.065	-0.007	-0.007
	(0.106)	(0.112)	(0.191)	(0.200)	(0.011)	(0.011)
Instrument Definition: Football County Only	X	X				
Instrument Definition: Includes Adjacent Counties			X	X		
Controls for Rep. Vote Shares in Past 4 Elections	X	X	X	X	X	X
Control for 2015 Unemployment Rate		X		X		X
Number of Observations	3,098	3,098	3,098	3,098	3,098	3,098

Notes: Observations are weighted by the total population of the county. Robust standard errors clustered at the DMA level are shown in parentheses.
*** p<0.01, ** p<0.05, * p<0.1