

The Impact of Media on Political Polarization: Evidence from the United States *

September 29, 2022

JEL Codes: C15, E60, H12, I18.

Keywords: Political Polarization, text-analysis, machine learning

Acknowledgement: I am grateful for the suggestions and insights offered by Salvador Ortigueira, Jinhui Bai, and Jia Yan.

*Michael Mahoney

WSU School of Economic Sciences. Pullman, WA 99163

Email: michael.mahoney@wsu.edu

Phone: (509) 551 8711

ABSTRACT:

Over the last thirty years, political polarization has risen sharply in the United States. Yet, it remains unclear the extent to which political polarization can be explained by changes in the American media environment. In this analysis, we quantify both political and media polarization using millions of lines of text data from political reporting and the proceedings of Congress. We utilize the estimator proposed in Gentzkow, Shapiro, and Taddy (2019) to construct an index representing the divergence in speech patterns between Fox News and CNN; we build a similar index to measure political polarization using speech data captured on the floor of Congress. We next use an Instrumental Variables approach to demonstrate that changes in media polarization causally effect changes in political polarization.

1 Introduction and Literature Review

That political polarization has risen in the United State over the last thirty years is undeniable. A host of theories exist to explain the phenomena, but there has been no comprehensive, quantitative analysis of the role media plays in driving political division. In this analysis we assess the causal impact of media on political polarization. We find a strong causal effect of media polarization on political polarization. Indeed, the relationship is highly significant and robust to the inclusion of economic and political controls, suggesting that changes in media coverage profoundly affect the political environment.

In the subsequent analysis, we apply a novel textual data set to the model developed in Gentzkow, Shapiro, and Taddy (2019). Their model allows us to quantify media polarization and political polarization in each month between September 2013 and September 2021 using text data. Our text data is divided into two components. The first is drawn from articles published by Fox News and CNN. The second is drawn from the proceedings of Congress. We pivot these raw text data into two matrices whose columns represent counts of phrases used by media personalities and members of Congress. We then use this count information to construction a measure of dispersion in language use between each group pair – Fox and CNN and Democrats and Republicans. This metric is our measure of polarization.

It is important to note the advantages of this approach over a more conventional measure of partisanship, such a roll-call votes. First, this technique allows us to determine which phrases are most polarizing, a result we include in the appendix. Second, this approach allows us to measure polarization in media, a context where there is no tangible measure of partisanship.

Once we have constructed our two indices, we conduct a causal analysis between the two polarization series. We use a two stage estimation approach, instrumenting media polarization with both television viewing hours and measures of coverage satisfaction reported by viewers of each media outlet. This analysis yields a positive, highly significant effect of media polarization on political polarization. The result is robust to the addition of economic and political controls.

The remainder of the paper is structured as follows: we first review the existing literature on political and media polarization, specifically examining

papers that utilize language processing techniques. We then elaborate on our data, and offer an expansive discussion of our collection and refinement processes. In section three, we introduce our model of speech and our measure of partisanship. Finally, we assess causality between our two series using an instrumental variables estimator. We discuss the validity of our instruments and present results.

1.1 Literature Review

There is a wealth of literature concerning political polarization. The question of the media’s influence on politics is less well studied, thus creating a gap in the existing literature. Below we examine a handful of key papers in this field and show that our analysis represents a unique contribution to the study of American politics.

Gentzkow, Shaprio, and Taddy (2019) measure political polarization across two-year sessions of congress from 1870 to 2016. They find heightened political polarization following the Civil War, a marked decrease in partisanship for much of the twentieth century, and then a sharp uptick in divisiveness beginning in the 1990s. In our analysis, we adopt their methodology with a few key modifications. First, we estimate our polarization series on a monthly basis to produce a much finer data set. Second, we use their approach to quantify media polarization, not just political partisanship.

Goet (2019) uses a similar approach to measure partisanship in the UK House of Commons over the period 1811-2015. His analysis provides substantial motivation for the estimator we present below: it affords us a method to measure polarization that does not require subjective classification. That is, our model is unsupervised (to borrow a phrase from computer science).

Prior (2013) examines the role that partisan media plays in driving political polarization and finds results quite different from our own. Namely, he finds no causal link between partisan media exposure and upticks in political polarization, concluding that most American politicians are fundamentally moderate. We find two issues with Prior’s analysis. First, there is no comprehensive statistical model. Prior’s analysis is a qualitative synthesis of polling data, not a properly designed econometric study.

Our second issue with Prior’s analysis is its timeframe, a problem admit-

tedly outside the author’s control. Prior examines the media/political environment of the late 1990s and early 2000s, a time period far removed from our analysis. It is thus possible that Prior’s results are correct for his time period. It is further possible that the effect media plays in political division has itself changed over time.

Waller and Anderson (2021) measure political polarization on Reddit (i.e. the dumbest place on the internet) during the period 2007-2021. They find that Reddit’s level of polarization responds to external events, with the 2016 election driving a large increase in partisanship – a result that fits neatly with our conclusions. They further find that much of the increase in divisiveness can be attributed to right-wing commentators operating on a few specific forums.

A more recent work by Ballard et. al. estimates partisanship in Congress using Tweets sent by congressmen between 2009 and 2020. They find that members of congress in “safe” districts tend to send the most polarizing tweets. There may be a connection to our analysis, if these “safe” districts tend to favor one news outlet over another.

A similar analysis by Pozen et. al. uses a large text dataset and natural language processing techniques to analyze constitutional debates between 1873 and 2016. They find a substantial divergence in language use surrounding these legal over the last forty years. Further, they find that much of this divergence has been driven by conservative politicians.

While the above papers use a host of techniques to quantify partisanship, only Prior (2013) attempts to assess the causal impact media plays in determining political polarization. As mentioned above, his analysis is lacking quantitative rigor and examines a time period very much unlike our own. The effect of media on political outcomes is therefore still an open questions, one we endeavor to resolve in the analysis below.

2 Data

Our raw data set consists of text drawn from 50,000 thousand documents published by Fox News and CNN between 2013 and 2021. From CNN we extract 16 million words written in opinion pieces and political reporting. From Fox News we collect another 13 million words, also from political and opinion pieces. For both sets of documents, we identify the month of publication and

the author.

In addition, we text mine the Congressional Record, a collection of 1,200 documents containing full text of House and Senate proceedings between 2013 and 2021. From these documents we extract roughly 80 million words. Below we discuss the intricacies of our collection methodology. The proceeding section is split into two subsections, one for media speech data and one for political speech data.

2.1 Media Speech Data

To obtain raw text inputs required for our analysis, we begin by extracting a set of URLs from CNN. These URLs will allow us to extract full text from their associated articles. CNN operates a comprehensive archive of all published articles running back to 2011, making this extraction process straightforward. We simply use python’s web-scraping tools to obtain a list of all URLs in the archive and then filter down for political and opinion pieces. In total this process yields 97,054 links to CNN political news and opinion pieces.

The Fox News website is possessed of no such archive. As such, URL extraction requires us to manipulate the site’s internal search engine. We use a webdriver interface program to automate this process, but search results of the Fox News website are still limited to 100 results. We therefore cannot simply find all articles categorized as ‘political’ or ‘Op-Ed’ by searching for these terms across the 2011-2021 time horizon. Instead we feed a series of search terms into the engine, filter for articles, and then extract all displayed URLs. We repeat this process for every month between July of 2011 and September of 2021. The result is 23,255 web addresses.

Once we have obtained lists of web addresses from both media websites, we use python’s text extraction tools to obtain full text documents from each webpage. The end result is a collection of roughly 23,000 articles from Fox News and 25,000 documents sampled from our CNN addresses. Each text document is then filtered to remove webpage formatting, to extract a publication date and an author. The filtered documents are then split into individual words, bigrams, and trigrams, which in turn are filtered for ‘stop words’, remaining webpage formatting, photograph captions, and author names.

The resulting dataset is far too large to read into the memory of a con-

ventional computer. We accordingly write each word, bigram, and trigram – along with the associated date, author, and news outlet – into a SQLite database. The database allows us to manage this staggering amount of information from a hard disk, rather than access memory. Memory management has proved a challenge throughout this analysis – and for both datasets we will need to restrict the total number of phrases to accommodate the constraints of a decidedly un-super computer.

We next pivot these data into a matrix of counts. Each row represents a media personality at a particular month. Each column represents counts of a particular phrase. For example, we have columns representing counts of the phrases “Gun Control Myths” and “President Barack Obama”. The rows represent the count of key phrases used by each author in a particular month. The matrix itself consists of 4,200 rows and 18,000 columns – 4,200 personality-month pairs, 18,000 phrases.

2.2 Political Speech Data

Our political speech data is obtained directly from both the Congressional Record and transcripts of Congressional Committee hearings. As above, we construct a dataset in which the columns represent particular phrases, and the rows represent counts of those phrases by speaker and month. Here, the speakers are members of the United States Congress (i.e. witless idiots), divided into Democratic and Republican parties.

We use python’s pdf reader tools to extract full text information from 1,200 daily records of the US Congress. We extract speakers by exploiting the formatting of these documents. Throughout the Congressional Record, speakers are delineated in all capital letters and prefaced by their title. We thus capture the speaker associated with each phrase by splitting each document into sections between title-speaker pairs. Title-speaker pairs that do not correspond to members of congress, such as those associated with committee witnesses, are discarded.

The result of this process is a collection of 80 million words, spoken by roughly 800 members of congress between 2013 and 2021. From there, it is a simple matter of pivoting the data into counts arranged by speaker and date. To conserve computer memory, we restrict ourselves to phrases spoken at least thirty times and consider only the 6,000 most commonly spoken phrases. The

end result is a data set with 25,000 rows and 6,000 columns.

3 Estimation

Our estimation strategy consists of two stages. First we quantify both media and political polarization following the methodology proposed by Gentzkow, Shaprio, and Taddy (2019).

3.1 Quantifying Polarization

Let i index authors and t index time. Each row of data is a J dimensional vector, c_{it} , of phrase counts, which we assume follow a multinomial distribution. For each observation, the multinomial density is parametrized by two objects:

1. A verbosity measure $m_{it} = \sum_j c_{jit}$, which measures the total number of phrases utilized by speaker i in time t
2. A J -dimensional vector of phrase probabilities, which is taken to depend on group affiliation. We denote this vector $q^{P(i)}(x_{it})$, where $P(i)$ represents group affiliation (Fox or CNN/Republican or Democrat) and x_{it} is a vector of speaker and time specific covariates.

Our subsequent formulation is not substantially different from a standard multinomial logitistic model. We assume that the elements of $q^{P(i)}(x_{it})$ are determined by a standard softmax functional form. That is, the individual choice probabilities are given by:

$$q_{jt}^{P(i)} = \frac{e^{\mu_{ijt}}}{\sum_l e^{\mu_{ilt}}} \quad (1)$$

Naturally, the μ terms are assumed to have a linear specification:

$$\mu_{ijt} = \alpha_{jt} + x'_{it}\gamma_{jt} + \varphi_{jt}I(Fox = 1) \quad (2)$$

The above formulation closely resembles a standard multinomial logit problem. The central difficulty in our analysis is the recovery of parameters α , γ and φ given that the number of choices is in the tens of thousands. Below we describe an estimator that allows for the recovery of these parameters in a

Mahoney – Polarization

high-dimensional setting. First, however, we will define our measure of polarization. Now, define

$$\rho_t(x) = \frac{q^{Fox}}{q^{Fox} + q^{CNN}} \quad (3)$$

We measure partisanship as,

$$\pi_t(x) = \frac{1}{2}q^{Fox}(x)\rho_t(x) + \frac{1}{2}q^{CNN}(x)(1 - \rho_t(x)) \quad (4)$$

This measure has a natural interpretation: it is the posterior probability that an observer with an even prior expects to assign to a speaker’s correct affiliation after hearing the speaker utter a single phrase. We then define average partisanship as,

$$\bar{\pi}_t = \frac{1}{N} \sum_i \pi_t(x_{it}) \quad (5)$$

Notice here that our goal is to recover estimates of $\pi_t(x)$, not to precisely estimate the causal terms α , γ , and φ . The conceptually simplest approach to form estimates of π would utilize maximum likelihood estimation. There are three limitations associated with direct maximum likelihood estimation. First, the multinomial formulation ignores important speech patterns in assuming independence between phrase counts. This means that the multinomial distribution is most certainly not the true distribution of phrase counts – and hence parameter estimates obtained via MLE will be biased.

Second, bias exists in MLE estimates due to sampling error. Our phrase count data is incredibly sparse, and so certain phrases may be spoken more by Fox News or CNN simply by chance. This will tend to increase our values of even prior $\hat{\rho}$, relative to the true value of ρ . As we’ll see the penalized estimator of Gentzkow, Shapiro, and Taddy (2019) allows us to account for this difficulty by assigning values of $\varphi_j = 0$ to least relevant phrases.

Third, and perhaps most important, the objective function associated with the maximum likelihood approach requires substantial computational power to evaluate. In our two applications below, the number of choice outcomes numbers well into the thousands. This will require the evaluation of denominator sums with thousands of terms, for each observation, at each iteration of

our optimization routine. Computation would require several months of run time on a conventional computer.

To account for this difficulty, we utilize the penalized estimator of Gentzkow, Shapiro, and Taddy (2019). The estimator uses a the Poisson approximation to a multinomial likelihood function first posited by Palmgren (1981) and extended by Baker (1994). This approximation saves us a considerable amount of computing power. The penalty terms also mitigate the sampling error caused by our data’s sparsity.

We minimize the following expression by our choice of α , γ and φ :

$$\sum_j \left[\sum_t \sum_i [m_{it} \exp(\alpha_{jt} + x'_{it} \gamma_{jt} + \varphi_{jt} I(P)) - c_{ijt} (\alpha_{jt} + x'_{it} \gamma_{jt} + \varphi_{jt} I(P)) \right. \\ \left. + \Psi(|\alpha_{jt}| + ||\gamma_{jt}||_1) + \lambda_j |\varphi_{jt}| \right] \quad (6)$$

Notice that this Poisson objective is separable across phrases and does not contain large sums of exponential terms. These two facts will make computation of optimal parameters substantially faster, as they allow for the utilization of distributed computing techniques. The penalty terms will ensure that our predicted probabilities have low errors. Since we are interested in prediction, and not the parameters themselves, the inclusion of these terms improves our measure of partisanship. Further, the penalty terms guarantee that many model parameters are set to zero. This will simplify the computation of the \hat{q} terms need to measure polarization.

The term λ_j , which limits the magnitude of our party effect, φ , is determined through an iterative process. We begin by setting λ_j to be large enough that the estimate of φ equals 0. For each phrase, we trace the value of lambda down to 0 in 100 steps, running the optimization routine over i and t each time. We then pick the value of lambda that minimizes a Bayesian Information Criterion.


For the penalty term Ψ , we use a value of 10^{-5} , as in GST (2019). This low penalty value ensures that solutions to the minimization problem always exist. The term also ensure that our matrix of estimates will be sparse, which proves computationally convenient.

3.2 Polarization Results

In our subsequent analysis we utilize trigram phrases. We compute the minimizing parameter values in equation (6) for both congressional speech data and media speech data. The resulting parameter vectors, α and γ are much too large to present within this paper, and are not of particular interest. In table 1 below we present selected φ terms obtained from both estimation contexts.

Once we have obtained estimates of all model parameters, we calculate analogs of equation (1). For each individual i in time t , we compute \hat{q}_i^{Fox} and \hat{q}_i^{CNN} , then calculate ρ and π for all individuals. Notice that we are comparing the divergence in phrase probabilities within individuals induced by a change in group affiliation. We then simply average over each speaker in a time period to obtain an overall measure of divergence in language use patterns. The result is two time series, one measuring media polarization, the other political polarization. We plot both series below:

Figure 1: Media Polarization 2013-2021




3.3 Causal Inference

The central purpose of this analysis is to determine the direction of causality between the two series recovered above. To assess causality between media polarization and political polarization we first employ an instrumental variables approach. To instrument for media polarization, we use two monthly series: surveys of viewer satisfaction for both media outlets and the average number of TV viewing hours in the US.

Coverage satisfaction is doubtless correlated with media polarization. Views.

Figure 2: Political Polarization 2013-2021



Similarly, Americans might consume more television media when media polarization is high. However, television viewing habits could only conceivably be correlated with political speech through its relationship to the overall media landscape.

Formally, we specify our first model as a standard two-stage least squares estimator, where the first stage is given by:

$$Media_t = \omega_0 + \omega_1 Sat_t^{fox} + \omega_2 Sat_t^{cnn} + \omega_3 TV_t + \epsilon_{1t} \quad (7)$$

Where the variable Sat represents the proportion of viewers of each outlet claiming satisfaction with network programming. The variable TV_t represents the number of hours of television media consumed by the average American household each day.

$$Pol_t = \beta_0 + \beta_1 \hat{Media}_t + \beta_P X_{tP} + \beta_E X_{tE} + \beta_T X_{tT} + \epsilon_{2t} \quad (8)$$

Where pol and $media$ represent indices of political and media polarization respectively. X_{tP} , X_{tE} , and X_{tT} represent the political, economic, and time covariates introduced above.

In estimating equation (8), we apply a basic Durbin-Watson test to detect the presence of autocorrelation in our error structure. The test yields convincing evidence that autocorrelation is indeed present, a fact we counter with robust standard errors. Throughout this section, all presented standard errors are inflated via White's methodology.

Mahoney – Polarization

The results of the two stage least squares analysis are given in table below. We report a number of specification results and include economic, political, and time controls. Our economic controls include the monthly unemployment rate, the quarterly GDP growth rate, the monthly CPI inflation rate, the quarterly share of output awarded to labor, and the monthly average of the S & P 500 index. Political controls include dummies representing the control of the House and Senate by the Republican party. Finally, we include time fixed effects.

In all cases, we find a positive, statistically significant coefficient on *media*, suggesting a causal relationship between media polarization and political polarization. In the full specification, we obtain a *media* coefficient of roughly 0.40, suggesting a one point increase in media polarization increases the political polarization index by 0.40. That this relationship is far from one-to-one is hardly surprising, media always exhibits a greater degree of partisanship than politics.

Table 1: Table One: IV Regression Results, Trigrams

Variable	(1)	(2)	(3)	(4)	(5)	(6)
Media	0.020 (0.070)	0.195* (0.107)	0.854*** (0.104)	0.338** (0.150)	0.129* (0.074)	0.338** (0.150)
Urate			-1.239*** (0.293)	-0.601*** (0.333)	-0.160 (0.261)	-0.601 (0.333)
Inflation			-1.452*** (0.286)	-0.944** (0.403)	-0.103 (0.236)	-0.944 (0.403)
GDP			-0.625*** (0.179)	0.051 (0.097)	0.038 (0.088)	0.051 (0.097)
Stocks			0.858*** (0.064)	0.161 (0.186)	0.215 (0.097)	0.161 (0.186)
Labor			0.203 (0.655)	0.032 (0.451)	-0.220 (0.341)	0.032 (0.451)
Rep Prez	12.376*** (1.247)	6.724*** (2.412)			11.389*** (1.180)	7.797*** (2.840)
Rep House	0.232 (0.891)	-3.373** (1.842)			1.499 (1.144)	-2.649 (3.406)
Rep Senate	0.250 (0.660)	2.614** (1.351)			0.806 (0.845)	1.060 (1.736)
Constant	53.542 (4.400)	43.685 (6.499)	-5.322 (7.906)	33.414 (12.325)	41.748 (6.175)	36.064 (12.234)
Time FE	No	Yes	No	Yes	No	Yes

3.4 Robustness Check

The first, and most obvious, robustness check is to repeat the entirety of the above analysis using counts of bigrams, rather than trigrams. Below we plot both polarization series obtained from this analysis, and show that the plots reveal similar trends to those in the trigram analysis. This implies that the polarization using less informative phrases mimics and supports the initial results uncovered using the trigrams as inputs.

Figure 3: Media Polarization 2013-2021 (Bigrams)





Figure 4: Political Polarization 2013-2021 (Bigrams)



We repeat the IV analysis of section 3.3 using the same set of controls. The results of this analysis are presented in table 2 below. Notice that once again, we obtain positive, statistically significant results in five of our six specifications.

Table 2: IV Regression Results, Bigrams

	(1)	(2)	(3)	(4)	(5)	(6)
Media	0.266 (0.234)	0.781* (0.235)	1.441** (0.181)	0.935*** (0.253)	0.432** (0.193)	0.935** (0.253)
Urate			-0.207 (0.673)	-0.085 (0.682)	0.299 (0.635)	-0.085 (0.682)
Inflation			-0.346 (0.456)	-0.313 (0.536)	0.543 (0.408)	-0.313 (0.536)
GDP			-0.598** (0.208)	0.122 (0.136)	-0.062 (0.140)	0.122 (0.136)
Stocks			0.573** (0.122)	-0.077 (0.264)	0.001 (0.186)	-0.077 (0.264)
Labor			-1.028 (1.079)	-1.262 (0.896)	-1.504 (0.855)	-1.262 (0.896)
Rep Prez	9.421*** (1.925)	5.649** (2.953)			8.990*** (1.690)	6.404* (3.302)
Rep House	-0.855 (1.050)	-4.440* (1.788)			-0.785 (1.570)	-5.206* (4.817)
Rep Senate	0.641 (0.716)	2.803* (1.003)			2.625* (1.193)	2.323 (1.866)
Constant	36.793 (14.450)	5.590 (14.255)	-44.819 (14.284)	-3.119 (12.325)	24.614 (14.524)	2.088 (15.890)
Time FE	No	Yes	No	Yes	No	Yes

4 Conclusion

In this analysis we have demonstrated that a causal effect exists between the American media environment and political polarization. Rising media polarization produces a positive, significant change in political polarization. This finding is in contrast to the results of Prior (2013), however there are a number of important distinctions between our analysis and his.

First, our paper examines a time period running from 2013-2021. This time horizon includes several important polarizing events – the Benghazi hearings, the election of Donald Trump, the Charlottesville rally – to name only a few. These events engendered a great deal of attention in both the media and in Congress. It is thus possible that the causal effect media plays in driving political polarization is itself a new phenomena.

Second, our analysis is driven primarily by textual information. That is, we capture political polarization by examining the divergence in language use between members of different political parties. The advantage to this approach is that it allows us to effectively disentangle the most polarizing phrases used by each group. In the appendix below, we present polarizing phrases in the media and congress for selected months in our sample. These are phrases that

Mahoney – Polarization

have values of φ that are large in magnitude – with positive values indicating a strong affiliation with republicans (Fox News) and negative values indicating a strong affiliation with democrats (CNN).

A final note on our analysis – we have only captured the role that traditional media plays in driving political polarization. During the time period in question, social media grew in importance and use. We have not captured or controlled for social media’s influence in politics in this analysis. Regardless, this potential correlation

5 Appendix

5.1 Predictive Phrases

In the tables below, we present that top 10 phrases associated with each group affiliation. Table three presents the phrases most highly associated with Democrats and Republicans, as measured by the φ_{jt} parameter in equation (6). We present these phrases for three selected months across our time horizon. Table four presents the phrases most associated with Fox and CNN, measured similarly.

Mahoney – Polarization

Congress March 2016				Congress July 2020					
Democrats		Republican		Democrats		Republican			
behavioral health care	φ	-2.72	College Program Act	φ	6.40	mandatory minimum sentences	-3.06	National Sheriffs Association	φ
public health emergencies	-2.30	-2.10	Cancer Awareness Month	5.19	Justice Policing Act	-3.05	Federal budget deficit	3.21	
mandatory minimum sentences	-2.10	-1.90	land port entry	3.60	massive tax cuts	-2.90	Lieutenant General Army	3.24	
nuclear power plants	-1.90	-1.64	United States Taiwan	3.47	York New York	-2.69	Army Civil Works	3.59	
committed United States	-1.64	-1.59	private health care	3.33	honest American people	-2.43	bipartisan budget agreement	3.68	
Agents Registration Act	-1.59	-1.42	Buy American Act	3.15	public health emergencies	-2.41	congressional delegation POMA	3.94	
Foreign Agents Registration	-1.42	-1.40	States Munitions List	3.13	health care plans	-2.40	Operation Warp Speed	4.08	
And even though	-1.40	-1.34	United States Munitions	3.13	Veterans Health Care	-2.31	private education loan	4.16	
Federal budget deficit	-1.34	-1.27	relief small businesses	2.94	Medicare Medicaid Social	-2.29	transpor research development	4.68	
behavioral health services	-1.27	-1.24	mental health condition	2.89	intimate partner violence	-2.13	child support enforcement	4.80	
Congress March 2018				Congress July 2021					
Democrats		Republican		Democrats		Republican			
death serious bodily	φ	-4.42	Service Main Street	φ	3.17	Endangered Threatened Wildlife	-4.37938	Federal Helium Reserve	φ
Medicare Medicaid Social	-3.43	-3.35	Peace Corps volunteers	3.18	Threatened Wildlife Plants	-4.27962	William National Defense	5.147469	
mental health professional	-3.35	-3.29	Lieutenant General Army	3.24	States Military Academy	-2.56797	Government Saudi Arabia	5.122877	
Americans deserve better	-3.29	-3.06	American Enterprise Institute	3.54	The First Amendment	-1.59529	The highest level	5.032817	
mandatory minimum sentences	-3.06	-3.05	Navy grade indicated	3.54	Public Health Emergency	-1.33022	nuclear power plants	4.724885	
Justice Policing Act	-3.05	-2.90	States Navy grade	3.54	President United States	-0.59694	Communist Party China	4.615472	
massive tax cuts	-2.90	-2.81	extraneous material currently	3.80	Chief Justice United	-0.35435	spare repair parts	4.523031	
western New York	-2.81	-2.65	material currently consideration	3.80	protect defend Constitution	-0.32685	construction defense activities	3.856043	
median household income	-2.65	-2.43	Coast Guard Reserve	3.94	people come country	-0.17049	military construction defense	3.855651	
honest American people	-2.43		Guard Authorization Act	4.90	telling American people	-0.12329	unemployment insurance ben	3.846372	
Media June 2015				Media August 2018					
CNN		Fox News		CNN		Fox News			
gun control laws	Varphi	-3.00	state based exchanges	12.16	assault style rifles	-6.48	players taking knee	Varphi	
donors Clinton Foundation	-2.62	-2.60	world ever seen	10.86	Lt Gov David	-6.08	Barack Obama issued	10.72	
premium tax credits	-2.60	-2.51	cover birth control	10.80	emails containing classified	-5.96	Lewis Center Ohio	10.51	
free speech rights	-2.51	-2.10	star Donald Trump	10.44	President Pro Tem	-5.53	Kentucky Secretary State	10.40	
Former House Speaker	-2.10	-1.92	White House taken	10.30	claims presidential election	-5.51	issued temporary restraining	10.37	
Former Speaker House	-1.92	-1.56	small medium sized	9.85	national security tax	-5.47	Washington Jobs Freedom	10.34	
Donald Trump Hillary	-1.56	-1.43	White House worked	9.83	override presidential veto	-5.36	real estate agent	10.26	
Court United States	-1.43	-1.28	middle class voters	9.81	talks United States	-5.11	hour long conversation	10.14	
first time ever	-1.28	-1.24	White House used	9.67	Clean Water Act	-5.05	US Marine Corps	10.00	
across Middle East	-1.24			9.44	Donald Trump conservative	-5.02	California San Francisco	9.98	
Media July 2020				Media March 2021					
CNN		Fox News		CNN		Fox News			
assault style rifles	Varphi	-6.48	terrorist attack Benghazi	11.53	leadership Democratic Party	-5.345857745	Jim Crow South	Varphi	
Lt Gov David	-6.08	-5.96	Trump administration rule	10.46	majority Supreme Court	-4.591922303	marching White House	10.69166041	
emails containing classified	-5.96	-5.53	rape incest life	10.35	House Sergeant Arms	-4.465800971	anti war activist	10.50238576	
President Pro Tem	-5.53	-5.51	American Red Cross	10.31	full court press	-4.194415161	federal gun control	10.25851558	
claims presidential election	-5.51	-5.47	civil rights activist	10.28	including former President	-4.143127254	United States intelligence	9.914988697	
national security tax	-5.47	-5.36	rebuild war torn	10.09	asylum United States	-4.129867578	minority owned businesses	9.638554239	
override presidential veto	-5.36	-5.11	Journal Constitution via	10.02	Joe Biden proposed	-3.945098909	avoid paying taxes	9.600972811	
talks United States	-5.11	-5.05	Congress Racial Equality	9.89	restoring voting rights	-3.728343662	data Centers Disease	9.176746142	
Clean Water Act	-5.05	-5.02	running mate former	9.71	investigation Hillary Clintons	-3.513591497	Barack Obama State	9.149851793	
Donald Trump conservative	-5.02			9.40	gun control advocates	-3.325731472	announced presidential campaign	9.116722975	

Mahoney – Polarization

The astute reader may notice that the phrases used by CNN do not seem to change all that much over the 2017-2020 period. This is a feature of the penalized estimator posited in equation (6). Essentially, the penalty terms ensure that more non-zero φ weights are assigned to Fox News. The terms given non-zero weights for CNN must therefore be somewhat static over-time. We observe a similar phenomena between Democrats and Republicans, although it is less pronounced.

Outcomes like this are an admitted concern with machine learning methods: our model predicts well, but it is exceedingly difficult to determine precisely why Fox News receives more non-zero weighting terms than CNN.

Regardless, many of the terms listed above seem rather natural. In July of 2020 the phrase most indicative of democrats was 'Justice in Policing Act', while republicans seems content to speak on more mundane matters. Generally speaking, the divide between Fox and CNN seems to carry a natural bent. Fox News speaks often of the Bangazi attack – a fixation of the outlet for many years. CNN speaks often of gun control measures and seems to frequently distinguish between Republicans who support Donald Trump and those who do not.

Not all terms appear sensible, and indeed without examining the entire context of the new cycle in a particular month it is difficult to understand why some terms receive such high weight. Again, this is a central shortcoming of machine-learning methods. However, the tables above are sufficient to conclude that our estimator has behaved reasonably.

5.2 Alternate IVs

In this section, we repeat the analysis of section 4, using Twitter data to construct an alternate IV. In the two tables below, we use the frequency of tweets issued by each media network as an instrumental variable. While the content of the tweets may be associated with political polarization, their frequency will not be. However, media polarization is likely to be correlated with tweet frequency. We thus argue that tweet frequency satisfies the exclusion restriction required for causal inference.

Table 5 below repeats the trigram analysis, while table 6 presents results for bigrams. In both cases, we find a positive and statistically significant co-

efficient on *media* in four of our six model. While the magnitude of our *media* coefficient is somewhat attenuated, the results of this robustness check qualitatively corroborate the conclusions of section 4. Also, while our coefficients are indeed smaller in magnitude, they are still large enough to suggest that a substantial portion of political polarization can be explained by changes in the media environment. Further, the coefficients on most economic controls remain insignificant.

Table 3: IV Regression Results, Trigrams, Alternate IVs

	(1)	(2)	(3)	(4)	(5)	(6)
Media	0.253*	0.433**	1.247***	0.426**	0.139	0.426**
	(0.183)	(0.239)	(0.124)	(0.225)	(0.166)	(0.225)
Urate			-0.593	-0.626	-0.278	-0.626
			(0.402)	(0.549)	(0.480)	(0.549)
Inflation			0.102	-0.475	0.112	-0.475
			(0.424)	(0.707)	(0.449)	(0.707)
GDP			-0.658	-0.265	-0.127	-0.265
			(0.392)	(0.485)	(0.302)	(0.485)
Stocks			4.14E-04	3.42E-05	1.45E-04	3.42E-05
			(9.81E-05)	(1.54E-04)	(1.12E-04)	1.54E-04
Labor			-0.167	-0.369	-0.355	-0.369
			(0.757)	(0.660)	(0.512)	(0.660)
Rep Prez	11.127	5.747			11.416	8.218
	(1.869)	(2.628)			(1.901)	(2.929)
Rep House	0.478	-3.654			1.487	-1.173
	(1.023)	(1.591)			(1.236)	(3.410)
Rep Senate	0.336	2.032			0.731	-0.011
	(0.703)	(0.457)			(1.313)	(1.930)
Constant	38.703	30.339	-24.556	34.434	43.541	35.607
	(11.119)	(13.790)	(7.756)	(14.816)	(10.961)	(13.149)
Time FE	No	Yes	No	Yes	No	Yes

The table below presents the results of our bigram analysis, here repeated using our alternate IV. The results of this IV estimation is even more aligned with our initial results. We establish significance in five of six model specifications.

Table 4: IV Regression Results, Bigrams, Alternate IVs

	(1)	(2)	(3)	(4)	(5)	(6)
Media	0.268 (0.306)	0.953** (0.398)	1.312*** (0.151)	1.189*** (0.372)	0.417* (0.249)	1.189** (0.372)
Urate			-0.271 (0.785)	-0.115 (1.250)	0.132 (0.982)	-0.115 (1.250)
Inflation			-0.100 (0.491)	-0.207 (0.885)	0.167 (0.711)	-0.207 (0.885)
GDP			-1.006 (0.550)	-0.038 (1.138)	-0.440 (0.608)	-0.038 (1.138)
Stocks			3.31E-04 1.16E-04	-1.78E-04 2.96E-04	4.14E-05 2.32E-04	-1.78E-04 2.96E-04
Labor			-1.929 (0.983)	-1.558 (0.962)	-1.974 (0.808)	-1.558 (0.962)
Rep Prez	9.144 (2.672)	4.196 (4.540)			8.756 (2.696)	3.910 (4.442)
Rep House	-0.486 (0.998)	-4.171 (2.697)			0.162 (1.518)	-6.593 (6.684)
Rep Senate	0.619 (0.656)	1.708 (0.889)			1.648 (2.110)	1.505 (3.554)
Constant	37.010 (18.585)	-2.929 (22.917)	-29.147 (12.594)	-15.354 (26.636)	26.141 (20.150)	-8.761 (22.910)
Time FE	No	Yes	No	Yes	No	Yes

References

- [1] Matthew Gentzkow, Jesse Shapiro, Matt Taddy. *Measuring Group Differences in High-Dimensional Choices: Method and Application to Congressional Speech* Econometrica 2019
- [2] Markus Prior *Media and Political Polarization* Annual Review of Political Science 2013
- [3] Isaac Waller, Ashton Anderson *Quantifying social organization and political polarization in online platforms* Nature 2021
- [4] Andrew O. Ballard, Ryan DeTamble, Spencer Dorsey, Michael Heseltine, Marcus Johnson *Dynamics of Polarizing Rhetoric in Congressional Tweets* Legislative Studies Quarterly 2022
- [5] Niels D. Goet *Measuring Polarization with Text Analysis: Evidence from the UK House of Commons, 1811–2015* Cambridge University Press: 2019
- [6] David Pozen, Eric Talley, Julian Nyarko. *A computational analysis of constitutional polarization*. Cornell Law Review (2019).