Martin Tekeli Brogaard (Christ's College)
Part IIB Dissertation 2022–2023

Inside the Fourth Estate: An Empirical Analysis of Political Commentary in US Cable News

Abstract

The political influence of news media is widely recognised. However, there is little systematic evidence of how media strategically formulate messaging to viewers. This dissertation opens up this black box by asking if commentators adjust their emotional delivery depending on the audience's political views. By applying topic modelling and sentiment analysis to over 200,000 partisan cable news transcripts, I show that the same commentator adjusts messaging in response to the audience. Strikingly, the direction of adjustment reflects a consistent *stronghold effect*: commentators increase the intensity of emotive rhetoric when speaking to an audience whose bias aligns with their ideology.

Word Count: 7489

1 Introduction

News media play a vital role in shaping the political landscape and, consequently, policy outcomes. From an efficiency stance, the principal goal of news outlets is to disseminate information to improve decision-making of the public. Increasingly, however, economists have become interested in the power of mass media to manipulate audiences politically. Recent evidence from the United States corroborates this concern: studies using exogenous channel positions (Martin and Yurukoglu, 2017) and experimental changes to news diet (Broockman and Kalla, 2023) find that Fox News exposure strengthens conservative attitudes. Similar conclusions on the electoral implications of partisan news find support in comparatively less democratic societies like Russia (Enikolopov et al., 2011).

To elucidate these findings, it has been proposed that a crucial aspect of the media's sway on politics is their use of slanted language and, in particular, emotional appeals. A growing body of literature contends that emotive rhetoric can mobilise listeners into political action (Brader, 2005; Valentino et al., 2011) and change attitudes on public policy (Renshon et al., 2015). Recent economic research also reveals that slanted language in the media has direct political consequences: in an AER paper, Djourelova (2023) exploits an AP media ban on the term "illegal immigrant" as a natural experiment to demonstrate a statistically significant persuasive effect on voters' stance on immigration policy.

Although these studies reaffirm that rhetoric matters for *receivers* of news, there has been little systematic research into how *senders* strategically tailor messaging to maximise influence on viewers with different preferences and biases. The main obstacle to identifying senders' incentives is the inability to construct a reliable counterfactual for comparisons: actors in news media will self-select into communicating to certain audiences based on unobservables, making it difficult to separate strategic incentives from heterogeneity in commentator characteristics by simply comparing speech across individuals.

This dissertation makes a considerable advancement in overcoming this challenge with a novel fixed effects identification strategy. By constructing an extensive database of cable news transcripts for a group of prominent American commentators, I provide an unprecedented investigation into how the *same* cable news commentator adjusts messaging in response to the audience's political orientation. Central to this identification approach is the exceptional degree of polarisation in the American cable news market. Given the indisputable political divides in viewership between Fox News (strongly conservative) and

MSNBC (strongly liberal), the US cable news market is an ideal setting for understanding how the same media actors change their message delivery as the ideology profile of the audience varies.

The study concentrates specifically on the relationship between emotive rhetoric and commentators' political alignment with audiences. There are two reasons for this. Firstly, emotional appeals in political messaging are underexplored and have recently been shown to play a key role in propagandist news content (Karell and Agrawal, 2022). As I will argue, if media use emotive rhetoric as a substitute for providing valuable information, strategic appeals to emotions may have adverse consequences for information efficiency. Secondly, quantifying emotive rhetoric from transcripts is computationally feasible with current sentiment analysis techniques, making text analysis highly tractable.

The dissertation's results offer three major contributions to extant scholarship on political incentives in the media. Firstly, I find that the same commentator adjusts messaging in response to audience bias, suggesting the presence of strategic messaging. Secondly, the direction of this adjustment reflects a meaningful *stronghold effect*: commentators consistently increase the intensity of emotive rhetoric when faced with an ideologically like-minded audience. Thirdly, the findings provide a compelling case for using topic modelling and sentiment analysis complementarily to extract new behavioural insights from text data in political economy research.

To summarise, Section 2 introduces the institutional context and presents a theoretical argument for strategic messaging rooted in a simple persuasion model. Section 3 describes how transcripts and commentator data are extracted. Section 4 and Section 5 outline how transcripts are sorted into topics and scored for the intensity of emotive rhetoric, respectively. Section 6 provides summary statistics, and Section 7 presents the econometric model used to test for messaging tactics. Next, Section 8 presents regression results with a discussion, and Section 9 describes robustness checks. Finally, Section 10 concludes.

2 Context and theory

To understand the behaviour of news commentators, a natural first step is to model them as choosing their messaging to optimise an objective. One possible narrative is that commentators are politically motivated and seek to maximise their influence on the audience's beliefs through political persuasion. To illuminate their incentives, this section first presents necessary background knowledge about the US cable news market. Next, to motivate the empirical analysis, I build a simple theoretical argument for why commentators may choose to adjust messaging depending on the audience's political bias.

2.1 Institutional context

In the United States, three cable networks dominate the market for political commentary: Fox, CNN, and MSNBC. Each of these channels provides 24–7 news coverage to millions of Americans every day. Importantly, the cable news market is notoriously segmented. According to a recent survey by the Pew Research Center (2020), 93% of regular Fox News viewers identify as Republican-leaning. In contrast, 95% of regular MSNBC viewers are Democratic-leaning. CNN is claimed to provide a neutral middle ground, with a majority of regular viewers nevertheless being Democrats (79% against 17% Republicans).

In contrast to broadcasting services, *cable* news are legally outside the jurisdiction of the FCC, implying minimal levels of content regulation by the government (Federal Communications Commission, 2021). In light of this, political divides in news consumption likely reflect product differentiation in the form of partisan reporting. Rather than appealing to the median viewer, outlets optimally differentiate themselves to capture consumer masses on each side of the political spectrum (Anand et al., 2007; Gentzkow and Shapiro, 2010). If viewers are indeed captive consumers, the partisan orientation of each channel's audience can be taken as given in the short run.

When considering individual actors in this setting, it is critical to distinguish between *anchors* and *commentators*: anchors normally feature exclusively on one channel and are hired by the network to provide reporting and conduct interviews. Commentators, on the other hand, are not formally affiliated with any network and may be invited by *multiple* channels to offer commentary on a given issue, often in an interview with an anchor. Although my contributions are chiefly empirical, the next section proposes a theoretical hypothesis for why rational *commentators* may alter messaging between channels.

2.2 A theoretical argument

Consider a cable news environment with two actors: commentator C and audience A. The commentator has a partisan bias P_C and is either left-wing ($P_C = L$) or right-wing ($P_C = R$). Suppose the true state of the world s can be located on the left–right political spectrum, such that $s \in [-1, 1]$. After observing s, the commentator C provides a message M to A, knowing what the prior political beliefs of the audience are. The objective of C is to induce A to support their side of the political spectrum.

2.2.1 Audience's choice

Depending on the channel's slant, A has a prior belief about s given by $\theta_A \in [-1, 1]$. Suppose A trusts information provided in C's message since viewers have already selected into the channel. After receiving M, A chooses to support a side:

$$\text{Support}_A = \begin{cases} L & \text{if } \mathbb{E}_A[s|M] < 0 \\ R & \text{if } \mathbb{E}_A[s|M] > 0, \end{cases}$$

where A's posterior expectation $\mathbb{E}_A[s|M]$ depends on A's prior θ_A and the message M.

2.2.2 Commentator's choice

The commentator has two options. First, after observing s, C can choose an *informative* message ($M = I$) with information s_M revealing that $s_M = s$. This leads A to update beliefs partly away from θ_A. Alternatively, C can send an *emotional* message ($M = E$) that engenders no updating and thus simply reaffirms the audience's prior beliefs.

Formally, suppose A's posterior expectation of s after the message is

$$\mathbb{E}_A[s|M = E] = \theta_A \tag{1}$$

$$\mathbb{E}_A[s|M = I, s_M] = \underbrace{\left(\frac{1}{1+\mu}\right) s_M}_{\text{Belief updating}} + \underbrace{\left(\frac{\mu}{1+\mu}\right) \theta_A}_{\text{Bias toward prior}}, \tag{2}$$

where $\mu \in [0, 1]$ measures the degree of the audience's "confirmation bias". This form of biased updating is similar to the one introduced by Hagmann and Loewenstein (2018) in their model of "persuasion with motivated beliefs". Optimal message delivery will depend on whether A's prior beliefs are aligned with C's side (ignoring knife-edge case of $\theta_A = 0$).

2.2.3 Commentator's optimal messaging

When ideologies are aligned. Suppose that $\{\theta_A < 0 \wedge P_C = L\}$ *or* $\{\theta_A > 0 \wedge P_C = R\}$. For all values of $s \in [-1, 1]$, it is *weakly* optimal for C to choose an emotional message. By not revealing any information, A will always hold onto the prior and support C's side. For some values of s, an informative message may even push A to the other side, which would be strictly sub-optimal.

When ideologies are misaligned. Either (i) $\{\theta_A > 0 \wedge P_C = L\}$ *or* (ii) $\{\theta_A < 0 \wedge P_C = R\}$. In both cases, C's choice depends on the realisation of s.

Consider case (i) where C is left-wing and A has a right-wing prior. If $s \geq \theta_A$, it does not matter what C does since A will always support R. Note, further, that C could only strictly benefit from choosing $M = I$ if $s < 0$ since for $0 \leq s < \theta_A$, A's belief updating would not make the posterior expectation $\mathbb{E}_A[s|M]$ cross the midpoint of $[-1, 1]$.

Whether $M = I$ can be optimal depends on both s and μ. In case (i), an informative message will definitely persuade A to support L when

$$\mathbb{E}_A[s|M = I, s_M] < 0 \iff s_M < -\mu\theta_A. \tag{3}$$

Intuitively, for an informative message to be optimal, $s_M = s$ must be sufficiently dissonant with A's prior θ_A. In addition, a larger confirmation bias weight μ will restrict the possible values of s for which $M = I$ is optimal for C because A's belief updating is weaker. The argument is symmetric for case (ii).

Though simplistic, the model reveals two subtle mechanisms. Firstly, a commentator has no reason to be informative when they are politically aligned with the audience. Instead, they may resort to uninformative emotional rhetoric when speaking to political strongholds, producing echo chambers. When misaligned, assuming s takes on a range of values in $[-1, 1]$, stronger confirmation bias implies fewer informative messages, meaning that belief inertia caused by partisanship reduces information efficiency. Critically, the result is predicated on the assumption that emotional messaging is less informative.

While other viable frameworks for this environment may exist, this model provides a distinct prediction for the study's empirical results: on average, news commentators should provide more emotional messages to like-minded audiences. The following sections build up to exploring the veracity of this hypothesis.

3 Constructing the dataset

The identification problem of the study is to determine how the *same* commentator adjusts emotive rhetoric as the channel varies. For this purpose, closed caption transcripts from partisan channels will be analysed. Three important building blocks must be retrieved: transcript text from news segments sorted by commentator, the channel name of each segment, and background information on ideology for each of the commentators.

3.1 Facial detection data

To retrieve cable news transcripts, I download closed caption text from individual shows in 2010–2022 using the *Stanford Cable News Analyzer*. Developed by Hong et al. (2021a), the tool uses a neural network MTCNN face detector on video data from the *Internet Archive's TV News* project, which amalgamates nonstop recordings and closed caption text from CNN, Fox and MSNBC from 2010–now. In total, the database contains 13 years or around 280,000 hours of TV news content. The MTCNN face detector is implemented jointly with the *Amazon Celebrity Recognition API* to identify news personalities. An individual is identified when the model's confidence score is above 0.5. For each detected individual, the Stanford Analyzer provides a unique ID for the 15-minute programme and, critically, time intervals in which the detected person features, down to the millisecond.

3.2 Data on commentator ideology

Political ideology data for relevant news personalities are sourced from the *Database on Ideology, Money in Politics and Elections (DIME)*, which stores information on 130 million campaign donations made by elite Americans. Importantly, the dataset contains a *cfscore* calculated for each individual: this is a weighted score that indicates political ideology of people based on campaign contribution history. If $cfscore > 0$, records indicate the person is Republican; if $cfscore < 0$, they are likely to be a Democrat.

Fortunately, DIME *cfscores* have already been carefully matched with people featuring in the Stanford Cable News Analyzer by Kim et al. (2022a), though for a different research purpose. By using their replication data, I retrieve *cfscores* for 977 individuals (703 commentators and 274 news anchors, with the former including politicians). Finally, I divide each commentator into Democrat or Republican based on the *cfscore*'s sign.

3.3 Trimming text files from facial detection time stamps

To extract closed caption transcripts for 15-minute programme IDs in bulk, I use the *Internet Archive API* (with research permission granted by the founder of the TV News Archive, Roger MacDonald). For each individual included in the database, there is a json file of all programme IDs and time stamps where they are detected. By using the API, I extract all transcripts that correspond to the IDs in the json files.

To isolate the closed caption text where a given person's face is detected by the Stanford Analyzer, a simple model was created in Rust: first, the model parses the json files from the Stanford Analyzer linked to each individual in the database. For each individual, it initialises a new thread, which finds the relevant txt files of closed caption texts with the API. Finally, the model iterates over each line of the txt file. If the time stamp of a line is within an interval from the face detection data, the text is written into a new file. After running, the model returned just over 1,100,000 files. Note, due to the researcher agreement with the Internet Archive, transcript text cannot be shared.

3.4 Pre-processing transcripts

Firstly, since short documents do not contain sufficient information for conducting proper text analysis, transcripts with less than 25 words were discarded. The remaining transcripts were then loaded into a "corpus" (collection of text files) in R. Before analysing transcripts, I follow common text analysis practice of pre-processing the corpus with RStudio's "tm" package (Gentzkow et al., 2019):

1. Removing stop-words (e.g., "the", "of", "to"), as these don't add much meaning.
2. Removing punctuation, as this is largely irrelevant for spoken phrases.
3. Removing numbers and non-alphabetical characters.
4. Lemmatisation: removing grammatical suffices, as most terms take on several grammatical forms with the same semantic meaning (e.g., "thinking" and "thinks" both become "think").

In total, the corpus contains 617,726 text files (315,456 for commentators and 302,270 for anchors). Overall, the average transcript length is 339 words. The average commentator has 449 appearances, and the average anchor 1,103.

4 Document classification with topic modelling

Naturally, messaging in news commentary will differ significantly depending on the topic being discussed. Since topics may be correlated with the frequency of channel appearance for a given person, it is necessary to control for topic differences to a sufficient degree.

To this end, I use a simple but widely used model in economics known as the Latent Dirichlet Allocation model. Similar to factor models like principal components, it is an unsupervised machine-learning model that assigns documents to topics by clustering texts with frequently co-occurring words. Below, I provide a brief theoretical summary based on the seminal theory in Blei et al. (2003) and an outline of how I implement the model.

4.1 Latent Dirichlet Allocation model

Suppose the corpus D contains M documents and V unique words. A document $d \in D$ is a sequence of N_d words. Assume news content can be divided into K broad topics. Now, let θ_d be a random variable for the *topic distribution* over d. Let ϕ_k be a random variable for the *word distribution* for topic $k \in \{1, 2, ..., K\}$. Consider the generative process:

1. Choose $\theta_d \sim \text{Dir}(\boldsymbol{\alpha})$ where $\text{Dir}(\boldsymbol{\alpha})$ is a multivariate Dirichlet distribution with a K-dimensional parameter vector $\boldsymbol{\alpha}$.
2. Choose $\phi_k \sim \text{Dir}(\boldsymbol{\beta_k})$ where $\boldsymbol{\beta}$ is a $K \times V$ parameter matrix, with each row $\boldsymbol{\beta_k}$ being a V-dimensional vector for topic $k \in \{1, 2, \ldots, K\}$.
3. For each word position j in the N_d -long sequence of words d:
   (a) Choose a topic $z_{d,j} \sim \text{Multinomial}(\theta_d)$.
   (b) Choose a word $w_{d,j} \sim \text{Multinomial}(\phi_{z(d,j)})$.

We are interested in finding the objects $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$, where $\beta_{kv} = \Pr(w = v|z = k)$ represents the probability of observing the v-th word in V given topic k. The purpose of the LDA model, then, is to generate posterior topic probabilities for document d:

$$\Pr(\theta, \boldsymbol{z}|d, \boldsymbol{\alpha}, \boldsymbol{\beta}) = \frac{\Pr(\theta, \boldsymbol{z}, d|\boldsymbol{\alpha}, \boldsymbol{\beta})}{\Pr(d|\boldsymbol{\alpha}, \boldsymbol{\beta})}, \tag{4}$$

where $\boldsymbol{z}$ is an N_d-dimensional vector of topic assignments for each of the N_d words in d. The numerator in (4) is given by

$$\Pr(\theta, \boldsymbol{z}, d|\boldsymbol{\alpha}, \boldsymbol{\beta}) = \Pr(\theta|\boldsymbol{\alpha}) \prod_{j=1}^{N_d} \Pr(z_j|\theta)\Pr(w_j|z_j, \boldsymbol{\beta}), \tag{5}$$

which by summing over topics and integrating over θ gives a marginal probability of d:

$$\Pr(d|\boldsymbol{\alpha},\boldsymbol{\beta}) = \int \Pr(\theta|\boldsymbol{\alpha}) \prod_{j=1}^{N_d} \sum_{z(j)} \Pr(z_j|\theta)\Pr(w_j|z_j,\boldsymbol{\beta})\mathrm{d}\theta. \tag{6}$$

Given independence over documents, the probability of observing the entire corpus D is

$$\Pr(D|\boldsymbol{\alpha},\boldsymbol{\beta}) = \prod_{d=1}^{M} \int \Pr(\theta_d|\boldsymbol{\alpha}) \prod_{j=1}^{N_d} \sum_{z(d,j)} \Pr(z_{d,j}|\theta_d)\Pr(w_{d,j}|z_{d,j},\boldsymbol{\beta})\mathrm{d}\theta_d. \tag{7}$$

The corpus-wide objects $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are chosen to maximise this expression. However, due to the functional form of the Dirichlet density, obtaining an analytical solution is intractable.

Instead, I fit the LDA model to my entire corpus with "Gibbs sampling": essentially, this algorithm iteratively augments the topic distribution by repeatedly reassigning terms to different topics based on how likely they are to belong to each topic until the distribution becomes stable (Griffiths and Steyvers, 2004). The model was run and evaluated for different numbers of topics using the LDAvis package, which offers a flexible way to analyse the most frequently occurring terms and visualises topics in a principal-components map.

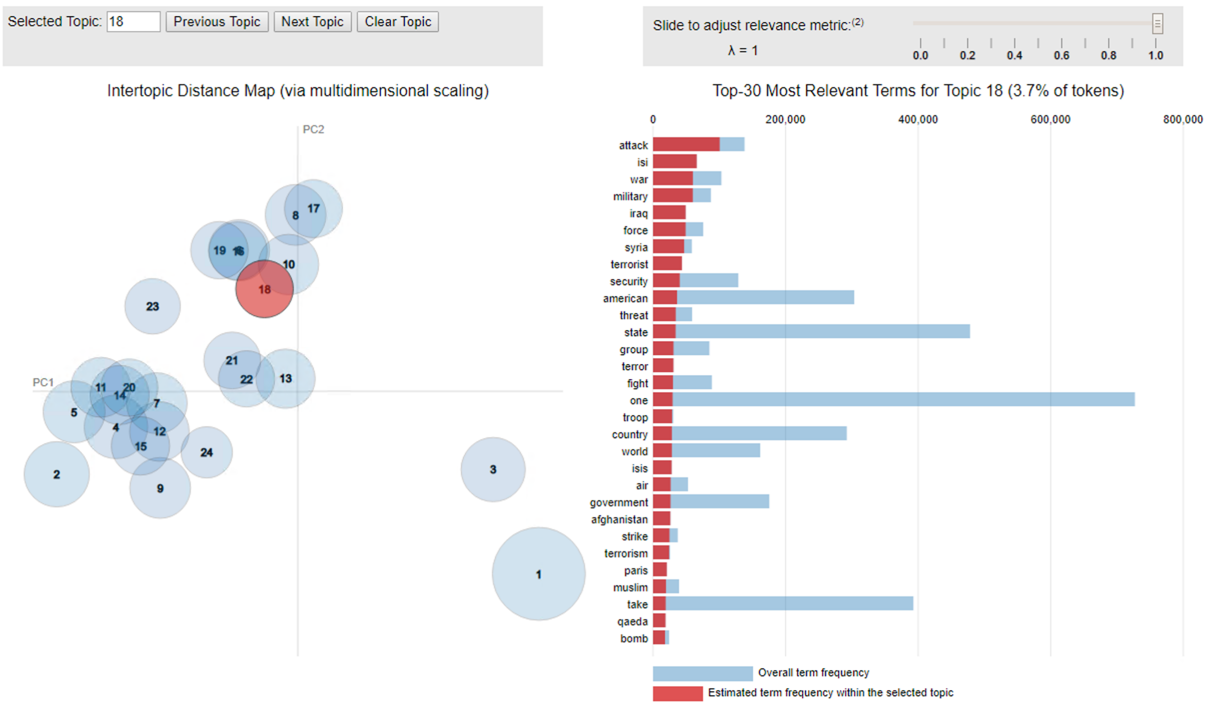

Figure 1: Exploration of LDA model using the LDAvis tool

Finally, posterior topic probabilities were calculated for all transcripts, and transcripts were classified into the topic with the highest posterior probability.

Unfortunately, the LDA model requires the statistician to specify the total number of topics *before* running it, which was unknown ex ante. A widely used measure of goodness-of-fit is the model "coherence score" (Röder et al., 2015), but estimating this would require immense computational power. Instead, different models were run with 12–26 topics and evaluated ex post. Since there is no prediction involved, this does not amount to data mining (Gentzkow et al., 2019). The number of topics was gradually increased until the topic patterns based on top terms did not change meaningfully.

4.2 Classifying and combining topic clusters

The majority of topics could be deciphered readily from top terms. After running several models, I chose one with 24 topics, as this led to granular clustering where most topics were clear. With a granular model, some topics could also be grouped ex post. For example, topics concerned with different elections were grouped into the single topic "elections". Apolitical topics like "sports" and "the weather" were also combined. A high number of transcripts were classified into the "no clear topic" category, but this has not proven to be problematic for the final results (see 9.1). Judging from Table 1, the model has done a defensible job of classifying transcripts despite its minimalist foundation.

Table 1: LDA topics for all 617,726 transcripts

Topic guess	**No. of transcripts**	**Important top terms**
Health emergency	21,987	hospital, crisis, disease, pandemic, vaccines
Middle East/terrorism	33,569	Iraq, Libya, Taliban, terror, Isis, islamic
Foreign policy	30,171	Iran, Korea, nuclear, diplomacy, treaty
Racism/social	21,100	racism, black, jew, bigotry, religion
Washington*	51,805	Senate, speaker, filibuster, president, oval
Family/school	15,852	mother, father, school, child, teacher
Russian interference	41,471	FBI, Russian, security, Mueller, justice
General/weather/sports*	90,182	sports, NFL, tonight, storm, weather
Guns/violence	24,239	gun, NRA, shooter, murder, suspect, police
Healthcare policy	24,467	Medicaid, Obamacare, tax, premium, insured
Elections*	69,181	Trump, Hillary, Romney, campaign, vote
Legal/supreme court	18,324	court, law, supreme, abortion, justice, rule
Economy	20,569	job, money, business, stock, wage, investment
Immigration	13,511	immigrant, illegals, border, asylum, citizen
No clear topic*	141,301	issue, country, say, think, good, essentially

*topic was combined from a set of smaller topics provided by the LDA model

5 Document scoring with sentiment analysis

5.1 Measuring emotive intensity

To quantify the presence of emotive rhetoric in transcripts, I employ lexicon-based sentiment analysis. I use the *NRC Emotion-Intensity Lexicon* by Mohammad (2018) to score transcripts. The lexicon was built by a diverse group of human annotators and provides an index of 10,170 words with two key descriptors: firstly, it specifies the emotion a given word is associated with out of four negative emotions (anger, disgust, fear, sadness) and four positive emotions (anticipation, joy, surprise, trust). I omit "trust", as its rhetorical interpretation may be considered ambiguous. Secondly, for each word, it provides an "intensity" score between 0 and 1 indicating the *strength* of the emotion the word conveys.

5.2 Scoring algorithm

Each document receives seven scores for each emotion according to the following algorithm. Let W_e be the set of identified words in document d that belong to emotion e in the lexicon. Each identified word $w \in \{1, 2, ..., |W_e|\}$ has an intensity score of $s_e(w)$ for this particular emotion. The intensity score for emotion e in document d is simply

$$S_d^e = \begin{cases} 0 & \text{if } W_e = \emptyset \\ \frac{1}{|W_e|} \sum_{w=1}^{|W_e|} s_e(w) & \text{otherwise.} \end{cases}$$

In essence, the score captures the average intensity of the emotion amongst the words linked to e. The approach is inspired by equivalent intensity scoring, such as the one used in Sharma et al. (2015). Note, setting the score equal to zero if no terms associated with e are detected is not innocuous: given the counting nature of the data, shorter documents are more likely to receive a score of zero (implications for inference addressed in 7.3.3). The tables below give summary statistics and examples of lexicon terms by emotion.

Table 2: Summary statistics for emotive intensity scores

	Negative				Positive		
	anger	disgust	fear	sadness	anticipation	joy	surprise
Mean	0.369	0.274	0.376	0.335	0.438	0.399	0.314
Std. Dev.	0.242	0.227	0.218	0.232	0.175	0.184	0.197
% zeros	21.2%	32.8%	15.5%	22.9%	12.1%	13.0%	19.9%
Transcripts	617,726	617,726	617,726	617,726	617,726	617,726	617,726

Table 3: Examples of NRC lexicon entries by emotion and intensity

Emotion	**Score**	**Example of high- and low-scoring words by emotion**
anger	Higher	outraged, brutality, hatred, furious, enraged, loathe, vicious, liar
	Lower	vexed, bitter, jealousy, conflict, frustration, annoying, irritated
disgust	Higher	cannibalism, perverted, massacre, slaughter, sewerage, filth
	Lower	dreadful, appalling, sneer, greedy, shame, cringe, impure, mess
fear	Higher	torture, horrific, terror, kill, holocaust, assassinate, doomed
	Lower	worry, risky, unsafe, precarious, anxious, warn, wary, careful
sadness	Higher	mourning, heartbreaking, tragic, depressing, bereaved
	Lower	lone, frown, unfavourable, adversity, negative, unlucky
anticipation	Higher	excited, eager, adventure, thrilling, hopeful, climax, urgent
	Lower	impatient, endeavour, plan, expected, fun, curious, await
joy	Higher	bliss, jubilant, elation, exuberance, love, dance, overjoyed
	Lower	good, comfy, decent, comfort, leisure, friend, chocolate
surprise	Higher	flabbergast, ambush, eruption, startle, alarmed, astonishing
	Lower	expect, fortunate, merriment, coincidence, unintentional

5.3 Evaluation of scoring algorithm

With the above approach, document scores are normalised by calculating the average intensity score of the terms associated with a given emotion. This is done, firstly, to ensure that longer documents are not scored systematically differently from shorter documents. Secondly, since the presence of one emotion mostly implies the absence of another, normalising by the total number of words would make scores negatively correlated, which would complicate making comparisons between emotion scores.

Some aspects of the scoring algorithm may give rise to reservations. The scoring method rests on the assumption that terms linked to a given emotion indicate the presence of that emotion in messaging. This is admittedly naive given that terms are not considered in syntactical context but as a mere "bag of words" (Gentzkow et al., 2019). More sophisticated NLP approaches have been developed to account for this, but due to their algorithmic complexity, the output is more difficult to interpret. This is the primary reason why I adopted a simple lexicon-based approach: the more complex the scoring methodology becomes, the more it will approach a black box with limited scope for external critique. Since I am principally concerned with drawing causal inference about messaging, the merits of adopting more sophisticated NLP tools had to be balanced against the degree to which they compromise methodological transparency.

6 Summary statistics

6.1 Transcript scores by speaker groups

For all *commentators* in the sample, Table 4 compares conditional means of emotion scores for each channel by commentator ideology.

Table 4: Mean scores for commentators by group and channel (with standard errors)

	MSNBC		CNN		Fox	
	Democrat	**Republican**	**Democrat**	**Republican**	**Democrat**	**Republican**
anger	0.359***	0.312	0.371***	0.333	0.339	0.348***
	(0.0010)	(0.0012)	(0.0010)	(0.0012)	(0.0012)	(0.0010)
disgust	0.272***	0.218	0.274***	0.235	0.247	0.260***
	(0.0009)	(0.0010)	(0.0009)	(0.0010)	(0.0011)	(0.0009)
fear	0.367***	0.327	0.373***	0.343	0.353	0.358***
	(0.0009)	(0.0011)	(0.0009)	(0.0011)	(0.0011)	(0.0009)
sadness	0.325***	0.271	0.334***	0.292	0.303	0.311***
	(0.0008)	(0.0009)	(0.0008)	(0.0009)	(0.0010)	(0.0008)
anticipation	0.442***	0.414	0.432***	0.413	0.339	0.423***
	(0.0007)	(0.0009)	(0.0007)	(0.0009)	(0.0009)	(0.0008)
joy	0.410***	0.387	0.401***	0.383	0.374	0.391***
	(0.0007)	(0.0010)	(0.0007)	(0.0009)	(0.0010)	(0.0008)
surprise	0.304***	0.286	0.304***	0.291	0.268	0.292***
	(0.0009)	(0.0011)	(0.0009)	(0.0011)	(0.0012)	(0.0010)
Transcripts	55,042	48,045	61,104	47,518	44,276	59,471
People	398	305	398	305	398	305

***For this channel, this political group has a higher mean score at 1% significance

We notice that average intensity scores are greater for Democrats on MSNBC and CNN. In contrast, scores are comparatively higher for Republicans on Fox. Without further controls, one cannot determine whether this pattern can be explained by a selection effect due to commentator heterogeneity or other confounding factors.

To provide further context, I also compare scores between anchors and commentators. For compactness, Table 5 displays means across *negative* emotions as defined by the NRC lexicon (anger, disgust, fear, sadness) and *positive* emotions (anticipation, joy, surprise). Notably, on all channels, anchors have higher average intensity scores. One explanation is that anchors may be chosen to report sensational stories whilst commentators are expected to provide sober analysis.

Table 5: Mean scores by speaker group and channel (with standard errors)

	MSNBC		CNN		Fox	
	Anchor	**Commentator**	**Anchor**	**Commentator**	**Anchor**	**Commentator**
negative	0.344***	0.308	0.367***	0.322	0.370***	0.315
	(0.0007)	(0.0006)	(0.0005)	(0.0006)	(0.0005)	(0.0006)
positive	0.401***	0.375	0.399***	0.372	0.396***	0.361
	(0.0005)	(0.0004)	(0.0004)	(0.0004)	(0.0004)	(0.0004)
Transcripts	67,491	103,087	121,224	108,622	113,555	103,747
People	274	703	274	703	274	703

***For this channel, this speaker group has a higher mean score at 1% significance

6.2 Transcript scores by news topics

Finally, the figure below demonstrates how average transcript intensity scores, grouped by positive and negative emotions, vary by topic across the entire sample. Whilst scores for positive emotions are rather uniform, there is greater variation for negative emotions: "gun violence", "racism", and "terrorism" have strong negative intensity. This aligns with expectations and thus gives further confidence in the scoring approach.

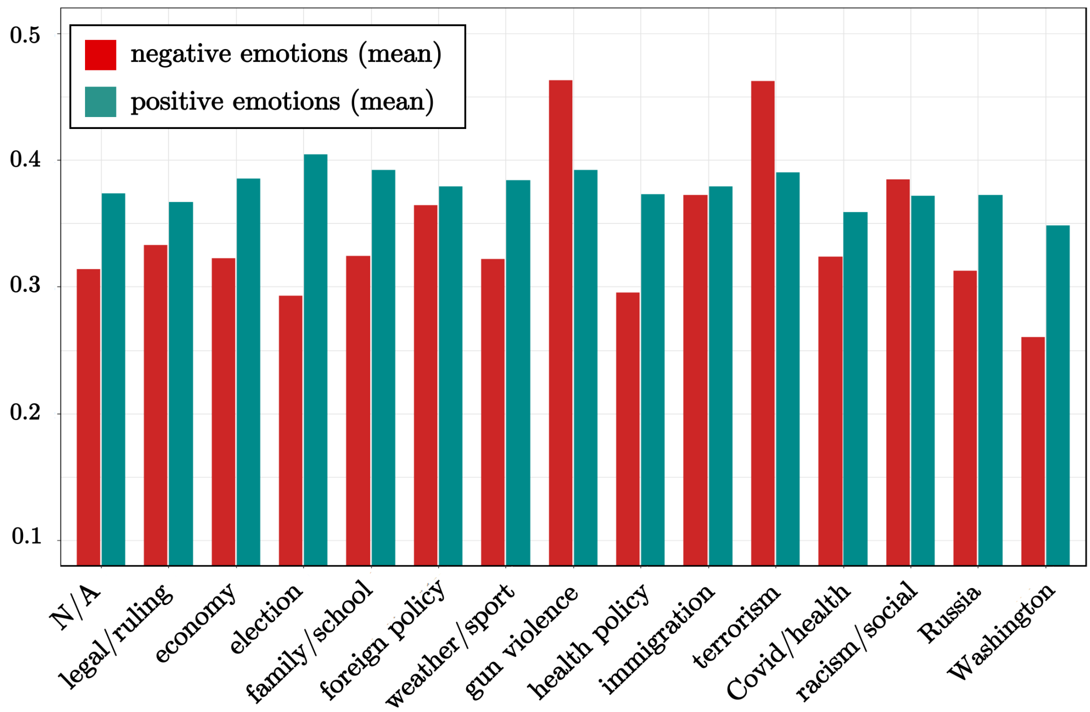

Figure 2: Emotive intensity scores by LDA topic clusters

To tackle confounding variables, the next section develops an econometric model to examine strategic communication, concentrating on the transcripts of commentators appearing on both Fox and MSNBC.

7 Empirical strategy

At the heart of this study's identification strategy is the fact that the *same* commentator features on both Fox and MSNBC. When testing for messaging tactics, this allows for the elimination of time-invariant selection bias. To ensure sufficient variation by individual, I restrict the sample of commentators to the 658 for whom I have at least 20 observations. For the main analysis, I also ignore CNN transcripts (analysed in 9.3 for completeness).

7.1 A naive specification

Consider a group of commentators who are frequently invited to deliver commentary on Fox *and* MSNBC. Suppose each commentator has an incentive to adjust the intensity of emotive rhetoric between the two channels and that this incentive may differ depending on their party affiliation. Formally, assume commentator i speaks on a programme p. Let $k = k(p)$ be the topic of programme p and $t = t(p)$ the date of transmission. For emotion e, commentator i on programme p speaks with an average emotive intensity given by

$$S^e_{i,p} = \beta_0 + \beta_1 Rep_i + \beta_2 Fox_p + \beta_3 Rep_i \times Fox_p + \overbrace{\alpha_i + \lambda_{k(p)} + \theta_{t(p)} + \varepsilon_{i,p}}^{u_{i,p}}, \tag{8}$$

where α_i is an unobserved fixed commentator effect, $\lambda_{k(p)}$ is a topic effect, $\theta_{t(p)}$ captures a time trend, and $\varepsilon_{i,p}$ is an idiosyncratic error. $Fox_p = 1$ indicates that the transcript is from a Fox segment. $Rep_i = 1$ means the commentator is a Republican, and $Rep_i = 0$ indicates they are a Democrat.

In (8), β_2 measures the change in the intensity score, all else equal, when a Democrat speaks on Fox compared to when they speak on MSNBC. For a Republican, the total channel effect is $\beta_2 + \beta_3$. If $\text{Cov}(Fox_p, u_{i,p}) = \text{Cov}(Rep_i \times Fox_p, u_{i,p}) = 0$, pooled OLS estimators of these two parameters of interest are consistent, incorporating both within- and between-commentator variation for efficiency gains.

However, contemporaneous exogeneity of $u_{i,p}$ is unlikely to hold. Firstly, the frequency with which a commentator is invited onto Fox is correlated with the topic effect $\lambda_{k(p)}$ for the news segment. Secondly, rhetoric and relative channel invitations may both exhibit time trends on average. To control for this, one should include dummies for the 15 LDA topics and calendar years. Clustered standard errors at the individual level should be applied to adjust for within-commentator autocorrelation (Bertrand et al., 2004).

7.2 Fixed effects model: testing for strategic messaging

Crucially, the above specification ignores the impact of commentator heterogeneity bias. Heterogeneity in unobserved commentator characteristics in α_i is likely associated with whom the network invites to appear on a segment. Estimators incorporating between-variation will be biased due to this selection effect. To establish a reliable test for strategic message delivery, the model should only focus on within-variation for commentators. For this reason, we consider the specification with individual commentator dummies

$$S^e_{i,p} = \gamma + \beta_2 Fox_p + \beta_3 Rep_i \times Fox_p + \underbrace{\sum_{k=2}^{15} \delta_k T_k}_{\text{topics}} + \underbrace{\sum_{t=1}^{12} \tau_t Y_{2010+t}}_{\text{years}} + \underbrace{\sum_{i=2}^{658} \omega_i D_i}_{\text{commentators}} + v_{i,p}, \tag{9}$$

which is equivalent to a fixed effects model. By controlling for commentator fixed effects in this way, estimating (9) amounts to measuring the "within-individual", "within-topic" effect of speaking on Fox. A Wald test for groupwise homoskedasticity is rejected at a 1% significance level (Table 8), implying the need for robust standard errors.

7.3 Further model considerations

7.3.1 Bias-efficiency trade-off

The fixed effects estimator requires strict exogeneity with the Fox dummy for unbiasedness, but there is little reason to believe that emotionality in messaging is related to future or past channel appearances. However, discarding between-variation is costly in mean-squared-error terms. A Hausman test for equality of FE and RE coefficients is rejected for all seven regressions at 1% significance, suggesting the presence of commentator heterogeneity bias (Table 8). Given the large sample size, efficiency is arguably less important, so the fixed effects estimates (Table 7) constitute my main results.

7.3.2 Z-scoring dependent variables

Given the arbitrary scaling of the intensity scores, interpretations of coefficient magnitudes are not readily meaningful. Therefore, to contextualise the intensity scores within this dataset, I standardise the dependent variables by subtracting their global means and dividing by their global standard deviations such that coefficient estimates are measured in terms of standard deviations of document scores across the whole corpus.

7.3.3 Potential bias from document length

As noted in 5.2, the counting nature of the scoring algorithm makes it slightly sensitive to differences in transcript length. This would distort estimates if the average transcript length differs systematically by channel–ideology combinations. Overall, MSNBC transcripts are slighly longer on average and relatively longer for Democratic commentators. In contrast, Republican transcripts are relatively longer than Democratic transcripts on Fox. Consequently, I control for document length to address this.

8 Results and discussion

Table 6: Pooled OLS regression results (z-scored dependent variables)

	anger	disgust	fear	sadness	anticipation	joy	surprise
controls⋆	...	...	...	...	...	...	...
Rep	–0.122***	–0.177***	–0.101***	–0.155***	–0.144***	–0.122***	–0.110***
	(0.032)	(0.023)	(0.024)	(0.022)	(0.022)	(0.020)	(0.020)
Fox	–0.071***	–0.109***	–0.056***	–0.084***	–0.188***	–0.183***	–0.171***
	(0.019)	(0.023)	(0.022)	(0.020)	(0.017)	(0.022)	(0.019)
Fox×Rep	0.170***	0.255***	0.144***	0.211***	0.240***	0.233***	0.222***
	(0.028)	(0.031)	(0.026)	(0.030)	(0.024)	(0.026)	(0.025)
constant	–0.490***	–0.350***	–0.307***	–0.258***	–0.261***	–0.148***	–0.410***
	(0.047)	(0.048)	(0.049)	(0.051)	(0.038)	(0.040)	(0.041)
N	206,461	206,461	206,461	206,461	206,461	206,461	206,461
R^2	0.126	0.149	0.104	0.122	0.049	0.061	0.068

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$ (clustered standard errors in parentheses)
⋆includes 14 topic dummies, 12 year dummies, and transcript length

Table 7: Fixed effects regression results (z-scored dependent variables)

	anger	disgust	fear	sadness	anticipation	joy	surprise
controls	...	...	...	...	...	...	...
Fox	–0.065***	–0.090***	–0.057***	–0.083***	–0.188***	–0.185***	–0.155***
	(0.016)	(0.016)	(0.015)	(0.016)	(0.014)	(0.014)	(0.014)
Fox×Rep	0.137***	0.193***	0.112***	0.171***	0.234***	0.227***	0.191***
	(0.024)	(0.022)	(0.022)	(0.024)	(0.021)	(0.020)	(0.018)
Average FE	–0.520***	–0.433***	–0.345***	–0.388***	–0.361***	–0.270***	–0.489***
	(0.028)	(0.030)	(0.028)	(0.029)	(0.032)	(0.035)	(0.030)
$\text{Fox}^{\star\star}_{Rep}$	0.071***	0.103***	0.055***	0.089***	0.046***	0.042***	0.036***
	(0.017)	(0.015)	(0.016)	(0.018)	(0.015)	(0.014)	(0.011)
N	206,461	206,461	206,461	206,461	206,461	206,461	206,461
Within R^2	0.084	0.102	0.072	0.080	0.036	0.040	0.047

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$ (robust standard errors in parentheses)
⋆⋆implied total channel effect for Republicans ($\hat{\beta}_2 + \hat{\beta}_3$)

Table 8: Diagnostics test statistics linked to fixed effects regressions

	anger	disgust	fear	sadness	anticipation	joy	surprise
Wald stat	22112***	11758***	21676***	19473***	29461***	30083***	20853***
Hausman stat	216.5***	208.9***	146.6***	197.3***	62.9***	144.1***	143.9***

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

The main results of the dissertation are the fixed effects estimates in Table 7, as these control for commentators' selection into channels on time-invariant unobservables. Note that the *total* channel effect for Republicans $\hat{\beta}_2 + \hat{\beta}_3$ had to be calculated separately with an equivalent regression. The results demonstrate a range of notable patterns regarding how commentators change their emotive rhetoric by channel.

Firstly, the Fox coefficient is statistically significant for both groups of commentators across all emotions. Even without observing the signs of the coefficients, this is an important result: statistical significance of the channel dummy indicates that the same commentator changes their rhetoric in response to the audience, controlling for the broad news topic. By itself, this finding offers evidence against any theory of news commentary that predicts no scope for strategic behaviour in messaging.

Secondly, we notice that, within each political group, the sign of the Fox dummy is the same for all regressions with different emotion scores as the dependent variable. This result is rather surprising and highlights the fact that it is not the emphasis on the *type* of emotive rhetoric that underlies the channel adjustment for a given commentator. One might expect that commentators would have an incentive to tone down the intensity of some emotions on a given channel in order to emphasise others. Instead, commentators adjust the intensity of *all* emotions depending on the audience being addressed.

Thirdly, and most strikingly, the direction of this adjustment mechanism is highly meaningful when broken down by party affiliation of commentators: the same Democratic-leaning commentator tones down the intensity of emotive rhetoric when speaking to a Fox audience relative to when speaking on MSNBC. In contrast, the same Republican-leaning commentator amplifies the intensity of emotive rhetoric on Fox, reflecting heterogeneity in incentives. On average across the seven emotions, Democratic commentators decrease their intensity by 12% of one global standard deviation on Fox whilst Republican commentators increase the intensity by 6% of one global standard deviation on Fox. We can interpret this as a *stronghold effect*: commentators consistently increase the intensity of emotive rhetoric when speaking to an audience whose ideological bias is congruent with their own political ideology. Given this systematic difference between the two commentator groups, the results are a strong indicator that commentators are politically motivated and alter their message delivery strategically in response to the audience.

A simple way to illustrate this stronghold effect is the following scatter plot. Each commentator represents a point with two coordinates, showing the average intensity scores (across all emotions) on MSNBC and Fox, respectively. Points situated on the 45-degree line represent commentators whose mean scores are identical across the two channels.

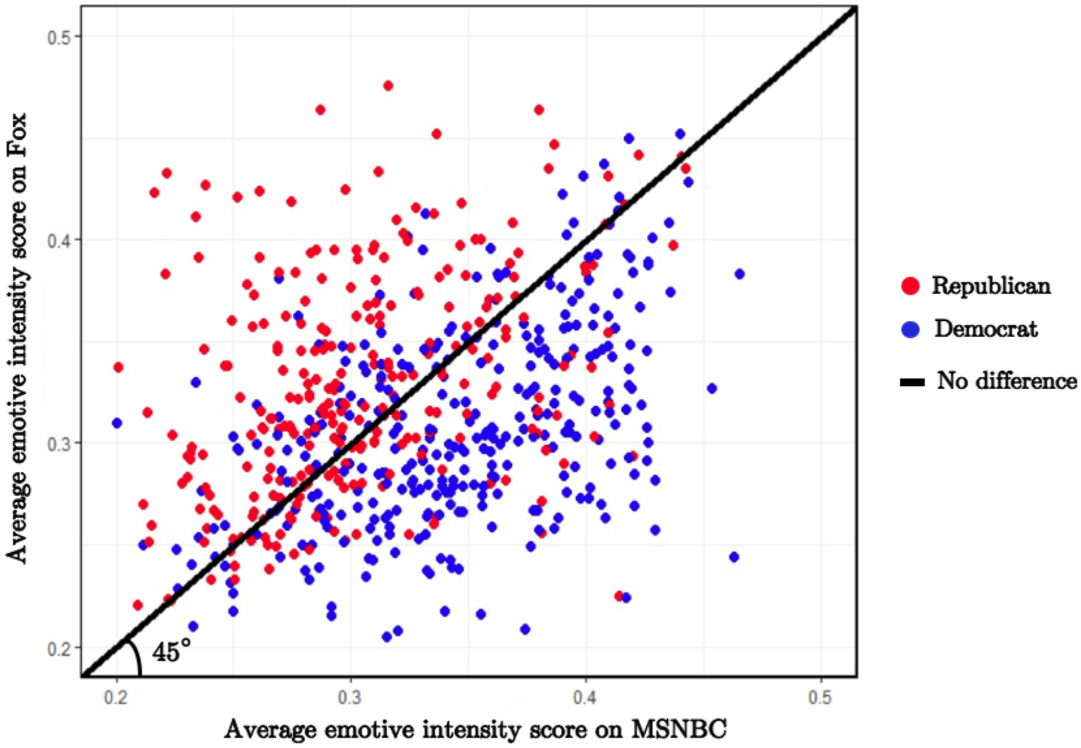

Figure 3: Scatter plot of intensity scores (averaged over all emotions)

The scatter plot effectively illustrates the stronghold effect. Firstly, the 45-degree line passes through the bulk of the points, suggesting that the stronghold effect is relatively symmetric across the two channels. Secondly, on average, Democrats clearly intensify their emotive rhetoric on MSNBC (blue points cluster below the line) whilst Republicans increase their emotive intensity on Fox (red points cluster above the line).

There are two plausible ways to frame this striking behaviour. The first explanation would emphasise strategic incentives of politically motivated commentators as discussed in 2.2: in this framework, commentators choose uninformative emotive rhetoric when speaking to strongholds to ensure viewers hold onto their prior beliefs. Insofar as emotional rhetoric is less informative, this would undermine information efficiency because commentators withhold information when it is unfavourable to their side. As captured by inequality (3) in 2.2.3, the more a given channel's audience insists on holding onto its prior beliefs, the larger is the inefficiency implied by commentators' strategic appeals to emotion. The combination of partisan divides and strategic messaging may therefore undermine the media's core mission of information provision. It should, however, be stressed that this theory is but one of several ways to conceptualise the amorphous role of emotive rhetoric. The results may be equally consistent with strategic frameworks where, for example, emotional messaging is persuasive even when it contains no information.

Alternatively, the stronghold effect may be related to systematic patterns in the interaction between partisan anchors and commentators during interviews: one possibility is that news anchors deliberately frame questions to stir an emotional reaction when commentator ideology matches channel slant. A counterargument to this critique is that a commentator's ideology is not always public knowledge. In this case, it was derived from private campaign contributions.

Fourthly, it is relevant to compare coefficient *magnitudes* by ideology. The table below provides p-values for a set of z-tests for equal coefficient magnitudes on the Fox dummies between Democrats and Republicans.

Table 9: Tests for differences in coefficient magnitudes by ideology

	Negative				Positive		
	anger	disgust	fear	sadness	anticipation	joy	surprise
p-value	0.797	0.553	0.927	0.803	0.000	0.000	0.000

H_0: Magnitudes of total effect of speaking on Fox are equal

We notice yet another consistent pattern: for all positive emotions as defined by the NRC lexicon (anticipation, joy, surprise), the magnitude of the total channel effect is greater for Democrats than for Republicans. For all negative emotions, the magnitudes are not statistically different at any reasonable level of significance. Thus, the nature of strategic messaging differs somewhat by ideology, with Democrats exhibiting a greater extent of message adjustment for positive emotions.

Taken together, the results support the view that emotional messaging is endogenous with the audience's political profile. The theoretical framework presented in 2.2 provides one explanation borne out by the data rooted in political persuasion, with definite efficiency implications. As the principal contribution of this dissertation is empirical, future work may wish to establish a more mechanistic understanding of the incentive structure underlying these salient findings.

9 Robustness checks

9.1 Restricting the sample of topics

For further robustness, I also estimate the fixed effects model where I discard documents that were classified into the "no clear topic" category in 4.2. This is to ensure that this classification procedure does not challenge the overall results.

Table 10: FE channel effects by group with restricted sample of topics

	anger	disgust	fear	sadness	anticipation	joy	surprise
Fox_{Dem}	–0.058***	–0.090***	–0.052***	–0.084***	–0.186***	–0.154***	–0.162***
	(0.015)	(0.015)	(0.015)	(0.016)	(0.015)	(0.014)	(0.014)
Fox_{Rep}	0.066***	0.102***	0.048***	0.083***	0.034**	0.038***	0.033***
	(0.017)	(0.015)	(0.016)	(0.018)	(0.015)	(0.015)	(0.011)
N	171,686	171,686	171,686	171,686	171,686	171,686	171,686

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$ (robust standard errors in parentheses)

Reassuringly, restricting the sample to transcripts outside this category does not lead to systematic differences with the estimates in Table 7, and coefficients all remain significant at the 5% level.

9.2 Day-of-the-week effects

The framing selected for a given topic in news commentary may vary considerably based on the day of the week: channels may choose to address lighter issues during the weekend and systematically invite commentators in light of this. As shown below, the results are robust when including day-of-the-week dummies, with negligible implied differences for directions and magnitudes of point estimates.

Table 11: FE channel effects by group with day-of-the-week controls

	anger	disgust	fear	sadness	anticipation	joy	surprise
Fox_{Dem}	–0.065***	–0.089***	–0.056***	–0.082***	–0.161***	–0.184***	–0.154***
	(0.016)	(0.016)	(0.015)	(0.016)	(0.014)	(0.014)	(0.014)
Fox_{Rep}	0.072***	0.104***	0.055***	0.089***	0.063***	0.042***	0.037***
	(0.017)	(0.015)	(0.016)	(0.018)	(0.013)	(0.014)	(0.011)
N	206,461	206,461	206,461	206,461	206,461	206,461	206,461

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$ (robust standard errors in parentheses)

9.3 Including CNN transcripts

Since CNN does not have an obvious partisan profile, CNN transcripts were left out of the main analysis focused on audience bias. Nevertheless, for transparency, total channel effects by group for a separate fixed effects model with CNN transcripts are shown below.

Table 12: FE channel effects by group (including CNN)

	anger	disgust	fear	sadness	anticipation	joy	surprise
Fox_{Dem}	–0.065***	–0.090***	–0.057***	–0.083***	–0.188***	–0.185***	–0.155***
	(0.016)	(0.016)	(0.015)	(0.016)	(0.014)	(0.014)	(0.014)
Fox_{Rep}	0.071***	0.103***	0.055***	0.089***	0.046***	0.042***	0.036***
	(0.017)	(0.014)	(0.016)	(0.018)	(0.015)	(0.013)	(0.011)
CNN_{Dem}	–0.012	–0.037***	–0.018*	–0.015	–0.056***	–0.060***	–0.021*
	(0.010)	(0.014)	(0.016)	(0.018)	(0.011)	(0.013)	(0.011)
CNN_{Rep}	0.028**	0.009	0.012	0.022*	–0.029**	–0.018	0.034***
	(0.012)	(0.011)	(0.011)	(0.012)	(0.013)	(0.013)	(0.012)
N	315,344	315,344	315,344	315,344	315,344	315,344	315,344

$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$ (robust standard errors in parentheses)

We notice a lack of consistent statistical significance for CNN, in contrast to the Fox effect. This is arguably not surprising given the channel's neutral profile. Point estimates on the CNN dummy for Democrats are, notably, all negative, and the majority of Republican coefficients are positive. Insofar as CNN is *relatively* more conservative than MSNBC and less conservative than Fox, this would still be consistent with the stronghold effect.

10 Conclusion

Understanding how the media formulate messaging to maximise political influence is of substantial democratic importance. This dissertation offers novel field evidence of how strategic messaging incentives operate in partisan cable news. By constructing a new transcript collection from facial detection data, the study makes three empirical contributions to extant scholarship on the political incentives of the media.

Firstly, the results show that the same commentator adjusts the intensity of emotive rhetoric between partisan channels. This provides a strong case for regarding actors in cable news media as strategic political agents. Secondly, the direction of adjustment represents a consistent *stronghold effect*: the same commentator increases the intensity of emotive rhetoric on the partisan channel that caters to viewers with the same ideology. The effect is consistent and statistically significant for all independent emotion dimensions considered. Thirdly, the study makes a methodological leap in political economy research by exemplifying how topic modelling and sentiment analysis can be applied jointly to uncover politically motivated behaviour.

Unlike existing studies, my identification strategy circumvents problems of selection bias. Whereas previous research has been confined to case studies of specific news topics, I significantly expand the sample by considering several topics and explicitly modelling them with unsupervised machine learning. The sample size, then, facilitates identification of channel differences for the same commentator, thus providing a credible basis for causal inference about strategic messaging.

When contextualised within current trends of deepening polarisation, my findings serve as a timely reminder of the need for a nuanced appreciation of the tactics media entities employ to wield their influence, and the profound implications such insights have for the health of our political institutions.

Bibliography

Anand, B., Di Tella, R. and Galetovic, A. (2007), 'Information or Opinion? Media Bias as Product Differentiation', *Journal of Economics and Management Strategy* **16**(3), pp. 635–682.
URL: *https://doi.org/10.1111/j.1530-9134.2007.00153.x*

Benchimol, J., Kazinnik, S. and Saadon, Y. (2022), 'Text mining methodologies with R: An application to central bank texts', *Machine Learning with Applications* **8**, 100286.
URL: *https://doi.org/10.1016/j.mlwa.2022.100286*

Bertrand, M., Duflo, E. and Mullainathan, S. (2004), 'How much should we trust differences-in-differences estimates?', *The Quarterly Journal of Economics* **119**(1), pp. 249–275.
URL: *https://www.jstor.org/stable/25098683*

Blei, D. M., Ng, A. Y. and Jordan, M. T. (2003), 'Latent dirichlet allocation', *The Journal of Machine Learning Research* **3**, pp. 993–1022.
URL: *https://dl.acm.org/doi/10.5555/944919.944937*

Bonica, A. (2016), 'Database on Ideology, Money in Politics, and Elections: Public version 2.0'. Accessed: 2023-01-11.
URL: *https://data.stanford.edu/dime*

Brader, T. (2005), 'Striking a Responsive Chord: How Political Ads Motivate and Persuade Voters by Appealing to Emotions', *American Journal of Political Science* **49**(2), pp. 388–405.
URL: *https://doi.org/10.2307/3647684*

Brambor, T., Clark, W. R. and Golder, M. (2006), 'Understanding interaction models: Improving empirical analyses', *Political Analysis* **14**(1), pp. 63–82.
URL: *https://doi.org/10.1093/pan/mpi014*

Broockman, D. E. and Kalla, J. L. (2023), 'Consuming cross-cutting media causes learning and moderates attitudes: A field experiment with Fox News viewers'. Preprint.
URL: *https://doi.org/10.31219/osf.io/jrw26*

DellaVigna, S. and Gentzkow, M. (2010), 'Persuasion: Empirical evidence', *Annual Review of Economics* **2**(1), pp. 643–669.
URL: *https://doi.org/10.1146/annurev.economics.102308.124309*

Dellavigna, S. and Kaplan, E. (2006), 'The Fox News Effect: Media Bias and Voting', *National Bureau of Economic Research* . Working paper.
URL: *https://doi.org/10.3386/w12169*

Djourelova, M. (2023), 'Persuasion through Slanted Language: Evidence from the Media Coverage of Immigration', *American Economic Review* **113**(3), pp. 800–835.
URL: *https://doi.org/10.1257/aer.20211537*

Enikolopov, R., Petrova, M. and Zhuravskaya, E. (2011), 'Media and political persuasion: Evidence from Russia', *American Economic Review* **101**(7), pp. 3253–3285.
URL: *https://doi.org/10.1257/aer.101.7.3253*

Federal Communications Commission (2021), 'The FCC and Freedom of Speech'. Accessed: 2023-02-19.
URL: *https://www.fcc.gov/consumers/guides/fcc-and-speech*

Gentzkow, B. M. and Shapiro, J. M. (2010), 'What Drives Media Slant? Evidence From U.S. Daily Newspapers', *Econometrica* **78**(1), pp. 35–71.
URL: *https://doi.org/10.3982/ECTA7195*

Gentzkow, M., Kelly, B. and Taddy, M. (2019), 'Text as data', *Journal of Economic Literature* **57**(3), pp. 535–574.
URL: *https://doi.org/10.1257/jel.20181020*

Griffiths, T. L. and Steyvers, M. (2004), 'Finding scientific topics', *Proceedings of the National Academy of Sciences of the United States of America* **101**, pp. 5228–5235.
URL: *https://doi.org/10.1073/pnas.0307752101*

Hagmann, D. and Loewenstein, G. (2018), 'Persuasion with motivated beliefs'. Working paper.
URL: *https://www.aeaweb.org/conference/2019/preliminary/paper/5ytibhd*

Hong, J., Crichton, W., Zhang, H., Fu, D. Y., Ritchie, J., Barenholtz, J., Hannel, B., Yao, X., Murray, M., Moriba, G., Agrawala, M. and Fatahalian, K. (2021a), 'Analysis of faces in a decade of US cable TV news', *Proceedings of the 27th ACM SIGKDD*

Conference on Knowledge Discovery & Data Mining, pp. 3011–3021.
URL: *https://doi.org/10.1145/3447548.3467134*

Hong, James and Crichton, Will and Zhang, Haotian and Fu, Daniel Y. and Ritchie, Jacob and Barenholtz, Jeremy and Hannel, Ben and Yao, Xinwei and Murray, Michaela and Moriba, Geraldine and Agrawala, Maneesh and Fatahalian, Kayvon (2021b), 'Stanford Cable News Analyzer'. Accessed: 2023-01-04.
URL: *https://tvnews.stanford.edu/*

Kamenica, E. and Gentzkow, M. (2011), 'Bayesian persuasion', *American Economic Review* **101**(6), pp. 2590–2615.
URL: *https://doi.org/10.1257/aer.101.6.2590*

Karell, D. and Agrawal, A. (2022), 'Small town propaganda: The content and emotions of politicized digital local news in the United States', *Poetics* **92**, 101641.
URL: *https://doi.org/10.1016/j.poetic.2021.101641*

Kim, E., Lelkes, Y. and McCrain, J. (2022a), 'Measuring dynamic media bias', *Proceedings of the National Academy of Sciences of the United States of America* **119**(32), pp. 3–5.
URL: *https://doi.org/10.1073/pnas.2202197119*

Kim, E., Lelkes, Y. and McCrain, J. (2022b), 'Replication Data for: Measuring Dynamic Media Bias', Harvard Dataverse.
URL: *https://doi.org/10.7910/DVN/I2K54B*

MacDonald, Roger (2009), 'TV News Archive'. Accessed: 2023-01-09.
URL: *https://archive.org/details/tv*

Martin, G. J. and Yurukoglu, A. (2017), 'Bias in cable news: Persuasion and polarization', *American Economic Review* **107**(9), pp. 2565–2599.
URL: *https://doi.org/10.1257/aer.20160812*

Mohammad, S. M. (2018), 'Word affect intensities', *LREC 2018 - 11th International Conference on Language Resources and Evaluation,* pp. 174–183.
URL: *https://doi.org/10.48550/arXiv.1704.08798*

Pew Research Center (2020), 'Americans' main sources for political news vary by party and age'. Accessed: 2022-10-18.

URL: *https://www.pewresearch.org/short-reads/2020/04/01/americans-main-sources-for-political-news-vary-by-party-and-age/ft_2020-04-01_newssources_04/*

Renshon, J., Lee, J. J. and Tingley, D. (2015), 'Physiological Arousal and Political Beliefs', *Political Psychology* **36**(5), pp. 569–585.
URL: *https://doi.org/10.1111/pops.12173*

Röder, M., Both, A. and Hinneburg, A. (2015), 'Exploring the space of topic coherence measures', *WSDM 2015 - Proceedings of the 8th ACM International Conference on Web Search and Data Mining,* pp. 399–408.
URL: *https://doi.org/10.1145/2684822.2685324*

Rozado, D. and Al-Gharbi, M. (2022), 'Using word embeddings to probe sentiment associations of politically loaded terms in news and opinion articles from news media outlets', *Journal of Computational Social Science* **5**(1), pp. 427–448.
URL: *https://doi.org/10.1007/s42001-021-00130-y*

Sharma, R., Gupta, M., Agarwal, A. and Bhattacharyya, P. (2015), 'Adjective intensity and sentiment analysis', *Conference Proceedings - EMNLP 2015: Conference on Empirical Methods in Natural Language Processing* (September), pp. 2520–2526.
URL: *https://doi.org/10.18653/v1/D15-1300*

Stigler, G. J. (1961), 'The economics of information', *Journal of Political Economy* **69**(3), pp. 213–225.
URL: *http://dx.doi.org/10.1086/258464*

Valentino, N. A., Brader, T., Groenendyk, E. W., Gregorowicz, K. and Hutchings, V. L. (2011), 'Election night's alright for fighting: The role of emotions in political participation', *Journal of Politics* **73**(1), pp. 156–170.
URL: *https://doi.org/10.1017/s0022381610000939*