Fight or Flight on Fox?

Partisan Fear Responses on U.S. Cable News Shows

August 30, 2022

Abstract

Empirical work on political communication has so far left out a potentially pivotal dimension – the unspoken emotional responses indicated by facial expressions. This paper shows how to measure these responses using deep-learning-based computer-vision algorithms in the context of U.S. cable news video. Using machine-generated metrics for expressed emotion, combined with mentions of politically divisive entities, we estimate the difference in emotion when cable-channel speakers hear mentions of entities that are from the opposing side of the political divide. We find that the most responsive emotion is *fear*: When Fox News personalities hear about Democratic entities, or when MSNBC personalities hear about Republican entities, the dominant expressed emotion in their faces is *fear*, which increases by at least 30% relative to mentions of neutral political entities or those from one's own partisan team. We find some evidence that this partisan fear response is stronger on Fox News than on MSNBC.

Electronic copy available at: https://ssrn.com/abstract=4204973

"Men are driven by two principal impulses, either by love or by fear."

– Machiavelli

1 Introduction

In media and politics, emotions are frequently used as communication or persuasion tools to "diplomatically" indicate an opinion, to motivate an audience to action, or to trigger an opponent into reaction. In U.S. politics, at least, there is evidence of intensification of the use of emotions in recent decades, before the advent of social media, and even before the increase in partisanship starting in the early 1990s [Gennaro and Ash, 2021, Gentzkow et al., 2019]. Given this increase in emotional intensity alongside partisan polarization and an increasing diversification of venues for political conversation, there is a renewed need to analyze the role of emotions in political discourse.

A rich and growing literature has begun to recognize and analyze the importance of emotions in political communication [e.g. Marcus, 2000, Lau and Rovner, 2009]. This literature has mostly measured emotion by the content of texts, for example transcribed political speeches [Rheault et al., 2016a, Kosmidis et al., 2019, Hargrave and Blumenau, 2020, Hobolt et al., 2021, Gennaro and Ash, 2021]. A smaller number of papers has analyzed the emotional content of audio [Dietrich et al., 2019a,b, Damann et al., 2022]. The most recent frontier has started to explore video content and visual representations of emotion [Boussalis et al., 2021, 2022]. It is in this last dimension that this paper makes a contribution.

In the context of U.S. cable news, we analyze facial expressions in response to the mention of partisan actors, which given their divisiveness might trigger strong emotions. We use a newly collected corpus of videos from three ideologically different cable news channels (CNN, MSNBC, and Fox News). We use audio analysis to detect mentions of partisan entities, such as "Democrats" or "Obama' on the left and "Republicans" or "Trump" on the right. Matching up these mentions to the associated video frames, we use a deep-learning-based computer-vision algorithm to extract measures of emotional responses by the individuals depicted on-screen.

We then analyze empirically the variation in emotional responses when entities from the

Electronic copy available at: https://ssrn.com/abstract=4204973

partisan opposition are mentioned, taking Fox News as representing Republicans and MSNBC as representing Democrats. Relative to mentions of entities from the same political side (e.g. Fox News mentioning "Trump"), when entities from the opposing side are mentioned (e.g. Fox mentioning "Obama"), the major emotional response is an increase in expressed *fear*. We observe no corresponding change for anger, disgust, happiness, sadness, or surprise.

The observed effect on fear is robust to a wide range of statistical specifications, including adjustments for overall variation across the networks, adjusting for the composition of selected entity words, or using variants on the emotion metric. We provide additional human validation of the machine classification of the fear response. The effect for mentions of opposing entities is large in magnitude – an increase of at least 30% in fear response relative to mentions of non-opposing entities. The effect is concentrated in Fox News, rather than MSNBC.

These results add to the literatures on emotions in politics, media bias, and political methodology. In the emotions literature, a large body of work has considered their importance in politics [Marcus, 2000]. In particular, some recent work has considered nonverbal behavior of political candidates and emotional responses (Bucy and Stewart 2018, Olivola and Todorov 2010), emotional contrasts between political candidates (Gong and Bucy 2016, Koppensteiner and Grammer 2010), and the effect of elite emotions on potential voters (Himelboim et al. 2016). The closest papers, by Boussalis et al. [2021] and Boussalis et al. [2022], look at gender differences in the facial expression of emotion. Our paper is the first to show a special role of fear responses, relative to other negative emotions, in partisan conflict, adding to the growing literature on affective polarization [Iyengar et al., 2019].

The second related literature is on media bias. This work includes a number of papers that have shown differences in the language used by news outlets associated with the right wing or the left wing [Gentzkow and Shapiro, 2010, Martin and Yurukoglu, 2017, Widmer et al., 2020]. The closest media bias paper is Boxell [2018], who shows using newspaper images that in the 2016 presidential election, Trump and Clinton were depicted with different facial expressions depending on the partisan slant of the outlet. That result is quite different from ours, which looks at the dynamic emotional response to mentions of partisan entities in video footage. Our paper shows that partisan disagreement is communicated not just

Electronic copy available at: https://ssrn.com/abstract=4204973

through textual content or facial expressions, but the interaction between the two.

Methodologically, this paper is notable in extending analysis of political communication into a new data modality - video. While a number of papers have used automated methods to study emotions in various political contexts (Rheault et al. 2016b; Boussalis and Coan 2021; Gennaro and Ash 2021; Dietrich 2021; Dietrich et al. 2019a; Dietrich et al. 2019b), those have used text or audio. It is only in the most recent papers that facial expressions have been mined for emotional expression [Boussalis et al., 2021, 2022]. Both of those papers look at level differences across groups (gender) in facial emotional expression. We take another step and look at how partisan groups react differently to the same partisan topic.

2 Measuring Emotion Expressions in TV News Video

This section describes the sequence of steps for taking our videos of cable news shows and building the dataset of emotion expressions associated with mentions of partisan entities. Appendix Table A.1 provides some summary statistics on this process.

Cable News Video Archive. The first step is to build a collection of videos from cable news shows. We obtained video footage and associated metadata (links, titles, and video descriptions) from all videos uploaded by the three cable news networks (FNC, MSNBC, CNN) to YouTube. Our dataset is built from 1,750 videos, containing 191 hours (11,454 minutes) of footage.[1] The raw dataset comes as video files, which consist of a continuous stream of images and accompanying segments of audio. A sample of image frames is shown in Figure 1.

Audio Transcription and Partisan Mention Detection. We first work with the audio data to detect mentions of partisan entities. The audio stream associated with each video file is extracted as mono-channel audio data. The audio files are converted to text with

[1]See Appendix Table A.2. The total length of all videos uploaded by each news channel is as follows: CNN (46 hours, 0 minutes, 21 seconds), MSNBC (50 hours, 15 minutes, 33 seconds), Fox News (94 hours, 41 minutes, 7 seconds). The earliest and the latest videos uploaded by each news channel are dated as follows: CNN (earliest: 05/02/2011, latest: 03/07/2020), MSNBC (earliest: 05/08/2017, latest: 03/07/2022), Fox News (earliest: 11/10/2016, latest: 03/06/2020).

Electronic copy available at: https://ssrn.com/abstract=4204973

Figure 1: Example Image Frames from Cable News Video Collection



Notes. Sample of image frames from Cable News Videos downloaded from YouTube for Fox News, CNN, and MSNBC. These examples are annotation pairs given to human annotators, as described in section "Validation of Fear Classifications" below. Left image is annotated as 'fear' by our machine classifier; right image is a randomly assigned comparison image.

Electronic copy available at: https://ssrn.com/abstract=4204973

automated transcription software,[2] with each word timestamped with a precision of one-tenth of a second.

We then searched the transcripts for sets of relevant partisan words. We started with the names of all presidents, presidential candidates, and the words 'president', 'democrat', 'democratic', 'liberal', 'conservative', and 'republican'. We then added the (sur)names of the top-100 most famous political actors, as indicated by a recent national survey.[3] These names and words, with associated partisanship tag (Democrat, Republican, or Neutral) are listed in Appendix Table A.3.[4]

Video Analysis and Emotion Detection. The image stream is converted into static frames by extracting 4 frames per second. We used an automated face detection algorithm to filter out any images not containing a face. As we care about facial responses to word mentions, we use a second algorithm to detect and drop faces that are speaking. Now we have a collection of images of non-speaking faces from the three cable networks.

Meanwhile, we build a classifier to extract emotions from faces depicted in video frames. We follow the dominant approach in social psychology, which focuses on six basic emotions: happiness, anger, sadness, fear, disgust, and surprise.[5] We have training data on face images with emotion labels from FER2013 [Goodfellow et al., 2013], a large and diverse dataset that is often used to train emotion classification models, including for social science applications [Tumen et al., 2017, Khaireddin and Chen, 2021, Minaee et al., 2021].[6] This dataset consists

[2]We use `librosa` (McFee et al. 2015), `ffmpeg` (Tomar 2006), and `pocketsphinx` (Zhang 2017), libraries written for `Python`.

[3]YouGov survey, administered Q1 of 2022.

[4]Note that "president" is the only "neutral" word in our list. We have also tagged it as a potentially "questionable" word that could lead to false positives. Appendix Table A.9 shows that our results are robust to dropping mentions of "president" and other such questionable words.

[5]See, e.g., Ekman [1992], Izard [1994], Ekman and Cordaro [2011], Levenson [2011], Panksepp and Watt [2011]. In the Ekman framework, these expressions are believed to be culturally universal and rooted in biological adaptive functions [Calvo and Nummenmaa, 2016]. In addition, there are the dimensional [Russell and Paris, 1994] and constructionist approaches [Barrett, 2006, Lindquist and Gendron, 2013], which believe that emotions are not produced in a discrete manner, but rather in a continuum on two orthogonal dimensions: valence (pleasantness versus unpleasantness) and arousal (calmness versus tension) [Calvo and Nummenmaa, 2016]. Analyzing these other dimensions is an important area for future work, but is not feasible at the moment since there are no image datasets labeled using that approach.

[6]The FER2013 dataset was originally compiled by Pierre Luc Carrier and Aaron Courville by using Google image search API by using emotion-related keywords. To diversify the range of faces in the dataset, these keywords were eventually combined with words related to gender, age or ethnicity to obtain facial image search queries. The first 4 of 1000 images returned as a result of the search query were kept for the next

Electronic copy available at: https://ssrn.com/abstract=4204973

of 35,887 images of faces, each labeled with one of the seven classes: anger (4953), disgust (547), fear (5121), happiness (8989), sadness (6077), surprise (4002), and neutral (6198). Figure 2 Panel A shows samples of images, with the associated classes, illustrating that the FER2013 dataset covers faces belonging to diverse genders, ethnicities, ages, compositions, and facial features. To help the model further generalize to a new dataset of video frames, we take the standard approach of augmenting the data with visually perturbed images (random flips, colorizing, rotations, etc.) containing the original label (see Appendix Figure A.2).[7]

Our classifier adapts recently developed approaches in computer vision that obtain high accuracy in identifying expressed emotions in images of human faces, among many other computer-vision tasks [e.g. Krizhevsky et al., 2012, Li et al., 2016, Yang et al., 2017, Rao et al., 2019, Deng et al., 2021]. We use a stacked convolutional neural network architecture (see Figure A.2), which uses a series of filters to scan over the image and learn relevant features for a prediction task. We train the model on the augmented FER2013 dataset with 80% used for training and 20% held-out for testing. The test-set accuracy of the trained model is 66.37%. As shown in the confusion matrix in Figure 2 Panel B, the model obtains good precision and recall across all seven categories. When fear is mis-classified, it is mostly mis-classified as sadness.

This trained classifier is then applied to the cable news images. For each image frame, the face features are extracted and processed to be in the same format as the classifier. We then obtain a probability distribution across emotions for each face, which we label with the top emotion class. The resulting dataset is a set of emotion labels for each face in each frame.

Merged Dataset. We then link the data on word mentions to the data on expressed emotions.[8] The relevant "events" in our data are mentions from the list of political words. We extract image frames from the pre- and post-event window, each 0.75 seconds long (corre-

stage of processing. In the next phase, OpenCV has been used to extract faces from collected images. Human labelers were then used to validate the labels of emotions, and incorrectly labeled images were discarded. All images were finally resized to be 48x48 pixels (for easier processing by neural network models) and converted to grayscale.

[7]Further, we tried enriching the FER2013 by using the Cohn-Kanade dataset (Lucey et al. 2010), but that did not improve accuracy results and the idea was discarded.

[8]For long words that last over multiple frames, the associated duplicate consecutive observations are removed.

Electronic copy available at: https://ssrn.com/abstract=4204973

Figure 2: Emotion Expression Classification

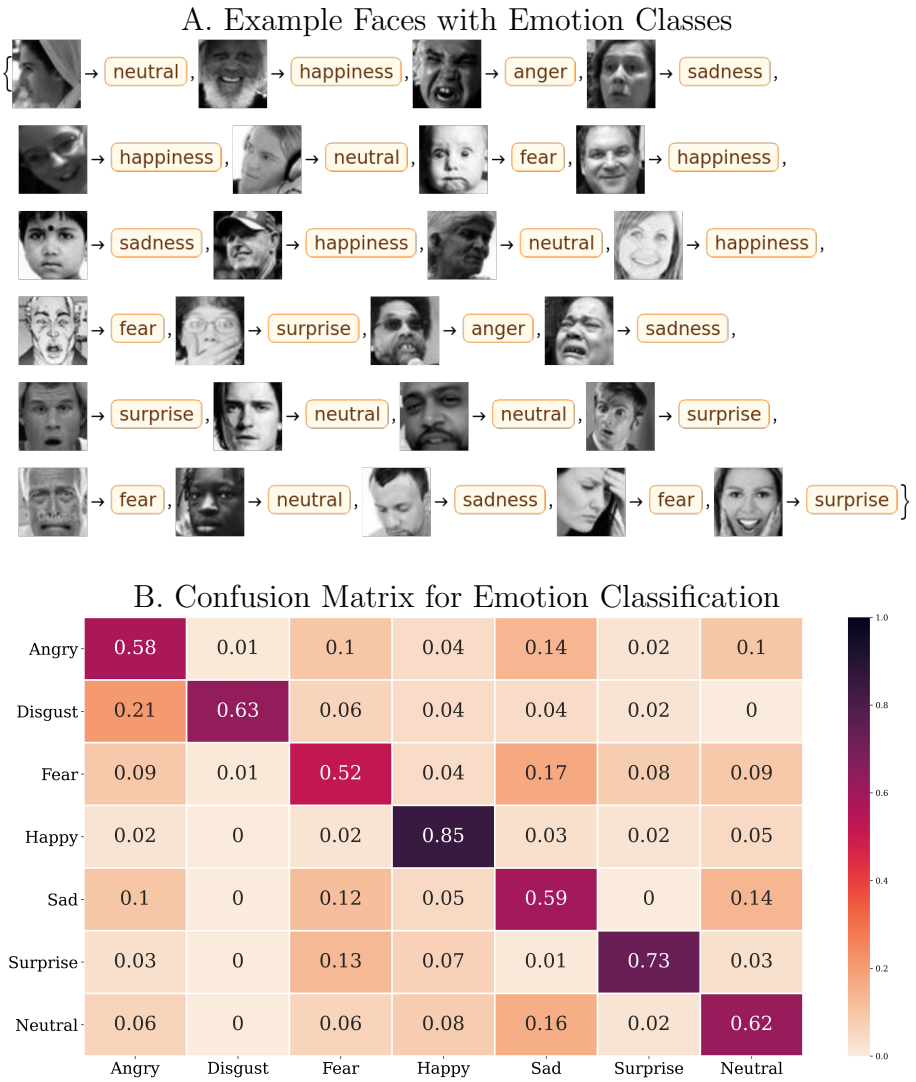

Notes. Panel A: Random sample of faces, with classes, from the FER 2013 Dataset [Goodfellow et al., 2013]. Panel B: Confusion matrix showing prediction frequencies across predicted classes (columns) by true class (rows) for our emotion classifier applied to FER 2013.

Electronic copy available at: https://ssrn.com/abstract=4204973

sponding to 3 frames each, 6 frames in total). Treatment windows that overlap with each other, or that are too closely located, are adjusted or discarded.

The appendix reports detailed summary statistics on this dataset. First, Appendix Table A.4 shows counts of events by network and word partisanship. Appendix Figure A.3 shows the shares by emotion class in their prevalence across these mentions.

Validation of Fear Classifications. As discussed further in Section 3 below, our empirical analysis of partisanship focuses on the emotion response of fear. Thus we wanted to gain additional confidence that our emotion classifier properly identified fear expressions in our new domain of cable news videos. To do so, we engaged a pair of human annotators to label images and assess our classifier's performance in the new domain.

Specifically, we sampled 1000 pairs of video frames from our collection that contained one face each. Of the pair, one face has a relatively high classifier probability of expressing fear, while the other face has a relatively low classifier probability of expressing fear. Examples of these image pairs are shown in Figure 1; in each pair by row, the left image's face was identified as expressing fear. For the human annotation, the ordering of the pair was randomized and the annotators were blinded to the machine labels.

For each of these image pairs, we asked the annotator to rank each pair along a four point scale: left image is definitely more fearful, left image is probably more fearful, right image is probably more fearful, right image is definitely more fearful. From a second annotator, we have a binary ranking (either left or right is more fearful). To make a fairer comparison and reduce noise, we limit our validation dataset to the set of 668 image pairs where the second annotator's binary ranking corresponds to the first annotator's. We normalized the annotator score and the classifier score to the same scale for comparison.

Electronic copy available at: https://ssrn.com/abstract=4204973

Figure 3: Validation of Classifier Fear Score Against Annotator Fear Score

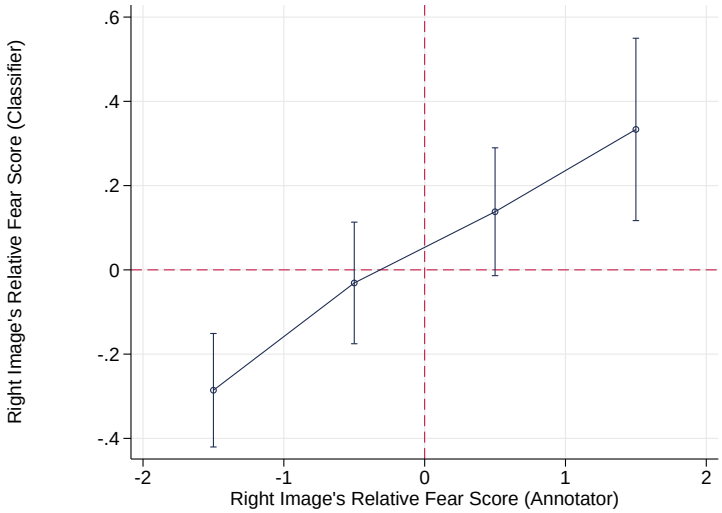

Notes. Average classifier fear score for the right image in a pair, plotted by the annotator's four-point fear score for the right image, relative to the left image. Data is limited to image pairs where the second annotator's binary ranking agrees with the first annotator's. Confidence intervals give standard errors of the mean.

Figure 3 plots the classifier's average relative fear score (vertical axis), binned by the annotator's relative fear score (horizontal axis). We can see that, reassuringly, the scores are highly correlated, with the classifier scores increasing with the annotator scores. In a univariate regression with the classifier score as the outcome and the annotator score as the regressor, the estimated coefficient is 0.206 ($p < .01$).[9] Hence, we have confidence that our classifier captures a dimension of fear expression in the cable-news videos that corresponds to human judgment.

3 Empirical Analysis

This section outlines our empirical approach and reports the results.

[9]The coefficient is similar, at 0.18 ($p < .01$), if we do not filter based on the second annotator's data and use all the annotated image pairs.

Electronic copy available at: https://ssrn.com/abstract=4204973

3.1 Empirical Approach

The goal of the empirical analysis is to measure the emotional framing of words with different partisan valence – Democrat, Republican, or Neutral. For each word w in these lists (Appendix Table A.3), we define an indicator variable P_{wc} equaling one when the partisan slant of a given network c goes against that entity. For example, $P_{wc} = 1$ for Democrats on Fox News, but $P_{wc} = 0$ for Republicans on Fox News. Conversely, on MSNBC the sign is flipped for Democrats and Republicans. For the more centrist network CNN, $P_{wc} = 0$ for all w. Summary tabulations on mentions by word are shown in Appendix Table A.3, while tabulations on mentions by network and category are reported in Appendix Table A.4.

Next, we have an emotional expression variable E_i^k, corresponding to the intensity that emotion k is expressed by depicted faces in a window of frames before and after mention i. There are seven emotions $k \in \{$angry, disgusted, fearful, happy, neutral, sad, surprised$\}$. We let $\boldsymbol{E}_i$ represent the vector of these variables. The baseline specification for each mention is a categorical variable, with E_i^k, equaling one if emotion k is the top or dominant emotion expressed in the window around i, meaning that the most faces were tagged with this emotion across frames. For robustness, we compute three alternative measures: (1) the share of faces on emotion k, (2) the total number of faces on emotion k, transformed by the inverse hyperbolic sine, and (3) an alternative categorical variable equaling one if emotion k is expressed at all in the window around i. Summary means on these outcomes by network and word partisanship, for each emotion, are reported in Appendix Figures A.4 (top emotion), A.5 (share of faces), A.6 (IHS face counts), and A.7 (any expression).

We estimate two regression specifications. First, to get at the joint relationship of all emotions on partisan opposition, we estimate

$$P_{iwc} = \boldsymbol{E}_i'\boldsymbol{\gamma} + \alpha_c + \alpha_w + \epsilon_{iwc} \tag{1}$$

where α_c include cable-channel fixed effects, and α_w includes word fixed effects, either by partisan category or by the specific word. This regression produces a vector of coefficients $\boldsymbol{\gamma}$ with seven items, corresponding to the seven emotion classes. In the second specification,

Electronic copy available at: https://ssrn.com/abstract=4204973

each emotion variable k is put on the left-hand side:

$$E^k_{iwc} = \gamma P_{wc} + \alpha_c + \alpha_w + \epsilon_{iwc} \tag{2}$$

with a separate regression for each of the seven emotion classes. Otherwise the terms are as in Equation (1). We estimate these models with ordinary least squares and robust standard errors, although we show robustness to clustering standard errors.

3.2 Results

The main results for how partisan opposition reflects emotional resonance are shown in Figure 4. Panel A shows the results from the joint-model specification with the emotions included together on the right-hand side of a single regression (Equation 1), while Panel B shows the results with the separate-model specification with each emotion included individually on the left-hand side in separate regressions (Equation 2). Each figure plots the estimated coefficients and confidence intervals summarizing how each of the seven emotion classes relate to partisan opposition – that is, FNC mentions of Democrat entities, and MSNBC mentions of Republican entities.

For both specifications, the result is the same. On cable news shows, mentions of the opposing political party is associated with a higher expression of fear. The effect on fear is relatively large in magnitude and statistically significant, while the effects for the other emotion classes are relatively small in magnitude and not statistically different from zero.

Now we dive deeper into the fear response to opposing entity mentions. Table 1 reports the statistical estimates from the main regression results across various specifications. We see that, consistent with the results from the coefficient plots, there is a positive and statistically significant effect. The coefficient does not change much without fixed effects (Column 1), adjusting for channel-level differences in fear response by adding cable channel fixed effects (Column 2), adjusting for overall variation across word categories (Columns 3 and 4), or even in the most conservative specification of adding word-specific fixed effects, meaning that the fear response is measured off emotion differences between the networks when the same word is mentioned. Overall, we can say that fear is the dominant emotion about 7-10

Electronic copy available at: https://ssrn.com/abstract=4204973

Figure 4: Emotional Expressions and Partisan Topics

A. Joint Specification, Regressing Opposition on All Emotions

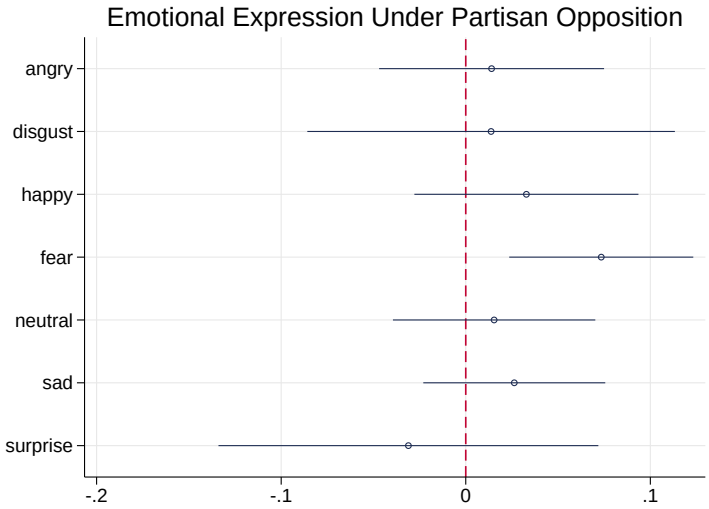

B. Separated Specification, Regressing Each Emotion Separately on Opposition

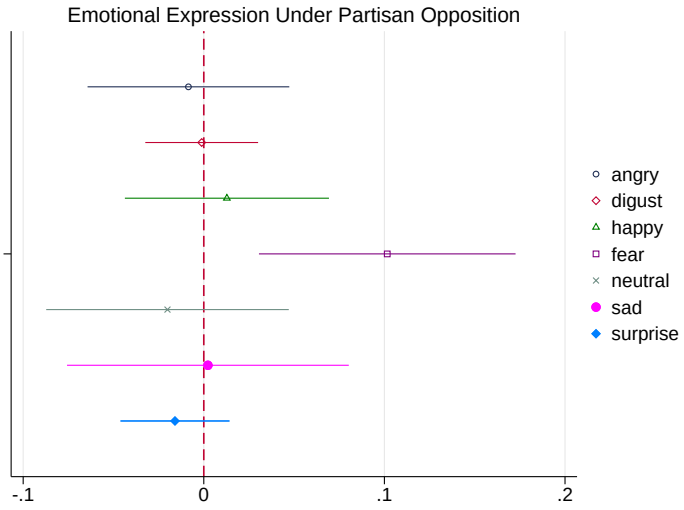

Notes. Coefficient plots for main regression results. Panel A is from a single "joint" specification where the outcome is "Partisan Opposition", equalling one for mentions of Democrat entities on FNC, or mentions of Republican entities on MSNBC. That is regressed on indicators for the top emotion being one of the seven indicated. Panel B is the separated specification; each coefficient comes from a separate regression, where the indicated top-emotion variable is regressed on Partisan Opposition. Fixed effects for the cable channel and for the detected word are absorbed. Error bars give 95% confidence intervals constructed using robust standard errors. For Panel A's associated regression table results, see Appendix Table A.5 Column 5; for Panel B, see Table 1 Column 5 and Appendix Table A.6 (all columns).

Electronic copy available at: https://ssrn.com/abstract=4204973

Table 1: Fear Expression and Partisan Topics: Additional Specifications

	(1)	(2)	(3)	(4)	(5)
	Effect on Prob(Fear is Top Emotion)				
Partisan Opposition	0.0809**	0.0868**	0.0669*	0.0781*	0.102**
	(0.0247)	(0.0248)	(0.0322)	(0.0325)	(0.0362)
Cable Channel FE		X		X	X
Word Category FE			X	X	
Word FE					X
N	1382	1382	1382	1382	1336
R-sq	0.008	0.012	0.009	0.014	0.052

Notes. Regression results for baseline specification, where the outcome is an indicator for whether this word mention's top emotion is fear. The treatment variable "Partisan Opposition" equals one for mentions of Democrat entities on FNC, or mentions of Republican entities on MSNBC. Cable Channel FE means fixed effect for the cable channel; Word Category FE means fixed effect for the word category (Republican, Democrat, Neutral); Word FE means fixed effect for the specific word (from Appendix Table A.3). Robust standard errors in parentheses. + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$.

percentage-points more of the time for opposing-party mentions, relative to a control mean of 23.3 percentage points. That is a proportionally large effect, an approximately 30% to 43% increase.

The appendix provides a wide range of supporting results and robustness checks. First, for comparison, Appendix Table A.5 provides the regression estimates for the joint specification, where P_{iwc} is regressed on $\boldsymbol{E}_i$ (Eq. 1). Across columns, there is a large effect for fear and no effect for the other emotions. Appendix Table A.6 reports the separate specification for the other emotions, with a similar story.

The results are not sensitive to how the outcome variable is specified. Appendix Table A.9 reports the regression results for fear using three other specifications: the share of emotions detected as fear, the inverse hyperbolic sine of the number of faces designed as fear, and a binary indicator for any face being designated as fear. Across these different outcome specifications, the message is the same that partisan opposition increases expressed fear. Comparing the response for the other emotions across these three outcome transformations, as done in Appendix Figure A.7, reinforces the point again.

There are three additional robustness checks. First, the results so far have just included

Electronic copy available at: https://ssrn.com/abstract=4204973

Table 2: Fear and Partisan Mentions: Fox News vs MSNBC

	(1)	(2)	(3)	(4)	(5)	(6)
	Effect on Prob(Fear is Top Emotion)					
Partisan Opposition	0.105**	0.0432	0.0968+	0.0552	0.105+	0.0485
	(0.0287)	(0.0481)	(0.0538)	(0.0726)	(0.0567)	(0.0814)
FNC Included	X		X		X	
MSNBC Included		X		X		X
CNN Included			X	X	X	X
Cable Channel FE			X	X	X	X
Word Category FE			X	X		
Word FE					X	X
N	935	447	1341	853	1297	812
R-sq	0.014	0.002	0.011	0.002	0.054	0.054

Notes. Regression results for comparison of FNC and MSNBC. Outcome is indicator for whether this mention's top emotion is fear. Treatment variable "Partisan Opposition" equals one for mentions of Democrat entities on FNC, or mentions of Republican entities on MSNBC. The rows "[Network] Included" indicate whether samples for the respective network are included. Cable Channel FE means fixed effect for the channel; Word Category FE means fixed effect for the word category (Republican, Democrat, Neutral); Word FE means fixed effect for the specific word. Robust standard errors in parentheses. + p < .0.1, * p < 0.05, ** p < 0.01.

FNC and MSNBC, because partisan slant (liberal or conservative) is needed to identify the effect. In Appendix Table A.8, we also add the relatively nonpartisan network CNN, which can affect the estimates by entering through the word fixed effects. The results are the same. Second, Appendix Table A.9 drops a few words in our entity list which could generate false positives (e.g. "president" or "bloomberg"), and that also does not affect the results. Third and finally, the statistical significance is not sensitive to how standard errors are clustered (Appendix Table A.10).

As a final analysis, we would like to know what cable network is driving our results. To this end, Table 2 compares the magnitude and significance of the results for FNC and MSNBC. We can see positive coefficients for both networks, and statistical significance goes down with the smaller sample. Overall, the estimates for FNC are larger in magnitude and more precisely estimated, and therefore more statistically significant. Hence, we have some evidence that Fox News speakers are more likely to express fear in response to mentions of the partisan opposition.

Electronic copy available at: https://ssrn.com/abstract=4204973

4 Conclusion

Media bias is not just a matter of text, in the topics covered or phrasing. This paper shows that bias is also communicated through how topics are framed emotionally, in the responses expressed in facial features. While this might seem obvious for those who watch political discussions live or on television, that has not been established quantitatively due to the technical difficulty in measuring emotional responses from video recordings of facial expressions.

Having filled that gap, this paper has also provided a generalizable method to measure emotional responses to particular terms in video recordings. We apply that method to a collection of cable news videos and show that emotional expressions in faces reflect partisan opposition. When a liberal (conservative) entity is mentioned on a conservative (liberal) cable network, the individuals onscreen tend to respond with more fearful facial expressions.

We started off with the hypothesis that mentions of opposing partisan groups would be reflected with more negative emotional responses. A priori, that could have been anger, disgust, or sadness. Or else it could have been a reduction in happiness toward a neutral emotion. Instead, we found that fear is the most responsive emotion in these televised political debates. That result is relevant to a broad literature in the political psychology of emotion [e.g. Elster, 1999, Brader, 2005, Gross, 2008], and more specifically to the literature on affective polarization [Iyengar et al., 2019]. In particular, a broad finding in emotion psychology is that fear can be reduced through exposure therapy [e.g. McNally, 2007]. If those insights apply in politics, then repeated exposure to partisan opponents in a safe environment may serve to reduce partisan animus.

These methods and results open up a new area for study, not just in media but in political communication more generally [Damann et al., 2022]. The method could be applied to campaign ads, electoral debates, parliamentary speeches, or courtroom videos, to name a few interesting domains. Going forward, studies of political slant or ideology cannot rely on text or audio, or even the video of the speaker, and confidently obtain a full picture. Based on our evidence, there is some additional partisan communication occurring through the unspoken facial responses of the interlocutors.

Electronic copy available at: https://ssrn.com/abstract=4204973

References

Gloria Gennaro and Elliott Ash. Emotion and reason in political language. *Economic Journal*, 2021.

Matthew Gentzkow, Jesse M Shapiro, and Matt Taddy. Measuring group differences in high-dimensional choices: Method and application to congressional speech. *Econometrica*, 87(4): 1307–1340, 2019.

George E Marcus. Emotions in politics. *Annual review of political science*, 3(1):221–250, 2000.

Richard R Lau and Ivy Brown Rovner. Negative campaigning. *Annual review of political science*, 12:285–306, 2009.

Ludovic Rheault, Kaspar Beelen, Christopher Cochrane, and Graeme Hirst. Measuring emotion in parliamentary debates with automated textual analysis. *PloS one*, 11(12):e0168843, 2016a.

Spyros Kosmidis, Sara B Hobolt, Eamonn Molloy, and Stephen Whitefield. Party competition and emotive rhetoric. *Comparative Political Studies*, 52(6):811–837, 2019.

Lotte Hargrave and Jack Blumenau. The declining importance of gender stereotypes for politicians' style in the uk. 2020.

Sara Hobolt, Moritz Osnabruegge, and Toni Rodon. Playing to the gallery: How politicians use emotive rhetoric in parliament. 2021.

Bryce J. Dietrich, Matthew Hayes, and Diana Z. O'brien. Pitch perfect: Vocal pitch and the emotional intensity of congressional speech. *American Political Science Review*, 113(4):941–962, 2019a. Publisher: Cambridge University Press.

Bryce J. Dietrich, Ryan D. Enos, and Maya Sen. Emotional Arousal Predicts Voting on the U.S. Supreme Court. *Political Analysis*, 27(2):237–243, April 2019b. ISSN 1047-1987, 1476-4989. doi: 10.1017/pan.2018.47. URL `https://www.cambridge.org/core/journals/political-analysis/article/emotional-arousal-predicts-voting-on-the-us-supreme-court/1047BF7D73A1B45BDB1C61A3A80E0F64`. Publisher: Cambridge University Press.

Taylor Damann, Dean Knox, and Christopher Lucas. More than words: How political rhetoric shapes voters' affect and evaluation. 2022.

Constantine Boussalis, Travis G Coan, Mirya R Holman, and Stefan Müller. Gender, candidate emotional expression, and voter reactions during televised debates. *American Political Science Review*, 115(4):1242–1257, 2021.

Electronic copy available at: https://ssrn.com/abstract=4204973

Constantine Boussalis, Travis Coan, and Mirya Holman. Mad, sad, but mostly glad: Hw men and women in politics communicate using emotions in images. 2022.

Erik P. Bucy and Patrick Stewart. The personalization of campaigns: Nonverbal cues in presidential debates. In *Oxford Research Encyclopedia of Politics*. 2018.

Christopher Y. Olivola and Alexander Todorov. Elected in 100 milliseconds: Appearance-based trait inferences and voting. *Journal of nonverbal behavior*, 34(2):83–110, 2010. Publisher: Springer.

Zijian Harrison Gong and Erik P. Bucy. When style obscures substance: Visual attention to display appropriateness in the 2012 presidential debates. *Communication Monographs*, 83(3):349–372, 2016. Publisher: Taylor & Francis.

Markus Koppensteiner and Karl Grammer. Motion patterns in political speech and their influence on personality ratings. *Journal of Research in Personality*, 44(3):374–379, 2010. Publisher: Elsevier.

Itai Himelboim, Kaye D. Sweetser, Spencer F. Tinkham, Kristen Cameron, Matthew Danelo, and Kate West. Valence-based homophily on Twitter: Network analysis of emotions and political talk in the 2012 presidential election. *New media & society*, 18(7):1382–1400, 2016. Publisher: Sage Publications Sage UK: London, England.

Shanto Iyengar, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra, and Sean J Westwood. The origins and consequences of affective polarization in the united states. *Annual Review of Political Science*, 22:129–146, 2019.

Matthew Gentzkow and Jesse M Shapiro. What drives media slant? evidence from us daily newspapers. *Econometrica*, 78(1):35–71, 2010.

Gregory J Martin and Ali Yurukoglu. Bias in cable news: Persuasion and polarization. *American Economic Review*, 107(9):2565–99, 2017.

Philine Widmer, Sergio Galletta, and Elliott Ash. Media slant is contagious. *Center for Law & Economics Working Paper Series*, 14, 2020.

Levi Boxell. Slanted images: Measuring nonverbal media bias. 2018.

Ludovic Rheault, Kaspar Beelen, Christopher Cochrane, and Graeme Hirst. Measuring emotion in parliamentary debates with automated textual analysis. *PloS one*, 11(12):e0168843, 2016b. Publisher: Public Library of Science San Francisco, CA USA.

Constantine Boussalis and Travis G. Coan. Facing the electorate: Computational approaches to the study of nonverbal communication and voter impression formation. *Political Communication*, 38 (1-2):75–97, 2021. Publisher: Taylor & Francis.

Electronic copy available at: https://ssrn.com/abstract=4204973

Bryce J. Dietrich. Using motion detection to measure social polarization in the US House of Representatives. *Political Analysis*, 29(2):250–259, 2021. Publisher: Cambridge University Press.

Brian McFee, Colin Raffel, Dawen Liang, Daniel P. Ellis, Matt McVicar, Eric Battenberg, and Oriol Nieto. librosa: Audio and music signal analysis in python. In *Proceedings of the 14th python in science conference*, volume 8, pages 18–25. Citeseer, 2015.

Suramya Tomar. Converting video formats with FFmpeg. *Linux Journal*, 2006(146):10, 2006. Publisher: Belltown Media Houston, TX.

Anthony Zhang. Speech recognition (version 3.8)[software]. In *Proc. ICCC*, 2017.

Paul Ekman. Are there basic emotions? 1992. Publisher: American Psychological Association.

Carroll E. Izard. Innate and universal facial expressions: evidence from developmental and cross-cultural research. 1994. Publisher: American Psychological Association.

Paul Ekman and Daniel Cordaro. What is meant by calling emotions basic. *Emotion review*, 3(4): 364–370, 2011. Publisher: Sage Publications Sage UK: London, England.

Robert W. Levenson. Basic emotion questions. *Emotion review*, 3(4):379–386, 2011. Publisher: Sage Publications Sage UK: London, England.

Jaak Panksepp and Douglas Watt. What is basic about basic emotions? Lasting lessons from affective neuroscience. *Emotion review*, 3(4):387–396, 2011. Publisher: Sage Publications Sage UK: London, England.

Manuel G. Calvo and Lauri Nummenmaa. Perceptual and affective mechanisms in facial expression recognition: An integrative review. *Cognition and Emotion*, 30(6):1081–1106, 2016. Publisher: Taylor & Francis.

James A. Russell and Faye A. Paris. Do children acquire concepts for complex emotions abruptly? *International Journal of Behavioral Development*, 17(2):349–365, 1994. Publisher: Sage Publications Sage CA: Thousand Oaks, CA.

Lisa Feldman Barrett. Solving the emotion paradox: Categorization and the experience of emotion. *Personality and social psychology review*, 10(1):20–46, 2006. Publisher: Sage Publications Sage CA: Los Angeles, CA.

Kristen A. Lindquist and Maria Gendron. What's in a word? Language constructs emotion perception. *Emotion Review*, 5(1):66–71, 2013. Publisher: Sage Publications Sage UK: London, England.

Ian J. Goodfellow, Dumitru Erhan, Pierre Luc Carrier, Aaron Courville, Mehdi Mirza, Ben Hamner, Will Cukierski, Yichuan Tang, David Thaler, and Dong-Hyun Lee. Challenges in representation

Electronic copy available at: https://ssrn.com/abstract=4204973

learning: A report on three machine learning contests. In *International conference on neural information processing*, pages 117–124. Springer, 2013.

Vedat Tumen, Omer Faruk Soylemez, and Burhan Ergen. Facial emotion recognition on a dataset using convolutional neural network. In *2017 International Artificial Intelligence and Data Processing Symposium (IDAP)*, pages 1–5. IEEE, 2017.

Yousif Khaireddin and Zhuofa Chen. Facial emotion recognition: State of the art performance on FER2013. *arXiv preprint arXiv:2105.03588*, 2021.

Shervin Minaee, Mehdi Minaei, and Amirali Abdolrashidi. Deep-emotion: Facial expression recognition using attentional convolutional network. *Sensors*, 21(9):3046, 2021. Publisher: MDPI.

Patrick Lucey, Jeffrey F. Cohn, Takeo Kanade, Jason Saragih, Zara Ambadar, and Iain Matthews. The extended cohn-kanade dataset (ck+): A complete dataset for action unit and emotion-specified expression. In *2010 ieee computer society conference on computer vision and pattern recognition-workshops*, pages 94–101. IEEE, 2010.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.

Jinpeng Li, Zhaoxiang Zhang, and Huiguang He. Implementation of EEG emotion recognition system based on hierarchical convolutional neural networks. In *International Conference on Brain Inspired Cognitive Systems*, pages 22–33. Springer, 2016.

Jufeng Yang, Dongyu She, and Ming Sun. Joint Image Emotion Classification and Distribution Learning via Deep Convolutional Neural Network. In *IJCAI*, pages 3266–3272, 2017.

Tianrong Rao, Xiaoxu Li, Haimin Zhang, and Min Xu. Multi-level region-based convolutional neural network for image emotion classification. *Neurocomputing*, 333:429–439, 2019. Publisher: Elsevier.

Zelin Deng, Qiran Zhu, Pei He, Dengyong Zhang, and Yuansheng Luo. A Saliency Detection and Gram Matrix Transform-Based Convolutional Neural Network for Image Emotion Classification. *Security and Communication Networks*, 2021, 2021. Publisher: Hindawi.

Jon Elster. *Alchemies of the Mind: Rationality and the Emotions*. Cambridge University Press, 1999.

Ted Brader. Striking a responsive chord: How political ads motivate and persuade voters by appealing to emotions. *American Journal of Political Science*, 49(2):388–405, 2005.

Kimberly Gross. Framing persuasive appeals: Episodic and thematic framing, emotional response, and policy opinion. *Political Psychology*, 29(2):169–192, 2008.

Electronic copy available at: https://ssrn.com/abstract=4204973

Richard J McNally. Mechanisms of exposure therapy: how neuroscience can improve psychological treatments for anxiety disorders. *Clinical psychology review*, 27(6):750–759, 2007.

Marcus D. Bloice, Christof Stocker, and Andreas Holzinger. Augmentor: an image augmentation library for machine learning. *arXiv preprint arXiv:1708.04680*, 2017.

Connor Shorten and Taghi M. Khoshgoftaar. A survey on image data augmentation for deep learning. *Journal of big data*, 6(1):1–48, 2019. Publisher: SpringerOpen.

Jason Wang and Luis Perez. The effectiveness of data augmentation in image classification using deep learning. *Convolutional Neural Networks Vis. Recognit*, 11:1–8, 2017.

Agnieszka Mikolajczyk and Michal Grochowski. Data augmentation for improving deep learning in image classification problem. In *2018 international interdisciplinary PhD workshop (IIPhDW)*, pages 117–122. IEEE, 2018.

Electronic copy available at: https://ssrn.com/abstract=4204973

A Appendix

Table A.1: Data Processing and Number of Observations

Steps	Description	CNN	MSNBC	Fox News	Total
1	#videos	558	597	595	**1,750**
1	#all frames extracted	723,732 (=100%)	662,484 (=100%)	1,363,468 (=100%)	**2,749,684 (=100%)**
2	#audio segments	787,339	1,065,759	1,144,690	**2,997,788**
3	#frames with emotion labels	117,447 (=16.23%)	129,702 (=19.58%)	287,672 (=21.10%)	**534,821 (=19.45%)**
4	#frames w/faces and w/no speech	53,730 (=7.42%)	59,202 (=8.93%)	149,448 (=10.96%)	**262,380 (=9.54%)**
5 / 6	#labeled frames in treatment windows	1320 (=0.18%)	1455 (=0.21%)	2642 (=0.19%)	**5,417 (=0.20%)**
	#all treatment windows	434	469	1015	**1,918**
	#'fear' images in treatment windows	331 (=0.05%)	374 (=0.06%)	753 (=0.06%)	**1,458 (=0.05%)**
	#treatment windows with 'fear' images	177	180	385	**742**
	#observations ('fear' images)	**331**	**374**	**753**	**1,458 (=0.05%)**

Notes. Supporting information on the data processing steps, with the number of videos and frames at various steps of the pipeline. #audio segments gives the count number of instances in the videos with continuous (uninterrupted) speech. A "treatment window" is a window of frames before/after the mention of a political entity. The percentages indicate the amount of data obtained (or lost) in each step of the extraction of fear emotions used in the analysis (specifically, 99.95% of all frames had to be discarded; 0.05% of all frames is usable).

Electronic copy available at: https://ssrn.com/abstract=4204973

Table A.2: Video Duration Information by Network

	MSNBC	CNN	FNC	**All**
Minimum Length	30S	29S	17S	**17S**
Maximum Length	5H:19M:18S	2H:20M:12S	9H:4M:3S	**9H:4M:3S**
Mean Length	5M:4S	4M:57S	9M:33S	**6M:33S**
Median Length	4M:20S	3M:29S	4M:53S	**4M:21S**
Total Length	**50H:15M:33S**	**46H:0M:21S**	**94H:41M:7S**	**190H:57M:1S**

Notes. H, M, and S stand for "hours", "minutes", and "seconds" respectively. The total analyzed video length across all networks is 190H:57M:1S.

Electronic copy available at: https://ssrn.com/abstract=4204973

Table A.3: List of Partisan Words with Frequencies

Left-Wing Words		Neutral Words		Right-Wing Words	
democrat	299	president*	731	republican	145
obama	71			trump	67
biden	41			west*	22
young*	37			conservative	20
clinton	34			paul	20
sanders	30			graham	17
bloomberg*	29			mcconnell	17
schiff	9			bush	14
kennedy	8			collins	9
liberal	7			king	8
nixon	7			brown	7
pelosi	6			warner	5
warren	6			alexander	4
polis	5			grassley	4
schumer	5			ridge	4
waters	5			ryan	4
blumenthal	4			sessions	4
frank	4			cruz	3
kerry	4			haley	3
omar	4			palin	3
cummings	3			angle	2
booker	2			blackburn	2
gore	2			cotton	2
reich	2			daniels	2
shaheen	2			delay	2
				murkowski	2
				romney	2
				walker	2

Notes. List of words coded as Democrat, Neutral, or Republican, according to column. Accompanying numbers give the number of mentions in the dataset. Includes all words mentioned at least twice. An asterisk (*) indicates that this is a questionable word that may produce false positives, which is dropped in the robustness check reported in Appendix Table A.9.

Electronic copy available at: https://ssrn.com/abstract=4204973

Figure A.1: Data Augmentation



Notes. We augment the FER2013 dataset by applying geometric transformations (random flips, colorings, rotations, sharpening and blurring) and color space transformations (changing RGB color channels and intensification of colors) to our images (similar techniques have been used widely in the field by using associated augmentation libraries (Bloice et al. 2017; Shorten and Khoshgoftaar 2019; Krizhevsky et al. 2012; Wang and Perez 2017; Mikolajczyk and Grochowski 2018).

Electronic copy available at: https://ssrn.com/abstract=4204973

Figure A.2: Convolutional Neural Net Architecture for Image Classification

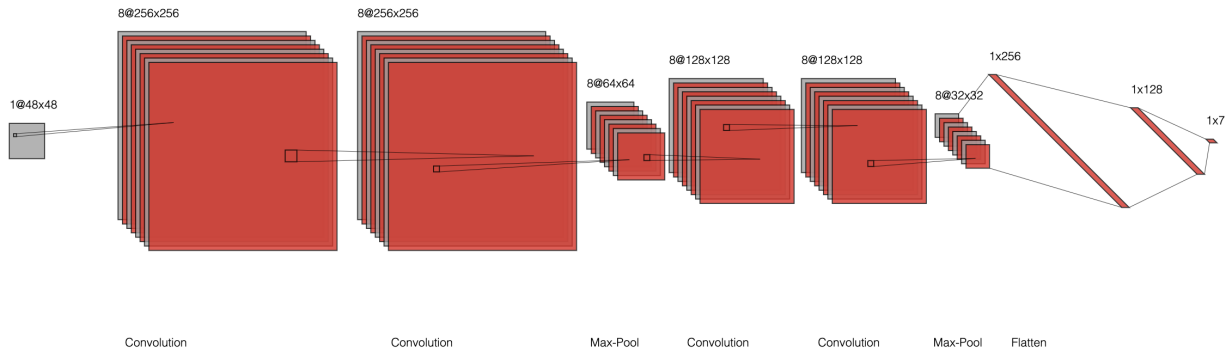

Notes. Diagram for neural net architecture for image classifier. The model uses convolution and pooling layers, both of which aim to reduce the input space (by extracting the key features from the image) and thus decreasing complexity and processing time. The flattening layer converts the output into a one-dimensional array and the classification layer calculates the final classification output.

Electronic copy available at: https://ssrn.com/abstract=4204973

Table A.4: Summary Tabulations on Mentions by Network and Category

	MSNBC	CNN	FNC	**Total**
Democrat	155	140	354	649
Neutral	165	179	387	731
Republican	127	87	194	408
Total	446	406	935	**1788**

Notes. Number of observations – meaning mentions in videos – by network (columns) and word category (rows).

Electronic copy available at: https://ssrn.com/abstract=4204973

Figure A.3: Summary Statistics: Rate of Top Emotions by Class

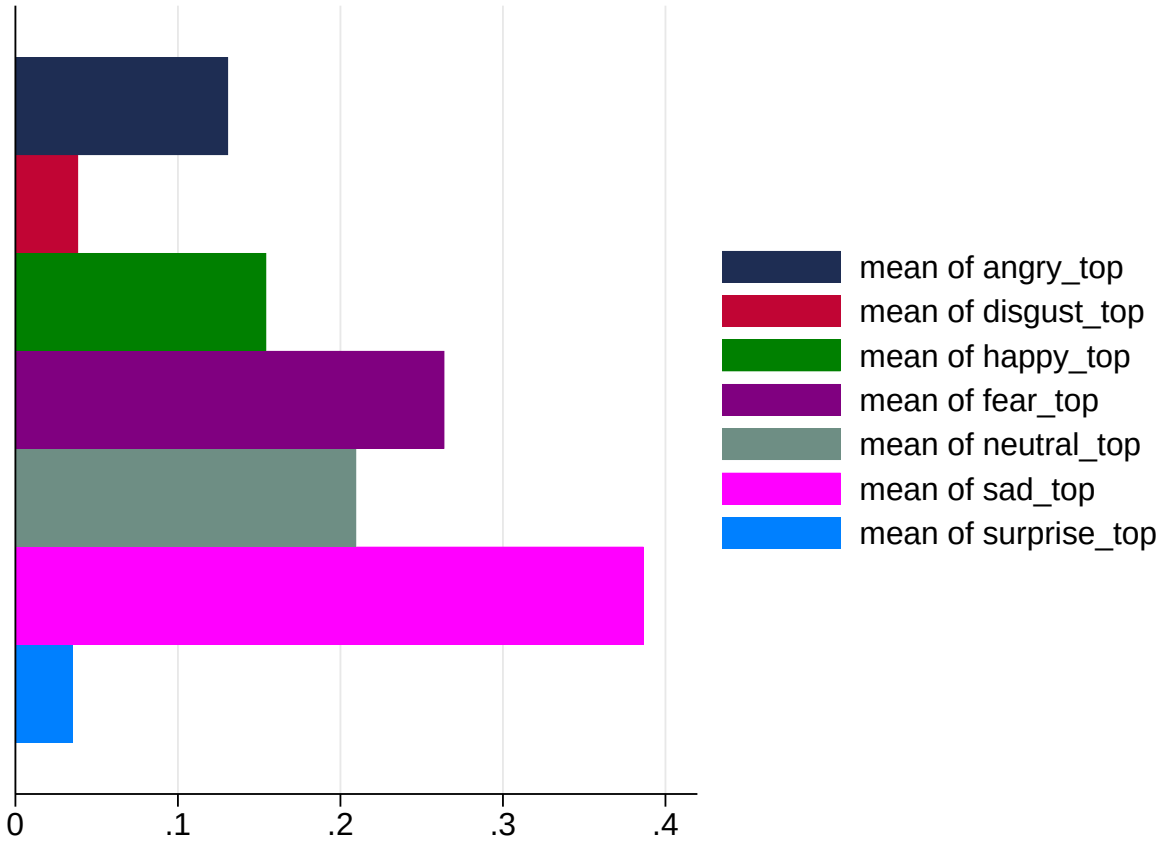

Notes. Summary statistics for all seven emotions for the main outcome – an indicator for whether that emotion is the top emotion in the present mention.

Electronic copy available at: https://ssrn.com/abstract=4204973

Figure A.4: Additional Summary Statistics: Top Emotions

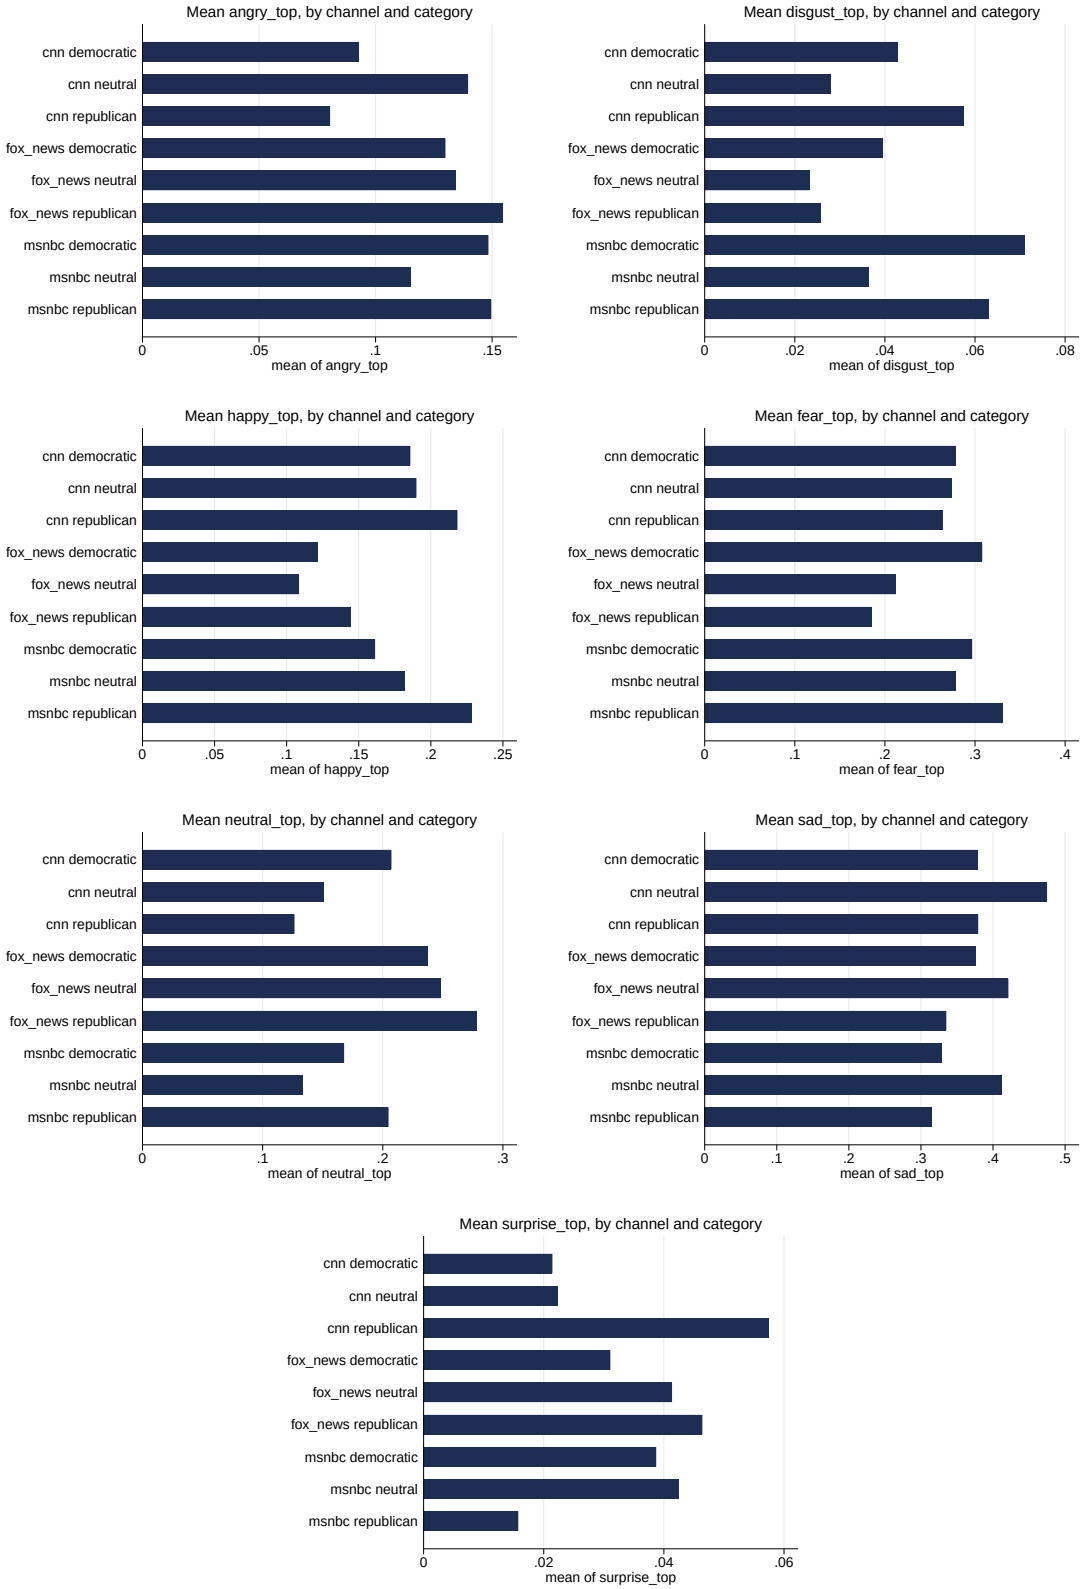

Notes. Mean rate as top emotion, by class, as indicated.

Electronic copy available at: https://ssrn.com/abstract=4204973

Figure A.5: Additional Summary Statistics: Share of Faces

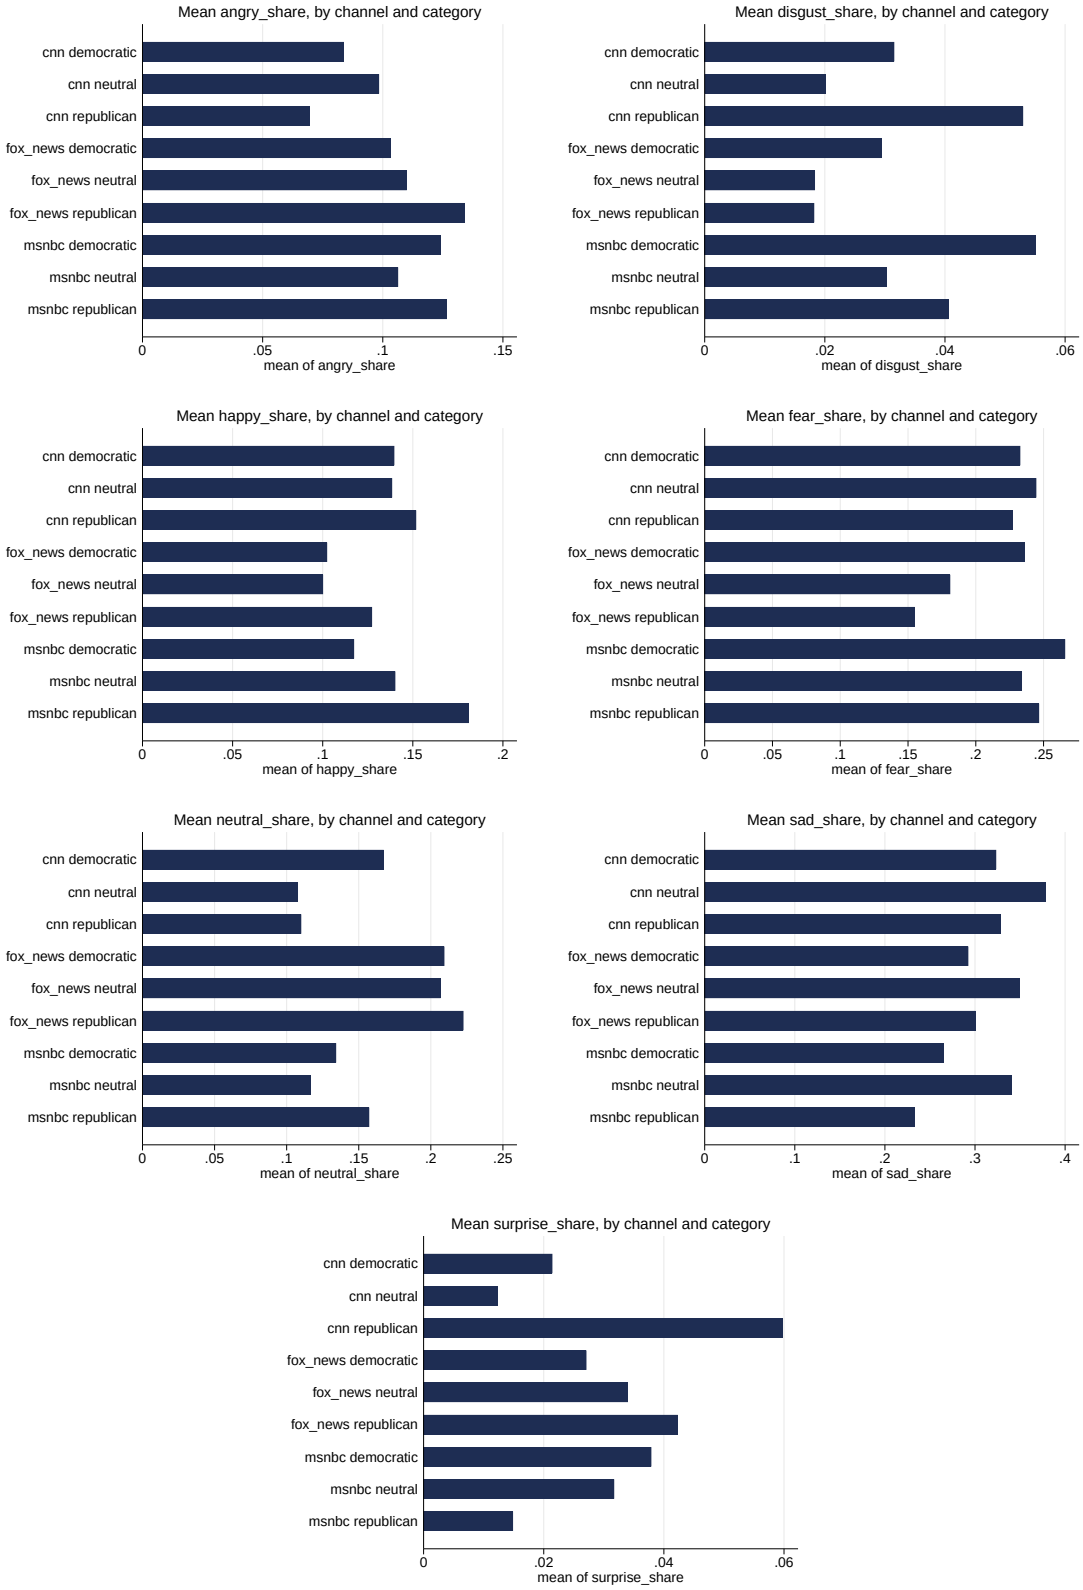

Notes. Mean share of faces on that emotion, by class, as indicated.

Electronic copy available at: https://ssrn.com/abstract=4204973

Figure A.6: Additional Summary Statistics: IHS Face Count Outcome

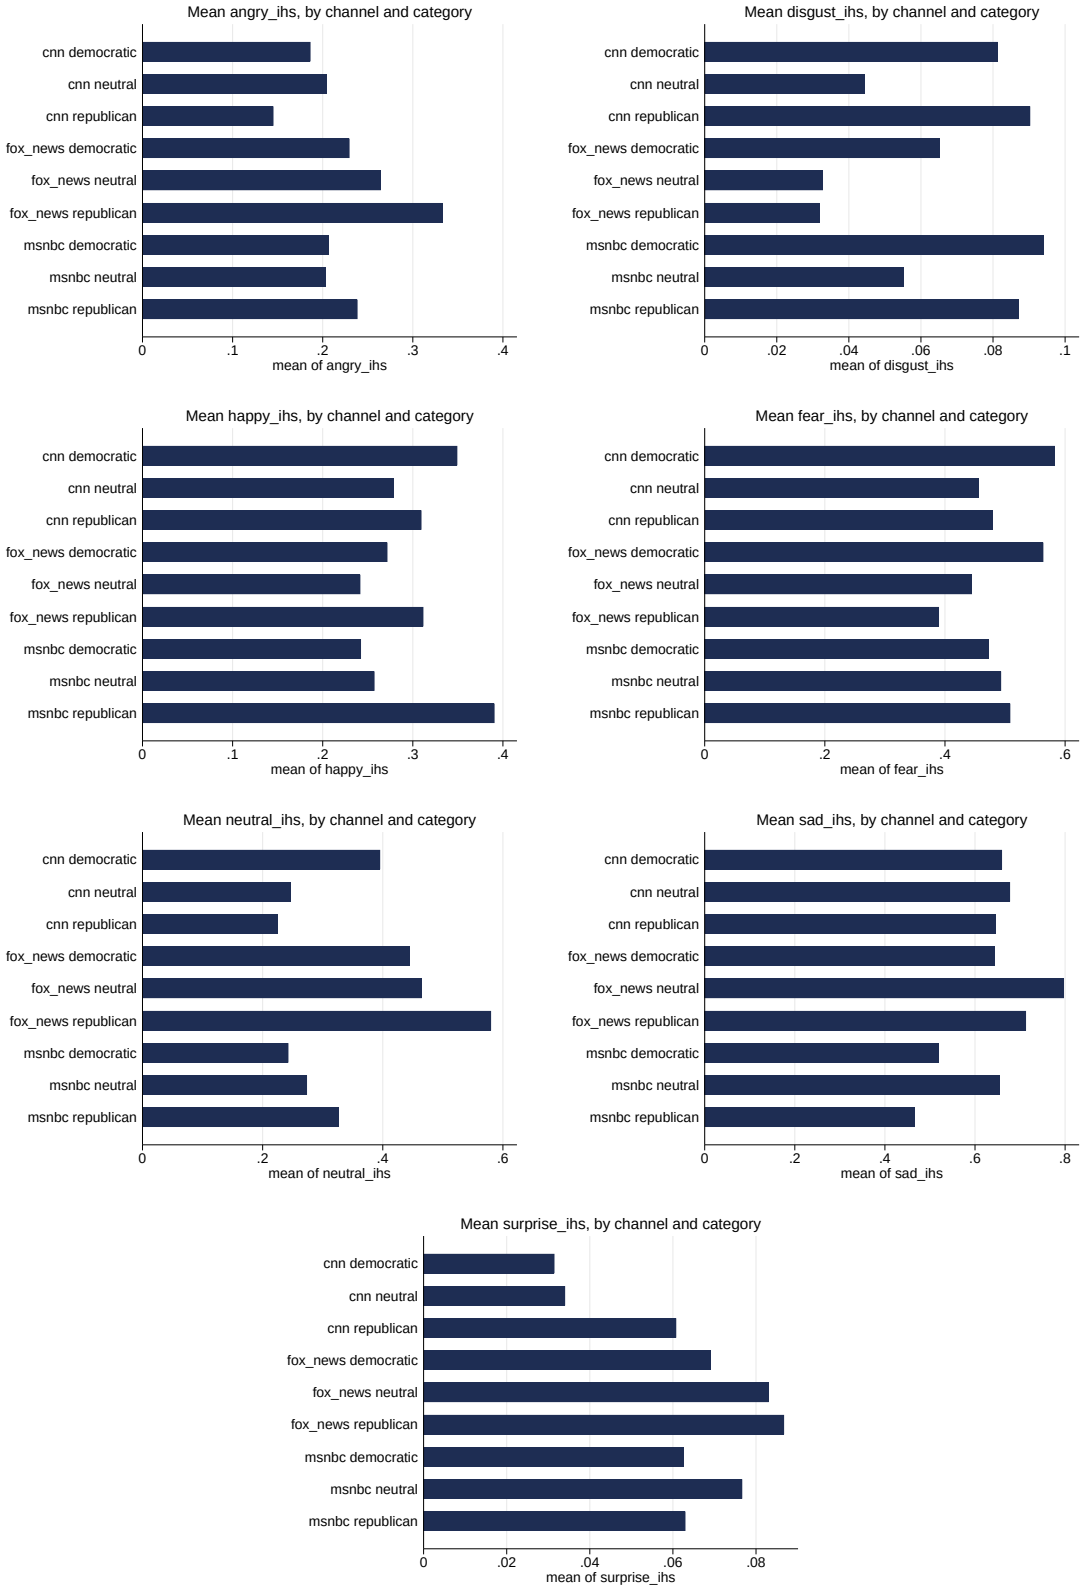

Notes. Mean IHS count of faces with that emotion, by class, as indicated.

Electronic copy available at: https://ssrn.com/abstract=4204973

Figure A.7: Additional Summary Statistics: Any Expression

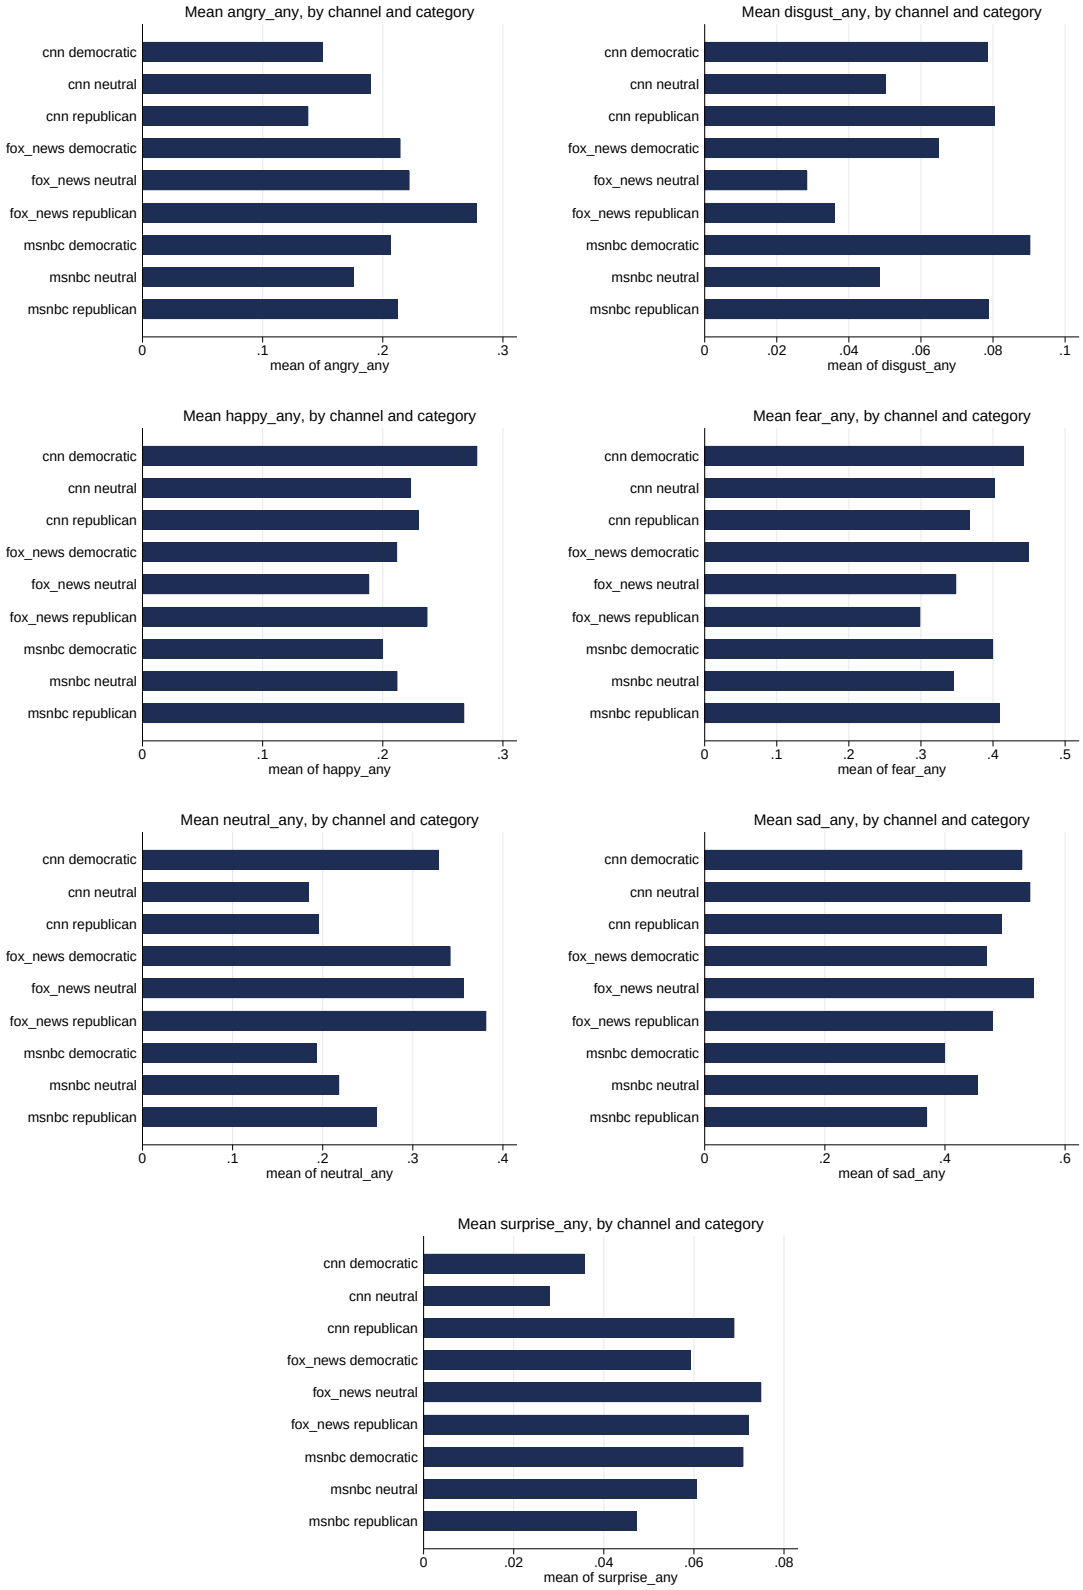

Notes. Mean rate that emotion is in any face, by class, as indicated.

Electronic copy available at: https://ssrn.com/abstract=4204973

Figure A.8: Main Summary Statistics on Fear Expression

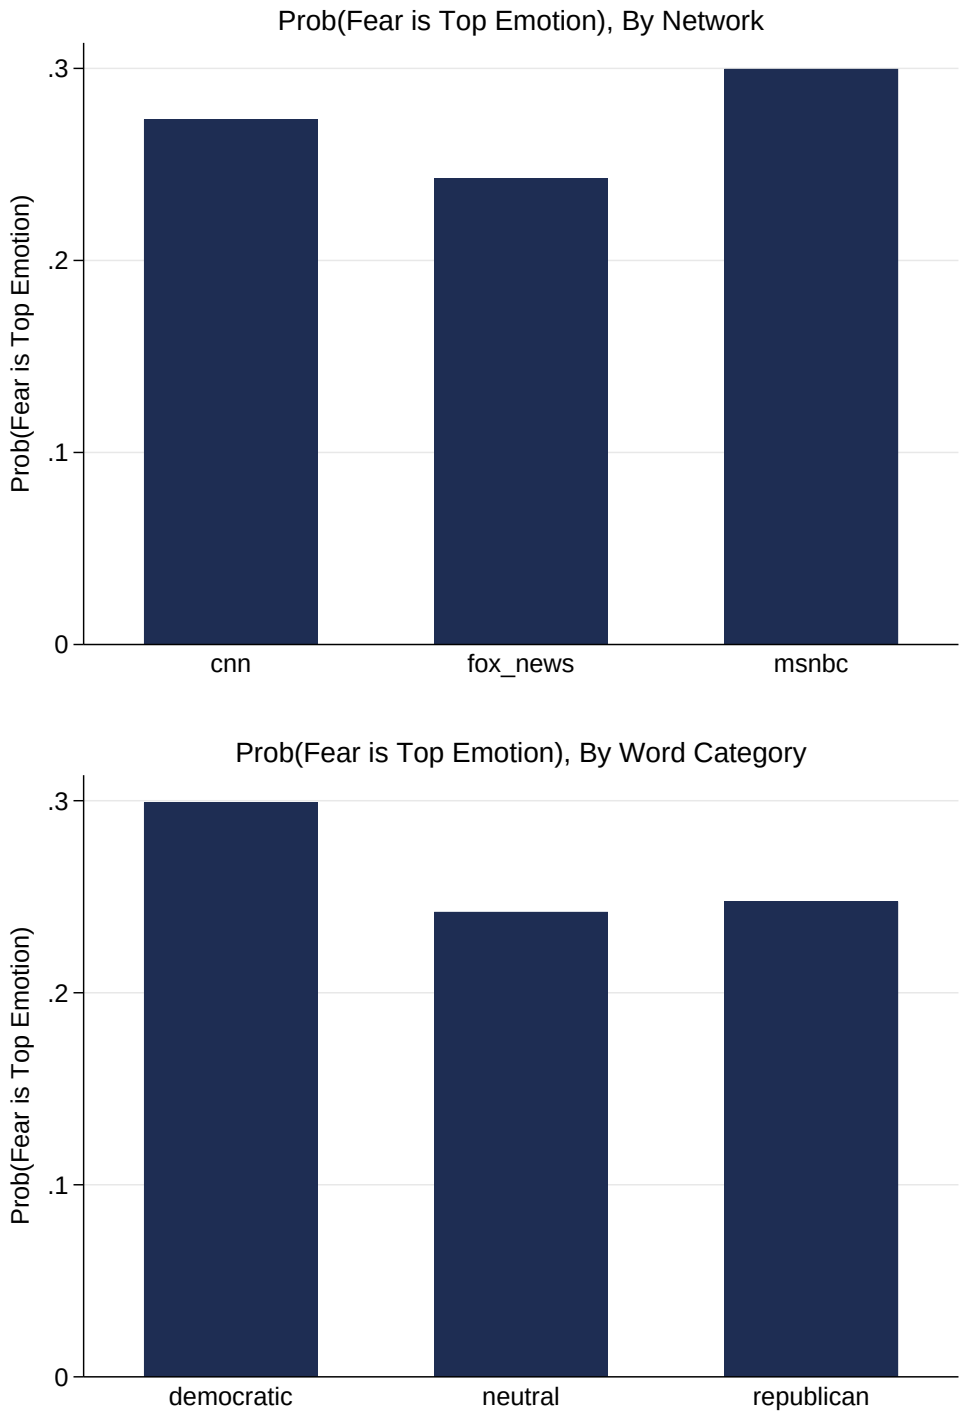

Notes. Summary statistics on main outcome variable – an indicator for whether fear is the top emotion in the present mention.

Electronic copy available at: https://ssrn.com/abstract=4204973

Table A.5: Emotion Expression and Partisan Topics, Joint Specification

	(1)	(2)	(3)	(4)	(5)
Top Emotion	Outcome: Partisan Opposition				
Anger	0.0311	0.0312	0.0166	0.0163	0.014
	(0.0416)	(0.0414)	(0.0320)	(0.0318)	(0.0311)
Disgust	0.1	0.116+	0.0139	0.03	0.0137
	(0.0681)	(0.0679)	(0.0525)	(0.0521)	(0.0508)
Happiness	0.0576	0.0689+	0.0416	0.0526+	0.0329
	(0.0414)	(0.0413)	(0.0319)	(0.0317)	(0.0309)
Fear	0.121**	0.127**	0.0691**	0.0754**	0.0734**
	(0.0341)	(0.0340)	(0.0263)	(0.0261)	(0.0254)
Neutral	0.0594	0.051	0.0426	0.0335	0.0154
	(0.0375)	(0.0373)	(0.0289)	(0.0287)	(0.0279)
Sad	0.0283	0.0266	0.0425	0.0409	0.0263
	(0.0335)	(0.0334)	(0.0258)	(0.0256)	(0.0251)
Surprise	-0.0535	-0.0557	-0.0353	-0.0375	-0.0311
	(0.0698)	(0.0695)	(0.0538)	(0.0533)	(0.0525)
Cable Channel FE		X		X	X
Word Category FE			X	X	
Word FE					X
N	1382	1382	1382	1382	1336
R-sq	0.012	0.022	0.415	0.426	0.515

Notes. Regression results for the joint specification, where the outcome is an indicator for "Partisan Opposition", equaling one for mentions of Democrat entities on FNC, or mentions of Republican entities on MSNBC. That is regressed on indicators for all seven of the emotions, equaling one when the respective emotion is the top/dominant emotion for the given word mention. Cable Channel FE means fixed effect for the cable channel; Word Category FE means fixed effect for the word category (Republican, Democrat, Neutral); Word FE means fixed effect for the specific word (from Appendix Table A.3). Robust standard errors in parentheses. + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$.

Electronic copy available at: https://ssrn.com/abstract=4204973

Table A.6: Emotions and Partisan Topics, Emotion-Outcome Specification

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Effect on Prob(Top Emotion)						
	Anger	Disgust	Happy	Fear	Neutral	Sad	Surprise
Partisan Opposition	-0.0085	-0.00115	0.0128	0.102**	-0.0201	0.00228	-0.0159
	(0.0285)	(0.0159)	(0.0288)	(0.0362)	(0.0342)	(0.0398)	(0.0154)
Cable Channel FE	X	X	X	X	X	X	X
Word FE	X	X	X	X	X	X	X
N	1336	1336	1336	1336	1336	1336	1336
R-sq	0.055	0.054	0.058	0.052	0.053	0.061	0.023

Notes. Regression results for separate specification, by emotion. Outcome is an indicator for whether this word mention's top emotion is as indicated in the column header. The treatment variable "Partisan Opposition" equals one for mentions of Democrat entities on FNC, or mentions of Republican entities on MSNBC. Cable Channel FE means fixed effect for the cable channel; Word FE means fixed effect for the specific word (from Appendix Table A.3). Robust standard errors in parentheses. + p < 0.1, * p < 0.05, ** p < 0.01.

Table A.7: Fear and Partisan Topics, Alternative Outcome Transformations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Effect on Fear Expression Outcome								
		Share			IHS			Any	
Partisan Opposition	0.0441*	0.0303	0.0433	0.104**	0.107*	0.116*	0.0937**	0.0776*	0.0641
	(0.0189)	(0.0247)	(0.0275)	(0.0397)	(0.0520)	(0.0580)	(0.0274)	(0.0359)	(0.0400)
Cable Channel FE	X	X	X	X	X	X	X	X	X
Word Category FE		X			X			X	
Word FE			X			X			X
N	1382	1382	1336	1382	1382	1336	1382	1382	1336
R-sq	0.009	0.012	0.047	0.005	0.007	0.042	0.008	0.011	0.047

Notes. Regression results for the baseline specification, but with different specifications for the fear outcome. "Share" means this is the share of measured emotions indicated as fear; "IHS" means it is the inverse hyperbolic sine of the number of frames indicated as fear; "Any" means that fear is indicated in any of the frames in the snippet. The treatment variable "Partisan Opposition" equals one for mentions of Democrat entities on FNC, or mentions of Republican entities on MSNBC. Cable Channel FE means fixed effect for the cable channel; Word Category FE means fixed effect for the word category (Republican, Democrat, Neutral); Word FE means fixed effect for the specific word (from Appendix Table A.3). Robust standard errors in parentheses. + p < 0.1, * p < 0.05, ** p < 0.01.

Electronic copy available at: https://ssrn.com/abstract=4204973

Figure A.9: Main Results – Alternative Emotion Specifications

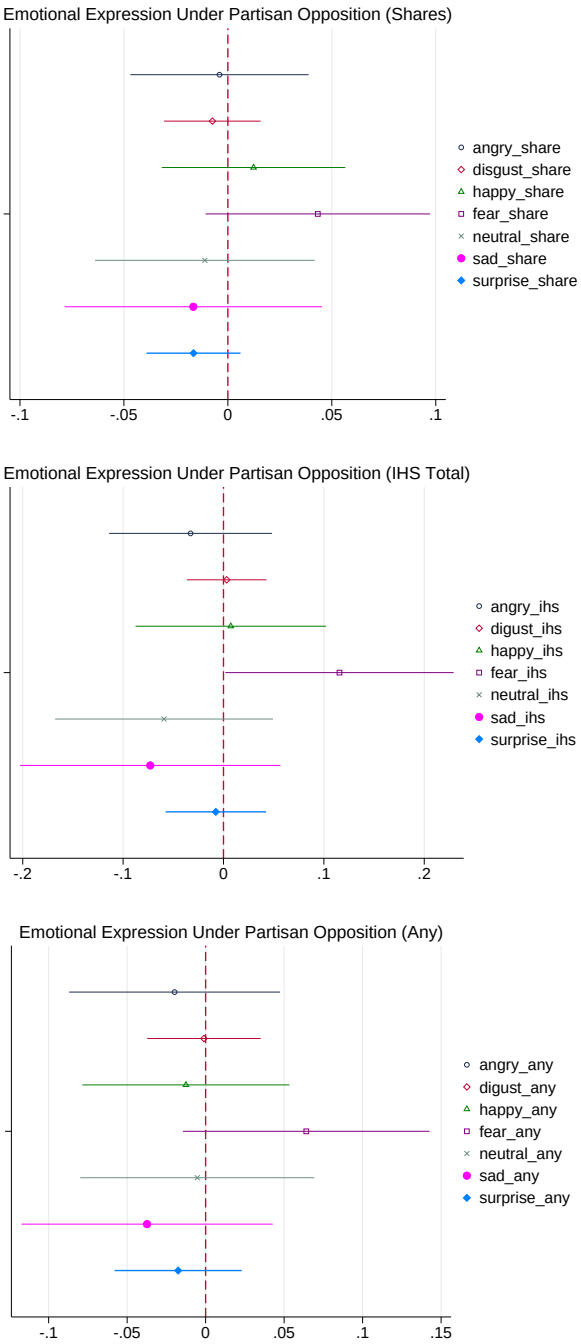

Notes. Coefficient plots from regression results for the separate specification, where the outcome is for one of the seven indicated emotions. "Share" means this is the share of measured emotions indicated; "IHS" means it is the inverse hyperbolic sine of the number of frames indicated; "Any" means that the emotion is indicated in any of the frames in the snippet. The treatment variable "Partisan Opposition" equals one for mentions of Democrat entities on FNC, or mentions of Republican entities on MSNBC. Fixed effects for the cable channel and for the detected word are absorbed. Error bars give 95% confidence intervals constructed using robust standard errors.

Electronic copy available at: https://ssrn.com/abstract=4204973

Table A.8: Fear Expression and Partisan Topics, including CNN

	(1)	(2)	(3)	(4)	(5)
	Effect on Prob(Fear is Top Emotion)				
Partisan Opposition	0.0683**	0.0868**	0.0553*	0.0802**	0.0955**
	(0.0235)	(0.0249)	(0.0279)	(0.0302)	(0.0328)
Cable Channel FE		X		X	X
Word Category FE			X	X	
Word FE					X
N	1788	1788	1788	1788	1734
R-sq	0.005	0.01	0.006	0.011	0.046

Notes. Regression results for baseline specification, where the outcome is an indicator for whether this word mention's top emotion is fear. Unlike the main table, the third "nonpartisan" network CNN is included and enters through the fixed effects. The treatment variable "Partisan Opposition" equals one for mentions of Democrat entities on FNC, or mentions of Republican entities on MSNBC. Cable Channel FE means fixed effect for the cable channel; Word Category FE means fixed effect for the word category (Republican, Democrat, Neutral); Word FE means fixed effect for the specific word (from Appendix Table A.3). Robust standard errors in parentheses. + p < 0.1, * p < 0.05, ** p < 0.01.

Table A.9: Fear Expression and Partisan Topics, Dropping Words with Potential False Positives

	(1)	(2)	(3)	(4)	(5)
	Effect on Prob(Fear is Top Emotion)				
Partisan Opposition	0.0950**	0.120**	0.130*	0.138*	0.0984**
	(0.0345)	(0.0390)	(0.0531)	(0.0597)	(0.0342)
Cable Channel FE		X		X	X
Word Category FE			X	X	
Word FE					X
N	769	725	769	725	969
R-sq	0.02	0.077	0.014	0.075	0.016

Notes. Regression results for baseline specification, where the outcome is an indicator for whether this word mention's top emotion is fear. Unlike main table, drops mentions of words that may generate false positives, as indicated in Appendix Table A.3. The treatment variable "Partisan Opposition" equals one for mentions of Democrat entities on FNC, or mentions of Republican entities on MSNBC. Cable Channel FE means fixed effect for the cable channel; Word Category FE means fixed effect for the word category (Republican, Democrat, Neutral); Word FE means fixed effect for the specific word (from Appendix Table A.3). Robust standard errors in parentheses. + p < 0.1, * p < 0.05, ** p < 0.01.

Electronic copy available at: https://ssrn.com/abstract=4204973

Table A.10: Robustness to Clustering Standard Errors

	(1)	(2)	(3)
	Prob(Fear is Top Emotion)		
Partisan Opposition	0.102**	0.102**	0.102**
	(0.0362)	(0.0376)	(0.0274)
Clustering	None	Video	Word
Cable Channel FE	X	X	X
Word FE	X	X	X
N	1336	1336	1336
R-sq	0.052	0.052	0.052

Notes. Regression results for baseline specification, where the outcome is an indicator for whether this word mention's top emotion is fear. Unlike main table, adds clustering of standard errors as indicated – by the video, or by the mentioned word. The treatment variable "Partisan Opposition" equals one for mentions of Democrat entities on FNC, or mentions of Republican entities on MSNBC. Cable Channel FE means fixed effect for the cable channel; Word Category FE means fixed effect for the word category (Republican, Democrat, Neutral); Word FE means fixed effect for the specific word (from Appendix Table A.3). Robust standard errors in parentheses. + $p < 0.1$, * $p < 0.05$, ** $p < 0.01$.

Electronic copy available at: https://ssrn.com/abstract=4204973