Additional Analysis and Methodology

CONTENTS

1 ADDITIONAL RESULTS

We report additional observations from the dataset that were omitted from the main paper, due to space considerations or the analysis being focused on text (section 1.2) rather than faces. While many of these patterns are not surprising, we use the following as examples to illustrate the types of questions that can be investigated with a decade of news video.

1.1 Who is in the news?

How much time is there when at least one face is on-screen in commercials? As observed in the main paper, the percentage of screen time when a face is on-screen in news content has risen from by 8.6%, from 72.9% in 2010 to 81.5% in 2019. This same percentage has only risen slightly in commercials in the same timespan (38% to 41%). Figure 1 shows the time-series chart that was omitted.

What is the average number of on-screen faces? The average number of faces visible on-screen is 1.38 in news content and 0.49 in commercials, and these figures vary little between channels. There is a rise in the number of faces over the decade, across all three channels, from 1.2 in 2010 to 1.6 in 2019, with much of the increase since 2015. By contrast, the average number of on-screen faces in commercials rises from 0.42 to 0.52 and is flat from 2012 onward. Figure 2 shows the time-series chart that was omitted.

What is the average size of faces? The average size of detected faces in news content, as a proportion of the frame height, has remained consistent over the decade, at roughly 33%, both in content and commercials (Figure 3). Note that some videos have black horizontal bars on the top and the bottom due a mismatch between the video resolution and the video aspect ratio (16:9 inside 4:3). We excluded these black bars from the frame height calculation.

We also investigate whether face size varies across data slices, such as by gender and presenter role, in news content and in commercials. Figure 4 shows the normalized distributions of face heights, by presented-gender, for news presenters and other faces in content; and all faces in commercials. In the news content, there are three dominant modes for both news presenters and other faces, likely corresponding to common visual formats in presentation.

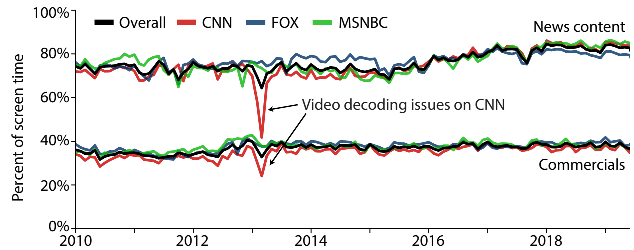

Figure 1: The percentage of time when faces are on-screen has increased for news content, but has remained static in commercials since 2013.

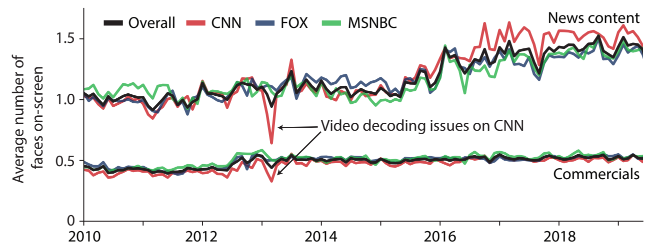

Figure 2: The average number of on-screen faces has increased on all three channels.

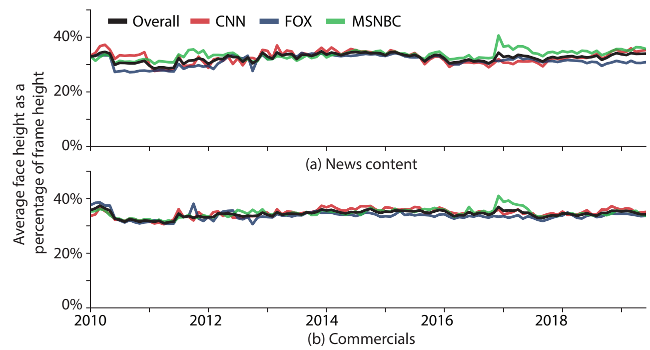

Figure 3: The average height of on-screen faces has remained mostly constant in both news content and commercials, but there is some variation within the decade. The average height of faces in news content and commercials is similar.

These modes are not present in commercials. Compared to news presenters, other faces in the news content are more likely to belong to the left most (smallest face) peak.

The peaks for female-presenting faces are slightly offset to the left of the peaks for male-presenting faces (Figure 4ab). Likewise, the lack of a disparity in commercials (Figure 4c) suggests that the disparity in news content is not caused by inherent bias in the face detector (for instance, in how it handles hair). While the cause is not obvious, one consequence of these distributional differences is that enforcing any face height threshold when analyzing the news content will disproportionately exclude female-presenting faces.

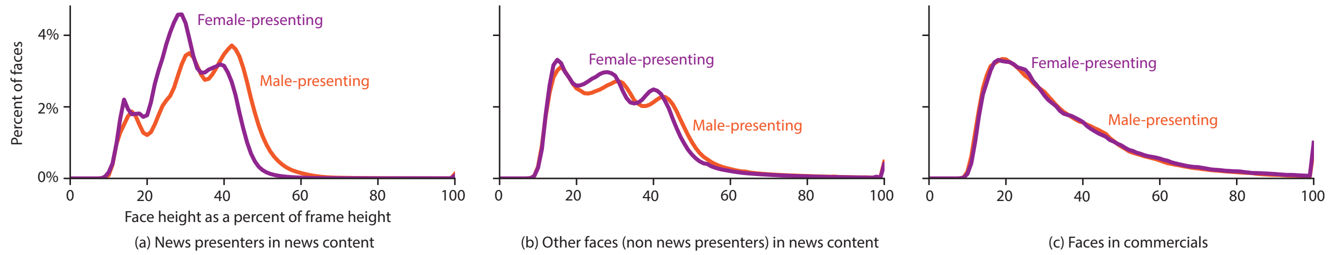

Figure 4: The distributions for the height of faces is tri-modal in news content, for both news presenters (a) and other faces (b), suggesting common visual formats in presentation. The peaks for female-presenting faces in news content are slightly to the left of the peaks for male-presenting faces. Interestingly, neither of these patterns are present in the commercials (c).

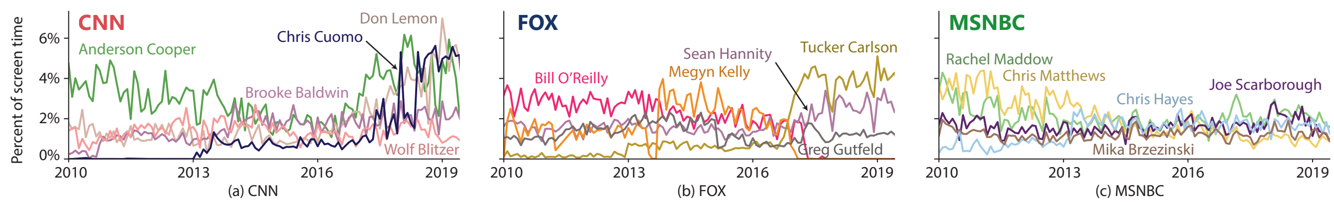

Figure 5: Screen time of the top five presenters on each channel. Since 2016, several of the top presenters on CNN have dramatically risen in screen time. Following Bill O'Reilly's firing and Megyn Kelly's departure from FOX in 2017, Sean Hannity and Tucker Carlson have risen in screen time. Since 2013, the variation in screen time among the top five hosts on MSNBC has been low compared to CNN and FOX.

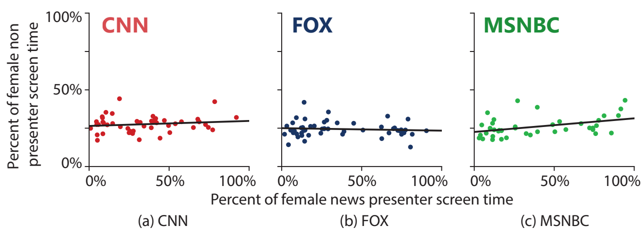

Figure 6: There is little correlation between shows that are predominantly presented by female-presenting news presenters and shows with the most screen time for female-presenting faces who are not news presenters.

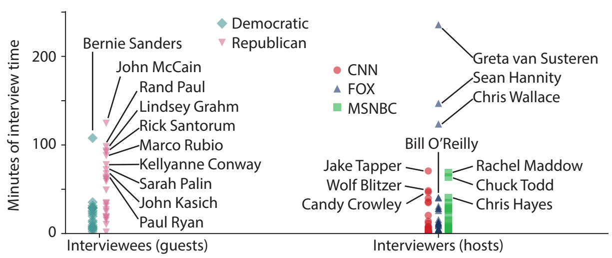

Figure 7: Interview time of the 44 politicians (interviewees) tested and hosts (interviewers). Note: Bernie Sanders is labeled Democratic due to his affiliation in the 2016 primary.

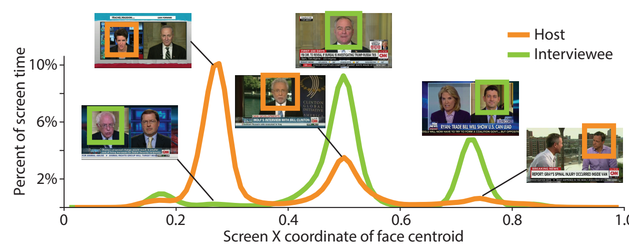

Figure 8: In interviews, the host appears overwhelmingly on the left or in the middle; interviewees appear in the middle or on the right.

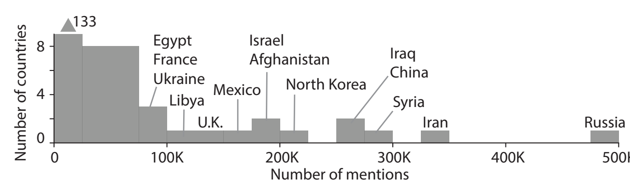

Figure 9: Not all countries receive equal attention in U.S. cable TV news than others. Russia, by a large margin, is covered the most, followed by Iran.

These patterns are present on all three channels; the shape and positions of the peaks vary slightly, while the degree of face height disparity between male- and female-presenting faces is similar.

Which news presenters receive the most screen time? The news presenters with the most screen time are Anderson Cooper (1,782 hours) on CNN, Bill O'Reilly (1,094) on FOX, and Rachel Maddow (1,202) on MSNBC. While the top presenter on each channel varies over the course of the decade (Figure 5), Cooper and O'Reilly hold the top spot for relatively long stretches on CNN and FOX, respectively. Maddow appears the most on MSNBC overall, but Chris Matthews holds the top spot for the early part of the decade (2010 to 2014). Since 2014, the top presenter on MSNBC has fluctuated on a monthly basis (Figure 5c).

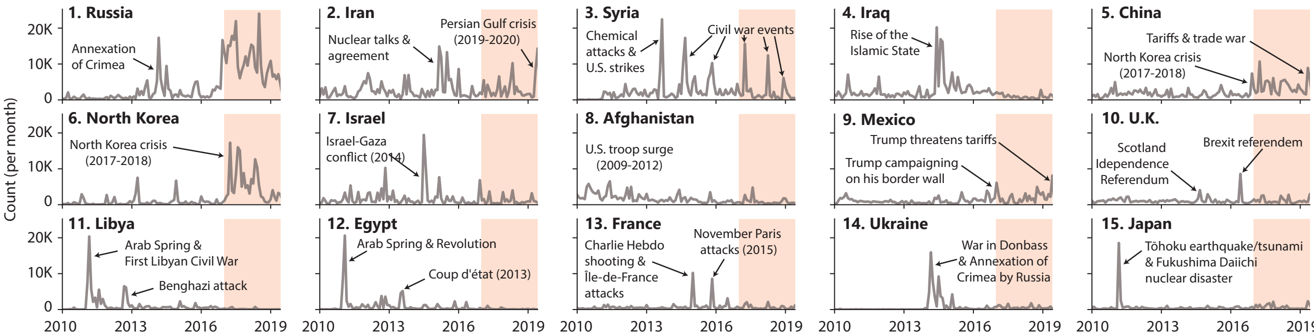

Figure 10: Major peaks in mentions of foreign countries occur around disasters and crises. Since the start of Trump's presidency, there has been an increase in coverage of Russia, China, and North Korea due to increased tensions and a marked shift in U.S. foreign policy (shaded).

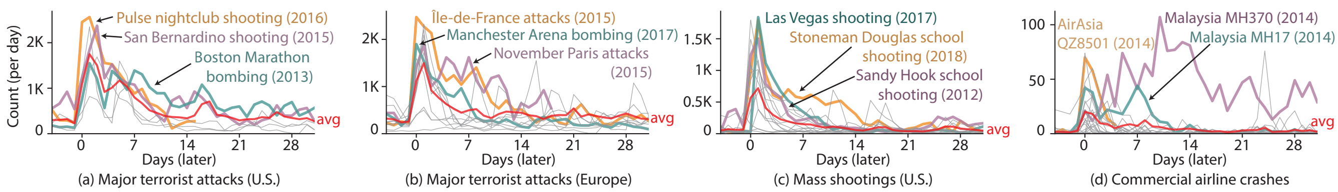

Figure 11: Following a major terrorist attack, mass shooting, or plane crash, usage of related terms increases and remains elevated for 2-3 weeks before returning to pre-event levels. A few plane crashes continued to be covered after this period as new details about the crash (or disappearance in the case of MH370) emerge. In the figure above, lines for individual events are terminated early if another unrelated event of the same category occurs; for example, the San Bernardino shooting (a terrorist attack) in December 2015 occurred three weeks after the November 2015, Paris attacks.

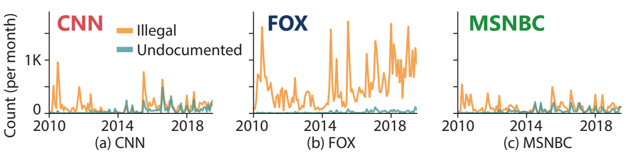

Figure 12: Counts of "illegal immigrant" and "undocumented immigrant" terminology in video captions, by month. `Illegal` is more common than `undocumented` on all three channels, but FOX uses it the most. `Undocumented` only comes into significant use from 2012 onward.

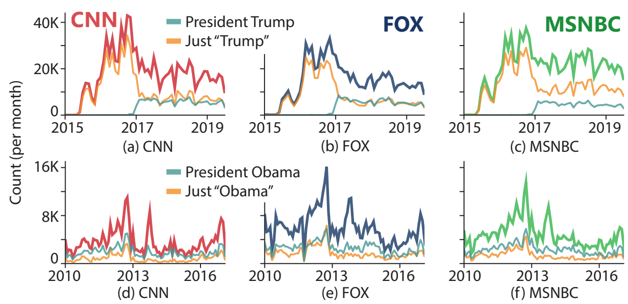

Figure 13: Counts of `Trump` and `Obama` peaked in election years (2016 and 2012). After his inauguration, `Trump` is referred to more often without `President` than with (MSNBC has the largest gap). In contrast, `Obama` is referred to with `President` more often than not. The channel color-coded lines represent the total counts of `Trump` and `Obama`, without exclusions such as the `Trump administration`. Note: most of these counts are captured by the n-grams that we identified as references to Trump and Obama's persons.

Do shows presented by female-presenting news presenters give more screen time to other female-presenting individuals overall? An individual show's gender balance is biased heavily by the gender of its host. For example, the show with the greatest female-presenting screen time is *Melissa Harris-Perry* on MSNBC and the show with the greatest male screen time is *Glenn Beck* on FOX.

We use the percentage of female-presenting news presenter screen time out of total news presenter screen time to measure the extent to which a show is female- or male-presented. As a measure of the gender balance for female-presenting individuals who are not presenters (non-presenter), we compute the percentage of female-presenting screen time for other faces (not identified as a news presenter) out of the screen time for all other faces. To evaluate whether shows that lean toward more female-presenting news presenter screen time also have more screen time for other female-presenting faces in general, we measure the linear correlation between the two percentages. We limited the analysis to shows with at least 100 hours of news content, to exclude short-lived shows and special programming.

We find no correlation on CNN (*slope* $= 0.03, R^2 = 0.02$) and FOX (*slope* $= -0.02, R^2 = 0.01$), and a weak positive correlation on MSNBC (*slope* $= 0.09, R^2 = 0.19$) (Figure 6). This suggests that shows hosted by female-presenting news presenters do not give proportionally more screen time to female-presenting subjects and guests. Our result contrasts with findings by the GMMP [5] that female journalists write disproportionately more articles about female subjects.

Which politicians get interviewed? Which presenters do interviews? Interviews are one of the ways that cable news networks bring on experts and provide politicians with a platform to express their views. We find interviews by looking for continuous segments of video when a presenter (interviewer) and interviewee are on-screen together and/or alternating back and forth (details in section 2.5.7). Empirically, we found that this approach identifies interview segments for 44 prominent American political figures that we tested (including 17 2016 US presidential candidates). (Note: we exclude Barack Obama, Mitt Romney, Donald Trump, and Hillary Clinton because they appear too frequently in non-interview contexts, leading to low precision in detecting interviews. Newt Gingrich and Mike Huckabee, who are both hosts and political figures, are also excluded.)

While the aforementioned exclusions and limited sample prevent this from being a true commentary on interview statistics for politicians in US news, the detection method and results are nonetheless interesting from a data exploration standpoint. In the interviews that we did detect (Figure 7), John McCain is featured the most. Many of the top interviewees shown are Republicans, due to our biased sampling toward 2016 presidential candidates and the relatively competitive and crowded Republican primary (compared to the Democratic primary that year). The top three interviewers are all hosts on FOX; Greta van Susteren (former host of *On the Record* on FOX) is the most prolific.

What is the visual layout of interviews? In the majority of interviews, the host appears on the left (split-screen) or in the middle, while the interviewee typically appears on the right (split-screen) or in the middle (Figure 8). This is in contrast to late night talk shows, which place the host on the right.

1.2 What is discussed in the news?

The amount of coverage that topics receive in the news can influence viewer perceptions of world events and newsworthy stories. To measure the frequency at which key topics are discussed, we count the number of times that selected words appear in the video captions.

How often are foreign countries mentioned? Foreign country names, defined in section 2.5.1, appear in the captions a total of 4.5M times. Most countries receive little coverage (Figure 9), and the eight countries with the highest number of mentions (`Russia`, `Iran`, `Syria`, `Iraq`, `China`, `North Korea`, `Israel`, and `Afghanistan`) account for 51% of all country mentions. `Russia` alone accounts for 11.2%. (If treated as a country, `ISIS` would rank 2nd after `Russia` at 8.4%.) Of these eight, five have been in a state of armed conflict in the last decade, while the other three have had major diplomatic rifts with the US. These data suggest that military conflict and tense US relations beget the prolonged coverage. No countries from Africa, South America, and Southeast Asia appear in the top eight; the top countries from these regions are `Libya`/`Egypt` (11th/12th), `Venezuela` (32nd), and `Vietnam` (25th). Bordering the US, `Mexico` is 9th, frequently appearing due to disputes over immigration and trade, while `Canada` is 21st.

Mentions of individual countries often peak due to important events. Figure 10 annotates such events for the 15 most often mentioned countries. For example, the Libyan Civil War in 2011, the escalation of the Syrian Civil War in 2012-2013, and the rise of ISIS (Syria, Iraq) in 2014 correspond to peaks. The countries below the top 10 are otherwise rarely in the news, but the 2011 tsunami and Fukushima Daiichi nuclear disaster; the 2014 annexation of Crimea by Russia; and the Charlie Hebdo shooting and November Paris attacks (both in 2015), elevated Japan, Ukraine, and France to brief prominence. Since the election of Donald Trump in 2016, however, there has a been a marked shift in the top countries, corresponding to topics such as Russian interference in US elections, North Korean nuclear disarmament talks, the Iran nuclear deal, and the trade war with China.

For how long do channels cover acts of terrorism, mass shootings, and plane crashes? We enumerated 18 major terrorist attacks (7 in the U.S. and 11 in Europe), 18 mass shootings, and 25 commercial airline crashes that occurred in the last decade, and we counted related n-grams such as `terror(ism,ist)`, `shoot(ing,er)`, and `plane crash` in the weeks following these events (section 2.5.2 gives the full lists of terms). Counts for `terrorism` and `shootings` return to the pre-event average after about two weeks (Figure 11a-c). Likewise, coverage of plane crashes also declines to pre-crash levels within two weeks (Figure 11d), though there are some notable outliers. Malaysia Airlines Flight 370, which disappeared over the Indian Ocean in 2014, remained in the news for nine weeks, and Malaysia Airlines Flight 17, shot down over Ukraine, also received coverage for four weeks as more details emerged.

Is it illegal or undocumented immigration? "Illegal immigrant" and "undocumented immigrant" are competing terms that describe non-US citizens who are residing in the US without legal documentation (e.g., a valid visa, green card, etc.), with the latter term seen as more politically correct [7]. Figure 12 shows the counts of when variants of these terms are said (section 2.5.3 gives the full list of variants). `Illegal` is used on FOX the most (59K times); FOX also has more mentions of immigration overall. From 2012 onward, `undocumented` has increased in use on CNN and MSNBC, though `illegal` still appears equally or more often on these channels than `undocumented`. Further analysis of who uses these terms on each channel is an interesting topic of future work.

How often are honorifics used to refer to President Trump and Obama? Honorifics convey respect for a person or office. We compared the number of times that `President (Donald) Trump` is used compared to other mentions of Trump's person (e.g., `Donald Trump`, just `Trump`). When computing the number of mentions of just `Trump`, we exclude references to nouns such as `the Trump administration` and `Melania Trump` that also contain the word Trump, but are not referring to Donald Trump (section 2.5.4 gives the full list of exclusions).

Our data suggests that although coverage of the incumbent president has increased since the start of Trump's presidency in 2017, the level of formality when referring to the president has fallen. Trump, in general, is mentioned approximately 3× more than Obama on a monthly basis during the periods of their respective presidencies in our dataset. As per the definition of being the president, the term `President Trump` only emerges on all three channels following his

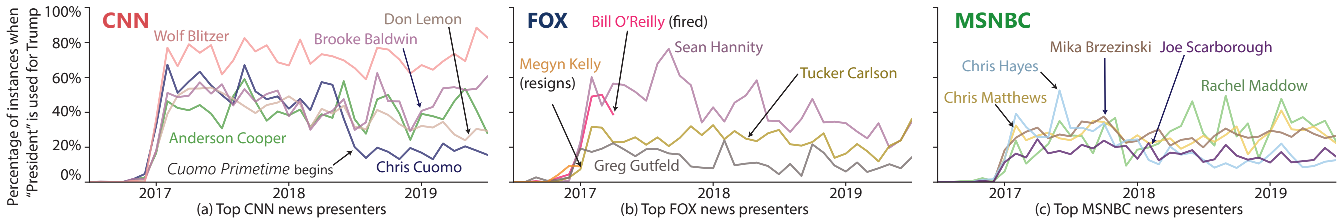

Figure 14: Percentage of time when the `president` honorific is said for `Trump` while a news presenter is on-screen increases after Trump's inauguration (top 5 presenters for each channel are shown). Chris Cuomo (CNN) drops from over 40% to under 20% in June 2018 with his transition from hosting *New Day* to *Cuomo Primetime*. Sean Hannity's (FOX) decline is more gradual over the course of Trump's presidency. From 2017 onward, Wolf Blitzer (CNN) is consistently above the other top hosts on any of the three channels (averaging 72%).

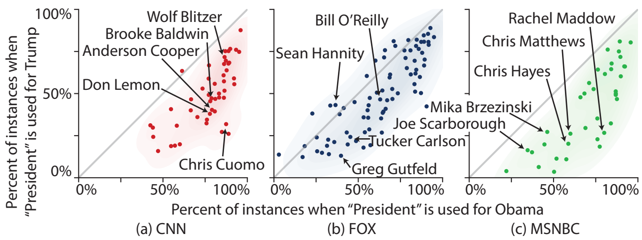

Figure 15: Percentage of mentions that use the `president` honorific for Trump (post-inauguration to January 20, 2017) and Obama (before January 20, 2017) by each news presenter (dots). A majority of presenters on all three channels use `president` a higher fraction of time when mentioning Obama than they do with Trump. The presenters with the highest screen time on each channel are annotated.

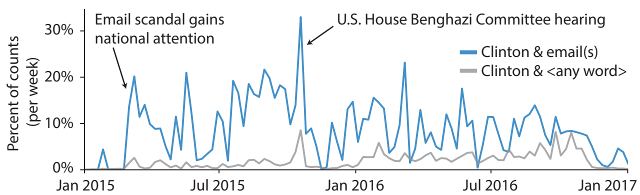

Figure 16: Hillary Clinton is on-screen up to 33% of the time when `email(s)` is mentioned (11% on average from 2015 to 2016). This is significantly higher than the percentage of time that Clinton is on-screen when any arbitrary word is said (1.9% on average in the same time period).

Person	Unique words ($Pr[person \mid word]$)
Bill O'Reilly (FOX)	opine (60.6), reportage (59.0), spout (58.6), urchins (57.9), pinhead[ed,s] (49.0, 51.5, 50.2)
Ed Schultz (MSNBC)	classers (71.2), beckster (61.6), drugster (59.9), righties (55.2), trenders (60.8), psychotalk (54.2)
Tucker Carlson (FOX)	pomposity (76.2), smugness (71.5), groupthink (70.5)
Sean Hannity (FOX)	abusively (76.1), Obamamania (53.3)
Glenn Beck (FOX)	Bernays (82.3), Weimar (62.2)
Rachel Maddow (MSNBC)	[bull]pucky (47.9, 50.7), debunktion (51.4)
Chris Matthews (MSNBC)	rushbo (50.5)
Kevin McCarthy (politician)	untrustable (75.9)
Chris Coons (politician)	Delawareans (63.8)
Hillary Clinton (politician)	generalistic (56.5)

Table 1: Unique words are often euphemisms or insults (`urchins` ≡ children, `beckster` ≡ Glenn Beck, `drugster`/`rushbo` ≡ Rush Limbaugh, `righties` ≡ conservatives, etc.). Others are the names of show segments or slogans. For example, `Psychotalk` is a segment of the *Ed Show*; Sean Hannity refers to the liberal media as `Obamamania media`; and Tucker Carlson brands his own show as the "sworn enemy" of `lying`, `pomposity`, `smugness`, and `groupthink`. Some rare words become unique due to being replayed often on the news; for example, Kevin McCarthy (U.S. representative) calls Hillary Clinton `untrustable` and Hillary Clinton uses `generalistic` in the same sentence as her infamous statement characterizing Trump's supporters as a "basket of deplorables".

inauguration to the office in January 2017 (Figure 13a-c), and, after this event, individuals on CNN and FOX use `President` nearly half of the time. By contrast, individuals on MSNBC continue to most commonly refer to him without `President`. We plot similar charts of `President Obama` over the course of his presidency from 2010 to January 2017 (Figure 13d-e) and find that, on all three channels, `President` is used more often than not.

1.3 What is said when people are on-screen?

Who uses unique words? We define vocabulary to be "unique" to a person if the probability of that individual being on-screen conditioned on the word being said (at the same time) is high. Table 1 lists all words for which an individual has a greater than a 50% chance of being on-screen when the word is said. (We limit analysis to words mentioned at least 100 times.) Political opinion show hosts (on FOX and MSNBC) take the most creative liberty in their words, accounting for all but three names in the list.

Which presenters are on-screen when the President honorific is said? A news presenter's use of the `President` honorific preceding Trump or Obama might set a show's tone for how these leaders are portrayed. As an approximation, we measure the visual association between news presenters' faces and the word `President`.

When a presenter is on-screen, we find that the honorific term `President` is used a greater percentage of time for Obama than for Trump, during the periods of their presidencies. On all three channels, most presenters lie below the parity line in Figure 15. However, the average FOX presenter is closer to parity than the average presenter on CNN or MSNBC in uses of `President` in reference to Trump and Obama (a few FOX presenters lie above the line). Furthermore, Figure 14 shows how the top hosts (by screen time) on each channel are associated with uses of `President` to refer to Trump over time. (Note: we cannot conclude from this data that news presenters say or do not say the president honorific, only that they are on-screen when it is said or not said. Refer to section 2.5.5 for details.)

How much was Hillary Clinton's face associated with the word* `email`*? Hillary Clinton's emails were a frequent news topic in 2015 and during the 2016 presidential election due to investigations of the 2012 Benghazi attack (on US diplomatic facilities) and her controversial use of a private email server while serving as the US Secretary of State. During this period, Clinton's face was often on-screen when these controversies were discussed, visually linking her to the controversy. We compute that during the time period from 2015 to 2016, Clinton's face is on-screen during 11% of mentions of the word `email(s)` (Figure 16), a significantly higher percentage than the 1.9% of the time that she is on-screen overall. (Refer to section 2.5.6 for details.)

2 EXTENDED METHODOLOGY

This section contains additional methodology and tables omitted from the condensed supplemental material in the main paper.

2.1 Commercial detection

Our heuristic algorithm is written using Rekall [4], an API for complex event detection in video, and is shown in Figure 17.

2.2 Identifying public figures

For important individuals who are not recognized by the Amazon Rekognition Celebrity Recognition API [1] or whose labels are known to be inaccurate, we train our own person identification models using FaceNet [8] descriptors. In the latter case, we determine a person's labels to be inaccurate if they are consistently missed or mis-detected on visual inspection of the videos. To obtain our own labels, we developed two human-in-the-loop labeling workflows optimized for people who are common (e.g., a President or news presenter who appears for hundreds of hours) and for people who are uncommon (e.g., a shooting victim or less-known public official). The approaches are described in 2.2.1 and 2.2.2, respectively. We determine which approach to use experimentally; if we can not find enough training examples for the common person approach, we switch to the uncommon person approach. The individuals for which we use our own labels are listed in Table 2.

Table 3 estimates the precision and recall of the labels for the individuals referenced in our paper analyses (e.g., important political figures and candidates). Precision is influenced by many factors, including the presence of individuals of similar appearance being prominent in the news. Because each individual represents only a small portion of overall face screen time, unbiased recall is difficult to compute without finding all instances of an individual. We make a best effort attempt to estimate recall by manually counting false negatives in frames sampled randomly from videos known to contain the individual (25 videos, 100 frames per video). We note that the number of samples per individual, found in these frames, varies due to the quantity and nature of an individual's coverage (e.g., appearances in interviews and the quality of their images).

2.2.1 Method for detecting uncommon individuals. To detect uncommon individuals (with less than $\approx$ 50 hours of screen time or 60,000 face detections), we use Google Image Search [6] to obtain initial images of the person. Next, we use FaceNet [8] to compute descriptors on these examples. We compute the L2 distances from these descriptors to the descriptors for all other faces in the dataset and display the faces visually by ascending L2 distance. We select instances of the faces that visually match the person, add them to the example set and repeat the process of computing L2 distances and displaying images until it becomes difficult to find additional examples (the top candidates are all images of other people). To make the selection process more time-efficient, we implemented range navigation and selection to label faces between L2 distance ranges at once if all or nearly all of the faces in the range are the correct person. Even so, the primary limitation of this approach is that the labeling time scales linearly with the frequency of the individual in the dataset.

2.2.2 Method for detecting common individuals. To detect common individuals, for whom it is impossible to browse all of their detections, we trained a simple linear classifier on the FaceNet [8] features. We use Google Image Search [6] to find initial examples, and we augment those by sampling faces from the dataset that are similar to the examples in FaceNet descriptor space. For the initial negative examples, we sample faces randomly and manually inspect the random samples that are most likely (based on L2 distance) to be positive examples. (This step is necessary because common individuals such as Donald Trump are likely to appear in the negative samples due to their high frequency in the dataset.) We then use these positive and negative examples to train a model. To further refine the model, we sample faces for which the model produces low confidence scores ($\approx$0.5) and label these as new examples or non-examples, repeating the training and labeling process until finding new positive examples becomes a challenge and the model precision is sufficient (evaluated by visually validating the faces that are labeled positive by the model).

2.3 Enumerating news presenters

Tables 4, 5, and 6 list the news presenters that we identified on each of the three channels (along with each person's estimated screen time). Note that the percentage of female-presenting individuals in the list is 52%, 42%, and 44% on CNN, FOX, and MSNBC, respectively.

2.4 Age for news presenters

We obtained birthdate information for 98% of the news presenters that we enumerated in section 2.3 using DBpedia [2] and manual Google and Wikipedia [9] search. For the birthdates queried from DBpedia, we manually verified the results to eliminate common errors such as the wrong birthdate due to the existence of another person of the same name. In a small number of cases (1%), only the birth year is available; for these individuals, we compute their age from January 1 of their birth year.

Our method does not estimate age directly from facial appearance, but instead derives age from a person's estimated identity. This avoids bias from visual factors such as makeup and is more precise than end-to-end methods that estimate age directly. It does, however, assume that the video was aired the same day that it was recorded for age to be accurate and does not account for old clips or still images. An obvious limitation is that individuals must first be identified to be included in the results.

Politicians	Notes
Donald Trump	Low recall from AWS
Hillary Clinton	Used for consistency to Trump
Barrack Obama	Used for consistency to Trump
Bernie Sanders	Used for consistency to Trump
Mitt Romney	Used for consistency to Trump
Dick Durbin	Not identified by AWS
News presenters	
Ana Cabrera	Not identified by AWS
Brian Shactman	Not identified by AWS
Bryan Illenas	Not identified by AWS
Dave Briggs	Not identified by AWS
David Gura	Not identified by AWS
Dorothy Rabinowitz	Not identified by AWS
Doug McKelway	Not identified by AWS
Ed Lavandera	Not identified by AWS
Griff Jenkins	Not identified by AWS
Jason Riley	Not identified by AWS
Jillian Mele	Not identified by AWS
Jim Pinkerton	Not identified by AWS
JJ Ramberg	Not identified by AWS
Lauren Ashburn	Not identified by AWS
Leland Vittert	Not identified by AWS
Louis Burgdorf	Not identified by AWS
Maria Molina	Not identified by AWS
Natalie Allen	Not identified by AWS
Nicole Wallace	Not identified by AWS
Pete Hegseth	Not identified by AWS
Richard Lui	Not identified by AWS
Rick Folbaum	Not identified by AWS
Rick Reichmuth	Not identified by AWS
Rob Schmitt	Not identified by AWS
Touré Neblett	Not identified by AWS
Trace Gallagher	Not identified by AWS
Yasmin Vossoughian	Not identified by AWS
Miscellaneous	
George Zimmerman	Used for consistency to Martin
Trayvon Martin	Not identified by AWS

Table 2: Individuals for whom we use our own labels. We use our own labels when no labels from AWS Rekognition Celebrity Recognition [1] are available; the AWS labels are known to have low precision or recall; or to be consistent on major comparisons between individuals labeled with our models and with AWS.

```
# Commercial Query
caption_words = rekall.ingest(captions, 1D)
histograms = rekall.ingest(database.table("hists"), 1D)
entire_video = rekall.ingest(database.table("video"), 3D)

# Find segments with >> delimiters
captions_with_arrows = caption_words
  .filter(word: '>>' in word)

# Find segments of black frames (where all of the pixels
# are black)
black_frame_segs = histograms
  .filter(i: i.histogram.avg() < 0.01)
  .coalesce(predicate = time_gap < 0.1s, merge = time_span)
  .filter(i: i["t2"] - i["t1"] > 0.5s)

# All segments between black frame segments in the
# video are candidates to be considered.
candidate_segs = entire_video.minus(black_frame_seqs)

# Candidate segments that contain >> delimiters are
# rejected
non_commercial_segs = candidate_segs
  .filter_against(
    captions_with_arrows,
    predicate = time_overlaps)

# Keep segments that were not rejected
commercial_segs = entire_video
  .minus(non_commercial_segs.union(black_frame_segs))

# Coalesce any overlapping intervals and filter intervals
# that are too short to be commercials
commercials = commercial_segs
  .coalesce(predicate = time_overlaps, merge = time_span)
  .filter(i: i["t2"] - i["t1"] > 10s)

# Find segments that have lowercase captions
lower_case_word_segs = caption_words
  .filter(word: word.is_lowercase())
  .coalesce(predicate = time_gap < 5s, merge = time_span)

# Find segments that have no captions
no_captions_segs = entire_video
  .minus(caption_words)
  .filter(i: 30 < i["t2"] - i["t1"] < 270)

# Compute the final commercial segments, coalesce
# nearby segments, and reject segments that are too long
commercials = commercials
  .union(lower_case_word_segs)
  .union(no_captions_segs)
  .coalesce(predicate = time_gap < 45s, merge = time_span)
  .filter(comm: comm["t2"] - comm["t1"] < 300s)
```

Figure 17: The Rekall [4] query for detecting commercials in each news video.

Name	Samples	Est. precision	Samples	Est. recall
U.S. political figures and candidates				
Amy Klobuchar	100	1.00	69	0.87
Barack Obama †	100	1.00	85	0.86
Ben Carson	100	0.99	132	0.85
Bernie Sanders †	100	0.99	42	0.83
Beto O'Rourke	100	1.00	50	0.58
Bill Clinton	100	0.89	59	0.90
Bill De Blasio	100	1.00	55	0.89
Bobby Jindal	100	0.99	133	1.00
Carly Fiorina	100	0.92	99	0.74
Chris Christie	100	0.98	118	0.87
Dick Durbin †	100	0.96	50	0.80
Donald Trump †	100	0.91	65	0.83
Elizabeth Warren	100	0.97	42	0.81
Gary Johnson	100	0.99	124	0.84
George W. Bush	100	0.72	71	0.80
Harry Reid	100	0.97	137	0.83
Herman Cain	100	1.00	100	0.90
Hillary Clinton †	100	0.89	136	0.84
Jeb Bush	100	0.96	79	0.92
Jim Gilmore	100	0.98	157	0.94
Jim Webb	99	0.99	158	0.89
Joe Biden	100	1.00	66	0.91
John Boehner	100	1.00	84	0.95
John McCain	99	0.99	196	0.91
Jon Huntsman Jr.	100	1.00	117	0.87
Kamala Harris	99	0.97	55	0.93
Kellyanne Conway	100	1.00	151	0.72
Kevin McCarthy	100	1.00	70	0.97
Lincoln Chafee	100	0.88	103	0.87
Lindsey Graham	100	1.00	107	0.88
Marco Rubio	100	1.00	93	0.85
Martin O'Malley	100	0.92	129	0.86
Michele Bachmann	100	0.91	104	0.92
Michelle Obama	100	1.00	107	0.76
Mike Huckabee	100	1.00	299	0.96
Mitch McConnell	99	1.00	81	0.83
Mitt Romney †	100	0.98	107	0.72
Nancy Pelosi	100	1.00	37	0.87
Newt Gingrich	100	0.98	226	0.94
Orrin Hatch	100	0.99	115	0.94
Paul Ryan	100	0.99	104	0.84
Pete Buttigieg	100	0.99	25	0.96
Rand Paul	100	1.00	140	0.94
Rick Santorum	100	1.00	168	0.92
Rick Perry	100	0.99	154	0.77
Ron Paul	100	1.00	185	0.96
Sarah Palin	100	1.00	126	0.85
Steve Scalise	100	0.97	109	0.94
Ted Cruz	100	1.00	102	0.85
Tim Kaine	100	0.99	185	0.92
Tulsi Gabbard	100	0.97	88	0.78
Miscellaneous				
George Zimmerman †	100	0.98	131	0.79
Trayvon Martin †	100	0.95	48	0.63

Table 3: Estimated precision is computed on ≈100 randomly sampled faces identified as each individual. Estimated recall is computed on actual instances of each individual's face found in a random sample of 2,500 faces, from 25 videos, known to contain each individual. († indicates our models.)

Ali Velshi (225.9 hours)	Alison Kosik (104.3)	Alisyn Camerota (271.1)	Amanda Davies (3.4)	Amara Walker (9.5)
Ana Cabrera (305.7)	Anderson Cooper (1782.3)	Andrew Levy (0.0)	Anthony Bourdain (110.8)	Arwa Damon (50.1)
Ashleigh Banfield (193.2)	Barbara Starr (156.5)	Becky Anderson (12.2)	Ben Wedeman (61.9)	Bianna Golodryga (16.0)
Bill Hemmer (0.2)	Bill Weir (16.0)	Brian Stelter (188.6)	Brianna Keilar (267.3)	Brooke Baldwin (898.6)
Campbell Brown (28.8)	Candy Crowley (140.7)	Carol Costello (311.4)	Chris Cuomo (678.0)	Christi Paul (84.1)
Christiane Amanpour (72.6)	Christine Romans (315.0)	Clarissa Ward (33.0)	Dana Bash (350.4)	Dave Briggs (91.7)
Deborah Feyerick (80.2)	Don Lemon (1098.8)	Drew Griffin (86.4)	Ed Lavandera (57.0)	Elizabeth Cohen (35.2)
Erica Hill (57.4)	Erin Burnett (539.6)	Errol Barnett (63.5)	Fareed Zakaria (230.3)	Frederik Pleitgen (71.4)
Fredricka Whitfield (477.8)	Gary Tuchman (37.4)	Gloria Borger (255.6)	Hala Gorani (28.6)	Howard Kurtz (39.0)
Jake Tapper (376.3)	Jamie Gangel (17.7)	Jean Casarez (35.5)	Jeff Zeleny (115.2)	Jeffrey Toobin (270.6)
Jessica Yellin (73.1)	Jim Acosta (220.9)	Jim Sciutto (282.3)	Joe Johns (118.3)	John Berman (584.2)
John King (377.0)	John Roberts (46.6)	John Vause (62.2)	John Walsh (20.6)	Kate Bolduan (322.6)
Kathleen Parker (21.7)	Kiran Chetry (54.5)	Kristie Lu Stout (4.2)	Kyra Phillips (105.1)	Kyung Lah (47.9)
Larry King (78.9)	Lisa Ling (25.2)	Lou Dobbs (0.3)	Lynda Kinkade (5.4)	Lynn Smith (0.2)
Martin Savidge (91.7)	Max Foster (34.4)	Michael Smerconish (177.6)	Michelle Kosinski (49.2)	Miguel Marquez (0.2)
Mike Galanos (2.4)	Mike Rogers (50.0)	Mike Rowe (4.8)	Morgan Spurlock (13.3)	Natalie Allen (75.5)
Nic Robertson (135.5)	Nick Paton Walsh (65.3)	Pamela Brown (110.2)	Paula Newton (17.3)	Piers Morgan (404.2)
Poppy Harlow (209.5)	Rachel Nichols (31.6)	Randi Kaye (148.0)	Richard Quest (90.5)	Richard Roth (7.4)
Robin Meade (2.0)	Rosemary Church (81.6)	S. E. Cupp (45.7)	Sanjay Gupta (200.1)	Sara Sidner (21.0)
Soledad O'Brien (91.6)	Stephanie Cutter (14.8)	Susan Hendricks (19.3)	Suzanne Malveaux (130.8)	T. J. Holmes (114.4)
Tom Foreman (44.0)	Van Jones (156.2)	Victor Blackwell (113.8)	W. Kamau Bell (43.9)	Wolf Blitzer (800.1)
Zain Asher (23.4)	Zain Verjee (24.2)			

Table 4: CNN. List of news presenters and their screen time in hours.

Abby Huntsman (51.3)	Ainsley Earhardt (211.9)	Alan Colmes (65.3)	Alisyn Camerota (141.3)	Andrea Tantaros (177.5)
Andrew Levy (160.1)	Andrew Napolitano (122.6)	Angela McGlowan (23.7)	Anna Kooiman (78.6)	Ari Fleischer (31.9)
Arthel Neville (108.9)	Bill Hemmer (383.0)	Bill O'Reilly (1093.8)	Bob Beckel (268.1)	Brenda Buttner (34.8)
Bret Baier (536.7)	Brian Kilmeade (638.4)	Brit Hume (171.7)	Bryan Llenas (33.6)	Byron York (77.8)
Cal Thomas (13.9)	Carol Alt (8.9)	Casey Stegall (26.2)	Charles Krauthammer (283.2)	Charles Payne (98.1)
Charlie Gasparino (45.9)	Cheryl Casone (33.1)	Chris Wallace (374.3)	Clayton Morris (217.4)	Dagen McDowell (44.4)
Dana Perino (437.2)	Daniel Henninger (53.1)	Dave Briggs (70.1)	David Asman (50.1)	David Hunt (1.0)
Dorothy Rabinowitz (7.2)	Doug McKelway (68.9)	Ed Henry (313.4)	Ed Rollins (32.1)	Elisabeth Hasselbeck (85.4)
Elizabeth Prann (25.0)	Ellis Henican (6.2)	Eric Bolling (394.9)	Eric Shawn (128.7)	Fred Barnes (10.8)
Geraldo Rivera (232.3)	Gerri Willis (27.8)	Glenn Beck (288.1)	Greg Gutfeld (782.5)	Greta van Susteren (487.5)
Gretchen Carlson (268.1)	Griff Jenkins (31.6)	Guy Benson (52.6)	Harris Faulkner (291.7)	Heather Childers (201.4)
Howard Kurtz (227.0)	James Taranto (4.6)	Jane Hall (0.1)	Janice Dean (41.6)	Jason Riley (25.3)
Jeanine Pirro (514.4)	Jedediah Bila (71.3)	Jehmu Greene (21.3)	Jennifer Griffin (57.9)	Jesse Watters (290.7)
Jillian Mele (118.9)	Jim Pinkerton (24.6)	John Fund (20.1)	John Roberts (65.5)	John Stossel (119.8)
Jon Scott (300.3)	Juan Williams (367.0)	Judith Miller (51.3)	Julie Banderas (98.2)	Karl Rove (252.4)
Katherine Timpf (60.2)	Katie Pavlich (83.4)	Kelly Wright (71.9)	Kevin Corke (40.1)	Kimberley Strassel (56.0)
Kimberly Guilfoyle (258.7)	Kristen Soltis Anderson (10.1)	Laura Ingle (31.0)	Laura Ingraham (498.0)	Lauren Ashburn (9.5)
Lauren Green (8.6)	Leland Vittert (136.6)	Leslie Marshall (73.3)	Manny Alvarez (12.2)	Mara Liasson (25.5)
Maria Bartiromo (81.0)	Maria Molina (67.1)	Mark Fuhrman (29.1)	Mark Levin (55.4)	Martha Maccallum (562.5)
Megyn Kelly (790.9)	Melissa Francis (84.6)	Michael Baden (19.5)	Mike Emanuel (99.7)	Molly Henneberg (28.4)
Molly Line (30.7)	Monica Crowley (89.3)	Neil Cavuto (737.6)	Paul Gigot (98.7)	Pete Hegseth (246.6)
Peter Doocy (87.9)	Phil Keating (35.9)	Rachel Campos-Duffy (11.2)	Raymond Arroyo (21.9)	Rich Lowry (37.0)
Rick Folbaum (42.8)	Rick Reichmuth (80.7)	Rob Schmitt (51.0)	Robert Jeffress (19.4)	Sandra Smith (71.7)
Sean Hannity (1071.8)	Shannon Bream (416.0)	Shepard Smith (360.2)	Steve Doocy (450.4)	Steve Hilton (81.2)
Stuart Varney (126.2)	Tammy Bruce (60.5)	Tom Shillue (145.5)	Tomi Lahren (15.8)	Trace Gallagher (131.7)
Trish Regan (44.5)	Tucker Carlson (865.3)	Uma Pemmaraju (48.3)	Walid Phares (28.2)	William Bennett (16.9)

Table 5: FOX. List of news presenters and their screen time in hours.

Abby Huntsman (29.0)	Al Sharpton (286.7)	Alec Baldwin (2.5)	Alex Wagner (174.8)	Alex Witt (261.3)
Ali Velshi (242.4)	Andrea Canning (4.9)	Andrea Mitchell (392.2)	Andrew Ross Sorkin (11.0)	Angie Goff (1.3)
Anne Thompson (10.1)	Ari Melber (395.0)	Ayman Mohyeldin (150.4)	Betty Nguyen (29.5)	Bill Neely (20.6)
Brian Shactman (213.3)	Brian Sullivan (13.3)	Brian Williams (282.5)	Carl Quintanilla (0.8)	Chris Hayes (839.5)
Chris Jansing (254.8)	Chris Matthews (1103.8)	Chuck Todd (550.3)	Contessa Brewer (49.8)	Craig Melvin (173.4)
David Faber (1.2)	David Gura (54.9)	Donny Deutsch (53.3)	Dylan Ratigan (109.7)	Ed Schultz (493.0)
Frances Rivera (44.2)	Greta van Susteren (21.7)	Hallie Jackson (105.0)	Jim Cramer (8.0)	Jj Ramberg (30.8)
Joe Scarborough (940.4)	John Heilemann (147.1)	Jose Diaz-Balart (88.3)	Josh Mankiewicz (13.1)	Joy-Ann Reid (337.1)
Kasie Hunt (112.6)	Kate Snow (51.6)	Katy Tur (187.1)	Kayla Tausche (2.1)	Keith Olbermann (109.7)
Kelly Evans (0.7)	Kelly O'Donnell (57.2)	Kerry Sanders (25.2)	Kristen Welker (212.2)	Krystal Ball (91.0)
Lawrence O'Donnell (688.0)	Lester Holt (13.1)	Louis Burgdorf (29.8)	Lynn Smith (28.0)	Mara Schiavocampo (18.5)
Mark Halperin (158.9)	Martin Bashir (114.7)	Matt Lauer (8.4)	Melissa Harris-Perry (197.9)	Meredith Vieira (1.2)
Miguel Almaguer (9.3)	Mika Brzezinski (696.7)	Mike Viqueira (46.9)	Natalie Morales (4.7)	Nicole Wallace (175.9)
Pete Williams (105.4)	Peter Alexander (97.5)	Rachel Maddow (1201.7)	Rehema Ellis (7.2)	Richard Engel (114.2)
Richard Lui (146.4)	Rick Santelli (1.3)	Ron Mott (16.7)	Ronan Farrow (31.4)	Savannah Guthrie (43.9)
Seema Mody (1.5)	Stephanie Gosk (14.0)	Stephanie Ruhle (111.5)	Steve Kornacki (358.6)	Steve Liesman (4.7)
Sue Herera (1.8)	Tamron Hall (200.5)	Thomas Roberts (198.8)	Tom Brokaw (29.2)	Tom Costello (24.5)
Touré Neblett (65.4)	Willie Geist (319.3)	Yasmin Vossoughian (66.6)		

Table 6: MSNBC. List of news presenters and their screen time in hours.

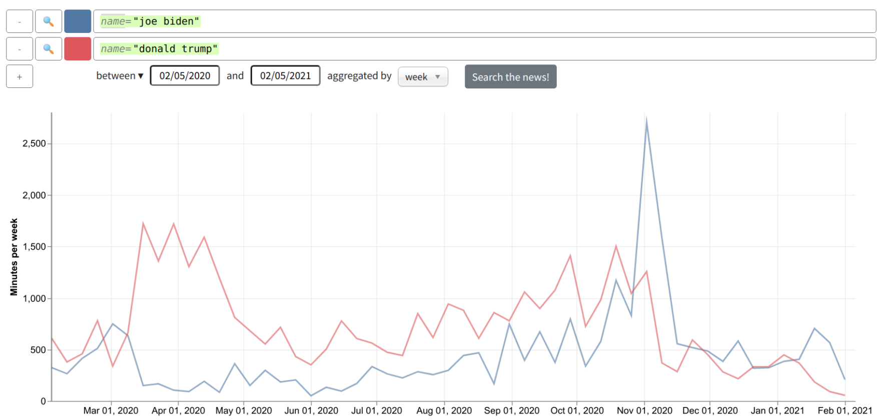

(a) Screen time

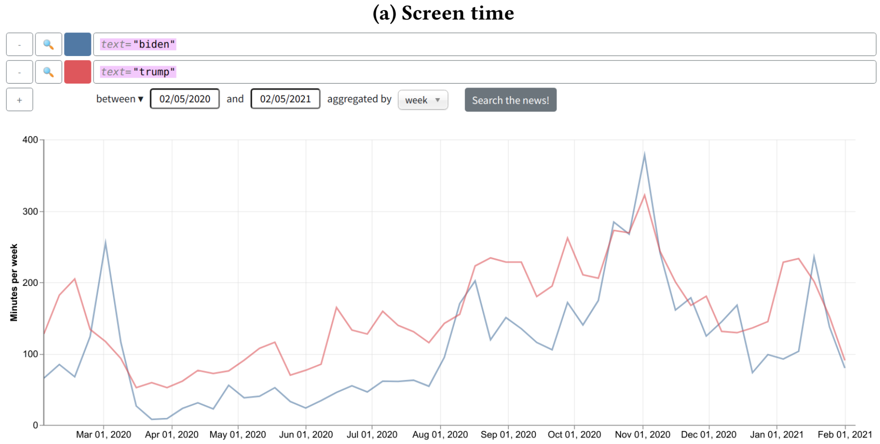

(b) Last name mentions (each word is dilated to a 1s time interval)

Figure 18: Comparison of screen time and name mentions between Donald Trump and Joe Biden from February 2020 to 2021, aggregated by week. Notice that there are differences in the patterns of both charts and that the screen time gaps are larger, in favor of Trump before the US presidential election on November 3, 2020. Both these charts can be viewed and edited in the public tool at https://tvnews.stanford.edu/.

2.5 Methodology for additional results

This section covers details that are not already contained in the methods section of the main paper or in section 2 of this document.

2.5.1 Counting mentions of foreign countries. To identify the set of most frequently mentioned countries, we constructed a list of country and territory names from [3], which includes all countries and territories with ISO-3166-1 country codes. We manually augment the list with country name aliases; for example, the `Holy See` and `Vatican` are aliases of one another and either term is counted as Vatican City. A few countries such as Mexico and Georgia are substrings of US state names, leading to over-counting in the results. To address this issue, we exclude occurrences of `Mexico` that are preceded by `New` and we omit `Georgia` entirely. (Mentions of Georgia in US cable news overwhelmingly refer to the US state and not the country.)

2.5.2 Counting mentions of terrorism, mass shootings, and major plane crashes. To measure how long the media continues to cover events after they take place, we counted the number of times that words related to terrorism, mass shootings, and plane crashes appear following an event. Table 7 and Table 8 show the events that were included in the analysis. For terrorism, we count instances of `terror(ism,ist)`, `attack`, `shooting`, `stabbing`, and `bombing`, which refer to the attack itself; for mass shootings, the list is `shoot(ing,er)`, which refers to the shooting or the mass shooter (searching more restrictively for instances of `mass shoot(er,ing)` yields a similar result, but sometimes mass is omitted in the news coverage); and for plane crashes the list is `(air)plane` or `airliner` followed by `crash` or `missing`. Because the keywords to measure news coverage are different between each category of event, the raw counts are not directly comparable across categories.

2.5.3 Counting mentions of illegal and undocumented immigration. We count the number of times that n-grams related to "illegal" and "undocumented" immigration appear in the captions to measure the prevalence of both terms in discussion around immigration. The n-grams used to measure uses of "illegal" are `illegal immigrant(s)`, `illegal immigration`, `illegals`, and `illegal alien(s)`. For "undocumented", the n-grams are `undocumented immigrant(s)`, `undocumented immigration`, and `undocumented alien(s)`.

2.5.4 Counting mentions of the president honorific in reference to Donald Trump and Barack Obama. We measure the number of times the "president" honorific is used when addressing each president. This requires classifying occurrences of the word `Trump` (and also `Obama`) in captions as having the "president" honorific, not having the honorific (e.g., `Donald Trump` or just `Trump`), or not referring to his person (e.g., Trump University).

For Donald Trump, we only count exact matches of `President Trump` or `President Donald Trump` as uses of "president". To count occurrences of without the honorific, we exclude occurrences preceded by `president` and instances followed by `administration`, `campaign`, `university`, and `care`, which are used in compound nouns with `Trump`. We also exclude occurrences preceded by `the` (e.g., to filter out other compound nouns of the form `the Trump ...`); note that this also removes `the Trump presidency`, which is

(a) A real interview.

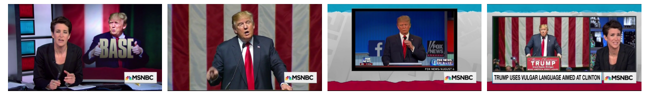

(b) Not an interview.

Figure 19: Example frames from a real and incorrectly detected interview. Note that both follow a pattern of a host and guest being on-screen, together and alone. The incorrectly detected interview contains videos and graphics of Donald Trump in lieu of his live person. As the presidents and leading candidates, Trump, Clinton, and Obama are discussed at length by hosts in visual contexts that appear similar to interviews.

not referring to his person, but his presidency. Finally, we exclude Donald Trump's immediate family: `Melania`, `Ivanka`, `Eric`, `Barron`, and `[Donald Trump] Jr`. These exclusions of nouns related to Trump (but not directed at his person) were selected by visual examination of the top 100 bigrams containing `Trump`.

The procedure for counting references to Barack Obama is identical, except that the excluded family members are `Michelle`, `Malia`, and `Sasha`.

2.5.5 Measuring visual association between news presenters and the president honorific. We extended the president honorific analysis (section 2.5.4) to when various news presenters are on-screen. The n-grams that are counted remain the same as in section 2.5.4. We start with the list of news presenters described in main text, but we only show news presenters with at least 100 total references to President Trump and 100 total references to President Obama to ensure that there are sufficient data for a comparison. This is to account for news presenters who retired before Trump became president or started after Obama stepped down.

2.5.6 Measuring visual association between Clinton and the word email. The Hillary Clinton email scandal and subsequent FBI investigation was a highly polarizing issue in the 2016 presidential election. To measure the degree to which Clinton's face became visually associated with the issue, represented by the word "email", we counted the number of times when "email(s)" was said, and the number of times it was said while Clinton is on-screen.

We count occurrences of `e mail(s)`, `email(s)`, and `electronic mail` as instances of email being said in the captions. There are 122K utterances of email in the captions between 2015 and 2017, while Hillary Clinton has 738 hours of screen time in the same time period. Clinton's face is on-screen during 14,019 of those utterances.

2.5.7 Detecting interviews. Our algorithm for finding interviews in TV news searches for interviews between a news presenter (the host) and a named guest X. We search for segments where the guest and the host appear together, surrounded by the guest appearing alone or the host appearing alone. Combining these segments captures an alternating pattern where a host appears, guest appears, ... that is indicative of an interview. The pseudo-code for this algorithm is shown in Rekall [4] in Figure 20.

We applied this interview detection algorithm to 44 people across our whole data set. These individuals are listed in Table 10. The algorithm is not perfect and, we exclude Barack Obama, Donald Trump, and Hillary Clinton due to those individuals appearing too often in video clips and still images. Their appearances, along with hosts, are often misclassified as interviews. For example, Donald Trump may be shown in a still image or giving a speech while the news content cuts back and forth to a host providing commentary (Figure 19). Events such as town-hall gatherings are also sometimes confused as interviews. As the leading candidates and presidents, Trump, Clinton, and Obama appear the most often in these challenging contexts.

We validated our interview detection algorithm by annotating 100 cable news videos, which contain interviews for three interviewees: Bernie Sanders, Kellyanne Conway, and John McCain. Table 9 shows the estimated precision and recall numbers for the three interviewees, as well as the total amount of interview screen time in the ground truth set for each interviewee.

Date	Event	Victims
Terrorist attacks (U.S.)		
4/15/2013	Boston Marathon bombing	286
12/2/2015	San Bernardino shooting	30
6/12/2016	Pulse Nightclub shooting	103
9/17/2016	2016 New York and New Jersey bombings	35
8/12/2017	Charlottesville car attack	29
10/31/2017	2017 New York City truck attack	20
8/3/2019	El Paso shooting	46
Terrorist attacks (Europe)		
4/11/2011	Minsk Metro bombing	219
7/22/2011	Norway attacks	396
7/17/2014	Malaysia Airlines flight 17 shootdown	298
1/7/2015	January 2015 Île-de-France attacks	42
11/13/2015	November 2015 Paris attacks	551
3/22/2016	Brussels bombings	375
7/14/2016	Nice truck attack	521
5/22/2017	Berlin Christmas market attack	68
6/3/2017	Manchester Arena bombing	273
8/17/2017	2017 London Bridge attack	59
2/19/2020	2017 Barcelona attacks	176
Mass shootings		
1/8/2011	Tucson, Arizona	21
7/20/2012	Aurora, Colorado	82
12/14/2012	Newtown, Connecticut	30
9/16/2013	Washington D.C.	21
5/23/2014	Isla Vista, California	20
5/17/2015	Waco, Texas	27
12/2/2015	San Bernardino, California	38
6/12/2016	Orlando, Florida	103
7/1/2017	Little Rock, Arkansas	28
10/1/2017	Las Vegas, Nevada	481
11/5/2017	Sutherland Springs, Texas	47
2/14/2018	Parkland, Florida	34
6/17/2018	Trenton, New Jersey	23
5/18/2018	Santa Fe, Texas	24
11/7/2018	Thousand Oaks, California	25
8/3/2019	El Paso, Texas	46
8/4/2019	Dayton, Ohio	37
8/31/2019	MidlandOdessa, Texas	33

Table 7: Major events included in the list of terrorist attacks and mass shootings.

Date	Plane crashes	Deaths
1/25/2010	Ethiopian Airlines Flight 409	90
5/12/2010	Afriqiyah Airways Flight 771	103
5/22/2010	Air India Express Flight 812	158
7/28/2010	Airblue Flight 202	152
11/4/2010	Aero Caribbean Flight 883	68
1/9/2011	Iran Air Flight 277	77
7/8/2011	Hewa Bora Airways Flight 952	74
4/20/2012	Bhoja Air Flight 213	127
6/3/2012	Dana Air Flight 992	159
11/17/2013	Tatarstan Airlines Flight 363	50
3/8/2014	Malaysia Airlines Flight 370	239
7/17/2014	Malaysia Airlines Flight 17	298
7/24/2014	Air Algérie Flight 5017	116
12/28/2014	Indonesia AirAsia Flight 8501	162
3/24/2015	Germanwings Flight 9525	150
8/16/2015	Trigana Air Flight 267	54
3/19/2016	Flydubai Flight 981	62
5/19/2016	EgyptAir Flight 804	66
11/28/2016	LaMia Airlines Flight 2933	71
2/11/2018	Saratov Airlines Flight 703	71
2/18/2018	Iran Aseman Airlines Flight 3704	66
3/12/2018	US-Bangla Airlines Flight 211	51
5/18/2018	Cubana de Aviación Flight 972	112
10/29/2018	Lion Air Flight 610	189
3/10/2019	Ethiopian Airlines Flight 302	157

Table 8: Plane crashes included in the analysis. This list includes all of the commercial airline crashes from 2010 to 2019 that involved at least 50 fatalities.

Interviewee	Hours	Precision	Recall
Bernie Sanders	3.5	91.7%	97.5%
Kellyanne Conway	2.2	91.8%	89.1%
John McCain	0.9	86.0%	99.5%

Table 9: Precision and recall for the interview query across 100 hand-annotated videos and the total amount of manually annotated interview screen time in the ground truth set for each interviewee.

```
# Interviews between a host and a named guest
faces = rekall.ingest(database.table("faces"), 3D)

# Select all faces (3s segments) identified as the
# guest and the faces of all hosts
guest_faces = faces.filter(
  face: face.name = guest_name)
host_faces = faces.filter(
  face: face.is_host)

# Coalesce adjacent segments since individuals are
# often on-screen for longer than the 3s sample rate
guest_segs = guest_faces.coalesce(
  predicate = time_gap < 30s,
  merge = time_span)
host_segs = host_faces.coalesce(
  predicate = time_gap < 30s,
  merge = time_span)

# Find segments when a host and the guest are on
# screen at the same time
guest_and_host_segs = guest_segs.join(
  host_segs,
  predicate = time_overlaps,
  merge = time_intersection)

# Find segments when the guest is on-screen without
# the host
guest_alone_segs = guest_segs.minus(
  guest_and_host_segs)

# Merge segments when the guest is on-screen alone
# with the segments when both the host and guest are
# on-screen and consider these to be segments of
# an interview
interview_segs = guest_and_host_segs.join(
  guest_alone_segs,
  predicate = before or after,
  merge = time_span)

# Merge the detected interview segments and return
# the ones that exceed a minimum interview duration
interviews = interview_segs
  .coalesce()
  .filter(interval:
     interval["t2"] - interval["t1"] >= 240s)
```

Figure 20: Rekall [4] query to retrieve interviews between a host and a named guest (e.g., Bernie Sanders).

Interviewee	Hours
John McCain	124.4
Bernie Sanders	107.8
Rand Paul	98.0
Lindsey Graham	93.3
Rick Santorum	91.9
Marco Rubio	87.9
Kellyanne Conway	77.7
Sarah Palin	72.0
Paul Ryan	67.5
John Kasich	63.5
Ted Cruz	61.5
Chris Christie	61.5
Mitt Romney	58.9
Ben Carson	49.1
Elizabeth Warren	35.4
Mitch McConnell	34.7
Carly Fiorina	33.7
Cory Booker	31.3
Kevin McCarthy	31.0
Tim Kaine	29.4
Chuck Schumer	28.9
Nancy Pelosi	28.9
Amy Klobuchar	28.5
Jeb Bush	26.8
Dick Durbin	25.8
John Boehner	24.6
Joe Biden	24.2
Bill Clinton	22.0
Bill De Blasio	19.6
George W. Bush	19.2
Steve Scalise	18.2
Bobby Jindal	17.3
Orrin Hatch	15.1
Martin O'Malley	14.6
Kamala Harris	12.9
John Cornyn	10.3
Tulsi Gabbard	9.6
Harry Reid	7.6
Pete Buttigieg	7.5
Jim Webb	6.1
Beto O'Rourke	5.3
Lincoln Chafee	4.4
Michelle Obama	2.3
Jim Gilmore	1.6
Newt Gingrich	185.3
Mike Huckabee	95.8

Table 10: Detected interview time for prominent U.S. political figures. Newt Gingrich and Mike Huckabee are listed separately because they are both hosts (news presenters) and politicians.

REFERENCES

[1] Amazon. 2020. Amazon Rekognition. https://aws.amazon.com/rekognition.

[2] Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. DBpedia: A Nucleus for a Web of Open Data. In *Proceedings of the 6th International The Semantic Web and 2nd Asian Conference on Asian Semantic Web Conference* (Busan, Korea) *(ISWC'07/ASWC'07)*. Springer-Verlag, Berlin, Heidelberg, 722–735.

[3] Duncalfe, Luke. 2020. ISO-3166 Country and Dependent Territories Lists with UN Regional Codes. https://github.com/lukes/ISO-3166-Countries-with-Regional-Codes.

[4] Daniel Y. Fu, Will Crichton, James Hong, Xinwei Yao, Haotian Zhang, Anh Truong, Avanika Narayan, Maneesh Agrawala, Christopher Ré, and Kayvon Fatahalian. 2019. Rekall: Specifying Video Events using Compositions of Spatiotemporal Labels. *arXiv preprint arXiv:1910.02993* (2019).

[5] Global Media Monitoring Project. 2015. Who Makes the News? https://whomakesthenews.org/gmmp-2015-reports/.

[6] Google LLC. 2020. Google Search. https://www.google.com

[7] Stephen Hiltner. 2017. Illegal, Undocumented, Unauthorized: The Terms of Immigration Reporting. *The New York Times* (Mar 2017).

[8] Florian Schroff, Dmitry Kalenichenko, and James Philbin. 2015. Facenet: A unified embedding for face recognition and clustering. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 815–823.

[9] Wikimedia Foundation. 2020. Wikipedia - The Free Encyclopedia. https://www.wikipedia.org