

Quotatives Indicate Decline in Objectivity in U.S. Political News

Tiancheng Hu¹, Manoel Horta Ribeiro², Robert West², Andreas Spitz³

¹ ETH Zürich, Zürich, Switzerland

² EPFL, Lausanne, Switzerland

³ University of Konstanz, Konstanz, Germany

tiancheng.hu@alumni.ethz.ch, manoel.hortaribeiro@epfl.ch, robert.west@epfl.ch, andreas.spitz@uni-konstanz.de

Abstract

According to journalistic standards, direct quotes should be attributed to sources with objective quotatives such as “said” and “told,” since nonobjective quotatives, e.g., “argued” and “insisted,” would influence the readers’ perception of the quote and the quoted person. In this paper, we analyze the adherence to this journalistic norm to study trends in objectivity in political news across U.S. outlets of different ideological leanings. We ask: 1) How has the usage of nonobjective quotatives evolved? 2) How do news outlets use nonobjective quotatives when covering politicians of different parties? To answer these questions, we developed a dependency-parsing-based method to extract quotatives and applied it to Quotebank, a web-scale corpus of attributed quotes, obtaining nearly 7 million quotes, each enriched with the quoted speaker’s political party and the ideological leaning of the outlet that published the quote. We find that, while partisan outlets are the ones that most often use nonobjective quotatives, between 2013 and 2020, the outlets that increased their usage of nonobjective quotatives the most were “moderate” centrist news outlets (around 0.6 percentage points, or 20% in relative percentage over seven years). Further, we find that outlets use nonobjective quotatives more often when quoting politicians of the opposing ideology (e.g., left-leaning outlets quoting Republicans) and that this “quotative bias” is rising at a swift pace, increasing up to 0.5 percentage points, or 25% in relative percentage, per year. These findings suggest an overall decline in journalistic objectivity in U.S. political news.

1 Introduction

Journalistic objectivity is the notion that news should contain accurate information and not convey the personal opinions or emotions of the writer (Ryan 2001; Calcutt and Hammond 2011). Historically, objectivity emerged alongside the conception of journalism as a profession (Schudson 1981) and has shaped many of the practices and norms in modern journalism (Boudana 2011). In the context of U.S. politics, with its two major political parties, this can also be interpreted as “equal treatment” of both parties (D’Alessio and Allen 2006). Bias in the news could affect public opinion (O’Connell 1999; Kahn and Kenney 2002) and lead to changes in voting behavior (DellaVigna and Kaplan 2007; Bernhardt, Krasa, and Polborn 2008).

Copyright © 2023, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

“Absolute” objectivity has been criticized as unattainable, as structural biases would creep into news production (Beharrell et al. 2009), or even as harmful, as the excessive balance of viewpoints could create an illusion of credibility for dubious or unsupported positions (Dixon and Clarke 2013). However, amidst the fragmented media ecosystem that emerged from the digitization of news outlets and the algorithmic serving of content (Thurman, Lewis, and Kunert 2019), journalism scholars have argued that objectivity has become ever more important to consumers of journalism (Boudana 2011; McNair 2017). This opinion is also held by the public worldwide, who, as of 2018, overwhelmingly agree that news media should be unbiased in its coverage of politics (Mitchell et al. 2018).

One of the concrete ways in which journalists have sought to report the news objectively is through the usage of direct quotes (Brooks et al. 2007; Stenvall 2008). Since journalists almost never directly observe the events they report, using quotes lends them more reliability and factuality than their own words (Van Dijk 1988). Furthermore, direct quotes would let people “speak for themselves,” following one of the golden rules of journalism (Ingram and Henshall 2012). However, even when using a direct quote, journalistic objectivity can still be compromised by the use of certain quotatives that relay the emotions of reporters to readers (Mencher and Shilton 1997) or the attempt to describe the speaker’s state of mind (Gidengil and Everitt 2003). For example, in the direct quote

“New York is not afraid of terrorists,” **boasted** Rep. Jerrold Nadler, a Democrat representing Manhattan,

the quotative (**boasted**) carries an illocutionary force from the reporter that influences how the reader perceives the quote itself, possibly distorting its original meaning (Caldas-Coulthard 1992). Objective quotatives, like “say” or “tell,” on the other hand, are considered neutral, as they imply little about the presumed intent or the fashion in which the quote was uttered (Sonoda 1997; Bell 1991).

Recent years were marked by increased political polarization (Abramowitz and Saunders 2008), mistrust in media (Brenan 2022), increased negative tone by politicians (Külz et al. 2022), and the perception that the public debate around politics has become less respectful and less fact-based (Doherty et al. 2019). Solutions to these issues are complex, but analyzing the bias and the departure from journalistic objectivity in political news coverage can help in-

form new practices and interventions that seek to improve the political news media ecosystem. Quotatives, in this context, are a powerful instrument to measure bias. Studying how journalists deviate from the standard usage of quotatives – e.g., “say” and “tell” (The Associated Press 2020) – allows researchers to quantitatively assess adherence to journalistic objectivity (Lee 2017) and reveal biases in journalistic coverage of politics (Gidengil and Everitt 2003).

Present work. This paper analyzes quotatives to study objectivity and media bias in political journalism. We ask:

- **RQ1** How has the usage of nonobjective quotatives evolved in U.S. political journalism?
- **RQ2** How do news outlets use nonobjective quotatives when covering politicians of different parties?

To answer these research questions, we developed a methodology to extract quotatives from a large-scale news corpus. We then performed a comprehensive study on how (and which) quotatives are used in direct quotes from U.S. politicians between 2013 and 2020, leveraging a large dataset of quotes from English-language media linked with relevant speaker metadata (Vaucher et al. 2021) and enriched with the political leanings of different U.S. outlets. By counting the usage of nonobjective quotatives like “shout” or “assert,” we analyze how U.S. news outlets of different political inclinations adhere to basic journalistic objectivity principles and how this adherence has evolved. Further, analyzing how outlets of different political inclinations use quotatives to talk about politicians of different parties, we study the evolution of quotative bias in news outlets.

Summary of findings. We find that the usage of nonobjective quotatives varies across different outlet categories. Overall, the more ideologically extreme an outlet is, the more nonobjective quotatives it uses. However, we also find that centrist outlets are experiencing a significant increase in the usage of nonobjective quotatives over the last years (about 0.6 percentage points, or 20% in relative percentage), suggesting that they may be “catching up” to the more biased outlets, which are not experiencing such significant increases (**RQ1**). We also find evidence of “quotative bias,” i.e., outlets tend to use nonobjective quotatives, especially when referring to politicians of opposing ideology. For instance, left and right-leaning outlets use nonobjective quotatives up to 2% more often when referring to Republican and Democrat politicians, respectively (**RQ2**). Last, we find that this quotative bias is increasing at a swift pace, increasing as much as 0.5 percentage points per year in absolute percentage, or 25% in relative percentage, for left-leaning outlets, suggesting a rapid increase in polarization (**RQ1** and **RQ2**).

Implications. Our findings indicate a decline in journalistic objectivity in U.S. political news, particularly from centrist outlets. This suggests that centrist outlets may play a role in the increasingly less respectful and fact-based debate around politics (Doherty et al. 2019). Further, we also find evidence of an increasing quotative bias, which could further erode trust in the media (Brenan 2022).

2 Background and Related Work

When a quote occurs in the news, three elements are typically involved: the source, i.e., the speaker who uttered the quote (underlined in the examples); the quoted content itself (in *italic*); and the quotative that introduces the quote (also known as a cue, reporting verb, speech verb, or attribution verb; in **bold**). Quotes can be classified as either direct, indirect, mixed, or pure (Cappelen and Lepore 1997), where only the former three types are typically of concern to journalists. We give an example of a direct quote in the introduction and of mixed and indirect quotes below.

Indirect quote: Sen. Ron Wyden of Oregon, the chairman of the Senate Finance Committee, **indicated** that *in 2019, about 100 to 125 corporations reported financial statement income greater than 1B USD.*

Mixed quote: Catsimatidis **said** he’d serve for 99 cents “*because I’m a grocer.*”

In direct and mixed quotes, a pair of quotation marks are used, and we can infer that the speaker uttered the quoted words, whereas indirect quotes may paraphrase the speaker’s words. Therefore, journalists have the most freedom in word choice in indirect quotes, as they can, to some extent, rewrite what the speaker said. In contrast, in direct quotes, journalistic norms require them to report the quoted words verbatim (Harry 2014).

In our work, we focus on direct quotes for the following reasons: Quotebank does not contain indirect quotes (Vaucher et al. 2021); automatically extracting the quotative in mixed quotes is technically challenging and, in some instances, impossible as there is no quotative, e.g., John will not help as he has “*done more for this house than all of us combined*”.

Measuring media bias. Previous work has studied media biases: how journalists’ and editors’ personal opinions, beliefs, and financial incentives shape what is considered newsworthy (McCombs and Shaw 1972) and how issues are covered (Iyengar 1994). Scholars argue that partisan media bias can harm democracy by distorting citizens’ political knowledge and increasing polarization (Bernhardt, Krasa, and Polborn 2008; Boudana 2011; McNair 2017). Thus, measuring media bias is the first step to improving our information ecosystem (Watts, Rothschild, and Mobius 2021).

Early studies in media bias required extensive manual annotation. For instance, Kobre (1953) studied how the press in Florida covered the U.S. 1952 presidential campaign by coding the number of inches of text given to each party, the position of pictures, etc., across hundreds of newspapers. However, in recent times and with the digitization of news, various methods have been developed to *automatically* measure media bias (Hamborg, Donnay, and Gipp 2019). Some of these methods are *audience based*, measuring how segregated news consumers are across outlets, e.g., Zhou, Resnick, and Mei (2011) use votes on Digg, a social news aggregator, to classify political articles. Others are *content based*, quantifying media bias by analyzing published content directly. For instance, Gentzkow and Shapiro (2010) measured bias using the frequency at which outlets reproduce partisan phrases in congressional speeches.

According to Budak, Goel, and Rao (2016), both content and audience-based approaches suffer from distinctive limitations. On the one hand, audience-based approaches do not scale beyond outlets for which detailed readership information can be obtained. On the other hand, content-based approaches struggle to generalize well across different types of news and outlets, e.g., methods that try to match politicians’ speeches to news only apply to a minority of news articles, limiting the scope of the results obtained.

Quotatives and bias. Quotatives can impact how readers perceive a news story and the involved speakers (Geis 1987; Just, Crigler, and Buhr 1999). For instance, Cole and Shaw (1974) carried out an experiment in which they changed “objective” quotatives like “said” for stronger verbs like “argued” or “insisted” and asked participants to rate stories across a variety of criteria. They found that in the modified versions, stories were perceived as more exciting and less objective, and speakers were perceived as more rash. Through quotatives, journalists can “paint reports on speech with any brush they like” (Geis 1987), which would not only reveal the beliefs and preferences of the writer (Gidengil and Everitt 2003) but also subtly influence the reader (Cole and Shaw 1974). In this context, quotatives have been used to measure political bias, sometimes referred to as “attribution bias.” This line of work dates from the 1960s when Merrill (1965) studied how Time magazine used quotatives (among other things) when referring to U.S. Presidents Kennedy, Truman, and Eisenhower. More recently, Gidengil and Everitt (2003) analyzed differences in quotative usage between male and female party leaders on Canadian television, finding that female leaders’ speech was reported with more negative and aggressive quotatives. With a similar methodology, Lee (2017) studied differences in nonobjective quotatives between offline and online newspapers, finding the former to adhere better to journalistic standards.

Quote attribution and analysis. Previous work in natural language processing has studied the problem of *quote attribution* (see Vaucher et al. (2021) for a review), an important task in understanding dialogue structure and developing better conversational agents. For each quote, the goal is to extract the speaker of the quote, either at the mention or entity level. This task is challenging as the speaker could be mentioned implicitly or require anaphora resolution. The task can be further combined with entity linking to extract unique IDs of speakers (Čuljak et al. 2022). Most prior work, however, has not dealt with the problem of quotative extraction.

Nonetheless, several datasets annotated for attributional relationship exist (Pareti 2012, 2016; Newell, Margolin, and Ruths 2018) that could be considered in this context. These datasets contain labels for the content, source, and cue for each attributional relationship. They can be viewed as an extension of The Penn Discourse TreeBank 2.0 (Prasad et al. 2008) that provides annotation of discourse relations and argument structures. While these datasets can potentially be used as resources for training a supervised model for quotative extraction, the attributional relationship they considered is much broader than quotation and thus not suitable for our study.

Existing work has also analyzed quotes from different perspectives. Niculae et al. (2015) found a systematic pattern in the outlets’ quoting behavior when covering the exact same event. Lazaridou, Krestel, and Naumann (2017) found that a machine learning classifier could reliably predict one of two news outlets based solely on the quotes they report, demonstrating media bias. Tan, Peng, and Smith (2018) showed a declining trend of bipartisan quote coverage with a bipartite graph of media outlets and the sentences they quoted. Külz et al. (2022) analyzed the quotes of U.S. politicians between 2008 and 2020 and found a decrease in negativity during Obama’s tenure and a sudden increase starting from Trump’s presidential primary campaign in 2015.

Relationship with prior work. In this paper, we set out to investigate how the usage of nonobjective quotatives evolved in U.S. political journalism (**RQ1**) and how it is modulated by media biases (**RQ2**). We do so by using dependency parsing to extract quotatives from a large dataset (see Sec. 3). Our method is related to *quote attribution*, a problem widely studied in natural language processing, with the key difference that previous methods aim to attribute quotes to speakers instead of finding the quotative used. Further, our approach is similar to previous work that derives automated media bias measurements (Budak, Goel, and Rao 2016). However, in contrast to previous work, we automate the measurement of *quotative* bias instead of relying on traditional manual annotation (Cole and Shaw 1974). Due to the scalability of our approach, we obtain results that help further understand the political news ecosystem (see Sec. 5). Namely, while previous work often attributes the decrease in journalistic objectivity to the rise of partisan media (McNair 2017), we find that centrist outlets in our dataset have systematically departed from journalistic standards.

3 Materials and Methods

3.1 Data and Data Processing

To study quotative usage across various news outlets, we use the Quotebank dataset (Vaucher et al. 2021), a web-scale corpus of quotes. Quotebank contains over 235 million unique quotes, extracted from 196 million English news articles from 377 thousand web domains between September 2008 and April 2020. We additionally obtain a list of current and former U.S. politicians with their party affiliations from Wikidata, in the same fashion as Külz et al. (2022). We filter Quotebank to consider the period containing the best-quality speaker attributions (May 2013 to 2020) and retain only quotes from politicians on this list. In cases where quotes are attributed to more than one speaker in Quotebank (which happens to 8.13% of speakers in 12.25% of the quotes), we heuristically attribute the quote to the speaker with the alphanumerically smallest Wikidata identifier. We validate speaker attribution in our filtered dataset on a manually annotated sample of 100 quotes and find that combining the speaker names provided by Quotebank with this heuristic yields 86% accuracy in identifying the correct ID.

To ensure the validity of our findings, we preprocess Quotebank as depicted in Figure 1. We 1) use heuristics to retain only direct quotes; 2) extract quotatives and remove

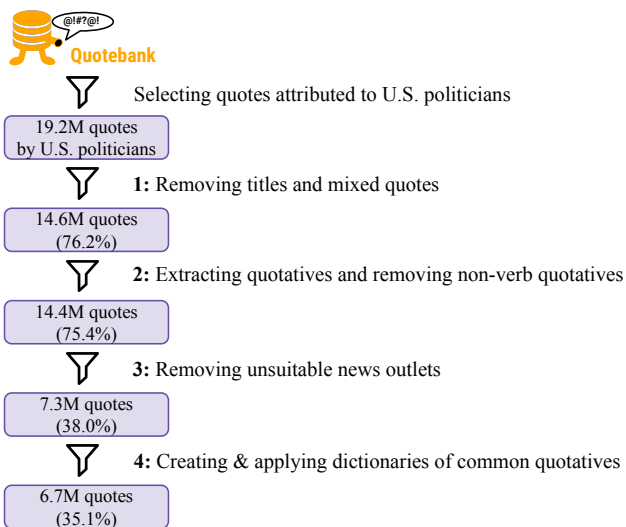


Figure 1: Data processing pipeline. We outline the key steps in our data processing pipeline and the percentage of retained data after filtering.

quotes without quotatives in the verb form; 3) filter quotes, keeping only those from U.S.-based outlets with human-verified bias ratings; and 4) create dictionaries of common quotatives, removing quotes with rare quotative verbs for which quotative extraction performs poorly. We detail each of these steps in the following paragraphs.

Step 1: removing titles and mixed quotes. To remove titles from the dataset that are erroneously recognized as quotes (e.g., movies), we apply a filter using the percentage of words in a quote whose first letter is capitalized. Afterward, to remove mixed quotes, we employ a sentence recognition filter that combines constituency parsing¹ and dependency parsing². We retain only quotes that can be parsed as a full sentence at the root level by constituency parsing and contain a subject and a predicate (root) in dependency parsing. These heuristics improve data quality (e.g., extracted quotatives in Step 2 are much more accurate) while retaining 76.2% of the dataset.

Step 2: extracting quotatives and removing non-verb quotatives. In the next step, we adopt a three-stage approach to extract the quotative from each quote using dependency parsing. First, we run dependency parsing and acquire a distribution of quotatives from the root node of each parsed quote. Second, we add a condition to ensure that in cases where one verb is identified as the root and another verb exists in a parallel node³, we choose the verb with the higher

¹Constituency parsing breaks down sentences into phrases and identifies their grammatical roles, e.g., in “I eat a big apple,” “a big apple” is a noun phrase. See Jurafsky and Martin (2022) for details.

²Dependency parsing extracts dependency relationships between words, with verbs typically being in the structural center, e.g., in “I eat a big apple,” “big” is an adjectival modifier of “apple.” See Kübler, McDonald, and Nivre (2009) for details.

³`csubj, ccomp, xcomp, advcl, acl, parataxis, conj,`

probability as the quotative (according to the distribution of verbs extracted in the first stage). Finally, we take the lemma of each extracted quotative and remove quotatives that are not in verb form. After this step, we retain around 75.4% of the original data.

Step 3: removing unsuitable outlets. We obtain a list of media bias ratings from mediabiasfactcheck.com (hereinafter MB/FC) and classify outlets into five categories based on the bias rating: left, left-center, least-biased, right-center, and right. We refer to left-center, least-biased, and right-center outlets as centrist outlets in the following. We remove quotes from outlets without a bias rating, from outlets that are not from the U.S. (also according to MB/FC data), and from outlets that have very few quotes (which may suggest data quality issues), only keeping outlets with more than 20 quotes over a period of 12 months. After this step, around 38.0% of the original data remains, all from relevant U.S. media outlets with human-verified bias ratings. Manual inspection of the removed data confirms that the removed outlets are predominantly non-news websites, small local newspapers, radio stations, and non-U.S news outlets.

Step 4: creating dictionaries of common quotatives. Inspired by Lee (2017) as well as the recommendations laid out in Reuters (2008) and The Associated Press (2020), we define quotatives as *objective* if they refer to the direct speech action and do not involve any subjective judgment of the action (e.g., like “say” and “tell”); and as *nonobjective* if they refer to some additional action or conduct and with subjective judgments (such as “boasted,” “rasped,” “taunted,” or “hailed”). To optimize for precision, we exclude common verbs with many non-quotative senses, such as “go.” Using this definition, we manually annotate the most frequent 99.5% of quotatives overall and the 98.0% of the most frequent quotatives per month. We consulted a professional journalist throughout this process, who suggested that the verbs “opine,” “pen,” and “utter” are only sometimes used nonobjectively. Since it would be infeasible to create a separate category just for these verbs, we excluded them. In the end, we curated a list of 32 objective and 152 nonobjective verbs (see Appendices A.1). We use this list to remove rare verbs (i.e., those not on the list), obtaining a final dataset with 6.7M quotes (35.1% of the original data) from 14,031 politicians in 989 outlets.

Data summary: Table 1 summarizes the most frequent quotatives and speakers in the final dataset. Consistent with prior literature (Gidengil and Everitt 2003), “say” and “tell” are the most commonly used quotatives (over 80% of the time), and Twitter is a common source for quotes (Bane 2019). We find minor differences in the coverage of each speaker across outlet categories, with Donald Trump being the most quoted speaker. Table 2 depicts the number of outlets per category and the fraction of quotes belonging to the ten top outlets in each category. No single outlet dominates an entire outlet category, but the distribution is more concentrated in the left and right categories, given the relatively smaller number of outlets in these categories.

`cc, relcl`, see <https://universaldependencies.org/en/dep/>

Quotative	%	Speaker	%	% Left	% Left-Center	%Least-Biased	%Right-Center	%Right
say	72.58	Donald Trump	15.64	19.51	15.59	15.18	13.85	16.93
tell	8.82	Barack Obama	3.50	4.36	3.90	2.60	3.39	4.36
write	3.82	Hillary Clinton	1.44	2.31	1.47	0.98	1.28	2.35
tweet	2.97	Nancy Pelosi	1.35	1.26	1.16	1.43	1.26	1.89
add	2.56	Bernie Sanders	1.24	2.67	1.17	1.00	0.99	1.60
ask	1.29	Joe Biden	1.23	1.60	1.16	1.13	0.94	1.89
continue	0.72	Mitch McConnell	1.04	1.03	0.98	1.11	0.99	1.13
respond	0.59	Chuck Schumer	0.90	0.73	0.85	0.98	0.85	1.03
declare	0.59	Lindsey Graham	0.88	1.13	0.78	0.83	0.78	1.30
state	0.45	Elizabeth Warren	0.86	1.43	0.83	0.77	0.69	1.08

Table 1: Statistics on top speakers and quotatives. Frequency of the top quotatives and speakers in the entire dataset after filtering. We also include the speaker frequency in each outlet category for reference. Among the listed quotatives, only “declare” is nonobjective.

Left	%	Left-Center	%	Least-Biased	%	Right-Center	%	Right	%
$n = 80$		$n = 249$		$n = 467$		$n = 142$		$n = 51$	
$k = 0.46M$		$k = 2.45M$		$k = 2.12M$		$k = 0.75M$		$k = 0.94M$	
CNN	20.17	Yahoo	6.15	The Hill	6.41	Washington Times	11.50	Breitbart	17.68
Raw Story	5.17	MSN	5.15	Roll Call	1.66	NW AR Democrat-Gaz.	3.42	Fox News	11.77
Salon	5.16	SFGATE	3.44	KVIA-TV	1.38	AR Democrat-Gaz.	3.10	Wash. Examiner	11.62
The Week	5.02	WaPo	3.28	UPI	1.29	Chicago Tribune	2.74	Newsmax	9.36
TPM	4.78	Politico	2.72	WTOP-FM	1.17	Laredo Morning Times	2.52	Daily Caller	4.52
Daily Beast	4.11	CBS	2.69	CBS Local	1.07	Boston Herald	2.40	Free Beacon	3.31
AlterNet	3.61	NBC News	2.50	KSL News	1.02	MyNorthwest	2.37	The Epoch Times	3.07
NY Magazine	3.40	NY Times	2.46	KTVQ-TV	0.95	Daily Herald	2.25	WorldNetDaily	2.86
Vox	2.69	Newsweek	2.03	WFMZ-TV	0.95	The Spokesman-Review	2.22	TheBlaze	2.76
Daily Kos	2.68	LA Times	1.74	WTHR-TV	0.94	Albuquerque Journal	2.14	CNSNews	2.64

Table 2: Statistics on outlets. Frequency and website names of the ten top outlets in each category. n : number of outlets in each category; k : number of quotes in each category.

3.2 Validation Error Analysis

To validate our verb extraction approach, we divide our data based on the detected verb type (objective and nonobjective) and manually annotate 100 randomly sampled quotes per year per verb type; the results are shown in Table 3. Our method achieves high accuracy in objective and nonobjective verb extraction, and this accuracy is stable over time, with a combined accuracy of over 90% across all years.

Before removing quotatives that are not contained in our dictionary of common quotatives (see **Step 4**), we also analyze the instances in which quotatives with rare verbs were retrieved by our model. We find that in many cases: 1) the actual quotative is not a verb (e.g., “according to” is a common quotative phrase); 2) the quote is a mixed quote without quotative; or 3) the method is applied to incomplete or noisy data and identifies an incorrect quotative.

3.3 Regression Model

Throughout the results section, we fit linear probability models (LPMs) specified as

$$y_q = \mathbf{X}_q \boldsymbol{\beta} + \varepsilon_q \quad (1)$$

using ordinary least squares, where y_q is a binary variable indicating whether a specific quote is nonobjective (1) or objective (0), \mathbf{X}_q is an array with explanatory variables associated with quote q , $\boldsymbol{\beta}$ is the array of coefficients we estimate, and ε_q is the error term. Coefficients obtained in this fashion are unbiased and robust (if the predicted probabilities lie between 0 and 1 (Horrace and Oaxaca 2003), which happens for all our analyses). We estimate standard errors and t -statistics with cluster robust standard errors (at the outlet level), accounting for autocorrelation between quotes as well as heteroskedasticity (Cameron and Miller 2015). More details about linear probability models can be found in Wooldridge (2010).

Year	Objective (%)	Nonobjective (%)	Combined (%)
2013	95	87	91
2014	98	86	92
2015	98	83	90.5
2016	95	92	93.5
2017	96	90	93
2018	96	90	93
2019	98	86	92
2020	98	91	94.5

Table 3: Validation of the proposed quotative extraction method. We take 100 random quotes from each year that our dataset covers, manually annotate the correct quotative, and show the accuracy. Note that the incorrect cases include those in which a quotative cannot be identified. This is most often because no clear quotative is present (due to mixed quotes or because the article text is incomplete).




Figure 2: Usage of nonobjective quotatives across outlets of different political leaning. For each media bias category (on the y-axis), we depict the usage of nonobjective quotatives per outlet (each represented by a circle \circ) and the overall average usage pooled across outlets (\times). Note that the x-axis is on a logarithmic scale. Pairwise differences between averages are statistically significant under the Wilcoxon Rank-Sum Test with Bonferroni correction.

Recent work has often favored LPM instead of logistic regression or other non-linear models due to the ease of interpretation and of incorporating fixed effects (Gomila 2021), e.g., see Dai et al. (2021). In our case, we use the LPM since our main purpose is to approximate the partial effects of the explanatory variables (Wooldridge 2010). We report all results in this paper with significance level $\alpha = 0.05$.

4 Results

4.1 RQ1: How Has the Usage of Nonobjective Quotatives Evolved?

Across the study period, we find that the usage of nonobjective quotatives produces a sensible ordering of the media bias categories considered, with the more partisan outlets using the most nonobjective quotatives and the less partisan outlets using the least. We depict this order in Fig. 2, where each circle (\circ) represents the average usage of nonobjective quotatives in one of the outlets considered, and crosses (\times)




Figure 3: Percentage yearly rate of change of nonobjective quotative usage for each outlet category. A solid circle denotes a significant effect ($p < 0.05$), and a hollow circle denotes an insignificant effect. Note that the reported trends correspond to the estimated β coefficients in Eq. (2).




Figure 4: Percentage yearly rate of change of nonobjective quotative usage for all outlets combined. We show the percentage of nonobjective quotatives after performing centering and plot a regression line showing the β coefficients estimated in our fixed effects model.

indicate the average usage pooled across each media bias category. When quoting politicians, U.S. least-biased outlets use nonobjective quotatives the least, followed by the left-center and right-center outlets, and finally, right and left outlets. Outlets considered more partisan by MB/FC used nonobjective quotatives more. Pairwise differences between averages are statistically significant under the Wilcoxon Rank-Sum Test with Bonferroni correction.

To study how the usage of nonobjective quotatives evolved, we use a fixed effects linear probability model. For each quote q , let $o[q]$ be the outlet in which q was reported, $b[q]$ be the bias category of the outlet, and $t[q]$ be the time when it was reported in months relative to the starting period of our dataset (May 2013). We define the model as

$$y_q = \alpha_{o[q]} + \gamma_{b[q]} + \beta_{b[q]}t[q] + \varepsilon_q, \quad (2)$$

where the dependent variable y_q equals 1 if the verb used in the quote q is nonobjective and 0 otherwise, $\alpha_{o[q]}$ is an outlet-level intercept, $\gamma_{b[q]}$ is a category-level intercept, and $\beta_{b[q]}$ is a category-level trend in the usage of nonobjective quotatives – the effect we are interested in estimating. Since we are modeling time-series data (one per outlet), autocorrelation may shrink the confidence intervals of model estimates (see Bertrand, Duflo, and Mullainathan (2004) for details). To address this, we estimate the model us-




Figure 5: Trends in the usage of nonobjective quotatives across outlets of different political leaning. For each media bias category (one per column), we show the percentage of nonobjective quotatives after performing outlet-level centering and plot the regression line showing the β coefficients estimated in our fixed effects model.

Quotative	Percentage Change	Quotative	Odds Ratio
say	-10.18↓	tweet	1 → 17.04
tweet	4.157↑	<i>falter</i>	1 → 11.88
tell	1.843↑	caption	1 → 10.86
write	1.555↑	<i>restate</i>	1 → 6.772
add	1.020↑	<i>remark</i>	1 → 5.614
respond	0.3538↑	<i>punctuate</i>	1 → 4.855
continue	0.3454↑	<i>blurt</i>	4.891 → 1
<i>declare</i>	0.2325↑	<i>disclose</i>	4.911 → 1
<i>remark</i>	0.2159↑	<i>enthuse</i>	6.335 → 1
<i>claim</i>	0.2028↑	<i>exult</i>	15.65 → 1

Table 4: Changes in quotatives used. We report the most changed quotatives between our dataset’s first and last 12 months in absolute change and odds ratio. *Italic* highlighting denotes nonobjective quotatives.

ing cluster robust standard errors, clustering on the outlet level (Cameron and Miller 2015).

We depict the estimated trends in nonobjective quotative usage in Fig. 3 (i.e., the estimated $\beta_{b[q]}$ in Equation (2)). Although the least-biased outlets used nonobjective quotatives less on average (Fig. 2), we find that their usage of nonobjective quotatives increases over time. We estimate that least-biased outlets increase their usage of nonobjective quotatives by 0.08% per year and that right-center and left-center outlets increase their usage by 0.10% and 0.06% per year, respectively. If we compare the level of nonobjective quotative usage from 2013 to 2020 (beginning and end of our study), these changes translate to relative increases of 19.9% for least-biased outlets, 21.3% for right-center and 13.6% for left-center outlets. In contrast, left outlets experienced a statistically insignificant decrease in their usage of nonobjective quotatives by 0.12% per year, and right outlets experienced a smaller, statistically insignificant increase in usage of nonobjective quotatives (of roughly 0.01%). We also experimented with an added seasonality factor shared across all outlets and politicians, either on a monthly or yearly basis. We exclude these results here since 1) accounting for seasonality does not substantially change the results, 2) ef-

fect sizes decrease only slightly when seasonality is considered, and 3) with only two four-year cycles, there is insufficient data for a robust analysis of seasonality.

We further illustrate the results obtained in the fixed effects model in Figure 4 and 5. In Figure 4, we center the overall quotative usage around 0 and plot the month-level nonobjective quotative usage, along with a regression line capturing the trend. The increase in nonobjective quotative usage across all outlets aggregated in percentage per year (i.e., the slope) is 0.079% ($p < 0.001$). In other words, we find that the overall rate of nonobjective quotative usage among all outlets is increasing, thereby supporting the argument that changes in the usage of quotatives indicate a steady decline in objectivity in U.S. political news.

In Figure 5, we plot at the outlet category level: we center each outlet time series around 0 and then report the month-level (demeaned) usage of nonobjective quotatives per outlet category, along with a regression line capturing the trend in each time series. Here, we observe that the usage of nonobjective quotatives increases for centrist outlets, decreases for outlets on the left, and only slightly increases for outlets on the right, although the two latter results were not statistically significant according to the model. In other words, the overall trend in decreasing objectivity can be explained by the trends in centrist outlets.

Another way to understand the change in quotative usage is to consider the extremes. We compare how quotative usage changes between the first 12 months (May 2013 - April 2014) and the last 12 months (May 2019 - April 2020) of our dataset. In Table 4, we report the quotatives that experienced the largest changes in terms of absolute percentage points (on the left) and odds ratio (on the right). We find that the usage of the quotative “say,” typically considered the gold standard of quotatives, fell by more than 10% percentage points. At the same time, we see an increase in other objective quotatives (e.g., tell), but this increase does not account for the entire ten percentage points. Lower in the list, we see that nonobjective quotatives like “claim,” “remark,” and “declare” are used more often. Finally, we highlight that quotatives reveal changes in where journalist source their quotes, with both “tweet” and “caption” (usually employed when the speaker uploads a picture or video on social media) experiencing large relative increases in usage.




Figure 6: Differences in nonobjective quotative usage between Democratic and Republican speakers. All estimates are significant. Reported differences correspond to the contrasts $\sigma_{\text{democratic}} - \sigma_{\text{republican}}$ in Eq. (3) (in percentage points).

4.2 RQ2: How Do News Outlets Use Nonobjective Quotatives When Covering Politicians of Different Parties?

Next, we investigate whether the outlets are biased in their quotative usage when they cover politicians from ideologically similar vs. opposing political parties.

Quotative bias across outlet categories. For each quote q , let $p[q]$ be the party of the politician who uttered the quote. Keeping with the notation in Eq. (2), we again use a fixed effects linear probability model

$$y_q = \gamma_{b[q]} + \eta_{p[q]} + \sigma_{b[q],p[q]} + \varepsilon_q, \quad (3)$$

where the dependent variable y_q equals 1 if the quotative used in the quote q is nonobjective and 0 otherwise, $\gamma_{b[q]}$ is a category-level intercept, $\eta_{p[q]}$ is a party-level intercept, and $\sigma_{b[q],p[q]}$ captures the interaction between pairs of outlet bias category ($b[q]$) and speaker party ($p[q]$). We emphasize that outlet and outlet category are outlet-level attributes, while the political party is a politician-level attribute. We again cluster standard errors on the outlet level to address autocorrelation. Note that here, we are particularly interested in the contrasts between different combinations of outlet categories and speaker parties, e.g., the difference between how left outlets quote Democratic and Republican speakers (in the model $\sigma_{\text{left,democratic}} - \sigma_{\text{left,republican}}$).

For each media bias category, we show the estimated percentage difference in the usage of nonobjective quotatives for Democratic and Republican speakers in Figure 6. For every outlet category, there is a significant difference in quotative usage between Democratic and Republican speakers. Notably, this difference is nearly 2%, around a third of the overall nonobjective quotative usage, for both left and right media outlets, which use more nonobjective quotatives when referring to politicians from opposing political parties. For centrist outlets, we see a Democratic bias in the usage of quotatives, with Republicans being quoted with nonobjective quotatives around 1% more for least-biased and left-center outlets and nearly 0.5% more for right-center outlets.

Matched analysis. A possible explanation for what we observe in Figure 6 is that outlets of different political leanings cover different quotes (Tan, Peng, and Smith 2018) and




Figure 7: Differences in nonobjective quotative usage between Democratic and Republican speakers for matched quotes. A solid circle denotes a significant effect ($p < 0.05$), and a hollow circle denotes an insignificant effect. Reported differences correspond to the contrasts $\sigma_{\text{democratic}} - \sigma_{\text{republican}}$ in Eq. (3) (in percentage points).

that these quotes lend themselves more or less to being attributed to speakers through nonobjective quotatives. Even if this were the case, one could still make a case against media bias, as journalistic textbooks and guidelines instruct the usage of objective quotatives regardless of the quote (Mencher and Shilton 1997; Brooks et al. 2007; Reuters 2008; The Associated Press 2020; Rich 2015). Nevertheless, we entertain this hypothesis by performing a matched analysis. Specifically, we identify quotes covered by both the left and right media outlets, merging left/left-center and right/right-center outlets for ease of comparison (we refer to these merged bias categories as left and right “combined”). Using only this subset of quotes ($n = 1.02\text{M}$, 15.13% of all quotes in our data), we fit the fixed effects model defined in Eq. (3).

For the two collapsed media bias categories and considering only matched quotes, we show the estimated difference in usage of nonobjective quotative for Democratic and Republican speakers in Figure 7. For right/right-center outlets, we find a small non-significant positive difference (0.01%). Note that in the non-matched scenario, in Figure 6, these two types of outlets behave differently – right outlets use more nonobjective quotatives when quoting democrats, whereas right-center outlets do so when quoting republicans. Since we aggregate right and right-center, it may be that these heterogeneous effects cancel each other. For left/left-center outlets, we find a significant difference of around -0.75% in the usage of nonobjective quotatives. This suggests that quote selection alone cannot explain the quotative bias previously observed and that quotative selection forms an additional source of bias on top of quote selection.

Trends in quotative bias. Finally, we investigate if quotative bias has evolved during the study period, using a fixed effects linear probability model:

$$y_q = \alpha_{o[q]} + \gamma_{b[q]} + \eta_{p[q]} + \lambda_{b[q],p[q]}t[q] + \varepsilon_q, \quad (4)$$

where the dependent variable y_q equals 1 if the quotative used in the quote q is nonobjective and 0 otherwise, $\alpha_{o[q]}$, $\gamma_{b[q]}$, and $\eta_{p[q]}$ are outlet, category, and party-level intercepts, and $\lambda_{b[q],p[q]}$ is the trend in the usage of nonobjective quotatives for each party/bias category combination.

For each media bias category, we depict the difference in the trends of nonobjective quotative usage for Democrats and Republicans in Figure 8. For left and centrist outlets,




Figure 8: Differences in percentage yearly rate of change of nonobjective quotative usage between Democratic and Republican speakers for each media outlet category. A solid circle denotes a significant effect ($p < 0.05$). Reported trends correspond to the contrasts $\lambda_{\text{democratic}} - \lambda_{\text{republican}}$ in Eq. (4) (in percentage points).

the gap between how nonobjective quotatives are used to quote Democrats and Republicans is increasing in the study period. These increases are statistically significant and substantial compared to the existing level of quotative bias observed in our data. For example, the estimated contrast of the trend is around 0.33% for left-center outlets, and the existing quotative bias is 1.08%. Thus, the annual relative increase of quotative bias is above 30%. Left outlets exhibit the most increase in quotative bias in absolute terms, at 0.5% per year. For right outlets, the difference in trend leans Republican, but the effect is not statistically significant.

5 Discussion

In this work, we analyzed quotatives to study political journalism, answering the following questions: How has the usage of nonobjective quotatives evolved in U.S. political journalism (**RQ1**)? How do news outlets use nonobjective quotatives when covering politicians of different parties (**RQ2**)? To answer these questions, we proposed a method to identify quotatives for direct quotes using dependency parsing. We then extracted quotatives from a large dataset of speaker-attributed quotes, resulting in over 6.7 million quotes over eight years, from 2013 to 2020. By counting the usage of objective and nonobjective quotatives, we analyzed the static and dynamic trends of quotative usage.

We find that the more partisan outlets use more nonobjective quotatives (Figure 2). However, during the study period, centrist outlets (classified as least-biased, left-center, and right-center by MB/FC) experienced a significant increase in the usage of nonobjective quotatives, suggesting that they may be “catching up” to the more biased outlets (Figure 5). We further observe that outlets tend to use more nonobjective quotatives when covering politicians of the opposing ideology, thereby exhibiting “quotative bias” (Figure 6). Even when we control for quotes by matching outlets on the quote level, we find that this bias still exists for left and left-center outlets (Figure 7). Finally, we find a rapid increase in quotative bias for most outlet categories over time, which may indicate that U.S. political news is becoming increasingly polarized (Figure 8).

These findings suggest that two simultaneous processes are at play: outlets are adopting more nonobjective quotatives overall *and* the usage of nonobjective quotatives is increasingly “mediated” by the party affiliation of quoted politicians. Both processes indicate a measurable decrease in journalistic objectivity. While detecting bias often requires some level of human judgment to determine neutrality, and while it is debatable how a neutral or balanced view can be presented in any specific context, quotative usage can be regarded as an easily quantifiable form of bias due to its prominence within journalism. There are clear and established rules for the usage of quotes on which journalists have historically agreed, as is evident from textbooks (Brooks et al. 2007; Mencher and Shilton 1997; Rich 2015) and editorial guidelines (Reuters 2008; The Associated Press 2020). Although objective journalism is a 20th-century invention and could be considered an anomaly throughout history, it is regarded as central to today’s democratic process. In this context, our results indicate a decrease in the level of objective quotative usage in U.S. political news coverage, which can be seen as a devolution of journalism as a profession.

An interesting question is to which extent Donald Trump, the most quoted speaker in our dataset, influences the findings observed in this paper. To investigate this, we re-run all fixed effects models in the paper in a filtered dataset from which we removed Trump’s quotes (see Appendices A.2). We find that, without Trump, there is no significant increase in the overall usage of nonobjective quotatives. However, we still observe quotative bias in all outlet categories except left-center outlets.

The ways in which the observed increase in nonobjective quotatives relates to broader trends in U.S. politics and the news ecosystem remains unclear. On the one hand, the observed trend may merely reflect the reality of the news business. As newspapers struggle to retain subscribers and attract clicks (Thurman, Lewis, and Kunert 2019), outlets (including the least-biased ones) may have succumbed to nonobjective quotatives as they adapt to the fast-paced style of Web-first publishing and try to produce engaging content. Alternatively, journalists themselves may be subject to trends of increasing polarization in the general public (Abramowitz and Saunders 2008), becoming more prone to Freudian slips when reporting the speech of politicians they (dis)like. On the other hand, the increase in quotative bias may influence people’s opinions about politicians (Cole and Shaw 1974) or erode the reader’s trust in the media outlet, as they might disagree with the opinions subtly embedded in the news piece by the writer (Gunther 1992).

Limitations. We highlight three limitations of our work. First, while Quotebank covers a large number of outlets over time, the number of quotes from each outlet does not necessarily reflect the actually published amount of content per outlet, nor does it represent the relative popularity of each outlet. Therefore, our findings should not be interpreted as the perception that an average news consumer in the U.S. may have. Second, since we focus on quotative verbs and do not consider adverbs, we may not capture the complete picture of quotative usage. For example, the hypothetical quo-

tative “say aggressively” should be categorized as nonobjective, but we still categorize it as objective since we do not consider adverb usage. Empirically, however, this combination of an objective quotative verb and a nonobjective adverb is uncommon (as it can be considered even more unprofessional than a nonobjective quotative verb) and is unlikely to affect our results. The occurrence of non-verb quotatives is rare as well. Also, despite our best effort, a few quotatives may be misidentified or not included in our dictionary. While this should not affect our overall analysis, analyzing these rare quotatives might yield additional insights. Third, while we carefully removed most indirect and mixed quotes from our dataset, some are expected to remain in the data. Thus, we cannot completely eliminate the effect of indirect quotes on the obtained results.

Future Work. Future work could use and extend our methodology to investigate trends in nonobjective quotative usage after the Trump presidency or in other countries and languages. Further, an even more comprehensive investigation of the landscape of quotative usage in political journalism could be obtained by extending our methodology to include indirect and mixed quotes and/or considering nonverb phrases and adverbs as quotatives. Last, future work could examine if nonobjective quotatives reflect changes in polarization or the media ecosystem and to what extent they impact readers’ opinions of politicians and news outlets.

A Appendices

A.1 List of Objective and Nonobjective Quotatives

Objective Quotative: say tell write tweet add ask continue respond state explain note read reply quote announce recall conclude post begin describe answer cite testify recount close email summarize finish caption inform accord preface

Nonobjective Quotative: declare warn claim argue insist remark joke suggest acknowledge urge promise comment assert quip proclaim admit share complain vow praise stress predict charge shoot observe emphasize boast reiterate pledge remind fire counter lament shout concede caution assure retort confirm exclaim contend advise laugh blast hit indicate yell press reflect tout fume mock muse interject gush apologize brag clarify thunder challenge hail interrupt snap elaborate chide chime plead lash intone confess disagree protest crow boom tease cry scold laud hint affirm crack implore scoff bellow chuckle rail lecture smile speculate scream bemoan reassure shrug marvel rip underscore decry commend gripe object confide jab pronounce taunt instruct enthuse admonish roar chastise whine rant reminisce reaffirm concur recite disclose beam whisper deflect posit rebuke pile falter articulate deride channel sneer blurt persist grumble ratchet punctuate forecast sigh sketch exhort explode burst preach cede interrogate diagnose gloat tee shy wax mourn exult goad backpedal restate howl

A.2 Analysis of Quotes Without Donald Trump

In Figure 9, we report the coefficients of interest for the fixed effects models depicted in Equations (2), (3), and (4) on a filtered dataset containing no quotes by Donald Trump.

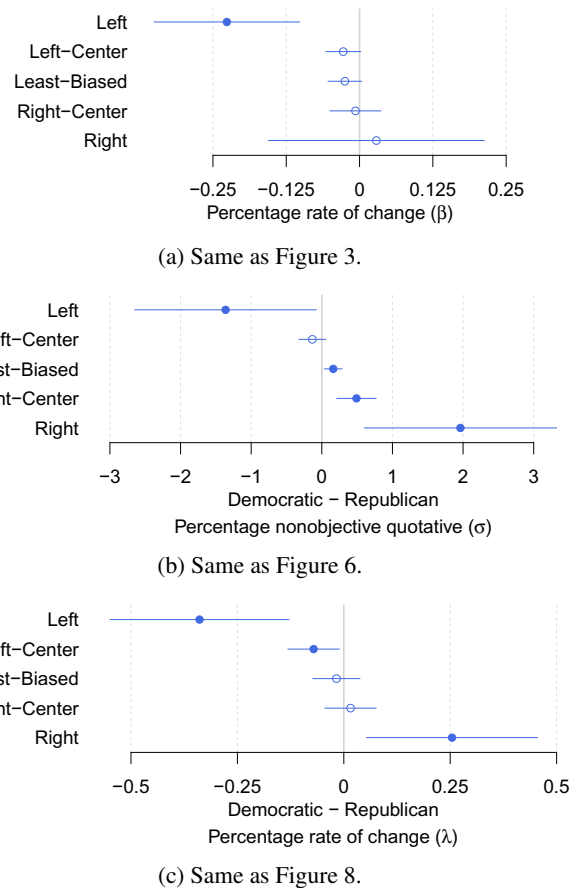


Figure 9: These figures correspond to the same analysis from Figures 3, 6, and 8 in the main paper but without any quotes from Donald Trump.

Ethical Statement

This work uses publicly available data to analyze quotes from U.S. politicians to study media bias. We match politicians’ names between Quotebank and Wikidata. Our study utilizes one identity characteristic – the political party of the involved speakers – at an aggregated level in order to investigate the effect of ideology in quotative usage. We do not perform any individual-level inferences. Additionally, we manually validate both speaker disambiguation and quotative extraction methods to minimize the risk of identifying the wrong individuals and ensure the veracity of our findings. Given that politicians are public figures and the importance of research to better understand the language used in political journalism and its implications, we believe that our work is in line with reasonable expectations of privacy (Doherty 2007). We do not foresee potential negative societal impacts coming from this research. On the contrary, we believe that a better understanding of our political media ecosystem is essential to improve it. We confirm that we have read and abide by the AAAI code of conduct.

(a) Coefficients shown in Fig. 3.

Outlet	β	SE	t	p
All	0.08	0.02	4.14	<0.001
Left	-0.12	0.07	-1.86	0.063
Left-Center	0.06	0.02	3.36	0.001
Least-Biased	0.08	0.02	4.77	<0.001
Right-Center	0.10	0.03	3.44	0.001
Right	0.01	0.09	0.07	0.943

(b) For coefficients shown in Fig. 6 (top) and Fig. 7 (bottom).

Outlet	$\sigma_{\text{dem}} - \sigma_{\text{rep}}$	SE	t	p
Left	-1.89	0.37	-5.10	<0.001
Left-Center	-1.09	0.09	-12.26	<0.001
Least-Biased	-0.88	0.08	-11.36	<0.001
Right-Center	-0.40	0.11	-3.58	0.013
Right	1.75	0.42	4.15	0.001
Left (Combined)	-0.72	0.09	-8.00	<0.001
Right (Combined)	0.10	0.14	0.73	0.883

(c) For coefficients shown in Fig. 8.

Outlet	$\lambda_{\text{dem}} - \lambda_{\text{rep}}$	SE	t	p
Left	-0.50	0.07	-7.67	<0.001
Left-Center	-0.33	0.03	-12.50	<0.001
Least-Biased	-0.29	0.02	-12.18	<0.001
Right-Center	-0.22	0.03	-7.82	<0.001
Right	0.17	0.06	2.71	0.170

Table 5: Regression result details. For key coefficients discussed throughout the paper, we present standard errors, t statistics, and p values.

Acknowledgements

We would like to thank Jonathan K ulz and Marko  uljak for their help in data preprocessing and Chris Stokel-Walker for having a look at our list of quotatives. This project was partly funded by the Swiss National Science Foundation (grant 200021_185043), the European Union (TAILOR, grant 952215), the Microsoft Swiss Joint Research Center, and the Deutsche Forschungsgemeinschaft under the Excellence Strategy of the German federal and state governments (DFG, grant EXC-2035/1 - 390681379). We also acknowledge generous gifts from Facebook and Google supporting West’s lab.

References

Abramowitz, A. I.; and Saunders, K. L. 2008. Is Polarization a Myth? *The Journal of Politics*, 70(2): 542–555.

Bane, K. C. 2019. Tweeting the Agenda. *Journalism Practice*, 13(2): 191–205.

Beharrell, P.; Davis, H.; Eldridge, J.; Hewitt, J.; Hart, J.; Philo, G.; Walton, P.; and Winston, B. 2009. *More Bad News (Routledge Revivals)*. Routledge.

Bell, A. 1991. *The Language of News Media*. Wiley-Blackwell.

Bernhardt, D.; Krasa, S.; and Polborn, M. 2008. Political Polarization and the Electoral Effects of Media Bias. *Journal of Public Economics*, 92(5): 1092–1104.

Bertrand, M.; Duflo, E.; and Mullainathan, S. 2004. How Much Should We Trust Differences-in-differences Estimates? *The Quarterly Journal of Economics*, 119(1): 249–275.

Boudana, S. 2011. A Definition of Journalistic Objectivity as a Performance. *Media, Culture & Society*, 33(3): 385–398.

Brenan, M. 2022. Media Confidence Ratings at Record Lows. *Gallup*.

Brooks, B. S.; Kennedy, G.; Moen, D. R.; and Ranly, D. 2007. *News Reporting and Writing*. Macmillan.

Budak, C.; Goel, S.; and Rao, J. M. 2016. Fair and Balanced? Quantifying Media Bias through Crowdsourced Content Analysis. *Public Opinion Quarterly*, 80(S1): 250–271.

Calcutt, A.; and Hammond, P. 2011. *Journalism Studies: A Critical Introduction*. Routledge.

Caldas-Coulthard, C. R. 1992. Reporting Speech in Narrative Discourse: Stylistic and Ideological Implications. *Ilha do Desterro A Journal of English Language, Literatures in English and Cultural Studies*, (27): 067–082.

Cameron, A. C.; and Miller, D. L. 2015. A Practitioner’s Guide to Cluster-robust Inference. *Journal of Human Resources*, 50(2): 317–372.

Cappelen, H.; and Lepore, E. 1997. Varieties of Quotation. *Mind*, 106(423).

Cole, R. R.; and Shaw, D. L. 1974. ‘Powerful’ Verbs and ‘Body Language’: Does the Reader Notice? *Journalism Quarterly*, 51(1): 62–66.

 uljak, M.; Spitz, A.; West, R.; and Arora, A. 2022. Strong Heuristics for Named Entity Linking. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies: Student Research Workshop, SRW@NAACL-HLT’22*.

Dai, H.; Saccardo, S.; Han, M. A.; Roh, L.; Raja, N.; Vangala, S.; Modi, H.; Pandya, S.; Sloyan, M.; and Croymans, D. M. 2021. Behavioural nudges increase COVID-19 vaccinations. *Nature*, 597(7876): 404–409.

D’Alessio, D.; and Allen, M. 2006. Media Bias in Presidential Elections: A Meta-Analysis. *Journal of Communication*, 50(4): 133–156.

DellaVigna, S.; and Kaplan, E. 2007. The Fox News Effect: Media Bias and Voting. *The Quarterly Journal of Economics*, 122(3): 1187–1234.

Dixon, G. N.; and Clarke, C. E. 2013. Heightening Uncertainty Around Certain Science: Media Coverage, False Balance, and the Autism-Vaccine Controversy. *Science Communication*, 35(3).

Doherty, C.; Kiley, J.; Tyson, A.; and Johnson, B. 2019. Public highly critical of state of political discourse in the US Pew Research Center.

Doherty, M. 2007. Politicians as a species of ‘Public Figure’ and the Right to Privacy. *Humanitas Journal of European Studies*, 1(1): 35–56.

Geis, M. L. 1987. *The Language of Politics*. Springer.

- Gentzkow, M.; and Shapiro, J. M. 2010. What Drives Media Slant? Evidence From U.S. Daily Newspapers. *Econometrica*, 78(1): 35–71.
- Gidengil, E.; and Everitt, J. 2003. Talking Tough: Gender and Reported Speech in Campaign News Coverage. *Political Communication*, 20(3): 209–232.
- Gomila, R. 2021. Logistic or linear? Estimating causal effects of experimental treatments on binary outcomes using regression analysis. *Journal of Experimental Psychology: General*, 150(4): 700.
- Gunther, A. C. 1992. Biased Press or Biased Public? Attitudes Toward Media Coverage of Social Groups. *The Public Opinion Quarterly*, 56(2).
- Hamborg, F.; Donnay, K.; and Gipp, B. 2019. Automated Identification of Media Bias in News Articles: An Interdisciplinary Literature Review. *Int. J. Digit. Libr.*, 20(4): 391–415.
- Harry, J. C. 2014. Journalistic Quotation: Reported Speech in Newspapers From a Semiotic-linguistic Perspective. *Journalism*, 15(8).
- Horrace, W. C.; and Oaxaca, R. L. 2003. New wine in old bottles: A sequential estimation technique for the LPM. Available at SSRN 383102.
- Ingram, D.; and Henshall, P. 2012. *The News Manual*, volume 1, chapter 8.
- Iyengar, S. 1994. *Is Anyone Responsible?: How Television Frames Political Issues*. University of Chicago Press.
- Jurafsky, D.; and Martin, J. H. 2022. *Speech and Language Processing*. online, 3rd ed. draft edition.
- Just, M.; Crigler, A.; and Buhr, T. 1999. Voice, Substance, and Cynicism in Presidential Campaign Media. *Political Communication*, 16(1): 25–44.
- Kahn, K. F.; and Kenney, P. J. 2002. The Slant of the News: How Editorial Endorsements Influence Campaign Coverage and Citizens' Views of Candidates. *American Political Science Review*, 96(2): 381–394.
- Kobre, S. 1953. How Florida Dailies Handled the 1952 Presidential Campaign. *Journalism Quarterly*, 30(2): 163–169.
- Kübler, S.; McDonald, R.; and Nivre, J. 2009. Dependency Parsing. *Synthesis Lectures on Human Language Technologies*, 1(1): 1–127.
- Külz, J.; Spitz, A.; Abu-Akel, A.; Günemann, S.; and West, R. 2022. United States Politicians' Tone Became More Negative with 2016 Primary Campaigns. *CoRR*, abs/2207.08112.
- Lazaridou, K.; Krestel, R.; and Naumann, F. 2017. Identifying Media Bias by Analyzing Reported Speech. In Raghavan, V.; Aluru, S.; Karypis, G.; Miele, L.; and Wu, X., eds., *IEEE International Conference on Data Mining, ICDM'17*.
- Lee, G. 2017. Verb Objectivity and Source Qualification: Comparison of Quotation Attributions in Offline and Online Newspapers. *Journalism*, 18(7): 890–906.
- McCombs, M. E.; and Shaw, D. L. 1972. The Agenda-Setting Function of Mass Media. *The Public Opinion Quarterly*, 36(2): 176–187.
- McNair, B. 2017. After Objectivity? *Journalism Studies*, 18(10).
- Mencher, M.; and Shilton, W. P. 1997. *News Reporting and Writing*. Brown & Benchmark Publishers.
- Merrill, J. C. 1965. How Time Stereotyped Three US Presidents. *Journalism Quarterly*, 42(4): 563–570.
- Mitchell, A.; Simmons, K.; Matsa, K. E.; and Silver, L. 2018. Publics Globally Want Unbiased News Coverage, but are Divided on Whether Their News Media Deliver.
- Newell, E.; Margolin, D.; and Ruths, D. 2018. An Attribution Relations Corpus for Political News. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation, LREC'18*.
- Niculae, V.; Suen, C.; Zhang, J.; Danescu-Niculescu-Mizil, C.; and Leskovec, J. 2015. QUOTUS: The Structure of Political Media Coverage as Revealed by Quoting Patterns. In Gangemi, A.; Leonardi, S.; and Panconesi, A., eds., *Proceedings of the 24th International Conference on World Wide Web, WWW'15*.
- O'Connell, M. 1999. Is Irish Public Opinion towards Crime Distorted by Media Bias? *European Journal of Communication*, 14(2): 191–212.
- Pareti, S. 2012. A Database of Attribution Relations. In *Proceedings of the Eighth International Conference on Language Resources and Evaluation, LREC'12*.
- Pareti, S. 2016. PARC 3.0: A Corpus of Attribution Relations. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation, LREC'16*.
- Prasad, R.; Dinesh, N.; Lee, A.; Miltsakaki, E.; Robaldo, L.; Joshi, A.; and Webber, B. 2008. The Penn Discourse TreeBank 2.0. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation, LREC'08*.
- Reuters. 2008. Reuters Handbook of Journalism.
- Rich, C. 2015. *Writing and Reporting News: A Coaching Method*. Cengage Learning.
- Ryan, M. 2001. Journalistic Ethics, Objectivity, Existential Journalism, Standpoint Epistemology, and Public Journalism. *Journal of Mass Media Ethics*, 16(1): 3–22.
- Schudson, M. 1981. Discovering the News: A Social History of American Newspapers.
- Sonoda, K. 1997. Subject-verb Inversion Before a Quotation in the Media Discourse. *Current English Studies*, 1997(36).
- Stenvall, M. 2008. On Emotions and the Journalistic Ideals of Factuality and Objectivity—Tools for Analysis. *Journal of Pragmatics*, 40(9): 1569–1586.
- Tan, C.; Peng, H.; and Smith, N. A. 2018. "You Are No Jack Kennedy": On Media Selection of Highlights from Presidential Debates. In *Proceedings of the 2018 World Wide Web Conference*.
- The Associated Press. 2020. *The Associated Press Stylebook: 2020-2022*. Basic Books.
- Thurman, N.; Lewis, S. C.; and Kunert, J. 2019. Algorithms, Automation, and News. *Digital Journalism*, 7(8): 980–992.
- Van Dijk, T. A. 1988. *News as Discourse*. Lawrence Erlbaum Associates.
- Vaucher, T.; Spitz, A.; Catasta, M.; and West, R. 2021. Quotebank: A Corpus of Quotations from a Decade of News. In *The Fourteenth ACM International Conference on Web Search and Data Mining, WSDM'21*. ACM.
- Watts, D. J.; Rothschild, D. M.; and Mobius, M. 2021. Measuring the News and Its Impact on Democracy. *Proceedings of the National Academy of Sciences*, 118(15).
- Wooldridge, J. M. 2010. *Econometric analysis of cross section and panel data*. MIT press.
- Zhou, D. X.; Resnick, P.; and Mei, Q. 2011. Classifying the Political Leaning of News Articles and Users from User Votes. In *Proceedings of the Fifth International Conference on Weblogs and Social Media, ICWSM'11*.