Shades Of Neutrality: Detecting Media Slant In A Political Monopoly Using Quotation Accuracy

LUCAS SHEN, Institute for Human Development and Potential, A*STAR

A natural starting point in the study of political media bias has been to examine differences in coverage intensity. This study develops a notion of quotation accuracy as a dimension of media slant. Using natural language processing (NLP), I operationalize three measures of quotation accuracy—substring alignment, bag-of-words overlap, and a two-stage transformer-based semantic measure. I combine the measures with applied econometrics to analyze coverage of parliamentary speeches in Singapore's flagship daily newspaper, a setting where coverage intensity shows no slant. Across all three measures, opposition speeches are quoted less accurately than those of the ruling party, demonstrating accuracy as a crucial yet underappreciated dimension of slant. NLP and machine learning features, and bounding analyses rule out competing explanations such as topical composition, ignorable tokens, or linguistic coherence. I then contextualize the findings in the government-media machinery, where transcript circulations (or lack thereof) could explain lapses in quotation accuracy. I theorize that these partisan differences in media engagement embed private information about slant.

CCS Concepts: • **Computing methodologies** → **Natural language processing**; **Information extraction**; *Causal reasoning and diagnostics*; • **Applied computing** → **Law, social and behavioral sciences**; • **Social and professional topics** → Governmental regulations.

Keywords Political Media Slant, Natural Language Processing, Quotation Accuracy, Semantic, Transformers, Singapore, Applied Econometrics, Parliamentary Speeches

1 Introduction

In an age where the media is omnipresent and shapes public opinion [9, 11, 26, 45, 46, 71] and electoral outcomes [1, 13, 18, 21, 25, 28, 33, 53], detecting political biases of the media has never been more pressing. Quantifying media bias is challenging because editorial decisions that shape what and how stories are covered are invisible to news consumers. One common starting point in detecting media bias has been to examine coverage intensity [44, 59–61, 63], but such measures overlook subtler forms of (mis)representation that occur through quotations of politicians. Moreover, intensity-based measures are less useful when the media is cognizant of the metric's public scrutiny. This study considers a notion of coverage accuracy, offering traction in contexts where intensity-based measures reveal no bias and, more generally, serves as a gauge of reliability.

Author's Contact Information: Lucas Shen, lucas@lucasshen.com, Institute for Human Development and Potential, A*STAR.

To examine richer forms of media bias, research has increasingly turned to natural language processing (NLP) and text processing techniques [30, 37, 54, 56, 73]. These computational methods quantify bias by detecting systematic differences in language use, sentiment, and framing that typically elude human readers [51, 73]. Still, a computational application often overlooked concerns coverage accuracy—are certain politicians quoted more accurately than others?

Accuracy matters because distortions can subtly expose readers to a version of reality that differs from the source [20, 34, 56]. Direct quotes carry significant weight as they are perceived as more authoritative and credible than paraphrased reports or summaries [20, 34]. A concrete example from the sample is the quote, "the alternative will be too painful to bear," which originated from the speech

> …So, let us not take our harmonious social fabric for granted because **the alternative *may* be too painful to endure**. This is one pillar of success that we must continue to invest in…

Though the quote and its originating substring look reasonably similar, the quote uses *will be* where the speech uses *may be*—a subtle modulation in assertiveness. Even such minor differences misrepresent a politician's stance and, when repeated over time, might distort political dialogue and sway public discourse in the political process [1, 9, 11, 13, 18, 18, 21, 25, 26, 33, 45, 46, 53, 71, 71, 73, 80].

The main computational problem in this study is to quantify quotation accuracy using NLP by comparing media quotations with their original speeches. I consider three complementary measures: two lexical and one semantic. First, the substring accuracy metric identifies the longest continuous sequence of matching words between the speech and the quote, reflecting how faithfully specific speech excerpts are reproduced. Second, the bag-of-words accuracy metric ignores word order and focuses on the proportion of overlapping words, enabling flexible comparisons of textual content even when the same content is paraphrased [47]. The third metric measures semantic alignment, where I use a two-stage transformer model based on a retrieval-reranking pipeline to measure how well the speech semantically supports the quote [43, 47, 52, 57, 65]. Together, these measures represent the first systematic application of modern retrieval and matching techniques to the problem of quotation fidelity in political reporting.

The extent of verbatim and semantic alignment between a quote and its speech constitutes quotation accuracy. In turn, systematic partisan differences in quotation accuracy constitute political media slant. I apply these accuracy measures to Singapore's flagship daily, where I find no partisan differences in coverage intensity. To estimate partisan differences in accuracy, I regress the quotation accuracy measures on an opposition indicator using ordinary least squares. I use NLP and machine learning (ML) to engineer linguistic features as controls. First, unsupervised topic models recover issue domains (e.g., climate or crime), ruling out partisan gaps driven by topic ownership. Second, I recompute the verbatim accuracy measures after removing stopwords to mitigate the possibility that partisan gaps arise from ignorable tokens and their usage. Third, I add controls for linguistic coherence, competency, and tone. To further address endogeneity bias, I use two bounding analyses (Section 5). Overall, the partisan gap in accuracy remains robust to these controls. This integration of NLP-based accuracy scoring with econometric methods bridges computational methods and statistical inference, with the analyses uncovering a previously undetected dimension of slant in an ostensibly neutral media system.

Finally, drawing on qualitative interviews, I connect the empirical findings to the real-world media-politician dynamic. From these insights, the observed partisan gap in accuracy may not stem from a deliberate editorial slant but from logistical asymmetries, in which early pre-circulation of speech transcripts is mainly exercised by the ruling party. However, I theorize that this logistical asymmetry nonetheless embeds privately-held information about the existence of media bias.

2 Background

2.1 Political Background

Singapore has a unicameral parliament based on the Westminster system. Each parliament has a maximum term of five years, after which a general election is held using plurality voting.

Group and Single-Member Constituencies. Since the late 1980s, most parliament seats have been contested under group representation,[1]where voters in each group constituency vote for a group of candidates instead of any single candidate.[2] Groups comprise 3 to 6 candidates, with at least one from a minority ethnic group. I account for group size and ethnicity in the analyses.

Non-Elected Members of Parliament. By the early 1990s, two additional schemes were introduced. The Non-Constituency Member of Parliament scheme awards seats to the best-performing opposition candidates who lost. The Nominated Member of Parliament scheme appoints non-partisan members through a selection committee. Both types are accounted for in the models.

Political Dominance of the Ruling Party. The People's Action Party has dominated Singapore's politics as the governing party since 1959, facing little opposition. The 2011 election, however, marked a turning point. For the first time, an opposition party, the Workers' Party, won a group constituency en route to winning 6 of 87 seats. They secured the group with a 9.4% margin and a geographically adjacent single-member constituency by 29.6%. The 2005–2016 sample thus coincides with this political shift. Accordingly, this study focuses on the ruling-versus-opposition difference [32, 76], in contrast to other contexts that emphasize left-right divides [30, 37].

[1] Also known as multi-member districts elsewhere.

[2] For instance, in the 2011 general election, only 12 of the 87 parliament seats were contested as single members.

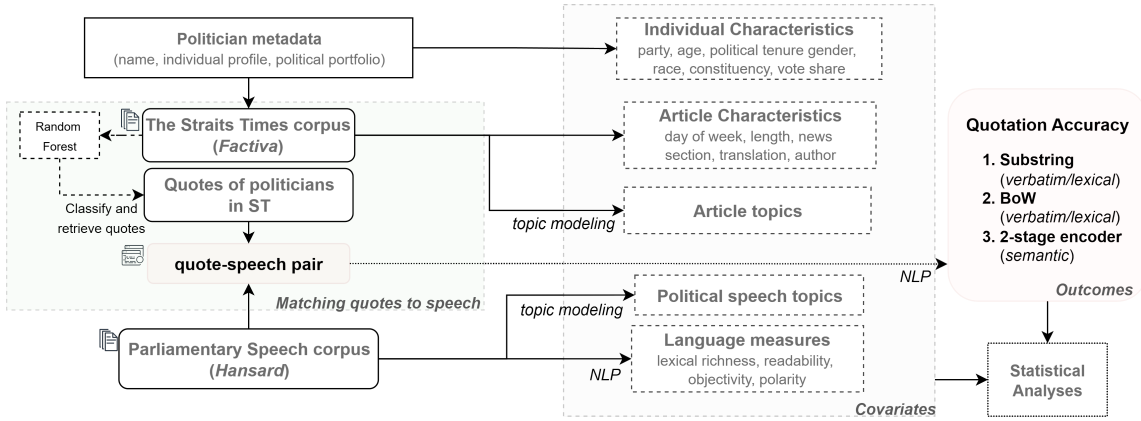

Fig. 1. Summary of pipeline to measure partisan coverage gap.

2.2 Mainstream Media

Mainstream media here primarily refers to daily newspapers, with *The Straits Times* standing out as Singapore's flagship daily. Owned by Singapore Press Holdings (SPH), it commands the largest readership, far exceeding that of other local dailies (Table A2) and is therefore the primary media of focus in this study.

The ruling party's dominance (Section 2.1) has allowed it to set media regulations. For instance, the Newspaper and Printing Presses Act (NPPA) requires that no newspaper be printed or published without a permit. Newspaper companies must be publicly listed with two classes of shares: ordinary and management. Holders of management shares are entitled to 200 votes compared to the one vote per ordinary shareholder [70]. These shares can only be transferred with government approval, creating perceptions of government-vetted nominees.[3]

SPH is publicly listed, with management shares held mainly by financial institutions and individuals with government ties [32]. Thus, while SPH is publicly listed, the NPPA enabled government influence over key appointments [32]. Moreover, trusted figures from state institutions have often held senior management and editorial leadership roles, including editor-in-chief, at The Straits Times [32]. This regulatory and ownership structure raises questions about whether profit maximization and media neutrality are compromised by ideological alignment [3, 8, 23, 32, 68]. It suggests that the media, while outwardly neutral and profit-driven, may reflect political connections and regulatory constraints that shape slant [32].

3 Methods

3.1 News and Speech Corpora

For *The Straits Times* news articles, I query Factiva using the names of politicians who served in parliament from 2005 to 2016. This period, spanning 12 years and four parliaments, including two full terms, encompasses the landmark 2011 election (Section 2.1) and offers a continuous window for examining quotation fidelity across a major turning point in Singapore's parliamentary politics. The query returns 62,132 unique news articles (title–date tuples) that mention at least one active politician during this period. Parliamentary transcripts are obtained from the Hansard repository, from which I automatically extract speaker–speech segments by parsing HTML tags (Appendix A).

3.2 Classification of Relevant Articles

Of the 62,132 articles containing mentions of the politicians, only a fraction contain relevant direct quotations from parliamentary speeches. To identify these, I train a *random forest* model on an initial training set of 1,419 manually labeled articles to predict which articles contain quotations from parliament. Each article is vectorized using n-grams ($n = 1$ to 5), supplemented with indicators for quotation marks and context-specific (e.g., parliament, bill, budget, debate, GRC). To address class imbalance, I apply weighted penalties to misclassifications of the minority class (articles with quotations). The model achieves 83% recall and 70% precision, yielding 3,421 articles (5.5% of the total) with at least one quotation from a parliamentary speech.

3.3 Matching Quotes to Speeches

Quotes are strings within single or double inverted commas or strings after speech colons. From the 3,421 identified articles, I extract quotations and match each to its originating parliamentary speech. Let f be the Ratcliff-Obershelp edit distance metric. For a quote q and a set of speeches $S = \{s_1, \ldots, s_L\}$, the matched speech is $s^* \in S$ which satisfies:

$$s^* = \arg\max_{s_\ell} \left\{ \mathcal{F}_1(q, s_\ell) \;\middle|\; s_\ell \in \{s_1, \ldots, s_L\} \right\},$$

where $\mathcal{F}_1$ is a composite function of f and is the substring accuracy measure (Equation (1)) defined in the following subsection.[4] This matching process yields 14,901 matched quote–speech pairs from 5,129 speeches by 204 politicians over 2005–2016.

[3] This concept of separate share classes is not unique. *The New York Times*, for instance, has class B shares for family owners with more voting power than class A shareholders.

[4] I manually oversee this matching by reviewing the best matches. Virtually all quotes are matched to a parliamentary speech the day before the print. This is consistent with news being *new* and explains why I do not include a control for the time gap between speeches and media pieces. Please see Niculae et al. for a better workflow to match quotes to speeches at scale [56].

3.4 Substring Quote Accuracy

The first measure, substring accuracy, scores accuracy according to the quote-length substring in the speech that best matches the quote. This substring measure merely formalizes how direct quotations work. If a quote is verbatim, there should be a quote-length substring in the speech that perfectly matches the quote (perfect score of 100).[5] Even if the quote is not verbatim, there should still be a quote-length substring in the speech that closely matches the quote.

Let q be the shorter string and s be the longer string, with lengths m and n, respectively. The substring accuracy locates the m-length substring within s that best matches q, scoring accuracy by this best partial match. Formally, the substring accuracy measure is:

$$\mathcal{F}_1(q,s) = \max_i \left\{ f\left(q, s_{(i:i+m-1)}\right) \,\middle|\, i \in \{1, \ldots, n-m+1\} \right\}. \tag{1}$$

The subscript $(i : i + m - 1)$ indicates the location of the substring of length m. $\mathcal{F}_1$ goes from 0 to 100, increasing in accuracy. Concretely, if q = "*a fundamental relook*" (20-character long) and the originating speech is s = "*… is taking a more fundamental relook at this regulation framework and see how best we can support this strategy …*", the measure searches all possible 20-character substrings in s that best matches q, which is [mor]"*e fundamental relook*", and then computes the edit distance between "*a fundamental relook*" and "*e fundamental relook*". This score forms the quotation accuracy score between the quote q and the speech s. Table 1 provides further examples.

3.5 Bag-of-Words Quote Accuracy

A different concern arises when quotes have words taken from other parts of the same speech, or when word order is not preserved. A concrete example is the quote "*… can afford their education in a public university*" from the speech "*… can afford to attend Government-funded universities ….*" Here, misquotation arises because *"Government-funded"* is replaced by *"public university"*, a term used elsewhere in the same speech that otherwise has the same meaning in context.

To allow an accuracy score that is forgiving of such cases, I define a second quote-accuracy measure: bag-of-words (BoW) accuracy. This measure is more tolerant of reordered or substituted phrases within the same speech that potentially retain equivalent meaning. I first define a new (third) string $c = q \cap s$ (with a slight abuse of notation) as the string containing words common to both q and s. Sorting the strings alphabetically yields $\widetilde{q}, \widetilde{s}$, and $\widetilde{c}$. The BoW accuracy score is then the maximum over all paired combinations of these strings:

$$\mathcal{F}_2(q,s) = \max \left\{ f\left(\widetilde{q}, \widetilde{s}\right), f\left(\widetilde{q}, \widetilde{c}\right), f\left(\widetilde{s}, \widetilde{c}\right) \right\}. \tag{2}$$

[5] Discounting punctuations that cannot be heard directly. I remove punctuations in preprocessing, as is standard.

Table 1. Examples of Quotes, Speeches, and Accuracy

Quote-speech pair		Accuracy score		
Quote	Originating parliamentary speech	Substring	BoW	Semantic
a fundamental relook	…is taking **a more fundamental relook** at this regulation framework and see how best we can support this strategy…	95	100	96
there will be a need for us to make sure we have regular fare increases of the right quantum	…So **there will be a need for us to make sure that we have regular fare increases of the right quantum**…	95	100	100
likely to go back down the slippery slope	…considering that some of them may be drop-outs or expelled from school, they are **likely to go back down the slippery road** …	90	92	99
core Singapore values	…Rather it is an acknowledgement that the **core Singaporean values** of multi-racialism and meritocracy can and should co-exist with each other…	90	100	99
resilience in response to ground reaction	…commended the PAP's **resilience in response to the ground's reaction** after the election, and I said this augurs well for Singapore…	85	91	100
too employer-focused	…and increase in absentee payroll are **rather employer-focused**…	85	89	99
the alternative will be too painful to bear	…So, let us not take our harmonious social fabric for granted because the alternative may be too painful to endure. …	84	75	0
That is the purpose of these amendments	…That is in fact **one of the purposes of this amendment** which we are bringing to Parliament…	80	100	95
deviates from the concept of free market	…this Bill **deviates from the idea of the concept of a free market**, where supply of services by companies is set by market demand …	80	100	100
the high end	…we noted that the rates were at **a higher end**, but we had the rates that were charged …	75	74	1
They (said) there are two purposes	…They gave the reasons that they wanted the two budget hotels **to serve two purposes**…	75	88	100
look them in the eye	…I want to be able to **look these men and others in the eyes** and say to them…	70	89	4
to help ensure all Singaporeans can afford their education in a public university	…To **ensure that all Singaporeans can afford to attend Government-funded universities**…	70	100	100
temporary or permanent solutions for this important issue	…to propose certain solutions, whether **temporary or permanent, to help resolve this, to me, a very important issue**…	65	100	93
completely unwarranted, alarmist, and show fundamental lack of understanding about the law	…I would venture to suggest that such statements are **alarmist and reveal a fundamental misunderstanding as to what this Bill and the law** is all about…	65	84	97

Besides edits for syntactic flow and grammatical demands, the BoW accuracy measure is more forgiving when the quote is not verbatim (Table 1). This could occur because of genuine lapses in the reporting journalist's attention or because the transcript contains a slightly different version of the speech.[6]

[6] As is customary, the computations remove punctuations and normalize all strings to lowercase (since punctuations and casings can neither be seen nor heard in a speech).

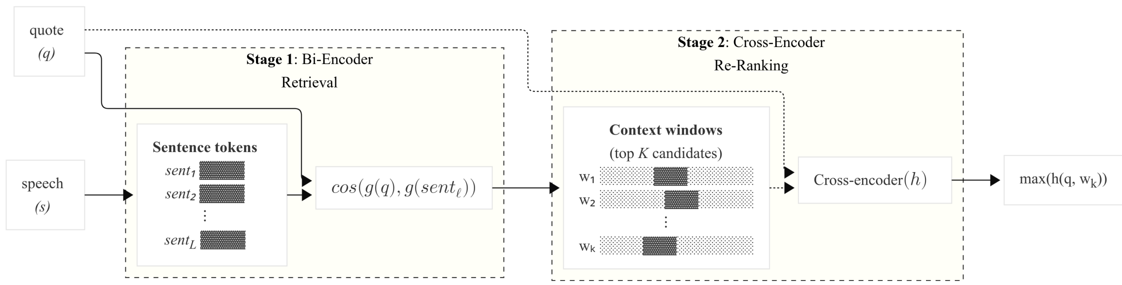

Fig. 2. Two-stage quote-to-speech semantic scoring.

3.6 Semantic Quote Accuracy

A third measure assesses whether a quote q semantically conveys the meaning of its originating speech s. I compute this using a two-stage transformer pipeline for candidate-passage retrieval and reranking [43, 47, 52, 57, 65].

Specifically, in the first stage, speeches are split into sentences $\mathcal{S} = \{\text{sent}_\ell\}_{\ell=1}^{L}$ (Figure 2). Each sentence and quote is (separately) mapped into 768-dimensional L2-normalized embeddings using the MPNET-based sentence-BERT biencoder g, which is pre-trained for semantic textual similarity and retrieval with high recall (*all-mpnet-base-v2*, ~109M parameters, [65, 72]). I then compute the cosine similarity, $\cos\left(g(q), g(\text{sent}_\ell)\right)$, of the quote and sentence embeddings, retaining the top ten candidates as seed sentences for reranking in the second stage.

Each candidate seed sentence at index ℓ is then expanded into a broader context window w_ℓ of up to ten neighboring sentences [47], subject to the constraint that the (q, w_ℓ) pair remains within the 512-token budget (including special tokens) of the cross-encoder model:

$$w_\ell \;=\; \text{concat}\left(\text{sent}_{\max(1,\ell-10)}, \ldots, \text{sent}_{\min(L,\ell+10)}\right).$$

I apply a cross-encoder h pre-trained and fine-tuned on the Microsoft Machine Reading Comprehension passage ranking task (*ms-marco-MiniLM-L12-v2*, ~33.4M parameters, [6, 65, 66, 79]). Its joint encoding enables token-level interactions that yield more precise judgments of contextual relevance over the narrowed search space [43, 47, 57].[7] The cross-encoder score is passed through a sigmoid layer σ, and rescaled to $[0, 100]$ (for consistency with the substring and BoW scores). The final score is the max score from among the ten candidates:

$$\mathcal{F}_3(q, s) \;=\; \max\left\{\sigma\big(h\left(q, w_\ell\right)\big)\right\}. \tag{3}$$

[7] The cross-encoder achieves near-state-of-the-art performance on MS MARCO passage reranking benchmarks with MRR@10 of 39.02 on the MS MARCO development set and NDCG@10 of 74.31 on TREC DL 2019, outperforming a comparable BERT-large model (*nboost/pt-bert-large-msmarco*: NDCG@10 = 73.36, MRR@10 = 36.48) and placing it at the top of its peers according to official cross-encoder model benchmarks (while being ~10*x* faster; https://web.archive.org/web/20250702081837/https://huggingface.co/cross-encoder/ms-marco-MiniLM-L12-v2#performance).

3.7 Content of Parliamentary Speeches and News Articles

To rule out coverage differences arising from underlying content differences, I apply a Latent Dirichlet Allocation (LDA, [12, 41, 64]) topic model, widely used in computational social science applications [77, 80]. With this, every parliament speech in the sample has a probabilistic association with one of K topics, where $\sum_K \rho_k = 1$ for each speech. The Dirichlet parameter α is set to low values, consistent with the fact that parliamentary speeches typically have only one dominant topic. The topic association of quotes is derived from the parliamentary speech topic model. Similarly, I train a separate topic model for the news article corpus. I preprocess the text to allow for joint terms. Topic number K is determined by maximizing topic coherence [17], evaluated across models from $K = 2, 4, \ldots, 100$, favoring lower values. The selected models are $K^* = 92$ for speeches and $K^* = 40$ for articles. Examples of topics (top five words per topic) learned from the speech corpus include:

(1) ⟨*cpf, retirement, minimum_sum, saving, cpf_saving*⟩
(2) ⟨*police, home_team, officer, crime, inquiry*⟩
(3) ⟨*premium, medishield_life, medishield, insurance, insurer*⟩
(4) ⟨*student, school, learn, education, teacher*⟩
(5) ⟨*fare, bus, public_transport, commuter, operator*⟩
(6) ⟨*sport, athlete, community, youth, sports*⟩

3.8 Additional Linguistic Measures

To further adjust for linguistic differences between politicians (Section 4.4), I consider measures of (i) objectivity, (ii) polarity, (iii) readability, and (iv) lexical richness. I compute objectivity and polarity measures for speeches using the *Pattern* sentiment analyzer (Python package `TextBlob`), which uses part-of-speech tagging. The readability measures (Python package `textatistic` [39]) are weighted averages of three textual information: (i) average words per sentence, (ii) average syllables per word, and (iii) the fraction of text made up of polysyllabic (three or more syllables) words—with readability decreasing in each of them. Lexical richness (Python package `lexicalrichness`) proxies for language sophistication, measured by the share of unique words—the higher this share, the richer the language [68, 69].

3.9 Dataset Statistics

The final dataset comprises 14,901 quote-speech pairs from 5,129 parliamentary speeches by 204 politicians (16 opposition, 44 non-partisan, 144 ruling party; Section 2.1) appearing in 3,421 articles over 2005–2016 (Appendix A). On average, quotes are 21 words long (standard deviation [SD] 39), with speeches 2,096 words long (SD 1,809), and 4.4 quotes per article (SD 3.1). The average quotation accuracy based on the three measures is: substring accuracy mean 91.5 (SD 12.8), BoW accuracy mean 96.7 (SD 10.2), and semantic accuracy mean 89.9 (SD 26.8). The 204 politicians include 47 women (23%), though female politicians account for a smaller share of coverage (17% of the quotes in the sample). Politicians average 49 years old (SD 8.3) with 7.5 years of tenure (SD 8.0).

4 Results

4.1 Detecting Partisan Gaps in Coverage

To estimate differences in coverage between politicians of the ruling and opposition parties, I estimate the ordinary least squares (OLS) model:

$$y_{iqrst} = \alpha + \beta\, opp_i + \sum_{k=2006}^{2016} \alpha_k\, year_{kt} + \sum_{\ell=11}^{13} \alpha_\ell\, parl_{\ell t} + \boldsymbol{\gamma}' \boldsymbol{X}_{iqrst} + \varepsilon_{iqrst}, \tag{4}$$

where y_{iqrst} is one of the three accuracy measures (Sections 3.4 to 3.6) for politician i, quote q from speech s reported in article r at year t. The coefficient of interest in β, which captures the partisan gap in quotation accuracy between opposition ($opp_i = 1$) and ruling-party ($opp_i = 0$) politicians. Standard errors are clustered at news articles.

β in Equation (4) does not carry an unambiguous causal interpretation since a candidate's choice of party is unlikely orthogonal to potential media coverage. To partially address endogeneity, Equation (4) includes covariates (Table 2). First, Equation (4) holds year and parliament constant to adjust for trends in media consumption or newsroom operations. To adjust for individual-level heterogeneity, Equation (4) includes politicians' gender and race, a quadratic in age and political tenure, and the interaction between politician type (e.g., minister or parliamentary secretary) and ministry portfolio (e.g., health or education). For news article controls, there are day-of-the-week dummies for variations in newsroom operations, news section (Table A3), a translation dummy for translated vernacular languages (as stated in the transcripts), and the (log of) speech, quote, and article word count. Finally, Equation (4) also adjusts for speech, quote, and article content (Section 3.7). A secondary model adjusts for election-related measures, including vote shares (Table 2), but for a smaller sample since some ruling-party politicians had walkovers (Section 2.1).

4.2 Quotation Intensity

Preliminary tabulations of the data by year suggest that opposition politicians in fact account for a larger share of quotations than their share of parliamentary seats (Figure C1a). For instance, in 2012, the opposition held 6 of 87 seats (6.9%) but accounted for 10.3% of the quotes. A similar pattern appears when aggregating by parliamentary term (Figure C1b).

Table 2. Description of Control Variables

A. Individual controls	
Gender	Gender of politician i
Race	Race of politician $i \in$ {Chinese, Indian, Malay, Eurasian/Others}
Age	Age of politician i at time t of speech s given
Tenure	Political tenure of politician i at time t of speech s given
B. Article controls	
Day of week	Dummies for article r publication day of week $\in$ {Mon, Tue, Wed, Thu, Fri, Sat, Sun}
Section	Dummies for article r's section $\in$ {*Singapore, Prime News, Top of the News, Home, ST, Insight, News, Money, Think, Review - insight, Sports, Opinion, World, Others* }
Translation	Dummy for a translation from vernacular to English in speech s
C. Topic distributions	
Speech topics	Topic distribution for speech s from model trained on speeches for $K = 92$
Quote topics	Topic distribution for quote q from model trained on speeches for $K = 92$
Article topics	Topic distributions for article r from model trained on articles for $K = 40$
D. Ministerial controls	
Type	Ministerial type of politician i at time t of speech s given
Portfolio	Ministerial portfolio of politician i at time t of speech s given
E. Electoral controls	
Group size	Number of politicians representing politician i's constituency $\in$ {1, 4, 5, 6}
Voters	Number of eligible voters in politician i's constituency
Votes (%)	Percentage of votes for politician i's constituency
Winner's share (%)	Percentage of valid votes for the winner in politician i's constituency

To formally test for partisan differences in coverage intensity, I estimate a model similar to Equation (4) at the article-speech level, comparing the length of the quotes per news article on a given speech. The estimates suggest that opposition and ruling-party politicians receive comparable levels of coverage. If anything, opposition quotes are longer, though the difference is not statistically significant (Table C1).

4.3 Quotation Accuracy

Table 3 reports the results for partisan gaps in accuracy. Panel A reports the substring accuracy outcome measure, which reflects the verbatim fidelity. The estimates in Panel A suggest that, on average, opposition politicians have quote snippets less accurate than those of ruling-party politicians. Panel B reports the bag-of-words accuracy measure, with the same conclusion–opposition quotes are less accurate. On average, opposition quotes have 1.4 to 2.1 points lower accuracy than ruling party quotes (10.7% to 21.0% of the standard deviations in accuracy). Relative to mean accuracies of 91.5 and 96.7 (out of 100), opposition quotes are 1.5–2.2% less accurate on the substring and BoW scales.

Panel C examines semantic accuracy, which quantifies whether the quote preserves the meaning of the speech beyond verbatim or surface-level similarity using a two-stage transformer model

Table 3. Difference in Quotation Accuracy, Opposition vs. Ruling-Party Politicians

	A. Substring accuracy		B. BoW accuracy		C. Semantic accuracy	
	(1)	(2)	(3)	(4)	(5)	(6)
Opposition	-1.858^{***}	-1.373^{*}	-2.149^{***}	-1.638^{**}	-4.053^{***}	-2.992^{**}
	(0.711)	(0.760)	(0.700)	(0.645)	(1.458)	(1.435)
Time fixed-effects	✓	✓	✓	✓	✓	✓
Individual controls	✓	✓	✓	✓	✓	✓
Article controls	✓	✓	✓	✓	✓	✓
Topic controls	✓	✓	✓	✓	✓	✓
Ministerial controls	✓	✓	✓	✓	✓	✓
Electoral controls	–	✓	–	✓	–	✓
R^2	0.180	0.204	0.110	0.117	0.230	0.270
F-statistic, time fixed-effects	4.69^{***}	2.99^{***}	4.24^{***}	2.07^{**}	3.94^{***}	2.5^{**}
F-statistic, individual controls	2.04^{**}	.82	1.88^{*}	.56	1.17	1.43
F-statistic, topic controls	1.66^{***}	1.55^{***}	1.25^{**}	1.26^{**}	1.26^{**}	1.2^{**}
F-statistic, ministerial controls	3.81^{***}	2.93^{***}	2.4^{***}	1.93^{***}	3.01^{***}	2.29^{***}
F-statistic, electoral controls	–	.89	–	1.82	–	3.23^{**}
Mean of dependent variable	91.5	91.9	96.7	96.8	89.9	90.3
N	14,901	10,898	14,901	10,898	14,901	10,898

Clustered standard errors in parentheses. ***, **, * indicate significance at the 1%, 5%, and 10% levels.

that assesses semantic support (Section 3.6). This semantics-based measure reinforces the same opposition disadvantage. On average, the opposition's quotes are 3.0 to 4.1 points less semantically faithful (11.2% to 15.1% of the standard deviation). Across all three measures, including electoral controls (even-numbered columns) does not change the findings.

All models include year and parliament trends (Section 4.1), so estimated partisan gaps are net of secular time trends. However, given that the sample spans over a decade, those average differences might drift over time. Interacting the party status with year shows that the partisan gap in accuracy was widest in the early half of the period and narrows thereafter, while persisting through the end of the sample (Figure 3).

4.4 Sensitivity Analyses

I run additional estimations to test the sensitivity of the results in Section 4.3. These include adjustments for journalists, clustering standard errors by speeches and journalists, and using alternative topic models with various K. I also recompute accuracy excluding stopwords to ensure journalists are not simply dropping filler words for opposition quotes. These do not change the finding that opposition quotes are less accurate (Table D1). Linguistic differences in opposition speeches may also hinder journalists' comprehension. I therefore add controls for: (i) objectivity, (ii) polarity, (iii) readability, and (iv) lexical richness (Section 3.8). If linguistic differences drive

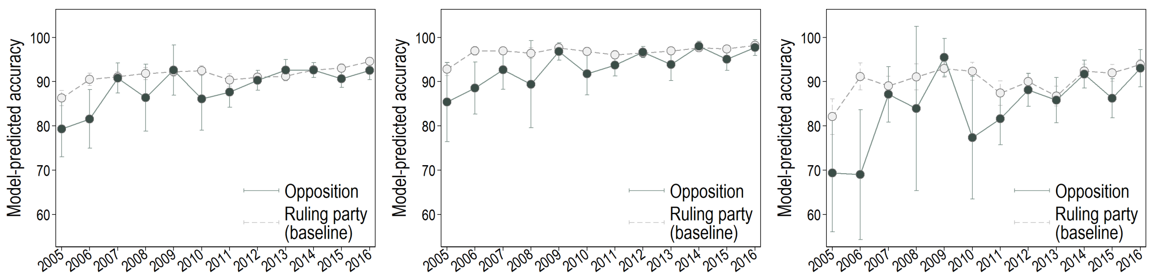

(a) Substring (b) BoW (c) Semantic

Fig. 3. Predicted Quote Accuracy by Party and Year. Points represent predicted marginal effects for the year from the interaction of party with year in the regression specification Equation (4). Models include controls for those listed in Table 3 (odd-numbered columns). Capped vertical lines are 95% confidence intervals.

the results, the estimated partisan gap should attenuate to zero with these controls. This is not the case (Table D2).[8] More generally, Figure 4 shows that the estimates from Section 4.3 for the substring measure are modest compared to alternative specifications, demonstrating that they are not artifacts of a single specification that yields large estimated differences in coverage. Figure D1 and Figure D2 present analogous results for the BoW and semantic accuracy measures.

5 Bounding Analyses

I turn to two bounding analyses to deal with compositional bias and selection bias in the OLS estimates. The aim is to demonstrate that the estimates from Section 4.3 likely understate the partisan gap in accuracy.

5.1 Ministerial Composition

First, there is a compositional bias even if party status was randomly assigned, since opposition status limits ministerial composition. As the opposition, politicians do not hold ministerial ranks and portfolios (e.g., Education or Health), meaning party status influences both media coverage and ministerial characteristics.[9]

Let the opposition dummy be $D_i \in \{0, 1\}$, and a rank dummy be $r_i \in \{0, 1\}$. y_{Di} and r_{Di} are the potential coverage and rank by party status. The observed difference in coverage between opposition politicians ($D = 1$) and ruling-party politicians ($D = 0$), at base (ministerial) rank ($r = 0$) can be decomposed (Appendix E, [5]):

$$\mathbb{E}[y_{1i} - y_{0i} | r_{1i} = 0] + \mathbb{E}[y_{0i} | r_{1i} = 0] - \mathbb{E}[y_{0i} | r_{0i} = 0]. \tag{5}$$

[8] However, accuracy degrades with poorer lexical richness and readability (Table D2).
[9] Unlike the typical Westminster system, there is no shadow cabinet.

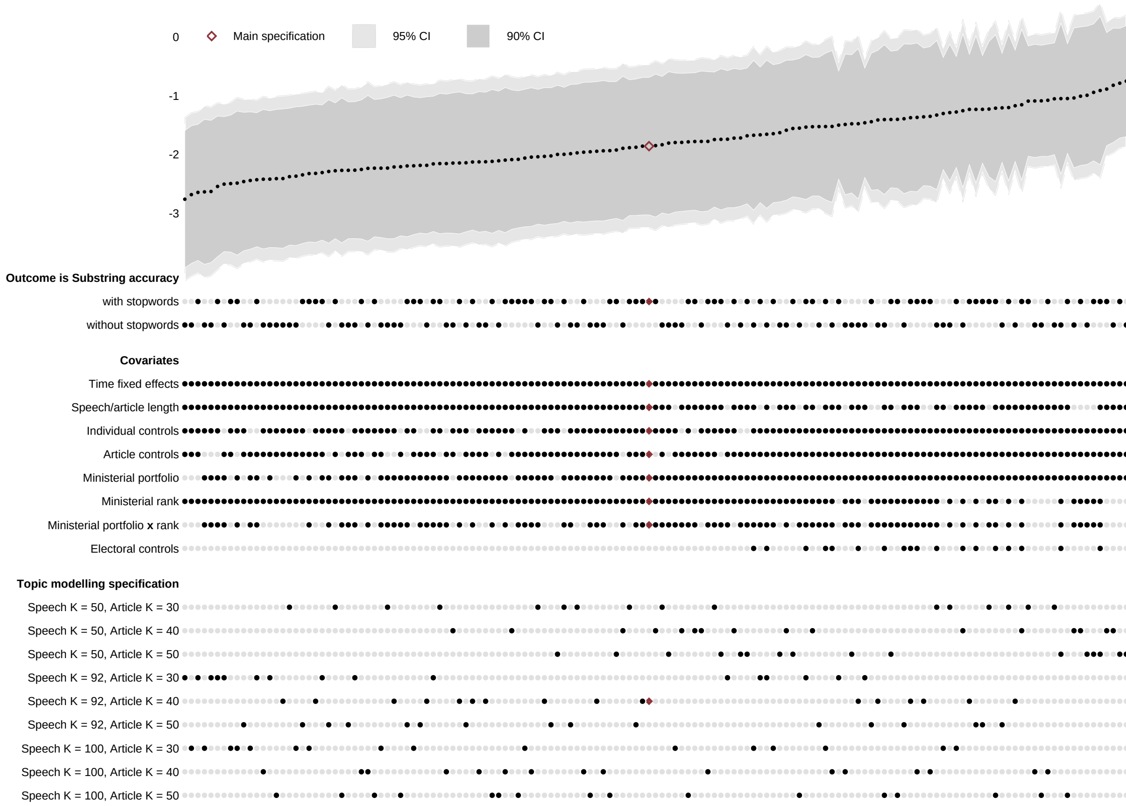

Fig. 4. Effect Sizes of Opposition Status on Substring Accuracy. Figure shows estimates, plotted in ascending order, across 144 specifications (2 substring accuracy measures, with/without substrings, × 8 covariate combinations, × 9 possible topic models). The chart indicates the above combinations. Combinations under the outcome and topic models are mutually exclusive, while covariate combinations are not. Shaded bands indicate 90% and 95% confidence intervals.

The first term reflects the true media slant between opposition and ruling-party politicians of the base rank. The compositional bias arises from the last two terms. The second term corresponds to the coverage of a ruling-party politician if assigned the base rank as an opposition. Since this applies to all ruling-party politicians (oppositions only hold base rank), it simplifies to the unconditional average: $\mathbb{E}[y_{0i}|r_{1i} = 0] = \mathbb{E}[y_{0i}]$. The third term represents a ruling-party politician's potential coverage at the base rank ($r = 0$), which arguably represents a lower-caliber or more junior politician. Hence, the last two terms evaluate to a positive bias (Appendix E; [68]).

The above implies that the OLS estimates from Section 4.3 constitute lower bounds on slant if:

(i) The true effect is negative, indicating media slant towards the ruling party; and

(ii) Lower-ranked ruling-party politicians receive less favorable coverage than the average ruling-party politician.

5.2 Selection on Observables and Unobservables

Second, if party assignment (ruling party vs opposition) is not conditionally random, a negative $\hat{\beta}$ may indicate selection bias rather than true slant. In particular, more capable candidates might select into the ruling party, leaving lower-caliber candidates in the opposition [76]. Consequently, opposition coverage may appear weaker simply because media attention tracks politicians' capabilities.

Since caliber is unobserved, the second bounding argument uses the proportional selection assumption [2, 58], bounding OLS estimates using how selection on observables relate to selection on unobservables [58]:

$$\beta^* = \tilde{\beta} - \left[\overset{o}{\beta} - \tilde{\beta}\right] \frac{R_{max} - \tilde{R}}{\tilde{R} - \overset{o}{R}}, \tag{6}$$

where $\tilde{\beta}$ and $\tilde{R}$ denote the coefficient and R^2 from the fully controlled regression, and $\overset{o}{\beta}$ and $\overset{o}{R}$ are from the uncontrolled model. R_{max} is set to $1.3\tilde{R}$ [58]. The idea is that the change in $\hat{\beta}$ when adding controls (observables) is informative on how unobservables might bias the estimate. Equation (6) scales the change in $\hat{\beta}$ by the relative change in $\hat{R}$, reflecting the potential explanatory power of the observables. Adjusting $\tilde{\beta}$ based on the above yields the bias-adjusted estimate β^*.

Assuming equal selection between observables and unobservables implies that the bias-adjusted slant is -2.88, -2.78, -7.04 (compared to -1.86, -2.15, -4.05 from Table 3). Moreover, assuming proportional selection between observables and unobservables, for the estimated partisan gap to be zero, the unobservables would need to be up to 16 times more important than the observables. Hence, adjusting for selection bias further supports that the estimated partisan gaps in accuracy are conservative in magnitude.

6 Discussion

This study bridges a computational approach with econometric analyses to examine the accuracy of political speech coverage. Applying the approach to Singapore—where political media bias is less overt, and coverage intensity shows no partisan gap—provides a useful testbed. I find no evidence that opposition politicians receive less coverage. This is consistent with accounts of how the media generally presents a veil of neutrality. However, as this study demonstrates, neutrality has gradations. While coverage intensity appears balanced, the accuracy measures developed in this study reveal a systematic partisan gap. This finding holds across verbatim and semantic measures, indicating lower accuracy for the opposition on both grounds of literal transcription and semantic faithfulness. These findings persist in a wide range of sensitivity analyses, including models that adjust for speech and linguistic qualities. Finally, two bounding analyses suggest the estimated partisan gaps in accuracy likely understate the true gap.

6.1 Related Work

These findings are consistent with media regulations that incentivize greater caution when quoting the ruling party, given its legal leverage. Journalists, especially in dailies, may trade off accuracy for timeliness [10], and inaccuracies in opposition quotes may stem from carelessness or synonym substitutions rather than intent. The results also fit demand-driven models of media bias, where bias reflects consumer preferences [9, 29, 30, 35, 44]. As media compete along a political spectrum for readership [42, 55], the rise of social media as an indirect competitor, with consumers more sympathetic to the opposition, may edge mainstream outlets towards the establishment.

This study connects to the political economy literature that estimates the size and direction of media bias without anecdotal evidence. It connects to research on the media in dominant single-party systems [25, 53, 62, 63], and to studies finding slight left-of-center bias in centrist US media [9, 37, 40, 75]. One major stream is coverage intensity, in which differences in the volume of coverage across parties and issues reveal media slant [44, 59, 61, 63]. Another influential approach focuses on language by measuring ideological proximity through the media-politician language overlap [30, 31]. By comparing the linguistic choices of politicians and news outlets, this line of work reveals that the media adopts partisan-coded language (e.g., Republican language such as death tax, war on terror vs. Democratic language such as estate tax, war in Iraq), capturing semantic patterns beyond coverage volume alone. These approaches move beyond anecdote and scale bias detection using observable coverage patterns and textual similarity.

This work aligns further with computational and network-based approaches to bias detection [7, 15, 48, 51, 54, 56, 73, 77, 81, 82]. These studies apply traditional NLP and transformer-based methods to capture bias in subtler forms [73], across diverse units of analysis, including political blogs [81], newspapers [7, 19], and news headlines [48]. Some also integrate computational measures into statistical tests of hypotheses [7]. Other studies specifically target framing bias, such as one showing that headlines after mass shootings in the US politicized gun-related issues in different ways (e.g., school safety, gun control, mental health), demonstrating how news outlets can cover the same events in different ways [48]. A related study on quotation patterns analyzes quoting patterns of President Obama using low-rank embeddings, revealing that different media outlets selectively cover different parts of the same speech, thereby influencing public perceptions without necessarily misquoting [56]. Relatedly, semantic-embedding pipelines have been used to quantify how much downstream social media discourse reflects press releases, demonstrating the use of embedding-based methods to capture similarity beyond word alignment [16].

Finally, LLMs are among the most recent wave of tools in media bias research [15, 73]. While promising, their reliability remains in question as a recent study finds that four widely used LLMs (GPT, Gemini, LLaMA, Claude) produce inconsistent ideological mappings of newspapers worldwide, hinting that bias detection should also account for context and institutional settings (Section 6.3) [15].

6.2 Coverage Fidelity

This study extends existing work by applying a two–stage retrieval–reranker pipeline [43, 47, 52, 57, 65] to quantify the semantic support of direct quotes from the source speeches. This application differs from standard neural retrieval frameworks, which are typically designed to find relevant documents in large corpora [38]. By contrast, this study re-purposes these models to assess semantic fidelity on known quote-source pairs. In a similar vein, existing fuzzy matching techniques are used to assess verbatim fidelity. By combining these accuracy measures with econometric methods, this approach uncovers statistical evidence of slant in an institutional setting where coverage intensity measures cannot.

Each of the three accuracy measures captures a different aspect of coverage fidelity. Substring accuracy is closest in spirit to verbatim quotation accuracy as traditionally defined: a match of contiguous words within the source speech. Exact wording is required, and even small, innocuous deviations (e.g., "college" $\longleftrightarrow$ "university") reduce the score. The bag-of-words measure is more forgiving when quotes are truthful but paraphrased or stitched from non-contiguous fragments. Semantic accuracy uses a transformer-based two-stage model to assess whether the speech supports the meaning of the quote. It will handle paraphrasing and synonyms through semantic representation. In particular, the semantic measure identifies quotes that are close to verbatim but could otherwise be interpreted as taken out of context as soundbites without substantive meaning (e.g., "look them in the eye"). In practice, the choice of the accuracy metric depends on how one weighs verbatim vs. semantic meaning, and the tolerance for paraphrasing. In this study, the consistency in findings across measures strengthens confidence that the partisan gap reflects systematic differences rather than artifacts of any single metric.

6.3 Logistics and a Separating Equilibrium

Here, I situate the findings within insights from primary and secondary research on a specific government-media communications machinery. The insights come from interviews with a senior journalist and two members of parliaments from recent terms. Together, their accounts suggest that the statistical evidence of a partisan gap in accuracy reflects a "separating equilibrium" where the stable equilibrium strategy in engaging the media differs by party (Figure 5).

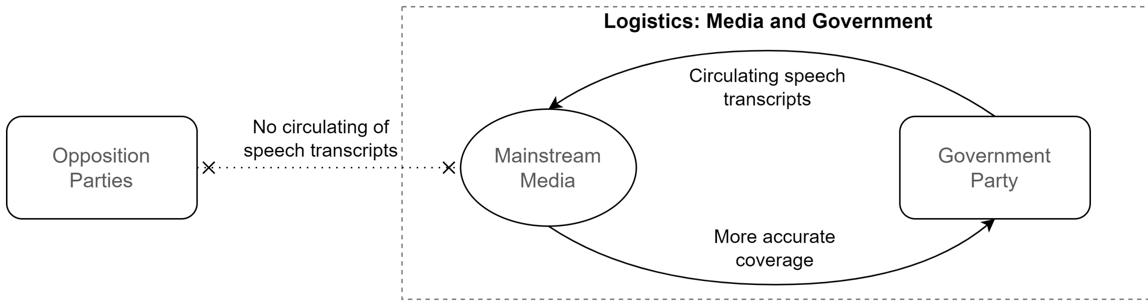

Fig. 5. Institutional Logistics and a Separating Equilibrium

Ruling-party politicians circulate their speech transcripts to the media in advance. Higher-ranking politicians introducing Bills or Motions not only share their speeches early but may also brief the media in advance. This pre-circulation gives media agents ample time to prepare. Non-senior ruling-party politicians also share their speeches in advance, or the media might request them.

By contrast, opposition politicians generally *do not* participate in this government-media machinery. They are less likely to voluntarily submit their transcripts after the sitting. If they do, it's often late. Since official transcripts are released only a week after the sitting, but news stories appear the same or the next day, journalists rely on shorthand notes, video recordings, or pre-circulated transcripts, mainly from ruling-party politicians (Appendix B). If early circulation improves accuracy, these partisan choices alone on media engagement could explain much of the observed partisan gap in quotation accuracy. Still, even if the entire partisan gap stems from differences in media engagement, the differences in media engagement reveal strong private beliefs about media slant as I discuss in Section 6.4 below.

6.4 Separating Equilibrium via Private Beliefs of Slant

Political agents care about getting covered accurately. They also care about spin in the framing of their messages [33, 48, 54, 77, 78]. As an illustration, a speech about "... raise taxes by x% to improve fiscal position ..." can be (mis)represented in different ways. One version is simply "raise taxes by x%". Another is "raise taxes by x% to stabilise fiscal position". The speech itself can have different spins. It can be framed as the government lacking fiscal prudence in recent years, or as a story about government foresight regarding future spending needs. The same message can also be threaded into a pre-existing narrative about social spending (or lack thereof). Spin increases with early circulation (Section 6.3) because it gives media agents time to enrich or distort the narrative. Thus, early circulation improves accuracy but also increases the risk of narrative distortion.

Being rational, politicians weigh the benefits of accuracy against the loss of narrative control, based on their private beliefs about the probability of spin, which differ by party (Appendix F). Ruling-party politicians often circulate their speeches early, implying the belief that the benefits of accuracy outweigh any potential negative spin. Opposition politicians typically do not, consistent with the reservation that giving the media early access will lead to unfavorable interpretations of the message.

Hence, even if differences in logistics and engagement fully explain disparities in accuracy (Section 6.3), the decision to not circulate still reveals strong privately held information of a media that slants.

6.5 Strengths and Limitations

A limitation of the study is that the analysis ends before more recent legislative periods, when opposition representation became more of a norm. Nevertheless, the sample spans twelve years and four parliaments, including two full terms, and crucially captures the landmark 2011 election when opposition representation made a breakthrough. The sample, therefore, provides a meaningful analytical window into partisan gaps in quotation fidelity during a major institutional turning point. The analyses also ensure that estimated partisan gaps do not hinge on any single year, as all models include year and parliament effects [27, 50]. Additionally, the period is sufficiently broad to capture the temporal dynamics of partisan gaps. What remains uncertain is how these gaps have evolved more recently, especially as the media landscape itself has shifted. Print readership has halved between 2017 and 2022, while television and social media have become more dominant news sources (Table A2, [24]), Examining partisan gaps in newer periods thus represents an important avenue of future work.

A second limitation is that the sample focuses on a single outlet, *The Straits Times*. This design is deliberate. As Singapore's flagship daily with the widest readership and well-documented leadership ties to the state [32], *The Straits Times* provides the clearest platform for evaluating journalistic behavior in political coverage. Smaller outlets exist, but with reach orders of magnitude lower. In this sense, any inaccuracies from *The Straits Times* would have the greatest potential impact on news consumers and the electorate at large.

At the same time, examining a single outlet in depth enables a closer study of institution-specific dynamics (Section 6.3). The observed partisan gap could stem from media slant or from a government-media machinery that excludes the opposition. The empirical approach cannot fully disentangle these explanations. However, I show that the choices made by ruling-party and opposition politicians still imply media slant (Section 6.4), and underscore that computational and applied statistics cannot substitute for attention to the institutional context [15].

Beyond *The Straits Times* and Singapore, the methodological approach is broadly applicable, as the accuracy measures can be computed wherever a media–source corpus exists (Figure 1). Extending it to other domestic outlets, including online and social media, where news consumption has shifted [24], would yield valuable comparisons. In more pluralistic political and media systems, the approach remains useful because the notion of an "opposition" is relative: for a left-of-center outlet, right-of-center politicians are the opposition, and vice versa [14]. This relativity is particularly salient in contexts where ideological leanings of outlets are already well documented [21, 30, 37, 56, 60, 82]. Moreover, the institutional machinery in Singapore (Section 6.3) also applies in other settings, where politicians cultivate long-term working ties with ideologically friendly journalists [78]. In such settings, identifying whether an outlet covers Party A less often than Party B, and, crucially, whether it does so with equal fidelity, provides a richer indicator of news reliability.

Computationally, the LDA topic modeling approach, while common and easily implemented [19, 56, 77, 80], has newer counterparts. This is less of a concern as the goal was to adjust for variation arising from differing speech content [27, 50, 80]. Nonetheless, newer topic modeling techniques can recover more coherent and contextually aware topics [4, 22, 36, 67, 74]. These will be useful when content is studied as an essential dimension [48] rather than a nuisance variation to partial out [27, 50].

Finally, while this study uncovers partisan gaps in quotation accuracy, it does not assess its impact on readers and the electorate [20, 34]. While out of scope of bias detection studies [9, 30, 37, 40, 44, 48, 49, 56, 59–61, 63, 75, 77], the impact of misquotation remains a critical avenue of research since news coverage affect perceptions [18, 71], media bias spreads to other outlets [80], and (mis)quotations can subtly distort reality [56]. Future studies should evaluate, through surveys or experiments, how misquotations or, more generally, coverage infidelity would affect public opinions.

6.6 Conclusion

Media slant operates through multiple dimensions, including fidelity in conveying political speeches. This study traces how political speeches change as they move from parliament to print. In an institutional setting widely regarded as neutral, the computational and statistical approach reveals subtle but systematic gradations in coverage accuracy along partisan lines. The framework generalizes to any other context with quote-to-source text pairs. More broadly, established computational methods can reveal hidden patterns when applied beyond their original domains.

Acknowledgments

I am grateful to a senior journalist from the national print media and a member of parliament for sharing relevant personal insights and experiences on parliamentary logistics and media engagement. For detailed comments, I thank Madhav Shrihari Aney, Wai Mun Chia, Yohanes Eko Riyanto, Chris Sakellariou, Gaurav Sood, You Yenn Teo, Walter Edgar Theseira.

To the late Giovanni Ko, my PhD supervisor, for seeding the idea to this study and imparting countless invaluable lessons along the way, particularly on the art and science of political economy and empirical inquiry.

Code and Data

Code and replication materials are available at https://github.com/LSYS/neutrality.

References

[1] Maja Adena, Ruben Enikolopov, Maria Petrova, Veronica Santarosa, and Ekaterina Zhuravskaya. 2015. Radio and the Rise of The Nazis in Prewar Germany. *The Quarterly Journal of Economics* 130, 4 (2015), 1885–1939. https://doi.org/10.1093/qje/qjv030

[2] Joseph Altonji, Todd Elder, and Christopher Taber. 2005. Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools. *Journal of Political Economy* 113, 1 (2005), 151–184. https://doi.org/10.1086/426036

[3] Simon P Anderson and John McLaren. 2012. Media Mergers and Media Bias With Rational Consumers. *Journal of the European Economic Association* 10, 4 (2012), 831–859. https://doi.org/10.1111/j.1542-4774.2012.01069.x

[4] Dimo Angelov. 2020. Top2Vec: Distributed Representations of Topics. arXiv:2008.09470 [cs.CL]

[5] Joshua D. Angrist and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton University Press, Princeton, NJ.

[6] Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Majumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, Mir Rosenberg, Xia Song, Alina Stoica, Saurabh Tiwary, and Tong Wang. 2018. MS MARCO: A Human Generated MAchine Reading COmprehension Dataset. arXiv:1611.09268 [cs.CL] https://arxiv.org/abs/1611.09268

[7] Ramy Baly, Georgi Karadzhov, Dimitar Alexandrov, James Glass, and Preslav Nakov. 2019. Multi-Task Ordinal Regression for Jointly Predicting the Trustworthiness and the Leading Political Ideology of News Media. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, Florence, Italy, 210–221. https://doi.org/10.18653/v1/N19-1216

[8] BBC News. 2013. Singapore profile - Media. http://www.bbc.com/news/world-asia-15966553

[9] Dan Bernhardt, Stefan Krasa, and Mattias Polborn. 2008. Political Polarization and the Electoral Effects of Media Bias. *Journal of Public Economics* 92, 5 (2008), 1092–1104. https://doi.org/10.1016/j.jpubeco.2008.01.006

[10] Fred C Berry. 1967. A Study of Accuracy in Local News Stories of Three Dailies. *Journalism Quarterly* 44, 3 (1967), 482–490. https://doi.org/10.1177/107769906704400309

[11] Timothy Besley and Andrea Prat. 2006. Handcuffs for the Grabbing Hand? Media Capture and Government Accountability. *American Economic Review* 96, 3 (2006), 720–736. https://doi.org/10.1257/aer.96.3.720

[12] David Blei, Andrew Y Ng, and Michael I. Jordan. 2003. Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3 (2003), 993–1022. https://doi.org/10.1162/jmlr.2003.3.4-5.993 arXiv:1111.6189v1

[13] Taylor C Boas and F Daniel Hidalgo. 2011. Controlling the Airwaves: Incumbency Advantage and Community Radio in Brazil. *American Journal of Political Science* 55, 4 (2011), 869–885. https://doi.org/10.1111/j.1540-5907.2011.00532.x

[14] Wayne Brittenden. 2000. The Political Compass. https://politicalcompass.org. Online political self-test positioning individuals on two axes: economic (left–right) and social (authoritarian–libertarian). First released December 20, 2000; active as of 2025.

[15] Alessio Buscemi and Daniele Proverbio. 2024. Large Language Models' Detection of Political Orientation in Newspapers. arXiv:2406.00018 [cs.CL] https://arxiv.org/abs/2406.00018

[16] Tristan J. B. Cann, Ben Dennes, Travis Coan, Saffron O'Neill, and Hywel T. P. Williams. 2025. Using semantic similarity to measure the echo of strategic communications. *EPJ Data Science* 14, 1 (2025), 20. https://doi.org/10.1140/epjds/s13688-025-00538-w

[17] Jonathan Chang, Jordan Boyd-Graber, Sean Gerrish, Chong Wang, and David M. Blei. 2009. Reading tea leaves: how humans interpret topic models. In *Proceedings of the 22nd International Conference on Neural Information Processing Systems* (Vancouver, British Columbia, Canada) *(NIPS'09)*. Curran Associates Inc., Red Hook, NY, USA, 288–296. https://doi.org/10.5555/2984093.2984126

[18] Chun-Fang Chiang and Brian Knight. 2011. Media Bias and Influence: Evidence from Newspaper Endorsements. *The Review of Economic Studies* 78, 3 (2011), 795–820. https://doi.org/10.1093/restud/rdq037

[19] Christopher Congleton, Peter van der Putten, and Suzan Verberne. 2022. Tracing Political Positioning of Dutch Newspapers. In *Disinformation in Open Online Media*, Francesca Spezzano, Adriana Amaral, Davide Ceolin, Lisa Fazio, and Edoardo Serra (Eds.). Springer International Publishing, Cham, 27–43.

[20] Dave D'Alessio. 2003. An Experimental Examination of Readers' Perceptions of Media Bias. *Journalism & Mass Communication Quarterly* 80, 2 (2003), 282–294. https://doi.org/10.1177/107769900308000204

[21] Stefano DellaVigna and Ethan Kaplan. 2007. The Fox News Effect: Media Bias and Voting. *Quarterly Journal of Economics* 122, 3 (2007), 1187–1234. https://doi.org/10.1162/qjec.122.3.1187

[22] Adji Bousso Dieng, Francisco J R Ruiz, and David M Blei. 2020. Topic Modeling in Embedding Spaces. *Transactions of the Association for Computational Linguistics* 8 (2020), 439–453. https://doi.org/10.1162/tacl_a_00325

[23] Ruben Durante and Brian Knight. 2012. Partisan Control, Media Bias, and Viewer Responses: Evidence from Berlusconi's Italy. *Journal of the European Economic Association* 10, 3 (2012), 451–481. https://doi.org/10.0.4.87/j.1542-4774.2011.01060.x

[24] Edson C. Tandoc Jr. 2022. Digital News Report 2022: Singapore. https://reutersinstitute.politics.ox.ac.uk/digital-news-report/2022/singapore.

[25] Ruben Enikolopov, Maria Petrova, and Ekaterina Zhuravskaya. 2011. Media and Political Persuasion: Evidence from Russia. *American Economic Review* 101, 7 (2011), 3253–3285. https://doi.org/10.1257/aer.101.7.3253

[26] Claudio Ferraz and Frederico Finan. 2008. Exposing Corrupt Politicians: The Effects of Brazil's Publicly Released Audits on Electoral Outcomes. *The Quarterly Journal of Economics* 123, 2 (2008), 703–745. https://doi.org/10.1162/qjec.2008.123.2.703

[27] Ragnar Frisch and Frederick V. Waugh. 1933. Partial Time Regressions as Compared with Individual Trends. *Econometrica* 1, 4 (1933), 387–401. https://doi.org/10.2307/1907330

[28] Matthew Gentzkow. 2006. Television and Voter Turnout. *The Quarterly Journal of Economics* 121, 3 (2006), 931. https://doi.org/10.1162/qjec.121.3.931

[29] Matthew Gentzkow and Jesse M Shapiro. 2006. Media Bias and Reputation. *Journal of Political Economy* 114, 2 (2006), 280–316. https://doi.org/10.1086/499414
[30] Matthew Gentzkow and Jesse M Shapiro. 2010. What Drives Media Slant? Evidence From U.S. Daily Newspapers. *Econometrica* 78, 1 (2010), 35–71. https://doi.org/10.3982/ECTA7195
[31] Matthew Gentzkow, Jesse M. Shapiro, and Matt Taddy. 2019. Measuring Group Differences in High-Dimensional Choices: Method and Application to Congressional Speech. *Econometrica* 87, 4 (2019), 1307–1340. https://doi.org/10.3982/ECTA16566
[32] Cherian George. 2012. *Freedom From the Press: Journalism and State Power in Singapore.* Singapore : NUS Press, c2012., Singapore.
[33] Alan S Gerber, James G Gimpel, Donald P Green, and Daron R Shaw. 2011. How Large and Long-lasting Are the Persuasive Effects of Televised Campaign Ads? Results from a Randomized Field Experiment. *American Political Science Review* 105, 1 (2011), 135–150. https://doi.org/10.1017/S000305541000047X
[34] Rhonda Gibson and Dolf Zillmann. 1993. The Impact of Quotation in News Reports on Issue Perception. *Journalism Quarterly* 70, 4 (1993), 793–800. https://doi.org/10.1177/107769909307000405
[35] Tim Groeling and Samuel Kernell. 1998. Is Network News Coverage of the President Biased? *The Journal of Politics* 60, 4 (1998), 1063–1087. https://doi.org/10.2307/2647731
[36] Maarten Grootendorst. 2022. BERTopic: Neural topic modeling with a class-based TF-IDF procedure. https://arxiv.org/abs/2203.05794
[37] Tim Groseclose and Jeffrey Milyo. 2005. A Measure of Media Bias. *Quarterly Journal of Economics* 120, 4 (2005), 1191–1237. https://doi.org/10.1162/003355305775097542
[38] Jiafeng Guo, Yinqiong Cai, Yixing Fan, Fei Sun, Ruqing Zhang, and Xueqi Cheng. 2022. Semantic Models for the First-Stage Retrieval: A Comprehensive Review. *ACM Trans. Inf. Syst.* 40, 4, Article 66 (March 2022), 42 pages. https://doi.org/10.1145/3486250
[39] Erin Hengel. 2022. Publishing While Female: are Women Held to Higher Standards? Evidence from Peer Review. *The Economic Journal* 132, 648 (05 2022), 2951–2991. https://doi.org/10.1093/ej/ueac032 arXiv:https://academic.oup.com/ej/article-pdf/132/648/2951/46859082/ueac032.pdf
[40] Daniel E Ho and Kevin M Quinn. 2008. Measuring Explicit Political Positions of Media. *Quarterly Journal of Political Science* 3, 4 (2008), 353–377. https://doi.org/10.1561/100.00008048
[41] Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. spaCy: Industrial-strength Natural Language Processing in Python. https://doi.org/10.5281/zenodo.1212303
[42] Harold Hotelling. 1929. Stability in Competition. *The Economic Journal* 39, 153 (1929), 41. https://doi.org/10.2307/2224214
[43] Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense Passage Retrieval for Open-Domain Question Answering. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Bonnie Webber, Trevor Cohn, Yulan He, and Yang Liu (Eds.). Association for Computational Linguistics, Online, 6769–6781. https://doi.org/10.18653/v1/2020.emnlp-main.550
[44] Valentino Larcinese, Riccardo Puglisi, and James M Snyder. 2011. Partisan Bias in Economic News: Evidence on the Agenda-setting Behavior of U.S. Newspapers. *Journal of Public Economics* 95, 9 (2011), 1178–1189. https://doi.org/10.1016/j.jpubeco.2011.04.006
[45] Horacio A Larreguy, John Marshall, and James M Snyder Jr. 2014. Revealing Malfeasance: How Local Media Facilitates Electoral Sanctioning of Mayors in Mexico. *National Bureau of Economic Research Working Paper Series*

No. 20697 (2014), 1–57. https://doi.org/10.3386/w20697

[46] Claire S H Lim, James M Snyder, and David Strömberg. 2015. The Judge, the Politician, and the Press: Newspaper Coverage and Criminal Sentencing across Electoral Systems. *American Economic Journal: Applied Economics* 7, 4 (2015), 103–135. https://doi.org/10.1257/app.20140111

[47] Gechun Lin. 2025. Using cross-encoders to measure the similarity of short texts in political science. *American Journal of Political Science* n/a, n/a (2025). https://doi.org/10.1111/ajps.12956 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajps.12956

[48] Siyi Liu, Lei Guo, Kate Mays, Margrit Betke, and Derry Tanti Wijaya. 2019. Detecting Frames in News Headlines and Its Application to Analyzing News Framing Trends Surrounding U.S. Gun Violence. In *Proceedings of the 23rd Conference on Computational Natural Language Learning (CoNLL)*, Mohit Bansal and Aline Villavicencio (Eds.). Association for Computational Linguistics, Hong Kong, China, 504–514. https://doi.org/10.18653/v1/K19-1047

[49] John R Lott and Kevin A Hassett. 2014. Is Newspaper Coverage of Economic Events Politically Biased? *Public Choice* 160, 1 (2014), 65–108. https://doi.org/10.1007/s11127-014-0171-5

[50] Michael C. Lovell. 1963. Seasonal Adjustment of Economic Time Series and Multiple Regression Analysis. *J. Amer. Statist. Assoc.* 58, 304 (1963), 993–1010. https://doi.org/10.2307/2282963

[51] Iffat Maab, Edison Marrese-Taylor, Sebastian Padó, and Yutaka Matsuo. 2024. Media Bias Detection Across Families of Language Models. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, Kevin Duh, Helena Gomez, and Steven Bethard (Eds.). Association for Computational Linguistics, Mexico City, Mexico, 4083–4098. https://doi.org/10.18653/v1/2024.naacl-long.227

[52] Ansel MacLaughlin, Tao Chen, Burcu Karagol Ayan, and Dan Roth. 2021. Context-Based Quotation Recommendation. *Proceedings of the International AAAI Conference on Web and Social Media* 15, 1 (May 2021), 397–408. https://doi.org/10.1609/icwsm.v15i1.18070

[53] Luke Miner. 2015. The Unintended Consequences of Internet Diffusion: Evidence from Malaysia. *Journal of Public Economics* 132 (2015), 66–78. https://doi.org/10.1016/j.jpubeco.2015.10.002

[54] Fred Morstatter, Liang Wu, Uraz Yavanoglu, Stephen R. Corman, and Huan Liu. 2018. Identifying Framing Bias in Online News. *Trans. Soc. Comput.* 1, 2, Article 5 (June 2018), 18 pages. https://doi.org/10.1145/3204948

[55] Sendhil Mullainathan and Andrei Shleifer. 2005. The Market for News. *The American Economic Review* 95, 4 (2005), 1031. https://doi.org/10.1257/0002828054825619

[56] Vlad Niculae, Caroline Suen, Justine Zhang, Cristian Danescu-Niculescu-Mizil, and Jure Leskovec. 2015. QUOTUS: The Structure of Political Media Coverage as Revealed by Quoting Patterns. In *Proceedings of the 24th International Conference on World Wide Web*. International World Wide Web Conferences Steering Committee, Florence, Italy, 798–808. https://doi.org/10.1145/2736277.2741688

[57] Rodrigo Nogueira and Kyunghyun Cho. 2019. Passage Re-ranking with BERT. arXiv:1901.04085 http://arxiv.org/abs/1901.04085

[58] Emily Oster. 2019. Unobservable Selection and Coefficient Stability: Theory and Evidence. *Journal of Business & Economic Statistics* 37, 2 (2019), 187–204. https://doi.org/10.1080/07350015.2016.1227711 arXiv:https://doi.org/10.1080/07350015.2016.1227711

[59] Riccardo Puglisi. 2011. Being The New York Times: the Political Behaviour of a Newspaper. *The B.E. Journal of Economic Analysis & Policy* 11, 1 (2011), n/a. https://doi.org/10.2202/1935-1682.2025

[60] Riccardo Puglisi and James M Snyder. 2011. Newspaper Coverage of Political Scandals. *The Journal of Politics* 73, 3 (2011), 931–950. https://doi.org/10.1017/s0022381611000569

[61] Nancy Qian and David Yanagizawa-Drott. 2017. Government Distortion in Independently Owned Media: Evidence from U.S. News Coverage of Human Rights. *Journal of the European Economic Association* 15, 2 (2017), 463–499. http://dx.doi.org/10.1093/jeea/jvx007

[62] Bei Qin, David Strömberg, and Yanhui Wu. 2017. Why Does China Allow Freer Social Media? Protests versus Surveillance and Propaganda. *Journal of Economic Perspectives* 31, 1 (2017), 117–140. https://doi.org/10.1257/jep.31.1.117

[63] Bei Qin, David Strömberg, and Yanhui Wu. 2018. Media Bias in China. *American Economic Review* 108, 9 (2018), 2442–2476. http://www.aeaweb.org/articles?id=10.1257/aer.20170947

[64] Radim Rehurek and Petr Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks (Proceedings of LREC 2010 workshop New Challenges for NLP Frameworks)*. University of Malta, Valetta, MT, 45–50. http://is.muni.cz/publication/884893/en

[65] Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics. https://arxiv.org/abs/1908.10084

[66] Nils Reimers and Iryna Gurevych. 2021. The Curse of Dense Low-Dimensional Information Retrieval for Large Index Sizes. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*. Association for Computational Linguistics, Online, 605–611. https://arxiv.org/abs/2012.14210

[67] Margaret E. Roberts, Brandon M. Stewart, Dustin Tingley, Christopher Lucas, Jetson Leder-Luis, Shana Kushner Gadarian, Bethany Albertson, and David G. Rand. 2014. Structural Topic Models for Open-Ended Survey Responses. *American Journal of Political Science* 58, 4 (2014), 1064–1082. https://doi.org/10.1111/ajps.12103 arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/ajps.12103

[68] Lucas Shen. 2021. Measuring Political Media Slant Using Text Data. https://www.lucasshen.com/research/media.pdf

[69] Lucas Shen. 2022. LexicalRichness: A small module to compute textual lexical richness. https://doi.org/10.5281/zenodo.6607007

[70] Singapore Statutes Online. 1974. Newspaper and Printing Presses Act.

[71] James M. Snyder and David Strömberg. 2010. Press Coverage and Political Accountability. *Journal of Political Economy* 118, 2 (2010), 355–408. https://doi.org/10.1086/652903

[72] Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2020. MPNet: Masked and Permuted Pre-training for Language Understanding. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 16857–16867. https://proceedings.neurips.cc/paper_files/paper/2020/file/c3a690be93aa602ee2dc0ccab5b7b67e-Paper.pdf

[73] Timo Spinde, Smi Hinterreiter, Fabian Haak, Terry Ruas, Helge Giese, Norman Meuschke, and Bela Gipp. 2024. The Media Bias Taxonomy: A Systematic Literature Review on the Forms and Automated Detection of Media Bias. arXiv:2312.16148 [cs.CL] https://arxiv.org/abs/2312.16148

[74] Dominik Stammbach, Vilém Zouhar, Alexander Hoyle, Mrinmaya Sachan, and Elliott Ash. 2023. Revisiting Automated Topic Model Evaluation with Large Language Models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, Houda Bouamor, Juan Pino, and Kalika Bali (Eds.). Association for Computational Linguistics, Singapore, 9348–9357. https://doi.org/10.18653/v1/2023.emnlp-main.581

[75] Daniel Sutter. 2012. Is the Media Liberal? An Indirect Test Using News Magazine Circulation. *Applied Economics* 44, 27 (2012), 3521–3532. https://doi.org/10.1080/00036846.2011.577024

[76] Netina Tan. 2014. Institutionalized Succession and Hegemonic Party Cohesion in Singapore. In *Party System Institutionalization in Asia: Democracies, Autocracies, and the Shadows of the Past*, Allen Hicken and Erik Martinez Kuhonta (Eds.). Cambridge University Press, Cambridge, 49–73. https://doi.org/10.1017/CBO9781107300385.003

[77] Oren Tsur, Dan Calacci, and David Lazer. 2015. A Frame of Mind: Using Statistical Models for Detection of Framing and Agenda Setting Campaigns. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*. Association for Computational Linguistics, Beijing, China, 1629–1638. https://doi.org/10.3115/v1/P15-1157

[78] Emma van der Goot, Toni G.L.A. van der Meer, and Rens Vliegenthart. 2021. Reporting on Political Acquaintances: Personal Interactions Between Political Journalists and Politicians as a Determinant of Media Coverage. *International Journal of Communication* 15, 0 (2021), 23. https://ijoc.org/index.php/ijoc/article/view/15313

[79] Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. MINILM: deep self-attention distillation for task-agnostic compression of pre-trained transformers. In *Proceedings of the 34th International Conference on Neural Information Processing Systems* (Vancouver, BC, Canada) *(NIPS '20)*. Curran Associates Inc., Red Hook, NY, USA, Article 485, 13 pages.

[80] Philine Widmer, Sergio Galletta, and Elliott Ash. 2023. Media Slant is Contagious. https://doi.org/10.2139/ssrn.3712218

[81] Tae Yano, Philip Resnik, and Noah A. Smith. 2010. Shedding (a Thousand Points of) Light on Biased Language. In *Proceedings of the NAACL HLT 2010 Workshop on Creating Speech and Language Data with Amazon's Mechanical Turk*. Association for Computational Linguistics, Los Angeles, California, USA, 152–158. https://aclanthology.org/W10-0723

[82] Yiyi Zhou, Rongrong Ji, Jinsong Su, and Jiaquan Yao. 2020. Uncovering Media Bias via Social Network Learning. *ACM Trans. Intell. Syst. Technol.* 12, 1, Article 12 (Dec. 2020), 12 pages. https://doi.org/10.1145/3422181

Supplementary Information: Table of Contents

§A. Data Appendix
§B. Background to Parliamentary Speeches
§C. Coverage Intensity
§D. Sensitivity Analyses
§E. Bounding Analysis Using Ministerial Composition
§F. Rational Choices in a Separating Equilibrium

A Data Appendix

News Article Corpus and Politicians. I start with the list of politicians who were active during the sample period, including many whose tenure began in the 1980s. For each politician, I collect metadata: ministerial rank and portfolio (where applicable), tenure dates, gender, and birth year. Birth year is unavailable for two non-partisan members; I impute these as 1950 and 1955 based on estimated age from available biographical information. I then query these politician names in Factiva, a database providing access to licensed news archives worldwide. This process returns over 112k news articles from The Straits Times. After deduplicating by title–date tuples, I end up with 62k unique articles for 204 politicians (Table A1).

Article Classification. Only a subset of articles that mention politicians will contain quotes from parliamentary speeches. To extract articles with parliament quotes, I label 1,419 articles from October 2011 to December 2012 (275 "positive" cases containing at least one quote) as a training set for use in a supervised machine learning workflow. I represent news articles as bag-of-words using n-grams (n = 1 to 5), supplemented by binary indicators for quotation marks and parliamentary keywords. The keywords include parliamentary terminology (parliament, debate, sitting, bill, budget), roles (MP, Minister, speaker), constituency types (SMC, GRC; Section 2.1), and quote indicators (said, colons, semicolons).

I compare six supervised learning algorithms (naive Bayes, decision tree, support vector machine, logistic regression, bagged decision trees, random forest) across multiple specifications addressing class imbalance through sample weighting and SMOTE (synthetic minority oversampling). Model hyperparameters are optimized via 5-fold cross-validation. The best, and eventual, classifier I used is a random forest with class weights, which achieves the best balance of recall (83%) and precision (70%), substantially outperforming the naive baseline of classifying all articles as negative (81% accuracy but 0% recall).

Classification Workflow. I adopt a classification workflow that iteratively improves the training set and classifier performance. Starting with the initial 1,419 labeled articles, I train a random forest and apply it to the remaining articles from 2012. The classifier predicts 471 articles contain parliamentary quotes. As I extract quotes from these articles (described below), I simultaneously verify each prediction—does the article truly contain parliamentary quotes? These verified labels convert the 471 articles from "predictions" into "ground truth", expanding my training set to 1,890 articles. I then retrain the classifier on this larger set and apply it to 2013 articles. I repeat this iteratively through the years of the sample period. Across all years, top predictive features naturally include parliamentary vocabulary (e.g., "parliament," "mps," "debate," "grc", "parliament yesterday") and those relating to question-and-answer (e.g., "asked", "responding").

Quote Matching and Extraction. For articles classified as positive, I extract candidate quotes using a semi-automated process. Candidate quotes are text strings within single or double inverted commas, or following speech colons. For each article, I extract all candidate quotes and identify potential source speeches, starting from the article date and the previous day (speeches typically appear in the news the same day or the next day). Each quote is matched against candidate speeches using the Ratcliff-Obershelp edit distance, which recursively identifies the longest common substring between quote and speech, then continues matching remaining portions. Text is preprocessed to lowercase with punctuation removed. The algorithm produces similarity scores and ranks candidate speeches. I review the top n matches for each quote and select the correct source speech, or mark the quote as false positive if no appropriate match exists. This process yields 14,901 quote-speech pairs with complete data from 5,129 speeches across 3,421 articles.

Topic Modeling. Topic number K is selected by maximizing coherence scores across models trained for $K = 2, 4, ..., 100$.[17] The Dirichlet parameter α is set to $1/K$, consistent with speeches typically having one dominant topic. I train separate LDA models for speeches ($K^* = 92$) and articles ($K^* = 40$) to control for content differences. For robustness checks, I train additional models at $K = 50, 100$ (speeches) and $K = 30, 50$ (articles). Text preprocessing includes stopword removal, lemmatization, and named entity removal (persons, dates, times, money, quantities). Importantly, I apply phrase detection using normalized pointwise mutual information across two passes over the corpus. Word pairs that frequently co-occur (exceeding a threshold of 0.55–0.7) are connected with underscores, allowing the model to treat multi-word expressions as single tokens. As a context-specific example, "minimum sum" becomes "minimum_sum" and "cpf saving" becomes "cpf_saving". A second pass identifies longer phrases like "early_childhood_education." This step is important because these phrases often have distinct meanings compared to their constituent words appearing separately. Other examples include: resale_flat, pre_school, civil_servant, unfair_dismissal, mental_capacity, integrated_shield_plans, group_representation_constituency, en_bloc_sale, progressive_wage_model, and transboundary_haze_pollution.

Table A1. Descriptive statistics

		Party		
	All	Ruling	Opposition	Non-partisan
	(1)	(2)	(3)	(4)
A. Sample				
Number of quotes	14,901	13,321	1,105	475
Number of MPs	204	144	16	44
#words in quote (SD)	20.9 (38.5)	20.9 (40.2)	20.1 (18.2)	23.7 (19.3)
B. Distribution of quotes (mean, SD)				
#quotes per MP	73.0 (130.8)	92.5 (147.6)	69.1 (92.1)	10.8 (9.1)
#quotes per article	4.36 (3.11)	4.23 (3.04)	6.18 (3.70)	4.92 (3.03)
#quotes per speech	2.91 (3.86)	3.01 (4.06)	2.41 (2.39)	1.96 (1.51)
#quotes per year	1,241.8 (402.3)	1,110.1 (349.6)	92.1 (50.6)	39.6 (14.7)
C. Quote accuracy (mean, SD)				
Substring accuracy	91.5 (12.8)	91.6 (12.9)	90.9 (12.7)	91.7 (11.5)
BoW accuracy	96.7 (10.2)	96.7 (10.0)	95.3 (13.2)	97.5 (7.3)
Semantic accuracy	89.9 (26.8)	89.8 (27.1)	89.3 (26.7)	93.8 (20.1)
D. Age, tenure, and gender				
Age (years, SD)	49.1 (8.3)	49.8 (7.9)	46.5 (10.6)	48.0 (8.7)
Political tenure (years, SD)	7.5 (8.0)	9.8 (8.2)	4.6 (6.6)	1.0 (0.7)
Female MPs (%)	23.0%	18.1%	18.8%	40.9%
Female quotes (%)	16.5%	14.7%	28.0%	40.4%

Table A2. Daily Newspaper Subscription

			2015 Subscriptions				2021 Subscriptions		
	Daily newspaper	Language	Unique Digital	Print	Digital	Total	Print	Digital	Total
1	Berita Harian (Berita Minggu)	Malay	911	44'600	2'500	47'100	21'200	7'300	28'500
2	The Business Times	English	6'658	29'200	18'500	47'700	14'200	26'000	40'200
3	Lianhe Zaobao	Chinese	13'727	148'600	39'300	187'900	85'200	51'700	136'900
4	Lianhe Wanbao	Chinese	1'137	82'500	9'100	91'600	35'700	27'700	63'400
5	The New Paper (The New Paper Sunday)	English	757	70'200	40'400	110'600	—	—	—
6	Shin Min Daily News	Chinese	—	—	—	120'200	62'700	27'700	90'400
7	The Straits Times (The Sunday Times)	English	60'871	304'300	177'400	481'700	145'700	315'100	460'800
8	Tamil Murasu (Tamil Murasu Sunday)	Tamil	—	—	—	12'800	5'900	1'900	7'800

Subscriptions of the eight daily newspapers wholly owned by Singapore Press Holdings. Parenthesis indicates the Sunday edition of the daily. For instance, *The Straits Times* is published as *The Sunday Times* on Sundays. In addition to total subscriptions, unique digital subscriptions, print subscriptions, and digital subscriptions are shown where available. 2015 subscription data comes from https://web.archive.org/web/20211003084400/https://corporate.sph.com.sg/system/misc/annualreport/2015/SPH_AR2015_DailyAverageNewspaperCirculation.pdf. 2021 subscription data comes from https://web.archive.org/web/20250516171119/https://mustsharenews.com/sph-media-trust-saga/, although numbers may have been inflated.

Table A3. Newspaper section by party

Newspaper section	Government	Opposition	Total
Home	963	85	1,048
Insight	50	7	57
Money	21	0	21
News	28	0	28
Opinion	11	0	11
Others	16	3	19
Prime News	1,929	152	2,081
Review - Insight	14	0	14
Singapore	8,213	623	8,836
Sports	14	0	14
ST	645	34	679
Think	12	3	15
Top of the News	1,871	198	2,069
World	9	0	9
Total	13,796	1,105	14,901

B Background to Parliamentary Speeches

This section synthesizes parliamentary business facts from the official source https://www.parliament.gov.sg/parliamentary-business/glossary, and from insights drawn from the written experiences of an anonymous reviewer, who is a member of parliament in the 13th parliament. Insights from the anonymous reviewer come mainly in the last two paragraphs.

Speeches are (mostly) scripted. Most of the speeches in the Parliament of Singapore are jointly planned by the Government in consultation with the Speaker of Parliament (Speaker), who, as a result, is rarely surprised by what Members of Parliament (Members) intend to say. Bills and Motions, for instance, require at least two days' notice before their introduction, or even longer depending on who introduces the bill. Extemporaneous and impromptu speeches exist but are rare and likely shorter. A Parliamentary Reporter records all these proceedings in shorthand before archiving them in the official repository.

True debates are rare. One common context in which Members speak is on Bills or Motions. The sponsor of the bill, for instance, usually a Political Office Holder, is allowed to speak twice, once to introduce the Bill/Motion and once more to close the debate. Other Members are allowed up to 20 minutes to debate the introduced Bill/Motion. Still, the interim response to the introduction and the response is likely prepared in advance. A different and rarer context is Ministerial Statements, given by a Minister regarding the Government's policy and decision. Although no notice is required for such statements, it is, in all likelihood, prepared in advance.

Deviations from scripted exchange. One context in which unscripted exchanges can arise is via Parliamentary Question Time, a period set aside at the beginning of every sitting, where Ministers or Members respond to Questions as filed and accepted by the Speaker until the end of the allocated Question Time. Since these Questions are filed in advance, responses are scripted. However, once a response is given, any other Member may spontaneously ask supplementary questions relating to the original Question. While these are conditional on having sufficient time, the Speaker may, in practice, allocate more time to the opposition than their proportional share based on seats held. There is room for unscripted responses to the extent that the supplementary questions have not been pre-empted or made known to colleagues in advance. Deviations from scripted exchanges among the incumbent party members may also arise when members want to engage further on points of clarification, but these usually occur near the end of a debate.

The main source of unscripted exchanges, therefore, relates to opposition Members who, unlike the incumbent party Members, do not circulate their speeches and questions in advance. As a result, it becomes difficult for most incumbent party Members to prepare scripted responses to speeches from opposition Members.

The extent to which speeches and their responses are predictable rests on the existing media-political machinery and differs by seniority and party status. Incumbent party Members are more likely to circulate both their speeches and questions in advance, both among their party colleagues and also to the media. The opposition Members, however, do not.

C Coverage Intensity

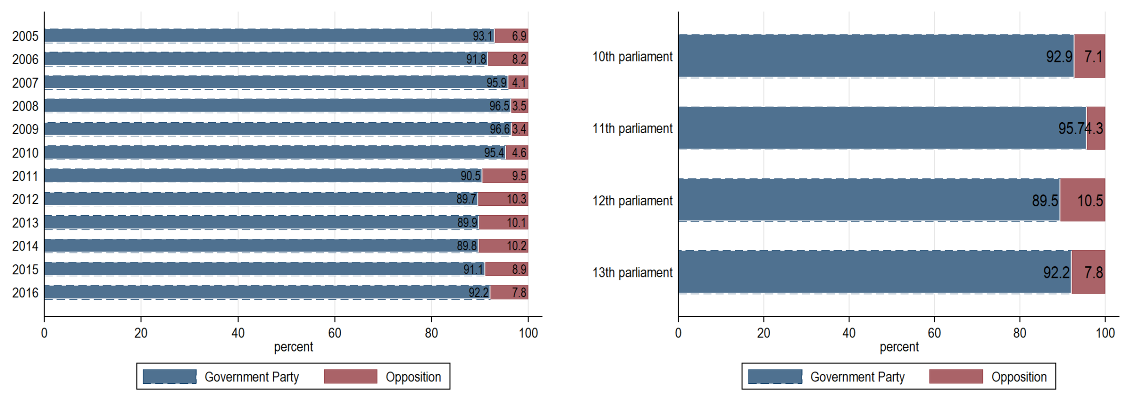

(a) Count of quotes over the years by partisanship (b) Count of quotes over parliaments by partisanship

Fig. C1. Count of quotes by partisanship over years and parliaments

Observations in Table C1 are at the article-speech level:

$$y_{irst} = \alpha + \beta\, opp_i + \sum_{k=2006}^{2016} \alpha_k\, year_{kt} + \sum_{\ell=11}^{13} \alpha_\ell\, parl_{\ell t} + \boldsymbol{\gamma}' \boldsymbol{X}_{irst} + \varepsilon_{irst}, \tag{C1}$$

similar to Equation (4) but aggregated to the article-speech level. The outcome is the log length of the number of quoted words, where the separate (direct) quotations that originate from the same speech and reported in the same news article are treated as a single observation. Time fixed-effects include both the parliament (10th, 11th, 12th, & 13th) and year (2005–16) fixed effects. Individual controls include the politicians' gender, ethnic (Chinese, Indian, Malay, and Eurasian/Others), and a quadratic in age and political tenure. Article controls include day-of-the-week dummies, newspaper section (e.g. Top of the News, Prime News), and a dummy for whether the quote was translated. Topic controls are vectors of probabilistic association (sum to one) of the newspaper article and parliament speech to topics uncovered using LDA. Ministerial controls include politician type (e.g. Deputy Prime Minister, Parliamentary Secretary) and political portfolio (e.g. Health, Education). Electoral controls include the electorate size that the politician represents, the group size of representation (group representation have sizes of 4–6), vote share, and winning margin in

Table C1. Difference in Coverage Intensity for Opposition vs. Ruling-Party Politicians, Article-Speech Level

	Outcome: Log of quote length by word count	
	(1)	(2)
Opposition	0.052	0.038
	(0.064)	(0.074)
Controls		
Time fixed-effects	✓	✓
Individual controls	✓	✓
Article controls	✓	✓
Topic controls	✓	✓
Ministerial controls	✓	✓
Electoral controls	–	✓
F-statistics		
F-statistic, year fixed-effects	2.682^{***}	1.613^{*}
F-statistic, individual controls	0.397	1.153
F-statistic, article controls	3.290^{***}	2.538^{***}
F-statistic, topic controls	2.762^{***}	2.413^{***}
F-statistic, ministerial controls	2.694^{***}	2.722^{***}
F-statistic, electoral controls		0.382
Mean of dependent variable	3.306	3.301
R^2	0.181	0.199
N	7,087	5,143

Clustered standard errors in parentheses. ***, **, * indicate significance at the 1%, 5%, and 10% levels.

the most recent general election. Politicians who had won by default (no opposition contest) in the most recent election have no electoral data. All regressions also include the intensity measures for the textual data—the (log) length of speech, speech paragraph, and news article. F-statistics report the test statistic for the null that the set of controls are jointly equal to zero. Robust standard errors in parentheses are clustered at news articles.

D Sensitivity Analyses

Table D1. Specification Checks for Quote Accuracy, Quote Level

					Clustering		Topic distribution of textual content					
	Baseline Regression	Journa-list FE	No ministerial controls	No translations	Cluster by speech	Cluster by journalist	Speech $K = 50$	Speech $K = 100$	Article $K = 30$	Article $K = 50$	Sentence topics	Parsimonious topics
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Panel A. Dependent variable is substring quote accuracy												
Opposition	−1.86***	−2.05***	−1.68***	−1.68**	−1.86**	−1.86**	−1.53**	−1.78**	−2.16***	−1.88***	−1.87***	−1.74**
	(0.71)	(0.77)	(0.55)	(0.70)	(0.75)	(0.73)	(0.72)	(0.72)	(0.70)	(0.71)	(0.71)	(0.70)
N	14,901	13,524	14,901	14,697	14,901	14,901	14,901	14,901	14,901	14,901	14,901	14,901
Panel B. Dependent variable is substring quote accuracy (No stop words)												
Opposition	−1.94***	−2.20***	−1.71***	−1.75**	−1.94**	−1.92**	−1.59**	−1.94***	−2.27***	−1.99***	−1.95***	−1.83***
	(0.72)	(0.78)	(0.56)	(0.71)	(0.76)	(0.77)	(0.72)	(0.73)	(0.71)	(0.72)	(0.72)	(0.71)
N	14,901	13,524	14,901	14,697	14,901	14,901	14,901	14,901	14,901	14,901	14,901	14,901
Panel C. Dependent variable is bag-of-words quote accuracy measure												
Opposition	−2.15***	−2.41***	−1.63***	−2.04***	−2.15***	−2.13***	−1.97***	−2.11***	−2.38***	−2.07***	−2.13***	−2.20***
	(0.70)	(0.75)	(0.55)	(0.70)	(0.71)	(0.78)	(0.72)	(0.72)	(0.70)	(0.69)	(0.70)	(0.71)
N	14,901	13,524	14,901	14,697	14,901	14,901	14,901	14,901	14,901	14,901	14,901	14,901
Panel D. Dependent variable is bag-of-words quote accuracy (No stop words)												
Opposition	−2.88***	−3.18***	−2.04***	−2.76***	−2.88***	−2.92***	−2.57***	−2.85***	−3.22***	−2.81***	−2.86***	−2.89***
	(0.91)	(0.99)	(0.70)	(0.92)	(0.93)	(1.10)	(0.93)	(0.94)	(0.90)	(0.89)	(0.92)	(0.91)
N	14,901	13,524	14,901	14,697	14,901	14,901	14,901	14,901	14,901	14,901	14,901	14,901
Panel E. Dependent variable is semantic quote accuracy												
Opposition	−4.05***	−4.68***	−2.97***	−3.69**	−4.05***	−4.23**	−3.25**	−3.71**	−4.87***	−4.06***	−4.01***	−3.92***
	(1.46)	(1.59)	(1.13)	(1.45)	(1.48)	(1.66)	(1.46)	(1.49)	(1.43)	(1.45)	(1.47)	(1.42)
N	14,901	13,524	14,901	14,697	14,901	14,901	14,901	14,901	14,901	14,901	14,901	14,901
Panel F. Dependent variable is *binarized* semantic quote accuracy (= 1 if Semantic accuracy ≥ 99)												
Opposition	−0.06***	−0.05**	−0.05***	−0.06***	−0.06***	−0.05**	−0.06***	−0.06***	−0.07***	−0.05**	−0.06***	−0.06***
	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.03)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)	(0.02)
N	14,887	13,513	14,887	14,682	14,887	13,513	14,887	14,887	14,887	14,887	14,887	14,887

Note: Observations are at the quote level; the regressions in this Table consider each quote as a separate observation, even if they originate from the same speech or are reported in the same news article. The set of controls is the same as in Table 3. Robust standard errors in parentheses are clustered at news articles.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

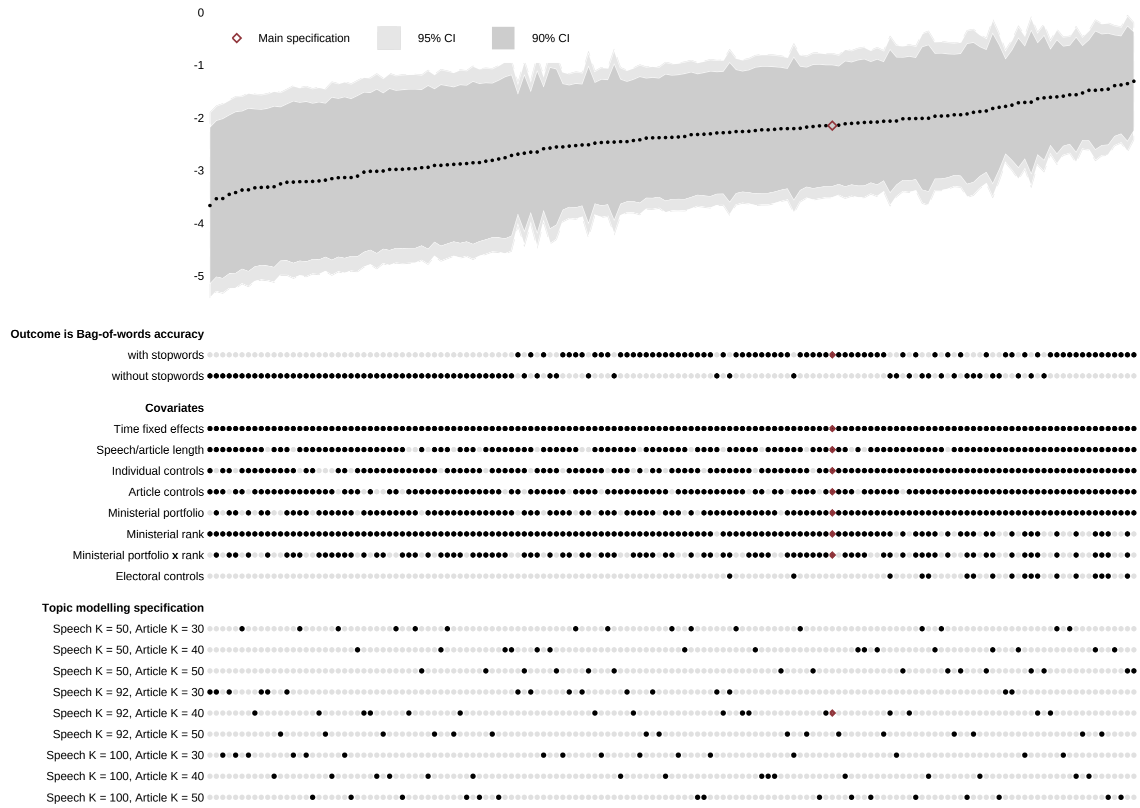

Fig. D1. Effect Sizes of Opposition Status on Bag-of-words Accuracy

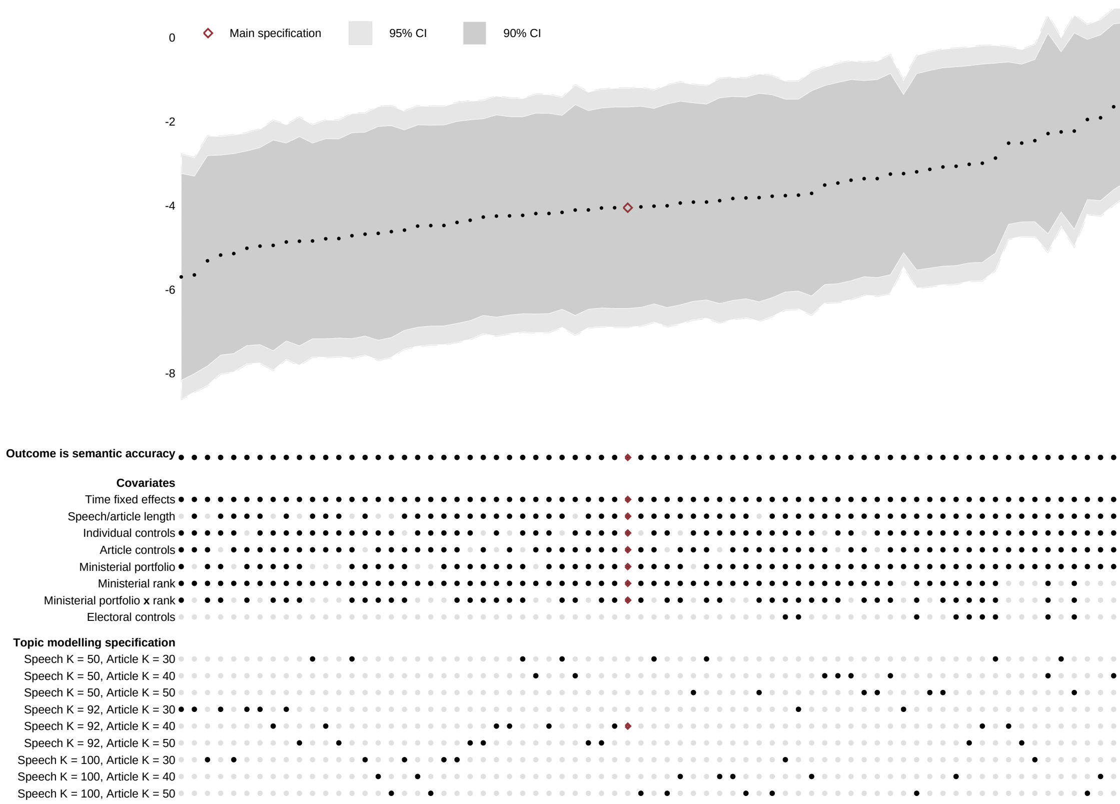

Fig. D2. Effect Sizes of Opposition Status on Semantic Accuracy

D.1 Additional linguistic controls

Table D2. Additional Language-based Measures as Controls

		Additional language and textual controls using 1st principal components of objectivity, polarity, readability, and lexical richness measures				
	Baseline results	Objectivity of textual content	Polarity of textual content	English grade/ readability	Lexical Richness	All
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Dep. Var. is substring quote accuracy						
Opposition	-1.86^{***}	-1.81^{**}	-1.86^{***}	-2.12^{***}	-1.76^{**}	-2.02^{***}
	(0.71)	(0.71)	(0.71)	(0.71)	(0.71)	(0.72)
Objectivity of speech and quote		-0.13^{**}				-0.13^{**}
		(0.06)				(0.06)
Polarity of speech and quote			-0.02			-0.03
			(0.06)			(0.06)
Grade/readability score of speech transcipt				0.65^{***}		0.64^{***}
				(0.08)		(0.09)
Lexical richness of speech transcipt					0.36^{***}	0.31^{***}
					(0.08)	(0.08)
Panel B. Dep. Var. is bag-of-words quote accuracy						
Opposition	-2.15^{***}	-2.12^{***}	-2.14^{***}	-2.24^{***}	-2.15^{***}	-2.23^{***}
	(0.70)	(0.70)	(0.70)	(0.70)	(0.70)	(0.70)
Objectivity of speech and quote		-0.08				-0.08
		(0.05)				(0.05)
Polarity of speech and quote			0.02			0.01
			(0.05)			(0.05)
Grade/readability score of speech transcipt				0.22^{***}		0.22^{***}
				(0.07)		(0.08)
Lexical richness of speech transcipt					0.14^{*}	0.12
					(0.08)	(0.08)
Panel C. Dep. Var. is semantic quote accuracy						
Opposition	-4.05^{***}	-4.10^{***}	-4.10^{***}	-4.08^{***}	-3.88^{***}	-3.97^{***}
	(1.46)	(1.46)	(1.46)	(1.46)	(1.46)	(1.46)
Objectivity of speech and quote		0.09				0.08
		(0.12)				(0.13)
Polarity of speech and quote			-0.10			-0.07
			(0.12)			(0.13)
Grade/readability score of speech transcipt				0.09		0.05
				(0.16)		(0.17)
Lexical richness of speech transcipt					0.42^{***}	0.41^{***}
					(0.15)	(0.15)

Clustered standard errors in parentheses. ***, **, * indicate significance at the 1%, 5%, and 10% levels.

E Bounding Analysis Using Ministerial Composition

The following shows that treating ministerial rank as a bad control suggests that the OLS estimate of the opposition dummy carries an attenuation bias (an upward bias when the true effect is negative), even if assignment into a party is random. Specifically, opposition status determines not only political media coverage but also the ministerial rank of politicians since the opposition politicians never get a higher than base (lowest) rank (Table E1). Hence, even conditioning on rank does not recover the true causal effect because of the politicians' composition and rank.

As a simplification, the following abstracts away from other regressors ($\boldsymbol{X}$), or more specifically, assumes that opposition status does not determine $\boldsymbol{X}$ other than rank. In reality, there is also a whole range of ministerial ranks/types, but these all collapse into a single dummy indicating a higher than base rank, with base rank as the omitted category. Let the opposition status indicator for politician i be

$$D_i = \begin{cases} 1 & \text{if politician } i \text{ is from an opposition party} \\ 0 & \text{otherwise,} \end{cases} \tag{E1}$$

and let the higher rank indicator for politician i be

$$r_i = \begin{cases} 1 & \text{if politician } i \text{ holds a higher than base (lowest) level rank} \\ 0 & \text{if base (lowest) rank.} \end{cases} \tag{E2}$$

Table E1. Ministerial Rank by party

Ministerial Rank/type		Party		
Rank	Description	Government	Opposition	Total
PM	Prime Minister	491	0	491
DPM	Deputy PM	1,228	0	1,228
Minister	MR1-4 rank	6,605	0	6,605
SMS	Senior Minister of State	1,020	0	1,020
MOS	Minister of State	531	0	531
Mayor	Mayor (of 1 of 5 districts)	200	0	200
Sps	Senior Parliamentary Secretary	103	0	103
Parl Sec	Parliamentary Secretary	102	0	102
Speaker	Speaker of Parliament	38	0	38
MP	Member of Parliament (base rank)	3,004	800	3,804
NCMP	Non-constituency MP	0	306	306
NMP	Nominated MP	475	0	475
Total		13,322	1,106	14,428

Following the setup of bad controls [5], opposition status (D_i) determines both media coverage (y_i) and the rank (r_i):

$$y_i = y_{0i} + (y_{1i} - y_{0i})D_i \tag{E3}$$

$$r_i = r_{0i} + (r_{1i} - r_{0i})D_i, \tag{E4}$$

where y_{1i} and r_{1i} (y_{0i} and r_{0i}) are the potential media coverage and potential ministerial rank of politician i as an opposition (ruling-party) politician. One can think of rank affecting media coverage as an omitted variable problem in that rank is increasing in some unobserved characteristics of a politician (e.g., competence, likeability, quotability, public image), and coverage is increasing in these characteristics. By the joint independence of $\{y_{1i}, r_{1i}, y_{0i}, r_{0i}\}$ and D_i, comparing opposition and ruling-party politicians conditional on the base level rank ($r_i = 0$) gives:

$$\begin{aligned}
&\mathbb{E}[y_i|D_i = 1, r_i = 0] - \mathbb{E}[y_i|D_i = 0, r_i = 0] \\
&= \mathbb{E}[y_{1i}|D_i = 1, r_{1i} = 0] - \mathbb{E}[y_{0i}|D_i = 0, r_{0i} = 0] \\
&= \mathbb{E}[y_{1i}|r_{1i} = 0] - \mathbb{E}[y_{0i}|r_{0i} = 0] \\
&= \mathbb{E}[y_{1i}|r_{1i} = 0] - \mathbb{E}[y_{0i}|r_{1i} = 0] + \mathbb{E}[y_{0i}|r_{1i} = 0] - \mathbb{E}[y_{0i}|r_{0i} = 0] \\
&= \mathbb{E}[y_{1i} - y_{0i}|r_{1i} = 0] + \mathbb{E}[y_{0i}|r_{1i} = 0] - \mathbb{E}[y_{0i}|r_{0i} = 0]
\end{aligned} \tag{E5}$$

The observed difference in political media coverage between opposition and ruling-party politicians can be decomposed into two parts. The first term in the last line of Equation (E5) is the true causal effect of opposition status on coverage, conditional on having the base ministerial rank.

The bias comes from the second and third terms. The second term is the potential coverage of a ruling-party politician had he been assigned a base ministerial rank as an opposition politician. The third term is the potential coverage of a ruling-party politician, given that he is assigned to a base ministerial rank.

In theory, the bias can go in either direction. But I propose here that a very plausible set of assumptions suggests that the bias is likely positive. First, the average coverage of a ruling-party politician who would have been assigned to the base rank had he been an opposition politician, is simply the average coverage of all ruling-party politicians, so that $\mathbb{E}[y_{0i}|r_{1i} = 0] = \mathbb{E}[y_{0i}]$.

The average coverage of a ruling-party politician assigned to a base ministerial rank, however, is likely below average if coverage increases with rank so that $\mathbb{E}[y_{0i}|r_{0i} = 0] < \mathbb{E}[y_{0i}]$. Together, the assumptions imply that $\mathbb{E}[y_{0i}|r_{1i} = 0] = \mathbb{E}[y_{0i}] > \mathbb{E}[y_{0i}|r_{0i} = 0]$, or that the bias term in the equation $\mathbb{E}[y_{0i}|r_{1i} = 0] - \mathbb{E}[y_{0i}|r_{0i} = 0] > 0$. Hence, the bias is positive. When the true causal effect of opposition status on coverage is negative, this becomes an attenuation bias towards zero, and the observed comparison of means from the OLS estimates constitutes a lower bound on the true magnitude of the opposition effect.

F Rational Choices in a Separating Equilibrium

To the extent that the detected partisan gaps in quotation accuracy fully reflect logistical differences—the choice to circulate speech transcripts in advance to media agents—the empirical findings point to privately-held information and beliefs about media slant as follows.

Allowing media agents early access to speeches leads to two things. First, it increases quotation accuracy ($A_i \in \mathbb{R}_+$). Second, and on the other hand, early access creates disutility from spin ($S_i \in \mathbb{R}_+$). Let c_1 indicate (early) circulation of speech and c_0 otherwise. A political agent i trades off between increasing quotation accuracy and spin by choosing whether to circulate. If the speech is not circulated, there is no spin, and the media reports the speech at face value, or that $S_i(c_0) = 0$. A political agent, therefore, gains from granting early access if

$$A_i(c_1) - S_i(c_1) > A_i(c_0), \tag{F1}$$

where the gain in accuracy is greater than the disutility from spin when circulating the speech in advance.

Political agents care about being quoted accurately, and they care about spin because they want control over the framing of their speeches. This precludes the framing pre-agreed upon by both parties. As an illustration, a quotation about "raising taxes by x% to improve fiscal position" can be misquoted in different ways. One is "raising taxes by x%", another is "raising taxes by x% to stabilize fiscal position".

The story itself can have different spins. It can be framed as the government lacking fiscal prudence in recent years or as a story about government prescience on future spending needs. It can also be threaded into a pre-existing narrative on social spending (or lack thereof). Spin increases with early circulation because it gives media agents time to develop or skew the narrative. To be precise, spin here has no value judgment but captures only magnitude for simplicity. To this extent, spin is undesirable because it takes narrative control away from political agents even if there is a non-zero probability that the political agent ends up 'looking good.' Without early access to develop a richer narrative, the media agent takes the "[improving] fiscal position" narrative at face value.

Here, I rely on the primary and secondary research insights. The institutional fact that a separating equilibrium exists, where the ruling-party politicians ($i = r$) choose to circulate their speeches but the opposition ($i = o$) does not, implies that

$$A_r(c_1) - S_r(c_1) > A_r(c_0) \tag{F2}$$

where the ruling-party politicians gain from circulating the speeches but

$$A_o(c_1) - S_o(c_1) < A_o(c_0), \tag{F3}$$

where the opposition politicians do not.

Moreover, if one takes the position where the detected differences in quotation accuracy in this paper are not reflective of slant but reflective entirely of logistical differences, implies that $A_r(c_\ell) = A_o(c_\ell) = A(c_\ell)$, where $\ell \in \{0, 1\}$. In which case, the above two equations (Equation (F2) and Equation (F3)) collapse into

$$A(c_1) - S_r(c_1) > A(c_0) > A(c_1) - S_o(c_1),$$

or that

$$S_o(c_1) > S_r(c_1), \tag{F4}$$

which says that even if observed differences in accuracy do not originate from slant but from choices in advance speech circulation, the separating equilibrium, at the very least, still reveals private beliefs about a media that slants towards the ruling party. These beliefs—to the extent that interactions between political and media agents are frequent and that media pieces involving political agents are frequent—are credible, given ample evidence for updating.