Essays on the political economy of mass media and representation

Shen, Yan Shun

2021

Shen, Y. S. (2021). Essays on the political economy of mass media and representation.
Doctoral thesis, Nanyang Technological University, Singapore.
<https://hdl.handle.net/10356/152696>

<https://hdl.handle.net/10356/152696>

<https://doi.org/10.32657/10356/152696>

This work is licensed under a Creative Commons Attribution-NonCommercial 4.0
International License (CC BY-NC 4.0).

Downloaded on 26 Apr 2023 22:54:24 SGT

Essays on the Political Economy of Mass Media and Representation

Shen Yan Shun

School of Social Sciences

A thesis submitted to the Nanyang Technological
University
in partial fulfillment of the requirement for the degree of
Doctor of Philosophy

2021

Statement of Originality

I certify that all work submitted for this thesis is my original work. I declare that no other person's work has been used without due acknowledgement. Except where it is clearly stated that I have used some of this material elsewhere, this work has not been presented by me for assessment in any other institution or University. I certify that the data collected for this project are authentic and the investigations were conducted in accordance with the ethics policies and integrity standards of Nanyang Technological University and that the research data are presented honestly and without prejudice.

9 July 2020

.....

Date

.....

Shen Yan Shun

Supervisor Declaration Statement

I have reviewed the content of this thesis and to the best of my knowledge, it does not contain plagiarised materials. The presentation style is also consistent with what is expected of the degree awarded. To the best of my knowledge, the research and writing are those of the candidate except as acknowledged in the Author Attribution Statement. I confirm that the investigations were conducted in accordance with the ethics policies and integrity standards of Nanyang Technological University and that the research data are presented honestly and without prejudice.

15 July 2020

.....

Date

Christos Sakellariou

Authorship Attribution Statement

This thesis does not contain any materials from papers published in peer-reviewed journals or from papers accepted at conferences in which I am listed as an author.

9 July 2020

.....

Date

.....

Shen Yan Shun

Acknowledgements

I am immensely grateful to Giovanni Ko for his support, guidance, and patience throughout. Through our discussions, I learned about political institutions, both elsewhere in the world and in my own country. The intellectual freedom he gave and the suggestions to use non-economic methods have been formative in my growth as a researcher.

I am very grateful to Madhav Aney, Yohanes Eko Riyanto, Christos Sakellariou, and Walter Theseira for their support and comments at various stages. I am especially thankful to Christos for his supervision and support in the later stages of my PhD, and also for copyediting the individual chapters. I also greatly appreciate the continuous support from Christos and Giovanni during the pandemic crisis.

The first chapter has benefited from comments from seminar participants at the Australasian Public Choice Conference and NTU. I am grateful to an anonymous journalist from *The Straits Times* for insights into parliamentary reporting and the mainstream media of Singapore in general, and an anonymous member of parliament for sharing experiences with the media and in parliament. This chapter also benefits from the written experience of an examiner who was a member of parliament.

For the last chapter, I thank Chris Youderian from SimpleMaps for clarifying how the USGS defines a primary U.S. county. For all three chapters, I have benefitted from the Python open-source community that made the data collection and machine learning implementations easy. Finally, I thank my cat for erasing only half a day's worth of coding work throughout my PhD.

Contents

Acknowledgements	i
Abstract	xi
Introduction	xii
1 POLITICAL MEDIA SLANT WITH TEXT DATA	1
1.1 Introduction	1
1.2 Background of Political and Media Institutions	5
1.2.1 Political Background	5
1.2.2 Mainstream media background	6
1.3 Data	8
1.3.1 Downloading and Matching of Quotes to Speeches	8
1.3.2 Quantifying Quotation Accuracy	10
1.3.3 Topic Distribution of Parliamentary Speeches and News Articles	13
1.4 Article-Speech Level Results	16
1.4.1 Empirical Strategy	16
1.4.2 Baseline Results, Article-Speech Level	22
1.4.3 Specification Checks, Article-Speech Level	24
1.5 Quote level Results	26
1.5.1 Baseline Results, Quote Level	26
1.5.2 Specification Checks	27
1.5.3 Bounds on OLS Estimates	28
1.6 Additional Text and Language-Based Measures	31
1.6.1 Speech Tone and Language Competency Controls	31
1.6.2 Representation of Speech Tone	33
1.7 Unpacking the Findings	33

1.7.1	Salience of Quotation Accuracy	33
1.7.2	Magnitudes	35
1.7.3	Contextualising via the Literature	36
1.7.4	Contextualising via Institutional Logistical Differences	37
1.7.5	Rational Choices in a Separating Equilibrium . .	39
1.7.6	Consequences of Quotation Inaccuracy	41
1.8	Conclusion	41
	Appendices	43
	A Appendix	44
A.0.1	Background to Parliamentary Speeches	44
A.0.2	OLS Estimate of the Opposition Dummy is a Lower Bound on the True <i>Magnitude</i>	45
A.0.3	Additional Tables and Figures	48
2	WOMEN ON BOARDS	64
2.1	Introduction	64
2.2	Background and Data	66
2.2.1	Data on Board Directors	66
2.2.2	Temasek Holdings and Government-Linked Companies	67
2.2.3	Share of Women in Parliament	70
2.3	Identification Strategy	72
2.4	First stage	78
2.4.1	Results	78
2.4.2	ADDITIONAL RESULTS AND ROBUSTNESS	79
2.5	Second stage	82
2.5.1	The Effect of Female Board Representation	82
2.6	Additional Results	87
2.6.1	Limits to Identification and External Validity . .	87
2.6.2	Alternative First-Stage Identification	88
2.7	Discussion	90
2.7.1	Inherent Limits to Causal Inference	90
2.7.2	A nudge is enough?	91
2.8	FINAL REMARKS	92

Appendices	94
B Data Appendix	95
3 WOMEN’S WAVE OR THE BLUE WAVE?	115
3.1 Introduction	115
3.2 Data	119
3.3 Context and Implications of MeToo	123
3.3.1 Determinants of Tweet Density	126
3.3.2 Selection of Women Candidates	127
3.4 Results	128
3.4.1 Empirical Strategy	128
3.4.2 Average Effect on Candidate Vote Share	130
3.4.3 Heterogeneous Effect/Backlash, by Existing Re- publican Support	131
3.4.4 Robustness	134
3.4.5 Back to the 2016 House Elections	135
3.5 Exploring Channels and Interpretations	135
3.5.1 Candidacy as Strategic Reaction	135
3.5.2 Turnout as a Channel	137
3.5.3 County-level Vote Changes	139
3.5.4 Marginal Districts	140
3.6 Discussion	142
3.6.1 What do the Tweets Measure?	142
3.6.1.1 Astroturfing?	142
3.6.1.2 Sexual harassment and disapproval of the Republican party (VOTER survey)	143
3.6.2 Changing Demographics	144
3.6.3 County-level Variation	145
3.7 Conclusion	146
3.8 Tables	147
Appendices	156
C Data Appendix	157
C.0.1 Data Details	157
C.0.2 Extra Figure and Tables	160

List of Tables

I	Summary Statistics	15
II	Description of Control Sets	17
III	BASILINE RESULTS FOR POLITICAL COVERAGE, ARTICLE-SPEECH LEVEL	23
IV	Baseline Results for Political Coverage, Quote-Level . .	26
V	Political Coverage of Backbenchers, Quote Level	29
VI	ADDITIONAL LANGUAGE-BASED MEASURES AS CONTROLS . .	32
VII	EXAMPLES OF SPEECHES AND QUOTES	34
A.0.1	DAILY NEWSPAPER SUBSCRIPTION	48
A.0.2	DIFFERENCES IN OBJECTIVITY AND POLARITY	49
A.0.3	POLITICAL COVERAGE DURING PRE-ELECTION PERIODS . . .	50
A.0.4	MINISTERIAL RANK BY PARTY	51
A.0.5	NEWSPAPER SECTION BY PARTY	51
A.0.6	SPECIFICATION CHECKS FOR QUOTE FRAGMENTS, ARTICLE- SPEECH LEVEL	53
A.0.7	SPECIFICATION CHECKS FOR QUOTE ACCURACY, ARTICLE- SPEECH LEVEL	54
A.0.8	SPECIFICATION CHECKS FOR QUOTE LENGTH, QUOTE-LEVEL	55
A.0.9	SPECIFICATION CHECKS FOR QUOTE ACCURACY, QUOTE-LEVEL	56
I	BREAKDOWN OF GLCs, BY YEAR AND SECTOR	69
II	SUMMARY STATISTICS	71
III	THE EFFECT OF WOMEN PARLIAMENT REPRESENTATION ON WOMAN BOARD REPRESENTATION	78
IV	THE EFFECT OF FEMALE BOARD REPRESENTATION ON ROE	83
V	THE EFFECT OF FEMALE BOARD REPRESENTATION ON FIRM VALUE	85

VI	THE EFFECT OF FEMALE BOARD REPRESENTATION ON SYSTEMATIC RISK	86
VII	SECOND STAGE RESULTS—ALTERNATIVE FIRST-STAGE	88
B.0.1	LIST OF DROPPED FIRMS	96
B.0.2	LIST OF FIRMS IN SAMPLE PERIOD 2000–17	97
B.0.3	SECTOR-INDUSTRY (GICS) BREAKDOWN	100
B.0.4	THE EFFECT OF GLC STATUS ON LOG OF FIRM VALUE	101
B.0.5	THE EFFECT GLC STATUS (BY OWNERSHIP %) ON FIRM VALUE	102
B.0.6	DESCRIPTION OF FIRM FINANCIAL CHARACTERISTICS FROM BLOOMBERG	103
B.0.7	TIER 1 GLCs MORE LIKELY TO HAVE AT LEAST 1 AND 2 WOMEN DIRECTORS	106
B.0.8	FALSIFICATION TEST—EFFECT OF WOMEN PARLIAMENT REPRESENTATION ON BOARD SIZE	106
B.0.9	TIER 1 GLCs HAVE HIGHER PERCENTAGE OF WOMEN DIRECTORS	107
B.0.10	FIRST-STAGE MAIN RESULTS, USING SHORTEST BOARD DURATION	107
B.0.11	FIRST-STAGE MAIN RESULTS, DROP MISSING INDIVIDUAL BOARD DATA	108
B.0.12	DIFFERENT MEASURES OF WOMEN PARLIAMENT REPRESENTATION	109
B.0.13	TIER 1 BOARD SEATS MORE LIKELY TO BE A WOMAN	110
B.0.14	ADDITIONAL ROBUSTNESS CHECKS FOR FIRST-STAGE	111
B.0.15	FALSIFICATION TEST—GOVERNMENT LINKAGE	112
B.0.16	THE (LAGGED) EFFECT OF WOMEN PARLIAMENT REPRESENTATION ON WOMAN BOARD REPRESENTATION	112
B.0.17	SECOND STAGE RESULTS—AT LEAST 1, 2, AND 3 WOMEN DIRECTORS	113
B.0.18	SECOND STAGE RESULTS—SECTOR-BY-YEAR TRENDS	114
I	SUMMARY STATISTICS	148
II	SELECTION OF TWEET DENSITY IN COUNTIES	149
III	THE EFFECT OF THE MeToo MOVEMENT ON CANDIDATE VOTE SHARE	150

IV	ROBUSTNESS	151
V	THE EFFECT OF THE MeToo MOVEMENT ON TURNOUT . . .	152
VI	CHANGE IN REPUBLICAN VOTE SHARE, DISTRICT-COUNTY LEVEL	153
VII	SELECTED DISTRICTS AND STATES	154
VIII	CORRELATION BETWEEN THE MeToo MOVEMENT AND INDIVIDUAL ATTITUDES (VOTER DATA)	155
C.0.1	Examples of Parsing Twitter User Geolocation	161
C.0.2	SELECTION OF WOMEN CANDIDATES INTO DISTRICTS	166
C.0.3	FULL REPORT OF INTERACTED COEFFICIENTS, FOR PARTY AND GENDER	167
C.0.4	ADDITIONAL ROBUSTNESS CHECKS	168
C.0.5	THE EFFECT OF THE MeToo MOVEMENT ON TURNOUT (FALSIFICATION— CHANGE IN TURNOUT PRESIDENTIAL ELECTION 2012–16) . . .	169
C.0.6	THE EFFECT OF MeToo ON INDIVIDUAL VOTING (VOTER DATA)	170
C.0.7	THE MEDIATED EFFECT OF THE MeToo MOVEMENT, 2016 HOUSE ELECTIONS	171
C.0.8	CHANGE IN REPUBLICAN VOTE SHARE, DISTRICT-COUNTY LEVEL (LOG TWEETS)	172
C.0.9	ADDITIONAL EFFECTS BY STATE AND DISTRICT	173
C.0.10	THE INTENSIVE AND EXTENSIVE MARGINS OF MeToo MOVE- MENT	174

List of Figures

I	Substring Accuracy Measure of Parliamentary Speeches	18
II	Effect Sizes of Opposition Status on Substring Accuracy	30
A.0.1	Graphics user interface with quote matching and extrac- tion	52
A.0.2	Effect Sizes of Opposition Status on Bag-of-words Accuracy	57
A.0.3	Count of quotes over the years by partisanship	58
A.0.4	Count of quotes over parliaments by partisanship	58
A.0.5	Quote length over time	59
A.0.6	Speech length over time	59
A.0.7	Article length over time	60
A.0.8	Bag-of-words quote accuracy measure over time	60
A.0.9	Distribution of quote length	61
A.0.10	Distribution of quote length at article-speech level	61
A.0.11	Distribution of accuracy measures	62
A.0.12	Distribution of speech and quote objectivity	62
A.0.13	Distribution of speech and quote polarity	63
I	INCREASES IN THE REPRESENTATION OF WOMEN IN PARLIAMENT	70
II	WOMEN DIRECTORS AND FIRM PERFORMANCE (BINNED MEANS)	74
III	INCREASES IN THE REPRESENTATION OF WOMEN ON CORPO- RATE BOARDS	76
IV	SUMMARY OF FIRST-STAGE RESULTS	80
B.0.1	Count of directors	104
B.0.2	At least 2 women directors	104
B.0.3	No parallel trends pre-2004	105
I	Caption for LOF	121
II	GEOGRAPHICAL DISTRIBUTION OF MeToo TWEETS IN 2018	122

III	GEOGRAPHICAL DISTRIBUTION OF WOMEN AND DEMOCRATIC VOTE SHARE	123
IV	WOMEN IN CONGRESS	125
V	REPUBLICAN VOTE SHARE AND MeToo MOVEMENT	126
VI	Caption for LOF	132
VII	Candidacy as Strategic Reaction	136
VIII	Caption for LOF	140
C.0.1	DISTRIBUTION OF MeToo TWEETS, BY COUNTIES	160
C.0.2	EXTENDED TIMELINE OF MeToo TWEETS, LEVELS	160
C.0.3	Candidacy as Strategic Reaction (Republican Women Challengers)	162
C.0.4	CORRELATIONS OF TWEETS AND REPUBLICAN VOTE SHARE, BY STATE (PART 1)	163
C.0.5	CORRELATIONS OF TWEETS AND REPUBLICAN VOTE SHARE, BY STATE (PART 2)	164
C.0.6	CORRELATION (BINNED SCATTERS) BETWEEN COUNTY DEMO- GRAPHICS, TWEETS, AND 2018 HOUSE ELECTIONS	165

Abstract

This thesis consists of three self-contained papers on the political economy of mass media and the economics of representation, all of which deal with the identification of subtle effects, if any.

The first paper explores the use of machine learning (ML) and natural language processing (NLP) in detecting political media slant in the mainstream media of Singapore. This paper finds robust evidence of slant towards the ruling party only when using newly defined measures of coverage *accuracy*—as opposed to the usual measures of coverage intensity. The additional methods from ML and NLP provide measures from the rich textual data, which help deal with identification.

The second paper tests the assertions that an increase in women representation on the corporate boards of directors improves firm financial outcomes. The identification strategy exploits the institution setting in Singapore where firms have varying linkages to the government, in the 2000-17 period where there is a substantial increase in women representation in politics. The main finding is that any observed link between female board representation and better financial performance can be attributed to permanent factors, such as corporate culture.

The last paper turns to the context of the U.S. elections right after the peak of the women's MeToo movement in 2018. Using returns from the House elections and county-level measures of the interest in the movement from Twitter data, this paper finds the expected advantage for Democratic women candidates and a disadvantage for the Republican men candidates, but only in those counties with high existing Republican support.

Introduction

This thesis comprises three discrete research questions related to the political economy of mass media and the economics of representation. The common element holding all three chapters together is the quantitative methods used to answer the questions.

In the first chapter, the question is whether politicians from different parties are represented differently by the mainstream media. A conventional approach in the social sciences and humanities discipline would be to ask stakeholders about their views and experiences on media slant. Such qualitative methods and case studies, however, may miss out on subtle effects. For instance, how can we meaningfully talk about differences in coverage accuracy, without first developing a notion of accuracy? How do we disentangle differences in media coverage arising from party affiliation versus differences arising from political office appointments? Differing representation of politicians in the media by coverage accuracy would be hard to demonstrate without compelling evidence from the quantitative approaches. Finding out the views of stakeholders is important, but additional value can be gained through quantitative methods that aim, as far as possible, to detect real effects.

As a contribution to the literature on media economics in the context of Singapore and more broadly to the political economy literature on media slant, chapter one adapts methods from the machine learning (ML) and natural language processing (NLP) toolbox. Chapter one uses supervised ML to scale data curation of media quotations of parliamentary speeches. To quantify quotation accuracy, chapter one adapts existing edit distance measures from the NLP to quantifying speech-to-quote accuracy. To account for differences in topics of parliamentary speeches, chapter one uses unsupervised ML to recover probabilistic associations

of speeches to clusters of speech topics. Finally, to account for how differences in speech characteristics, chapter one adapts computational linguistics on speech complexity, competency, and tone, which includes writing a custom and open-sourced programming package. This approach allows for objective quantification of coverage and accounting for confounders.

The quantitative methods complement richer qualitative approaches. To gain a better understanding of the media-political complex of Singapore for the first chapter, I conduct primary research through face-to-face interviews with agents from both sides of the media-political relationship: a senior political journalist and a Member of Parliament. The chapter also integrates written accounts provided by an examiner of a previous version of the thesis, based on his experience as a Nominated Member of Parliament. The combined insights substantially informed sections of the chapter that provide institutional context on the logistics involved in parliamentary procedures.

The second and third chapters also adopt a quantitative approach to answer their respective questions. The second chapter examines whether having more women representation on corporate boards materially improves firm financial performance. This chapter speaks to an established set of studies that is often criticised for being overly descriptive and not informative on the causal effects of increasing women representation. The nature of the endogeneity problem is that firms that appoint more women on boards likely have some other unobserved confounder that is also correlated with firm financial performance. One such confounder could be organisational culture, where firms that care more about fair hiring practices are also the firms that pay more attention to the type of best practices in business management that leads to a market advantage. Another confounder could be that firms that perform better in the market are firms that have a wider connection and pool of candidates to select from, and are thus in a better placed to appoint more women directors. Until the endogeneity problem is addressed, the results provide descriptive insights—on the differences in firm characteristics for firms with high women board representation and firms with low women board representation—rather than causal

insights. To approach causal inference, chapter two takes an instrumental variables approach. A plausibly exogenous instrument for women board representation that largely retains only the exogenous variation in women board representation can address the endogeneity problem and recover the causal effect of women board representation on firm performance.

The third and final chapter examines whether the grassroots MeToo movement benefited certain political parties and candidates in the 2018 midterm elections. The 2018 elections saw the largest representation of women in Congress, most of whom are from the Democratic party, and coincided with the buildup and peak of the MeToo movement on social media. A natural question arising here is whether the movement materially benefitted the electoral performance of certain parties or even certain candidates who fit the MeToo zeitgeist and led to the historic performance of women candidates. This chapter contributes to a small but nascent set of political economy literature concerning social media. The general question here is whether social media have the same impact as traditional mainstream media, which is well documented. While a connection between the MeToo movement and the performance of Democratic women candidates in the 2018 midterm elections appear obvious, the extent of a causal interpretation, if any, is not without the use of appropriate quantitative methods. The MeToo movement may be correlated with shifts in demographics that are also correlated with shifts in the party vote share. Any study that aims at documenting the impact of the movement must therefore account for the underlying demographics. It may also be the case that candidates that did well during the elections, would have done well anyway, if not for the movement because they inherently possess certain qualities.

Overall, the results reveal subtle insights and provide evidence against specious claims. In the first chapter, contrary to public perceptions, intensity in parliamentary speech coverage is similar for politicians of both parties. On the other hand, accuracy in parliamentary speech coverage is different for both parties, something more subtle and underappreciated in the discourse on media slant. This finding is

robust to controlling for the topics and characteristics of speeches derived from natural language processing methods. In the second chapter, any observed positive correlation between women board representation and firm performance systematically disappears once approaches that account for confounders are employed. This suggests a limited impact of women board representation on firm financial performance, whether positive or negative. The third chapter shows that the apparent advantage in votes that the Democratic women candidates get in the 2018 House elections also exists in the 2016 House elections before the peak of the movement. This points to the perceived advantage of individual candidates from the MeToo movement as artefacts of geographical variation in demographics and the MeToo movement. On the other hand, the MeToo movement conferred an advantage to the extent that the overall success of the Democratic women candidates is mediated by strategic candidate selection, especially against their male counterparts, and the overall advantage of the Democratic party is mediated by an increase in turnout.

Finally, all three chapters of the thesis layer qualitative discussion on top of the quantitative results by referencing richer descriptive accounts wherever appropriate, and all three chapters have been carefully contextualised and accurately positioned in their respective institutional settings.

Chapter 1

MEASURING POLITICAL MEDIA SLANT USING TEXTUAL DATA: EVIDENCE FROM SINGAPORE

1.1 Introduction

A common strategy in the literature that systematically measures political media slant is to compare differences in *intensity* of coverage (Larcinese et al. 2011; Puglisi 2011; Puglisi and Snyder 2011; Qian and Yanagizawa-Drott 2017; Qin et al. 2018). This paper develops a notion of coverage *accuracy* using a novel dataset on direct quotations of speeches from the Singapore Parliament, which I construct using methods from automated text processing, machine learning (ML), and natural language processing (NLP).

A rich literature already exists on the US media (e.g. in Groseclose and Milyo 2005; Gentzkow and Shapiro 2010) where they find a slight left-of-centre bias, but their findings cannot be directly extended to Singapore—where a monopoly supplies all daily newspapers and where a single party dominates the regular democratic elections. To test for political media slant in Singapore without relying on anecdotes, I assemble a new panel of direct quotation data by extracting the direct quotation of politician speeches in parliament, as reported in the flagship daily *The Straits Times*, and matching the quotes back to their originating speeches in

parliament. This data allows me to compare coverage accuracy of the ruling party and its opposition.

Specifically, this paper tests whether the *The Straits Times* quotes the opposition speeches from parliament with lower accuracy than their ruling-party counterparts. To quantify *coverage accuracy* any speech-quote pairing, I define two measures of quote accuracy using existing measures of edit distance.¹ The first accuracy measure this paper defines is a *substring* accuracy measure, which scores quote accuracy using the best partial-substring match—this is the quote-length substring within the speech that best matches the quote. The second measure, which is the *bag-of-words* accuracy measure, on the other hand, scores quote accuracy by using the subset of words common to both quote and speech. The scores go from 0 to 100 (increasing in accuracy).

The empirical approach is straightforward. With quotes as the unit of analysis, I regress quotation accuracy (and intensity) on an opposition dummy using the OLS specification. If there are systematic differences in coverage between ruling-party and opposition politicians, the opposition dummy will pick this up. The main finding is that the parliamentary speeches of the opposition politicians are quoted with lower accuracy relative to the speeches of the ruling-party politicians. Conditional on the observables, the opposition speeches are 1.5 to 2.4 points (by the substring and bag-of-words accuracy measures) less accurate than those of their ruling-party counterparts, which is 11.6% to 22.9% of the standard deviations. Compared to the average accuracy of 91.4 and 96.4 per quote, the opposition get quotes that are 1.6% to 2.5% less accurate.

To deal with identification issues and competing explanations, and as part of the methodological contribution in this paper, I draw on rich textual content. First, I use unsupervised ML to recover the topic distributions from the text data so that each parliamentary speech and each

¹Edit distance measures are solutions to the problem of quantifying similarity between two strings (of words in this case). A common measure is the Levenshtein distance. For example, the Levenshtein distance between *intention* and *execution* is five, because five operations (of insertion, deletion, or substitution of characters) is the minimum that is required to convert one string to the other. Applications of edit distance include NLP (e.g. spell-checking, machine translation, and mutual language intelligibility) and computational biology (e.g. similarities in DNA/RNA sequencing), among others.

news article have a probabilistic vector for topics (e.g. climate change or crime). Controlling for speech topics helps rule out the interpretation that observed partisan differences in coverage arise from partisan differences in speech content. Second, I use alternative construction of accuracy by removing stop words (e.g. "of", "the", "until") as part of the robustness checks. This helps rule out the case where the token distributions of opposition speeches skew towards the use of more stop words which can be ignored by journalists. A third alternative interpretation is that the opposition speeches are less coherent. To address this, I include controls of language competency from the quantitative linguistics literature, and the baseline conclusions do not turn on these controls.

To deal with any remaining bias, I turn to two bounding arguments. First, controlling for ministerial portfolio potentially introduces a compositional bias since it is also determined by the opposition status (as a case of bad controls Angrist and Pischke 2009). Here, the bias arises even if opposition status is random, and a set of institution-specific factors predicts that the OLS estimates are biased towards zero. The second bounding argument uses the proportional selection assumption (Altonji et al. 2005; Oster 2017)—and the data suggest that the unobservables drive the results in the same direction as the observables. This does not rely on the random assignment assumption and reinforces the OLS as conservative estimates on the extent to which opposition coverage is less accurate than that of the ruling party.

Overall, there is evidence of a subtle but systematic difference in coverage accuracy. At face value, the lower accuracy and higher fragmentation of quotations for the opposition suggest media slant towards the ruling party. Through primary research, I contextualise the statistical findings in institutional-specific media-political machinery where the ruling-party politicians grant media agents early access to their speech transcripts, but the opposition party does not. I finally argue that even if this institutional-specific machinery can explain away coverage accuracy, the machinery itself embeds private beliefs about a media that slants.

The proposed contribution of this paper is methodological. Using coverage accuracy to discern bias can be used for any media platform with a known transcript, and abstracts away from other media biases (e.g. gatekeeping) that are less testable. Moreover, even if assignment into political parties is determined by unobserved characteristics that also determine coverage, the confounding bias from accuracy as the basis of comparison is less severe than that using intensity, since direct quotation should be accurate in any case. In the media of Singapore, there would have been no systematic evidence of slant if the only measure is coverage intensity. To my best knowledge, this paper is the first that narrows down media bias to coverage accuracy using the text data from politician speeches.² Another contribution is the large-scale systematic evidence of political media slant in Singapore, which complements other richer descriptions of media and politics in the same context (e.g. George, 2012).

This paper otherwise relates most to the literature that estimates the size and direction of media bias without relying on anecdotes—in particular, to the literature that uses textual information from congressional speeches Gentzkow and Shapiro 2010, that takes an incumbent-government dimension (Larcinese et al. 2011; Durante and Knight 2012; Lott and Hassett 2014), that uses the intensity of coverage to discern bias (Larcinese et al. 2011; Puglisi 2011; Puglisi and Snyder 2011; Qian and Yanagizawa-Drott 2017; Qin et al. 2018), that looks at the interaction between the media and a dominant single-party (Enikolopov et al. 2011; Miner 2015; Qin et al. 2017, 2018), and the literature that finds a slight left-of-centre bias in an otherwise centrist US media (Groseclose and Milyo 2005; Bernhardt et al. 2008; Ho and Quinn 2008; Sutter 2012).

More broadly, this paper relates to the rich literature on the political economy of mass media—on how mass media can facilitate accountability (Besley and Prat 2006; Bernhardt et al. 2008; Ferraz and Finan

²A related contribution is on how machine learning has its own place in the applied econometric toolbox (as discussed in Mullainathan and Spiess 2017). Supervised machine learning in this study expedites the data collection, while unsupervised learning affords machine-annotated data free from a researcher's (my) subjective priors.

2008; Snyder and Strömberg 2010; Larreguy et al. 2014; Lim et al. 2015), have welfare effects (Besley and Burgess 2002; Strömberg 2004; Corneo 2006; Eisensee and Strömberg 2007), and impact electoral outcomes (Gentzkow 2006; DellaVigna and Kaplan 2007; Boas and Hidalgo 2011; Chiang and Knight 2011; Enikolopov et al. 2011; Gerber et al. 2011; Adena et al. 2015; Miner 2015).

Section 3.3 starts by providing a background to the political and media institutions of Singapore. Section 3.2 details the use of ML and NLP in constructing the panel and accuracy measures. Section 1.4 begins by discussing the empirical strategy before presenting the results at the article-speech level. Section 1.5 presents the results at the quote level. Section 1.6 rules out additional language-based explanations, Section 1.7 follows with discussion. Section 3.7 concludes.

1.2 Background of Political and Media Institutions

1.2.1 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 must be held within three months of the dissolution of parliament. Political candidates are elected into parliament based on plurality voting.

Group and Single-Member Constituencies. Since the late 1980s, most parliament seats are contested under group representation³—voters in group constituencies vote for a group slate instead of any single politician. In the 2011 general election, for instance, only 12 of the 87 parliament seats were contested as single members. Groups must be made of between 3 to 6 candidates, of which, at least one must be of an ethnic minority. In the sample, group sizes vary between 4 to 6. The regression analyses control for politician ethnic and group size.

³Also known as multi-member districts elsewhere.

Non-Elected Members of Parliament. Two other schemes were implemented by the early 1990s. One is the non-constituency member of parliament scheme, allowing the best-performing opposition losers to get seats. The other is the nominated member of parliament scheme, where non-partisan individuals are appointed by a selection committee. In the regression analyses, the ministerial controls include the politician type categorical variable, which includes the above two types of politicians in addition to other types (e.g. minister or parliamentary secretary).⁴

Political Competition in Recent Years. The dominant and ruling political party has always been the People's Action Party (PAP), which forms the government with little legitimate challenge from opposing parties since 1959. The election in 2011 however, cast a different hue. All but five seats were contested by the opposition. Moreover, the main opposition of the day—the Workers' Party—won a breathtaking 6 of the 87 seats in parliament. This marks the first time in Singapore's political history where an opposition party won a (5-member) group constituency, with the sixth seat won in a neighbouring single-member constituency. The margins of victory for the opposition were non-trivial: the group won with a 9.4% margin and the single member by a 29.6% margin. This political background provides a basis for the sample period of 2005–16, approximately five years before and after the landmark 2011 election.

1.2.2 Mainstream media background

Media-Related Regulations. The political dominance of the ruling party has allowed it to put in place media-related legislations, which directly regulate the publication of print media and potentially influence the content and tone of print media. The most direct regulation is the *Newspaper and Printing Presses Act (NPPA)*, first enacted in 1974. The NPPA requires that no newspapers are to be printed or published without permits. Moreover, the NPPA requires that newspaper companies

⁴Before April 2017, the non-constituency and nominated members have similar but less voting rights compared to elected members. From April 2017 onwards, non-constituency members of parliament will have the same voting rights as members of parliament, but this does not affect the 2005–16 sample.

be publicly listed with two classes of management shares—ordinary shares and management shares. Holders of management shares are entitled to 200 votes to the 1 vote allotted to the ordinary shareholder (Singapore Statutes Online 1974). These management shares can only be transferred to those with the approval of the government, potentially creating the perception of government-vetted nominees.^{5 6}

Indirect media-related legislations also potentially affect media coverage. The *Internal Security Act* bans subversive documents and publications. The *Official Secrets Act* bans publications containing state secrets. The *Sedition Act* encompasses a wide range of offences defined as crimes committed against the state. The *Defamation Act* covers libel. The interpretation of these laws however, leave plenty of room for ambiguity, lending weight to the perception that journalists practice self-censorship.⁷

Dailies in Singapore. Mainstream media in this paper mainly refers to the daily newspapers (or just dailies). The Singapore Press Holdings (SPH) company wholly owns eight of the nine local dailies in Singapore. The ninth daily—the *Today* newspaper—is jointly owned by SPH (40%) and MediaCorp (60%).⁸ The flagship daily in Singapore is *The Straits Times*, an English publication with the largest readership (Table A.0.1).⁹

As per the requirements of the Newspaper and Printing Presses Act, SPH is publicly listed on the Singapore Stock Exchange, with 99.9% of

⁵This concept of separate share classes is not unique. *The New York Times* for instance has class B shares for family owners with more voting privileges than class A shareholders.

⁶The Workers' Party however, have had no issues getting permits for their newsletter—*The Hammer*.

⁷Another media-related legislation implemented in 2019 is the *Protection from Online Falsehoods and Manipulation Act*, which gives the Executive the power to order for a correction of a falsehood shared online, with recourse through an appeal system either through the Executive or through the judicial system. This legislation is more directly related to the non-mainstream social media, with its first use on Facebook (see <https://www.channelnewsasia.com/news/singapore/brad-bowyer-facebook-post-falsehood-pofma-fake-news-12122952>), though it does impose a general and additional restriction on media content.

⁸From a 2004 press release: http://www.sph.com.sg/media_releases/150 (Retrieved: 30 May 2017).

⁹The English daily with the next highest readership—*The New Paper*—has a print subscription of an order smaller (approximately 70 thousand compared to 300 thousand).

ordinary shares held by the public. The remaining 0.1% are management shares mostly held by financial institutions. A relatively small number of management shares is held by the CEO and directors of SPH (less than 0.1% of management shares).¹⁰ In other words, most management shareholders (those with 200 votes per share) are profit-maximisers rather than owners with ideological agendas (in the vein of Anderson and McLaren 2012; Durante and Knight 2012), though the fact that notable senior management positions are occupied by those with links to the government has also been documented (BBC News 2013; George 2012). Section A.0.1 furnishes details on the background of Parliamentary procedures and speeches.

1.3 Data

1.3.1 Downloading and Matching of Quotes to Speeches

Downloading the Textual Data. For *The Straits Times* news articles, I query Factiva using the names of politicians who served in parliament in the years 2005–2016. This returns 62,132 unique news articles (as identified by the title-date tuple) which mention at least one of the active politicians in the sample period. The parliamentary procedure transcripts are from the publicly available and official repository. I automate the extraction of speaker-speech chunks by parsing the HTML tags. The resulting data, however, does not make a distinction speeches on different types of parliamentary procedures, such as bills or motions. Section A.0.1 provides some background.

Supervised Classification of Articles Containing Quotes. Of the 62,132 articles containing mentions of the politicians, only a fraction will contain relevant direct quotations from parliamentary speeches. To automatically and accurately identify these relevant news articles containing the quotes of politicians, I use the *random forest* algorithm—a supervised learning classifier which takes an initial learning set of 1,419

¹⁰http://sph.listedcompany.com/main_shareholder.html.

manually-labelled articles—to predict which of the 62,132 articles contain quotations from parliament.¹¹ The random forest classifies the articles with 89% accuracy, 83% recall rate, and 70% precision rate. This classification yields 3,425 news articles (5.5% of the 62,132) with at least one quotation of a parliament speech.^{12 13}

Matching Quotes to Parliamentary Speeches. I define *quotes* as strings within single or double inverted commas, or strings after speech colons. From the above 3,425 articles with parliament quotes, I automatically extract the quotations and match them to their originating speeches as follows (implementation is through a simple custom-written GUI, illustrated in Figure A.0.1).

Let f be the Ratcliff-Obershelp pattern recognition algorithm, an existing edit distance measure. Like other edit distances, it measures the minimum edits required to change a string, of words in this case, into another, providing a measure of string similarity or probability alignment between two sequences, and has the well-behaved properties of a metric (triangle inequality, non-negativity, and symmetry (Jurafsky and Martin 2000)).¹⁴ For a quote q and a set of speeches $S = \{s_1, \dots, s_L\}$, the matched speech is $s^* \in S$ which satisfies:

$$s^* = \arg \max_{s_\ell} \left\{ \mathcal{F}_1(q, s_\ell) \mid s_\ell \in \{s_1, \dots, s_L\} \right\},^{15}$$

¹¹The random forest classifier is a bag of decision trees, where individual trees are constructed by bootstrapping over model inputs and the learning sample for a good bias-variance trade-off, and where each tree carry a vote on how to classify an article. I use class weights to mitigate learning-biases because of the news articles' class-imbalance—that a disproportionately large fraction of the articles does not contain quotes. I tune and evaluate out-of-sample prediction performance using k -fold cross-validation.

¹²It is also possible to use a *dummy* or *naive* classification which classifies as *yes* if a news article contains *parliament**, and *no* otherwise. The naive classifier has an accuracy score of 0.78 (compared to 0.89 from the random forest, a 14% improvement), a better recall score of 0.93 (compared to 0.83), a precision score of 0.47 (compared to 0.70, a 49% improvement), and an F-score of 0.66 (compared to 0.93, a 41% improvement).

¹³For a sense of scale, if the 1,419 manually-labeled news articles took one week to do, parsing the 62,132 news article would have taken approximately 43 weeks to for a single person.

¹⁴The Ratcliff-Obershelp pattern recognition algorithm is one among many string metrics in the computational literature. This paper uses the Ratcliff-Obershelp because it is an efficient built-in method in the Python ecosystem and that there is no reason to believe that other base string metrics would substantially change the baseline conclusions in this paper.

¹⁵ $\mathcal{F}_1$ is the quote accuracy score (1.1) defined in the following subsection.

where $\mathcal{F}_1$ is a composite function of f .¹⁶

This matching process yields 14,903 pairs of quote and speech, from 5,227 speeches made by 204 politicians, over the sample period 2005–16.

1.3.2 Quantifying Quotation Accuracy

The immediate issue with using the baseline measure f directly is that it will likely return an accuracy (similarity) score of zero because the quote needs to be edited (or added to) substantially before it resembles the original speech. In what follows, I define two accuracy measures based on the pre-existing Ratcliff-Obershelp edit distance measure. For the two accuracy measures defined below, I am not aware of them as standard or pre-existing measures.¹⁷ Mnemonic names are given for ease of reading.

Substring Quote Accuracy Measure. The first way I deal with this is to score quote *accuracy* according to the quote-length substring in the speech that best matches the quote. Let q be the shorter string and s the longer string, with lengths m and n , respectively. The substring accuracy measure locates the m -length substring within s that best matches q and scores accuracy according to this best partial substring match. Formally, substring accuracy measure for two strings q and s is:

$$\mathcal{F}_1(q, s) = \max_i \left\{ f \left(q, s^{(i:i+m-1)} \right) \mid i \in \{1, \dots, n - m + 1\} \right\}, \quad (1.1)$$

where f is the Ratcliff-Obershelp measure, increasing in accuracy. The superscript $(i : i + m - 1)$ indicates the location of the substring of length m . $\mathcal{F}_1$ goes from 0 to 100, increasing in accuracy (because f also goes from 0 to 100).

Concretely, if there is a quote $q = \text{"a fundamental relook"}$ which is 20-character long and the originating speech recorded is $s = \text{"...is taking"}$

¹⁶In practice, I manually oversee this matching by looking at the best matches. Virtually all quotes are matched to a parliamentary speech that occurred a 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 speech and media report.

¹⁷They may, however, be developed ad-hoc and used in other private applications that relate to fuzzy searches.

a more fundamental relook at this regulation framework and see how best we can support this strategy ...", then substring accuracy measure searches for a 20-character string in s that best matches q , which is the 20-character substring [mor]"*e fundamental relook*", and then computes the edit distance between "*a fundamental relook*" and '*e fundamental relook*'. This score then forms the quotation accuracy score between the quote q and the originating speech s .

The substring accuracy measure merely formalises how direct quotations work. If a quote is verbatim, there should always be a quote-length substring in the speech that perfectly resembles the quote (perfect score of 100).¹⁸ Even if the quote is not verbatim, there should still be a quote-length substring in the speech that closely resembles the quote, as illustrated above.

Bag-of-words Quote Accuracy Measure. A different concern is when quotes have words taken from different parts of the same speech or when the order of words is not preserved. A concrete example is the quote "*...can afford their education in a public university*" which came from the speech "*...can afford to attend Government-funded universities ...*". Here, misquotation arises because a phrase that should appear in the verbatim quote ("*Government-funded*") is replaced with another phrase ("*public university*") that was used elsewhere in the same speech but otherwise refers to the same thing.

To allow an accuracy score more forgiving to type of error above, I define the second quote accuracy measure—bag-of-words accuracy measure. With slight abuse of notation, I first define a new (third) string $c = q \cap s$, as the string containing words common to both q and s . All three strings are then sorted according to alphabetical order to get $\tilde{q}$, $\tilde{s}$, and $\tilde{c}$. Accuracy score (1.2) is then the maximum of the scores from all paired

¹⁸Discounting for punctuations which cannot be heard directly. I remove punctuations in pre-processing, as is standard.

combinations of the three strings:¹⁹

$$\mathcal{F}_2(q, s) = \max \left\{ f(\tilde{q}, \tilde{s}), f(\tilde{q}, \tilde{c}), f(\tilde{s}, \tilde{c}) \right\}. \quad (1.2)$$

This second measure is useful for two reasons. Besides edits for syntactic flow and grammatical demands, the bag-of-words accuracy measure is more forgiving when the quote is not verbatim, because of genuine slips in attention by the reporting journalist during the sitting, or because the transcript contains a (slightly) different version of the speech. Another reason for the bag-of-words accuracy measure, is because the article-speech level analyses will concatenate all quotes coming from a speech and reported in an article as one observation. The substring accuracy measure is no longer sensible in this case, but the bag-of-words accuracy measure still offers a measure of quote accuracy by comparing the concatenated quotes to the originating speech using words common to both.²⁰

Comparison to Other Records of Quotation Accuracy. Based on the quote accuracy scores computed using the substring accuracy measure, 19.8% of the articles contain some objective form of misquotation ($\mathcal{F}_1 < 100$). The proportion of articles with misquotations drops to 14.9% if perfect accuracy is defined as either $\mathcal{F}_1$ or $\mathcal{F}_2$ having a perfect score of 100. These figures are comparable to the 13.1% to 18.2% found in journalism studies of direct quotation accuracy, such as (Berry 1967; Marshall 1977), where they typically send out surveys to persons (most of whom are not politicians) mentioned in the news article and ask if they were accurately quoted.

At the quote level of the data in this paper, about 59.6% of the quotes contain some objective form of misquotation. Even with errors in the recording and matching of speeches and quotes in the data (e.g. HTML tags not always consistently used in the backend), this figure is still

¹⁹The max of the three components is taken instead of a weighted average for two reasons. One is that a weighted average requires specification of the weight for each component, and there is no clear way to do this. The other is that the measure is meant to be lenient in the first place—hence the maximum.

²⁰As is customary, the computations under the hood remove punctuations and normalise all strings to lowercase (since punctuations and casings can neither be seen nor heard in a speech).

much lower than the 90% misquotations found through a check with tape recordings in Lehrer (1989) at the sentence level (which I interpret as synonymous to the individual quote fragments in this paper). The main explanation for this is likely the context—where coverage of political speeches is more accurate than general. Table VII lists concrete examples of quote accuracy score for a sample of speech-quote pairings as part of the discussion in Section 1.7. These examples further provide a sense of how the quote accuracy measures align with human intuitions of quotation accuracy.

1.3.3 Topic Distribution of Parliamentary Speeches and News Articles

Retrieving Latent Topics. I use the *Latent Dirichlet Allocation* (LDA, Blei et al. 2003) model, an unsupervised learning algorithm, to recover clusters of topics from the news articles and parliament speeches. A *topic* in the information retrieval literature refers to a collection of words that occur frequently together. Three examples learned by the topic model trained on the sample parliament speeches are:²¹

1. $\langle \textit{cpf}, \textit{retirement}, \textit{minimum_sum}, \textit{saving}, \textit{cpf_saving} \rangle$
2. $\langle \textit{police}, \textit{home_team}, \textit{officer}, \textit{crime}, \textit{inquiry} \rangle$
3. $\langle \textit{premium}, \textit{medishield_life}, \textit{medishield}, \textit{insurance}, \textit{insurer} \rangle$

Every parliament speech in the sample will have some probabilistic association to one of K topics, where K is a pre-defined total number of topics that a speech can draw from. The sum of the probabilistic topic association of each speech $\sum_K \rho_k$ must by definition be equal to 1. In practice, I set the parameter of the Dirichlet distribution α to be very low, embedding the assumption that parliament speeches are unlikely to have high associations to more than one topic. Similarly, I train a

²¹These are taken from the LDA model trained on the parliamentary speech corpus on $K = 92$ topics. The words/phrases in the angular brackets are the top 5 words for topic 4, 18, and 31, with the relevance metric at 0.5. The phrases are words joint by '_', these are word that co-occur frequently and "glued" before the model is trained. The visualisation of the LDA results from the sample textual data is available at <https://lsys.github.io/media-lda/ldavis.html>.

separate topic model for the news article corpus, attributing to each article a topic vector. The topic association of quotes comes from the parliament speech topic model.

No Human Input Required. One major advantage of the LDA is that it is unsupervised—no human input is needed to impose a topic structure before or after the discovery of topics. Besides saving resources on the manual classification of thousands of speeches and news articles, the classification avoids researcher-induced bias on the content of speeches and articles.

Choosing Optimal Number of Topics. A major disadvantage in practice is in choosing the number of topic clusters K —a model hyperparameter that is not optimised within the model. To deal with this, I automate evaluation using topic coherence scores that correlate well with human intuition of topic coherence (Chang et al. 2009). I train topic models with K starting from 2, increasing in steps of 2, up till 100, and then select K^* using the highest topic coherence score (favouring lower K 's). The baseline results use topic controls from the $K = 92$ parliament speech topic model, and the $K = 40$ news article topic model. There is some judgement involved in choosing K . In the specification checks, I show that the baseline conclusions do not change based on the choice of K .

Implementation. Pre-processing of the textual data and implementation of the LDA are done using the *SpaCy* and *gensim* libraries. In pre-processing, a simple algorithm is passed over the text looking for words that co-occur frequently, identifying phrases such as "minimum sum" in topic 1 above, where the meaning is very different if the two words are considered separately. The data appendix provides more details.

Summary Statistics. Table I presents the summary statistics. 1,106 of the 14,903 (7.4%) individual quote fragments come from the opposition. Opposition quotes are, on average, from politicians who are younger and have shorter political tenures. The quotes also come from a higher proportion of women. Notably, panel B shows that the opposition has shorter speeches. This might explain the perception of how the opposition seems to get less intense coverage, when those perceptions do not

TABLE I
Summary Statistics

	Full sample			Non-opposition			Opposition politicians			Non-opposition – Opposition
	N	Mean	Std Dev.	N	Mean	Std Dev.	N	Mean	Std Dev.	ρ -value
<u>Panel A. Main outcomes</u>										
Quote (word count)	14,903	20.94	38.45	13,797	21.01	39.63	1106	20.05	18.23	0.13
Log of quote	14,903	2.68	0.95	13,797	2.69	0.94	1106	2.57	1.05	0.00
Substring quote accuracy measure	14,903	91.41	12.85	13,797	91.45	12.86	1106	90.80	12.71	0.10
Bag-of-words quote accuracy measure	14,903	96.40	10.52	13,797	96.50	10.26	1106	95.11	13.26	0.00
<u>Panel B. Length controls (word count) for textual data</u>										
Paragraph (of speech)	14,903	103.50	60.79	13,797	104.23	61.69	1106	94.37	47.23	0.00
Speech	14,903	2112.39	1824.96	13,797	2188.47	1858.28	1106	1163.35	909.74	0.00
Article	14,903	625.36	263.42	13,797	619.65	255.48	1106	696.60	339.70	0.00
<u>Panel C. Other control variables</u>										
Age	14,887	51.47	6.96	13,781	51.69	6.70	1106	48.74	9.23	0.00
Age ²	14,887	2697.39	702.22	13,781	2716.36	676.53	1106	2460.97	935.37	0.00
Tenure	14,903	10.97	7.55	13,797	11.11	7.48	1106	9.18	8.10	0.00
Tenure ²	14,903	177.22	207.00	13,797	179.42	208.02	1106	149.77	191.88	0.00
Female	14,903	0.17	0.37	13,797	0.16	0.36	1106	0.28	0.45	0.00
Rank	14,903	5.39	3.44	13,797	5.00	3.27	1106	10.28	0.45	0.00
Translations	14,903	0.01	0.12	13,797	0.01	0.11	1106	0.04	0.19	0.00
Group size	14,428	4.77	1.26	13,322	4.87	1.14	1106	3.58	1.89	0.00
Voters	14,428	126,159.37	38,920.13	13,322	128,444.28	36,166.56	1106	98,637.13	56,458.30	0.00
Vote	11,206	74,347.06	28,957.89	10,100	77,327.63	27,531.60	1106	47,128.50	27,440.08	0.00
Vote (%)	11,206	64.42	7.92	10,100	65.73	6.88	1106	52.42	6.60	0.00
Margin	10,900	35,389.49	22,908.53	10,100	37,525.17	22,418.44	800	8426.44	4545.91	0.00
Margin (%)	10,900	30.02	14.37	10,100	31.48	13.78	800	11.52	7.07	0.00

account for differences in the base speech length. I address selection issues in the following section.

Overall, the final (unbalanced) panel data includes 14,903 quote level observations over 12 years, 4 parliaments, 3 general elections, 3 by-elections, with 3,425 newspaper articles, 204 politicians, and 5,130 parliamentary speeches.

1.4 Article-Speech Level Results

1.4.1 Empirical Strategy

The first level of analysis is at the article-speech level, where all quotes originating from a speech s reported in article r are mechanically treated as a single observation. For example, if an article contains two quotes (fragments) from the same speech, these two quotes are concatenated and taken as one observation. The combined length of these two quotes is the first (of three) outcome measure—the intensity of coverage. For the second outcome measure—the *count* of quote fragments from a speech—this is recorded as 2. The third is quote accuracy measure (1.2) based on the common set of words between the concatenation of the two quotes and their originating speech.²² At this level, there are 204 politicians, 3,425 articles, and 5,227 speeches over 12 years, for a total of 7,098 politician-article-speech level observations, of which, 640 are from the opposition.

To compare the above three coverage measures between the opposition and the ruling-party, I estimate the OLS model:

$$y_{irst} = \alpha + \beta opp_i + \sum_{k=2006}^{2016} \alpha_k year_{kt} + \sum_{\ell=11}^{13} \alpha_\ell parl_{\ell t} + \gamma' \mathbf{X}_{irst} + \varepsilon_{irst}, \quad (1.3)$$

where y_{irst} is one of the above three measures of media coverage for politician i 's speech s at time t , reported in article r . The estimand of

²²At the article-speech level, the quotation accuracy is as defined by accuracy measure (1.2) which is computed based on words common to both speech and quote(s), because using best partial string match of a concatenation of multiple quotes is meaningless.

TABLE II
Description of Control Sets

A. Individual controls	
Gender	Gender of politician i
Race	Race of politician $i \in \{\text{Chinese, Indian, Malay, Eurasian/Others}\}$
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 $\in \{\text{Mon, Tue, Wed, Thu, Fri, Sat, Sun}\}$
Section	Dummies for article r 's section $\in \{\text{Singapore, Prime News, Top of the News, Home, ST, Insight, News, Money, Think, Review - insight, Sports, Opinion, World, Others}\}$
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	Number of votes for politician i 's constituency
Votes (%)	Percentage of votes for politician i 's constituency
Margin	Number of winner's vote – number of loser's votes in the politician i 's constituency
Margin (%)	Ratio of <i>majority</i> to the number of valid votes

Notes—Sections appearing no more than 6 times in the sample are collapsed under *Others*. Table A.0.5 shows the distribution of sections. In the rare instances where a politician holds two different ranks for two different ministry portfolio, the higher of the two ranks is recorded. Portfolios under non-extant ministries are mapped to their modern day equivalents. For instance, the Ministry of Information, Communications and the Arts (MICA) portfolio is mapped to the current day Ministry of Communications and Information (MCI) portfolio.

interest is β , the coefficient of the opposition dummy variable that takes on value 1 if politician i is from *any* opposition party. Identification issues notwithstanding, a negative β estimate suggest a systematic political media slant towards ruling-party politicians.

β in equation (1.3) unfortunately does not carry an unambiguous causal interpretation, since a candidate's choice of party is unlikely to be orthogonal to potential media coverage. Nevertheless, I control for a set of covariates that likely determines how the media covers politician speeches and that are also correlated to partisanship, including controls that originate from the richer textual data (Table II).

First, equation (1.3) includes year and parliament fixed-effects. If there are trends in media consumption or newsroom operations, the year fixed-effect will capture them. Quotes, for instance, exhibit a small

FIGURE I
Substring Accuracy Measure of Parliamentary Speeches

but gradual rise in accuracy over time (Figure I). Parliament fixed-effects capture potential trends in political sentiments and changes in parliament composition. The 12th Parliament for instance, had an increase in women representation (even in the absence of binding gender quotas), and an increase in opposition representation.

β in equation (1.3) above is not identified with individual fixed-effects since no candidate in the sample switched from the opposition to the ruling party (and vice versa). To mitigate concerns about individual characteristics driving observed differences between opposition, I include instead the politicians' gender, race, and a quadratic for age and political tenure. X_{it} also includes the full interaction of politician-type (e.g. minister or parliamentary secretary) and ministry portfolio of politicians (e.g. health or education). Certain politicians, for instance, will simply get more coverage bandwidth because of seniority. The portfolio accounts for some of this effect. These are all recorded as at the date of the speech given.

For the set of news article controls, there are publication day-of-the-week dummies for variations in newsroom operation by day. News section dummies control for the section in which a news article appears in, and

a translation dummy controls for quotes translated from a vernacular to English (as stated in the transcripts).²³ Standard errors are adjusted to allow for clusters within each newspaper article r .²⁴

Using Text Data to Deal with Alternative Interpretations. Here I consider the "political speech content" interpretation, "trivial words" interpretation, and the "speech coherence" interpretation.

One concern is related to partisan ownership of political topics (in the sense of Petrocik 1996; Puglisi 2011). If $\hat{\beta}_{OLS}$ is negative, suggesting that opposition coverage is less favourable, perhaps this arises because of differences in the political focus of the parties. For example, the ruling party may tend to speak more about political issues of greater interest or immediate relevance, and thus get more favourable coverage. Even with quotation accuracy as an outcome measure, perhaps accuracy suffers because journalists take mental breaks during (opposition) speeches that are less interesting, and not because of party status.

To deal with this concern, I control for the topics of the speeches and news articles using the the output from the LDA (Section 1.3.3). Concrete examples of topics learned from the parliamentary speech corpus, represented by the top five words are:

1. $\langle \text{cpf, retirement, minimum_sum, saving, cpf_saving} \rangle$
2. $\langle \text{police, home_team, officer, crime, inquiry} \rangle$
3. $\langle \text{premium, medishield_life, medishield, insurance, insurer} \rangle$
4. $\langle \text{student, school, learn, education, teacher} \rangle$
5. $\langle \text{fare, bus, public_transport, commuter, operator} \rangle$
6. $\langle \text{sport, athlete, community, youth, sports} \rangle$

The baseline results always include the topic distributions. While the output from the LDA says nothing about partisan ownership of topics, it is sufficient to remove variations in coverage that are correlated with

²³Speeches of ruling-party politicians for example, may involve more discussions on geopolitical entities and affairs, and these are usually reported in certain newspaper sections (e.g. *World*). Table A.0.5 shows the distribution of newspaper sections by partisanship. Opposition politicians, for instance, never appear in *World*, *Money*, or *Opinion* sections in the sample.

²⁴In the robustness checks in Sections 1.4.3 and 1.5.2 also allows for standard errors to be clustered by parliamentary speech or the reporting journalist.

the content of speeches and news articles. This mitigates concerns that differences in the political speech content are driving the results.²⁵

The second concern has to do with the usage of words in the opposition speeches. Perhaps the opposition uses more trivial words in their speeches, and journalists can ignore these words for efficiency, or even clarity, while retaining the core meaning of the speech. Mechanically, this would imply lower accuracy scores even if the journalist had changed the words for the sake of clarity. To mitigate such concerns, I perform robustness tests (in Sections 1.4.3 and 1.5.2) using alternative constructions of the accuracy scores by removing stop words in the pre-processing stage. The baseline conclusions are unaffected.²⁶

A third concern is language competency. For example, it may be the case that the speeches of the opposition are less coherent than those of the ruling party. Hence journalists have greater difficulty following and quoting opposition speeches accurately. To address this concern (in Section 1.6), I tap on the quantitative linguistics literature that estimates English language level based on the distribution of words in a text. Controlling for readability of the speech transcript and language sophistication using lexical richness measures barely attenuates the baseline estimates.

Bounds Using Differences in Ministerial Composition. To deal with any residual bias in the OLS estimates, I turn to two bounding arguments. First is a compositional bias that arises even if party assignment is random. The compositional bias occurs because the opposition status limits ministerial profile. Opposition politicians for instance, never get a ministerial rank that is higher than the base rank. In addition, opposition politicians hold no ministerial portfolios (e.g. Education or Health).²⁷ Hence party status determines both media coverage and ministerial characteristics.

²⁵The clustering of the topic distribution of the speeches and news article is apolitical, so although a natural tendency is to see if there are additional partisan differences in contentious political topics, the output so the LDA says nothing about which topics are contentious and which are not.

²⁶It may be worth noting that direct quotations should be verbatim even in the presence of errors or language inefficiency, so usage of stop words should have been preserved anyway.

²⁷Unlike the typical Westminster system, there is no shadow cabinet.

Let the opposition dummy be $D_i \in \{0, 1\}$, and a rank dummy be $r_i \in \{0, 1\}$, and abstract away from the other covariates X_{it} . The observed difference in coverage between an opposition politician ($D = 1$) and a ruling-party politician ($D = 0$), with both of equal base rank ($r = 0$), can then be decomposed as:²⁸

$$\mathbb{E}[y_{1i} - y_{0i} | r_{1i} = 0] + \mathbb{E}[y_{0i} | r_{1i} = 0] - \mathbb{E}[y_{0i} | r_{0i} = 0], \quad (1.4)$$

where the first term is the true causal effect (or slant) as the difference between the potential coverage of a candidate assigned to the opposition party and the ruling party, conditional on potential assignment to the base rank as an opposition.

The compositional bias comes from the second and third terms, and institution-specific factors help sign the bias as follows. The second term is the potential coverage of a ruling-party politician had he been given assignment to the base rank ($r = 0$) as an opposition. Since this assignment would apply to all ruling-party politicians (the opposition only ever gets base rank), this is simply the unconditional average coverage of ruling-party politicians— $\mathbb{E}[y_{0i} | r_{1i} = 0] = \mathbb{E}[y_{0i}]$. The third term is the potential coverage of a ruling-party politician conditional on assignment to the base rank ($r = 0$), and this I argue implies an *inferior* candidate, which means the conditional expectation is below the unconditional one— $\mathbb{E}[y_{0i} | r_{0i} = 0] < \mathbb{E}[y_{0i}]$. Taken together, this decomposition in equation (1.4) implies an *attenuation bias* if:

- (i) the true effect (first term) is negative, $\mathbb{E}[y_{1i} - y_{0i} | r_{1i} = 0] < 0$, which is the case if there is media slant favouring the ruling party;
- (ii) the second term is the unconditional expectation of potential coverage of a ruling-party politician, $\mathbb{E}[y_{0i} | r_{1i} = 0] = \mathbb{E}[y_{0i}]$, since assignment to base rank conditional on opposition assignment applies for *all* ruling-party politicians; and
- (iii) the third term is lower than the unconditional expectation of potential coverage of a ruling-party politician, $\mathbb{E}[y_{0i} | r_{0i} = 0] < \mathbb{E}[y_{0i}]$, if on average, inferior ruling-party candidates are less likely to take on

²⁸The appendix in Section A.0.2 provides more details of this decomposition based on bad controls (Angrist and Pischke 2009).

senior positions, and coverage is increasing in quality of candidate. Hence the second and third terms are $\mathbb{E}[y_{0i}] - \mathbb{E}[y_{0i}|r_{0i} = 0] > 0$, which goes in the opposite direction of the true effect.

Bounds Using Proportional Selection. If party assignment is not conditionally random, one can interpret a negative $\hat{\beta}_{OLS}$ as a selection bias, where for whatever institutional reason, candidates that are more competent or hold themselves better in public select into the ruling party, leaving the opposition with an inferior set of candidates. If this is the case, then opposition coverage would naturally be less favourable if coverage is increasing in the calibre of the politicians.²⁹

Since calibre is unobserved, amongst other factors, the second bounding argument uses the proportional selection assumption (Altonji et al. 2005; Oster 2017) which provides a formal bound on OLS estimates based on how the movement in OLS estimates when additional observables are added is informative about selection on unobservables, *after* normalising for corresponding movements in the variation explained. The formal bound developed in Oster (2017), assuming equal selection on both observables and unobservables, is:

$$\beta^* = \tilde{\beta} - \left[\overset{\circ}{\beta} - \tilde{\beta} \right] \frac{R_{max} - \tilde{R}}{\tilde{R} - \overset{\circ}{R}}, \quad (1.5)$$

where $\tilde{\beta}$ is the OLS estimate of β with all controls included, and $\tilde{R}$ is the corresponding R^2 . $\overset{\circ}{\beta}$ and $\overset{\circ}{R}$ corresponds to the OLS results without controls. $R_{max} = 1.3$ as recommended in Oster (2017).³⁰ I show below in Section 1.5.3 that the bounds computed using equation (1.5) has a larger magnitude than the OLS estimates ($\beta^* < \tilde{\beta} < 0$), reinforcing the argument of an attenuation bias even if party selection is not random.

1.4.2 Baseline Results, Article-Speech Level

Table III presents the baseline results at the politician-article-speech level, where the three panels present results for the three outcome

²⁹Tan (2014) provides some background on the institution-specific features which confers the ruling party of Singapore an advantage in selecting candidates.

³⁰A larger R_{max} only further reinforces the findings.

TABLE III
BASILINE RESULTS FOR POLITICAL COVERAGE, ARTICLE-SPEECH LEVEL

	Log of quote length by word count		Count of quote fragments		Bag-of-words quote accuracy measure	
	(1)	(2)	(3)	(4)	(5)	(6)
Opposition	0.052 (0.064)	0.038 (0.074)	0.328*** (0.089)	0.318*** (0.109)	-2.499*** (0.785)	-2.554*** (0.832)
<i>Controls</i>						
Time fixed-effects	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Article controls	Yes	Yes	Yes	Yes	Yes	Yes
Topic controls	Yes	Yes	Yes	Yes	Yes	Yes
Ministerial controls	Yes	Yes	Yes	Yes	Yes	Yes
Electoral controls	—	Yes	—	Yes	—	Yes
<i>F-statistics</i>						
<i>F</i> -statistic, year fixed-effects	2.682***	1.613*	1.753**	2.127**	4.589***	2.033**
<i>F</i> -statistic, individual controls	0.397	1.153	1.322	2.775***	2.490**	1.855*
<i>F</i> -statistic, article controls	3.290***	2.538***	5.370***	4.487***	3.189***	2.248***
<i>F</i> -statistic, topic controls	2.762***	2.413***	2.314***	2.164***	1.563***	1.440***
<i>F</i> -statistic, ministerial controls	2.694***	2.722***	5.106***	4.987***	2.360***	1.826***
<i>F</i> -statistic, electoral controls		0.382		1.660		1.476
<i>Mean of dependent variable</i>	3.306	3.301	2.101	2.119	96.326	96.587
<i>R</i> ²	0.181	0.199	0.241	0.256	0.118	0.139
<i>N</i>	7,087	5,143	7,087	5,143	7,087	5,143

Notes—Observations are at the article-speech level, where 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 from a speech made in a vernacular to English. 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 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. Politicians who had won by default (no opposition contest) in the most recent election have no electoral data. *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.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

variables. Each panel shows two sets of results; without and with electoral controls.

The results suggest some slant, with indications of higher fragmentation and lower quotation accuracy for the opposition politicians. In column (1), conditional on the time fixed-effects (both year and parliament), speech and article lengths, and the individual, ministerial, article, and topic controls, the opposition politicians get quotes that are 5 log points longer, suggesting that the opposition politicians get more coverage, although this is not statistically significant at the 10 percent level. In column (3), opposition politicians get about a third more quote fragments from each article-speech observation (about 18% of the standard

deviation), which is significant at the 1% level.

In the third panel, the outcome variable is the bag-of-words accuracy measure. This measure quantifies accuracy based on the subset of words common to both speech and quote. From the result in column (5), opposition politicians get quoted approximately 2.5 points less accurately than ruling-party politicians at the 1 percent level of significance (about 26 percent of the standard deviation). The estimates barely change with electoral controls (even-numbered columns).

At the article-speech level where the concatenation of all quotes from a speech reported in a news article is the unit of observation, there is no evidence that the opposition gets quoted at a lower intensity than the ruling party from their speeches in parliament, conditional on the observables which include the length of the speeches and news articles. This does not square with those perceptions of slant where the ruling party gets more coverage.³¹ The opposition speeches however, are quoted using more fragments, and their quotes are less accurate relative to those of the ruling party, providing the first sign that their coverage is more fragmented and less accurate. I discuss the magnitudes of the OLS estimates and implications in section 1.7 after the quote level results.

1.4.3 Specification Checks, Article-Speech Level

Table A.0.6 presents specification checks for the article-speech level result where the opposition gets more quote fragments. Column (1) reproduces the baseline result for comparison (column (3) of Table III). Column (2) models quote fragments as an over-dispersed count variable using the negative binomial regression, and the result is similar. The baseline result where the opposition gets more fragments of quote is

³¹In Figures A.0.3-A.0.4, eyeballing the distribution of coverage intensity by year and parliament suggest that intensity for the ruling party is proportional to the share of seats held in parliament. Or, that the ruling party gets more coverage simply because they have more share of seats in parliament. The results at the politician-year level (untabulated) is also consistent with the finding where there is no difference in the coverage intensity. The opposition politicians get featured in the same number of articles (counted as the number of articles containing at least one quote from parliament) per year compared to the ruling-party politicians.

also robust to the inclusion of journalist fixed-effects and beat assignment dummies in column (3), excluding translations in column (5), and adjusting standard errors for clusters within speeches and journalists in columns (6) and (7).

In columns (8)–(13), the baseline results also survive different assumptions on the topic structure of the textual data. Columns (6) and (7) use the $K = 50$ and $K = 100$ speech topic model. Columns (8) and (9) use the $K = 30$ and $K = 50$ article topic model for article distributions. Column (10) uses the topic distribution of the entire sentence containing the quote, instead of just the quote. Column (11) uses the most parsimonious topic distribution specification available in the sample—the $K = 50$ speech topic model and the $K = 30$ news article topic model.

An exception however, is in column (4). I exclude ministerial controls on concerns that ministerial rank itself is an outcome of opposition status (as a case of bad controls Angrist and Pischke 2009), and the opposition status is no longer significant. I do not fully understand how to interpret this result—including ministerial controls introduces a compositional bias, and excluding them loses non-trivial heterogeneity among the politicians.

Table A.0.7 tests the sensitivity of the baseline finding in column (5) of Table III where opposition quotes are less accurate, using the same set of robustness test as before (but without the count model). The outcome variable in panel A is constructed in the same way as in Table III. Panel B uses the alternative construction of quote accuracy where stop words are removed in pre-processing. The results from both panels are robust and never fall below the 1 percent level of significance.

Panel B in particular, mitigates concerns that the opposition's lower quotation accuracy can be narrowed down to a quote's trivial contents—that journalists are less careful only with the usage of stop words when quoting the opposition while treating the non-stop words (or the substantial words) the same way for all politicians.

TABLE IV
Baseline Results for Political Coverage, Quote-Level

	Log of quote length by word count		Substring quote accuracy measure		Bag-of-words quote accuracy measure	
	(1)	(2)	(3)	(4)	(5)	(6)
Opposition	-0.138*** (0.047)	-0.145*** (0.055)	-1.455** (0.701)	-1.485** (0.753)	-2.434*** (0.707)	-2.271*** (0.675)
<i>Controls</i>						
Time fixed-effects	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
News article controls	Yes	Yes	Yes	Yes	Yes	Yes
Topic controls	Yes	Yes	Yes	Yes	Yes	Yes
Ministerial controls	Yes	Yes	Yes	Yes	Yes	Yes
Electoral controls	—	Yes	—	Yes	—	Yes
<i>F-statistics</i>						
F-statistic, year fixed-effects	3.090***	2.050**	5.229***	3.189***	5.125***	2.501***
F-statistic, individual controls	1.104	1.379	2.027**	0.902	2.055**	0.742
F-statistic, topic controls	2.412***	2.290***	2.106***	2.078***	1.608***	1.459***
F-statistic, ministerial controls	3.901***	3.997***	3.904***	3.964***	2.865***	2.180***
F-statistic, electoral controls		1.096		0.990		1.264
Mean of dependent variable	2.680	2.672	91.405	91.863	96.399	96.647
R ²	0.056	0.068	0.171	0.197	0.113	0.124
N	14,887	10,900	14,887	10,900	14,887	10,900

Notes—Observations are at the quote-level; the regressions in this Table considers each quote as separate observations, even if they originate from the same speech or are reported in the same news article. The set of controls are the same as in article-speech results in Table III. 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.

1.5 Quote level Results

1.5.1 Baseline Results, Quote Level

Here, the unit of analysis is individual quote fragments. At this level, there are 204 politicians, 3,425 articles, and 5,227 speeches over 12 years, for a total of 14,887 quote fragments (with complete controls) of which, 1,106 are from opposition speeches. The model I estimate is:

$$y_{qrst} = \alpha + \beta opp_i + \sum_{k=2006}^{2016} \alpha_k year_{kt} + \sum_{\ell=11}^{13} \alpha_{\ell} parl_{\ell t} + \gamma' \mathbf{X}_{qrst} + \varepsilon_{qrst}, \quad (1.6)$$

which is synonymous with (1.3) at the article-speech level, but with an additional indexing of individual quote fragments q , the smallest unit of analysis here. The outcome variable y is alternatively: (i) log of individual quote length by word count, (ii) substring accuracy measure, and (iii) bag-of-words accuracy measure. Standard errors are adjusted to allow for clusters within each news article.

Table IV reports the results. The first panel confirms the finding at the article-speech level where the coverage of the opposition is more

fragmented. Conditional on the time trends and observables, column (1) indicates that the opposition politicians get shorter individual quote fragments. On average, the quotes of opposition politicians are approximately 13% shorter than quotes of ruling-party politicians. Taken together with the article-speech level results, the evidence suggests that the quotes of the opposition are more fragmented—opposition politicians are covered at the same intensity from a speech in an article as their ruling-party counterparts, but the coverage comes from more but shorter quote fragments.

The second and third panels are the results for quotation accuracy. In the second panel, the outcome variable is the substring accuracy measure, which reflects exactly the verbatim nature of direct quotations. Like the bag-of-words measure, it ranges from 0 to a perfect 100. In the third panel, the outcome variable is the bag-of-words accuracy measure, which scores accuracy using the subset of words common to both speech and quote. Compared to substring accuracy measure, the bag-of-words measure allows for edits of quotes for syntactic flow and grammatical demands. Consistent with the article-speech level results, opposition politicians are quoted less accurately by both measures at the quote level. On average, opposition quotes have approximately 1.5 to 2.4 points lower accurate than quotes of the ruling party (11.6% to 22.9% of the standard deviations in accuracy).³²

1.5.2 Specification Checks

Table A.0.8 presents the specification checks for the results on log of individual quote length (along the same lines as Table A.0.6 in Section 1.4.3). Panel B, C, and D use alternative measures of quote length: (B) by character count, (C) by word count without stop words, and (D) by

³²Including electoral controls can result in either an upward or downward bias. If walkovers typically involve ruling-party politicians sufficiently established to ward off opposition challenge, then excluding these politicians would induce a bias towards zero if these politicians get better coverage because of their prominence. On the other hand, if walkovers are associated with politicians with lower visibility—since they do not need to campaign as much—then leaving them out would artificially accentuate the observed difference between the opposition and the ruling party. The results in columns (2) and (4) of Table IV are consistent with the latter explanation, though the differences are small in magnitude, which can be attributed to sampling variation.

character count without stop words. The results are mostly robust, with journalist fixed-effects removing most of the observed opposition effect only when stop words are removed.

Table A.0.9 presents a similar set of specification tests for the results on quote accuracy, which remain statistically significant. Unlike the result for quote length above, including journalists and beat fixed-effects does not move the estimates for quote accuracy towards zero, suggesting that the tendency to quote ruling-party politicians more accurately is more institutional than just the discretion of individual journalists.

1.5.3 Bounds on OLS Estimates

In Section 1.4.1 above, I argue that bias in the OLS estimates are attenuating ones, implying that the OLS estimates are conservative. First, an attenuation bias arises when the individual politicians' ministerial controls are included, even if party status is random (as discussed in Section 1.4.1). Where a selection issue persists, a second bounding argument can be made using the proportional selection assumption—the assumption that the proportion of selection on observables is equal to the selection on unobservables. This assumption implies that all unobserved factors, which affects political coverage and party status, are equal in importance to all the observables available on hand, including the length of speeches and articles, individual and ministerial characteristics, the topics of the speeches and articles, and additional language-based measures (detailed in the following section).

Assuming equal selection between observables and unobservables implies that the treatment effect of opposition status on quote accuracy (or slant) is -2.18 and -3.39, compared with the OLS regression estimates of -1.46 and -2.24 from Table IV. Moreover, for the estimated treatment effect of the opposition status to be zero, the degree of selection on unobservables relative to observables must be high—the unobservables must be 58 times and 16 times more important than the observables. Hence, the estimates using equal selection on observables and unobservables further supports the OLS estimates as conservative on the extent to which the opposition politicians receive less accurate coverage.

TABLE V
Political Coverage of Backbenchers, Quote Level

	A. Log of quote length by word count		B. Substring quote accuracy measure		C. Bag-of-words quote accuracy measure	
	(1)	(2)	(3)	(4)	(5)	(6)
Opposition	-0.116** (0.053)	-0.164** (0.064)	-1.242 (0.778)	-2.195*** (0.820)	-2.605*** (0.781)	-2.680*** (0.752)
Time fixed-effects	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Article controls	Yes	Yes	Yes	Yes	Yes	Yes
Topic controls	Yes	Yes	Yes	Yes	Yes	Yes
Electoral controls	No	Yes	No	Yes	No	Yes
F-statistic, time fixed-effects	1.623*	1.262	3.119***	2.268***	2.287***	1.951**
F-statistic, individual controls	1.641	1.872*	3.964***	3.450***	2.241**	0.788
F-statistic, topic controls	2.986***	3.635***	2.252***	2.771***	1.252**	1.476***
F-statistic, electoral controls		0.647		1.823		3.373***
Mean of dependent variable	2.681	2.661	91.553	91.785	96.415	96.533
R ²	0.130	0.161	0.236	0.281	0.190	0.251
N	3882	3091	3882	3091	3882	3091

Notes—Observations are at the quote level, but only for politicians who are backbenchers; the regressions in this Table considers each quote as separate observations, even if they originate from the same speech or are reported in the same news article. The set of controls are the same as in article-speech results in Table III. 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.

Nonetheless, to simplify the analysis, and to directly compare politicians of both parties that belong to the base rank, I repeat the regression analysis only for the backbenchers. Table V reports the results. The point estimates vary compared to the baseline regressions, but the sign and magnitudes are largely similar. Focusing on the substring accuracy measure and with electoral controls included in column (4), the point estimate is even larger, and to this extent is consistent with the bounding analyses discussed above.

More generally, Figure II plots the point estimates for the quote-level results with substring accuracy as the outcome, from 144 different specifications (= 2 substring accuracy measures as the outcome, with and without substrings $\times$ 8 covariates combinations $\times$ 9 possible topic modelling specifications) to show how sensitive the main results are. The main specification is from column (3) of Table IV. The estimates are ranked in ascending order, with the chart indicating the specific combinations. Combinations under the outcome and topic models are mutually exclusive, but the ones under the covariates are not. Certain specifications have weaker statistical significance, but all have the same signs and approximately the same magnitude. Figure A.0.2 shows the same chart with bag-of-words accuracy as the outcome measure.

FIGURE II
Effect Sizes of Opposition Status on Substring Accuracy

1.6 Additional Text and Language-Based Measures

1.6.1 Speech Tone and Language Competency Controls

One concern with the baseline results, is that the opposition speeches are less coherent than those of their ruling-party counterparts. A related concern is that the media prefers reporting on speeches of a particular tone (e.g. negative sounding, or subjective sounding). To rule these out, I test whether speech: (i) objectivity, (ii) polarity, (iii) readability,³³ and (iv) lexical richness can pick up the differences in media coverage between politicians. If speech competency and tone are driving the results, then once speech tone and competency are controlled for, the opposition estimate should attenuate to zero. This is not the case.³⁴

To generate objectivity (objective vs. subjective) and polarity (positive vs. negative) measures for speeches, I use the *Pattern* sentiment analyser implemented in the *TextBlob* library, which uses part-of-speech tagging so that words in different parts of a speech get different weights. The readability measures are weighted averages of three different pieces of textual information: (i) average word per sentence, (ii) average syllable 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 measures proxy for language sophistication using information on the occurrences of unique words—the more unique words used, the higher the language sophistication.^{35 36}

³³*Readability* in the literature refers to how easy it is to read a piece of text, the context in this paper deals with speeches, and so *understandability* (of a speech) might be a better word. However, I retain the use of the term *readability* since it is a fairly established term.

³⁴This section uses data at the quote level. I perform the same set of tests at the article-speech level, and the findings are qualitatively the same.

³⁵The readability measures are computed using the *textatistic* library at <http://www.erinhengel.com/software/textatistic>, while the lexical richness measures are computed using the *lexicalrichness* library I wrote and is hosted at <https://github.com/LSYS/lexicalrichness>.

³⁶Importantly, the measures of lexical richness are computed using measures robust to changes in text length (e.g. Maas, mean segmental type-token ratio (MSTTR),

TABLE VI
ADDITIONAL LANGUAGE-BASED MEASURES AS CONTROLS

	Additional language and textual controls using 1st principal components of readability, lexical richness, objectivity, and polarity 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 Log of quote length by word count						
Opposition	-0.138*** (0.047)	-0.121*** (0.047)	-0.132*** (0.047)	-0.136*** (0.047)	-0.135*** (0.048)	-0.113** (0.047)
Objectivity of speech and quote		-0.059*** (0.005)				-0.055*** (0.005)
Polarity of speech and quote			0.030*** (0.005)			0.017*** (0.005)
Grade/readability score of speech transcript				-0.005 (0.006)		-0.005 (0.006)
Lexical richness of speech transcript					-0.001 (0.005)	-0.001 (0.005)
Panel B. Dep. Var. is substring quote accuracy measure						
Opposition	-1.455** (0.701)	-1.452** (0.701)	-1.462** (0.701)	-1.608** (0.702)	-1.272* (0.703)	-1.448** (0.704)
Objectivity of speech and quote		0.030 (0.064)				0.021 (0.065)
Polarity of speech and quote			-0.077 (0.061)			-0.064 (0.062)
Grade/readability score of speech transcript				0.380*** (0.087)		0.355*** (0.089)
Lexical richness of speech transcript					0.473*** (0.074)	0.451*** (0.076)
Panel C. Dep. Var. is bag-of-words quote accuracy measure						
Opposition	-2.434*** (0.707)	-2.411*** (0.709)	-2.425*** (0.707)	-2.443*** (0.705)	-2.379*** (0.703)	-2.352*** (0.704)
Objectivity of speech and quote		-0.083 (0.052)				-0.074 (0.054)
Polarity of speech and quote			0.048 (0.049)			0.030 (0.051)
Grade/readability score of speech transcript				0.021 (0.087)		0.003 (0.091)
Lexical richness of speech transcript					0.206*** (0.078)	0.205** (0.080)
N	14,887	14,885	14,885	14,887	14,836	14,834

Notes—Observations are at the quote level. The regressions in this Table are the same as in the baseline specification in Table IV (replicated in column (1)), but with the additional language measures. Readability and lexical richness are concepts with various proposed measures in practice. The readability and lexical richness measures in this Table are the first principal component of the relevant measures from principal component analyses. The objectivity and polarity measures are first principal components of the different textual components (e.g. speech sentence, speech paragraph). Robust standard errors adjusted for clusters by newspaper article in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

Table VI replicates the baseline models at the quote level with the additional text and language controls. In panel A, objectivity and polarity explain some variation in quote length, with more subjective and

Measure of Textual Lexical Diversity (MTLD), and HD-D (McCarthy and Jarvis 2010; Torruella and Capsada 2013)), since the speeches have varying lengths. The data appendix provides further details.

positive-sounding speeches getting more coverage. The opposition coefficient for quote length, however, remains insignificant. Lexical richness also explains some variation in quote accuracy (panels B and C), with accuracy increasing in the lexical richness of a speech. The baseline conclusion remains otherwise unchanged with all four of the additional language controls, providing evidence against the concern that the opposition speeches are covered less favourable because of the tonality or understandability of their speeches. Another concern is that linguistic competencies may differ in terms of grammatical errors, and that the mechanical inaccuracies arises because of corrections either in the news article or the official transcripts. Direct quotes, however, should still be verbatim and so are the official transcripts, with corrections, if any, occurring in a different part of the text.

1.6.2 Representation of Speech Tone

Another way in which differences between the opposition and the ruling party might manifest is in the representation of the original speech tone. Specifically, I test whether the opposition dummy can predict how similar the tone of the speech is to the tone of its quotation(s). If there is bias in representation along the dimension of speech tone, then there should be a bigger change in tone for the opposition than the ruling party. While the descriptive statistics indicate that journalists from *The Straits Times* prefer reporting on speeches that sound neutral and positive (see Figures A.0.12 and A.0.13), the results from Table A.0.2 suggests no systematic differences in the representation of speech tone between the opposition and the ruling-party politicians.

1.7 Unpacking the Findings

1.7.1 Salience of Quotation Accuracy

One reason this paper focuses on quotation accuracy is that it is arguably a cleaner measure of coverage, suffering less from confounding issues.

TABLE VII
EXAMPLES OF SPEECHES AND QUOTES

Quote fragment	Originating speech	Quote Accuracy
a fundamental relook	...is taking a more fundamental relook at this regulation framework and see how best we can support this strategy ...	95
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
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
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
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
too employer-focused	...and increase in absentee payroll are rather employer-focused ...	85
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
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
the high end	...we noted that the rates were at a higher end , but we had the rates that were charged ...	75
They (said) there are two purposes	...They gave the reasons that they wanted the two budget hotels to serve two purposes ...	75
look them in the eye	...I want to be able to look these men and others in the eyes and say to them ...	70
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
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
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
<i>Notes</i> —The speech column shows the portion of the speech which contains the best partial substring match. The text in bold is an approximation of the best partial substring match. Quote accuracy is as defined by the substring accuracy measure.		

Direct quotes are also particularly salient in journalism since they potentially carry more weight in the public eye compared to other types of statements in a news article, such as indirect testimonies (Gibson and Zillmann 1993; D'Alessio 2003). Getting quoted from a speech with more fragments may carry a higher risk of the quotes being taken out of context (intentional or not); and getting quotes that are less accurate directly increase the risk of misrepresentation (again, intentional or not).

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 fairly similar, the quote accuracy score is 84 (out of 100 using the substring accuracy measure). In particular, the quote uses *will be* where the speech uses *may be*—a subtle modulation in assertiveness.

Media coverage of parliamentary speeches may also be important in an institution like Singapore where campaign periods are in practice very short—10 days or less.³⁷

1.7.2 Magnitudes

From the results at the article-speech level and the quote level analyses, the opposition politicians are covered just as intensely as their ruling-party counterparts from the parliamentary speeches, conditional on time trends and observables. This finding runs counter to those perceptions of media bias in Singapore where the ruling party gets higher coverage *rates*.

³⁷The *Parliamentary Elections Act* allows for a maximum of a 55 day campaign period, the three general elections in the sample period however, have campaign durations of only 10 days or less.

On average, opposition politicians get quoted from a third more fragments or, 18% of the standard deviation. Compared to the unconditional average of 2.1 quote fragments per politician in an article-speech observation, the opposition politicians are quoted using 14% more fragments. On average, the opposition politicians also get quotes that are less accurate by 1.5 to 2.4 points (by the substring and bag-of-words quote accuracy measures), which is 11.6% to 22.9% of the standard deviations. Compared to the unconditional accuracy of 91.4 and 96.4 (out of 100) per quote, opposition get quotes that are 1.6% to 2.5% less accurate.

Overall, the opposition gets less favourable coverage of their parliamentary speeches. The magnitudes however, are not especially large. First, though the opposition is quoted from more fragments, this is on average less than one full unit of a fragment. Or, that more of the opposition quotes are similar in fragmentation to the ruling party than they are different. Second, though the opposition politicians get quotes that are systematically less accurate, once the differences are accounted for, the opposition do get quotes of *decent* accuracy—conditional on the observables, the opposition still get quotes that are 90 (out of 100) points accurate. Table VII provides more speech-quote examples for a sense of how the accuracy scores translate into human evaluation of accuracy.

1.7.3 Contextualising via the Literature

Here, I contextualise the findings where the opposition politicians get quoted less accurately in the news articles in general. In the subsection that follows, I contextualise under a very institutional-specific logistics difference.

First, the evidence is not in favour of an overt media capture. While inequality (Corneo 2006; Petrova 2008) and media concentration (McMillan and Zoido 2004; Besley and Prat 2006) increase the risk of capture, and while it is a fact that high-ranking public officials have been placed in senior management positions of *The Straits Times* (George 2012), their effects may have been mediated by the news monopoly's listing on

a stock exchange with 99.9% of its shares owned by the public.³⁸

The findings are however consistent with how the heavy media-related regulations in Singapore prod journalists, in such a way that they are just a tad more careful when quoting the ruling party, which has greater resources to bring legal actions to bear. Journalists, especially those working in dailies (as opposed to weeklies), trade-off between accuracy and timeliness (Berry 1967). Less accurate quotations of the opposition through the allocation of time spent on checking need not be an intentional partisan slant. The examples in Table VII suggest that inaccuracies in the quotes come from carelessness (or liberties in words/synonyms); and not from altering quotes with nefarious intent.

Third, the findings are also consistent with the story of demand-driven bias within a spatial competition where: (i) news consumers drive media bias,³⁹ (ii) social media, with anti-establishment sentiments, offers a marginal substitution to mainstream media, and (iii) media outlets compete along a political space (as in Hotelling 1929; Mullainathan and Shleifer 2005), with the opposition and the ruling party on opposite ends. In the absence of competitors, the mainstream media monopoly locates in the centre to gain the largest market share possible. But social media, which locates near the extreme opposition end of the political spectrum, leaves a truncated space for mainstream media to optimise market share over, and the middle of this truncated space is slightly towards the establishment relative to the full competition space.

1.7.4 Contextualising via Institutional Logistical Differences

Here, I situate the findings in a context that emerges from primary and secondary research, from both media and political agents, that accounts for a particular government-to-media communications machinery. For

³⁸The remaining shares are owned by the senior management, which is consistent with incentive compatibility under the principal-agent problems.

³⁹Mullainathan and Shleifer (2005) and Gentzkow and Shapiro (2006) model bias as determined from the demand-side, and several empirical evidence supports a demand-driven bias (e.g. Groeling and Kernell 1998; Bernhardt et al. 2008; Gentzkow and Shapiro 2010; Larcinese et al. 2011).

primary research, I conduct a face-to-face interview with a senior journalist at The Straits Times and a member of parliament in the 14th parliament. Secondary research comes from a reviewer who recounts, in detail and writing, his/her experiences as a member of parliament in the 13th parliament.

First, ruling-party politicians circulate their speech transcripts in advance to the media. Politicians who introduce a Bill or Motion, usually of higher ranks or political office holders, not only circulate their speeches to the media in advance but may have also briefed the media agents in advance. This allows the media ample time to prepare their media pieces. Non-senior ruling-party politicians also circulate their speeches to the media in advance, or the media may even initiate the request for the transcript.

Opposition politicians, however, do not share the same practices mentioned above. Furthermore, opposition politicians are less likely to voluntarily send in their transcripts after the sitting, and if they do, send it in late. It also stands to reason that if they are unwilling to voluntarily share, they are also unlikely to respond to media agents who initiate the request, if any, for a copy of the transcript before the sitting begins.

Given that the official speech transcripts are released only a week after sitting, and that virtually all news stories are either same or next day, media agents have to rely on their shorthand notes, video recordings of the sitting, or the pre-circulated speech transcripts. If accuracy increases in access (and timeliness) to the transcripts, the above would suggest that the logistical difference between the ruling and opposition parties would completely explain away the differences in coverage accuracy. The same can be said about the fragmentation of coverage. This is a potential explanation that, unfortunately, cannot be ruled out with the data.

1.7.5 Rational Choices in a Separating Equilibrium

However, to the extent that the detected partisan difference in quotation accuracy fully reflects logistical differences—the choice on whether to circulate speech transcripts in advance to media agents—the findings in this paper still point to 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. By choosing whether to circulate, a political agent i trades off between increasing quotation accuracy and spin. 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

$$A_i(c_1) - S_i(c_1) > A_i(c_0),$$

or, that 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 kind of 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 $y\%$ ", another is "raising taxes by $x\%$ to stabilise 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 judgement for simplicity but captures only magnitude. 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 insights drawn from the primary and secondary research. From 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

$$A_r(c_1) - S_r(c_1) > A_r(c_0)$$

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

$$A_o(c_1) - S_o(c_1) < A_o(c_0),$$

where the opposition politicians do not.

Moreover, if one takes the position where the detected differences in quotation accuracy in this paper is not reflective of slant, but reflective entirely of logistical differences⁴⁰—this 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 collapse into

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

or that

$$S_o(c_1) > S_r(c_1),$$

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 reveal 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.

⁴⁰This position, which is taken by the anonymous journalist and an anonymous member of parliament from the primary and secondary research, implies that the detected quotation accuracy difference in this paper is capturing $A_r(c_1) > A_o(c_0)$, which is not a fair indication of media slant because of logistical differences (c).

1.7.6 Consequences of Quotation Inaccuracy

Finally, even if media slant manifests in quotation inaccuracy, it is an agnostic error in the context of this study in that there is no clear implication of inaccuracy. In Section 1.7.1, I draw from the communications literature which suggests that direct quotations carry non-trivial weight in a media piece. Table VII, however, which shows some examples of misquotations, suggests that the inaccuracies in the sample are likely innocuous errors and one, therefore, should not expect significant consequences to the slight misquotations.

Moreover, from the above Section 1.7.5, the institutional fact that a separating equilibrium exists—where the opposition politicians are willing to trade away quotation accuracy to avoid spin—suggest that the former is less important than the latter. One way forward to address the consequences of misquotation would be to get participants in a randomized controlled trial to evaluate how their judgement of a speech is affected by misquotation. This is a future avenue of research.

1.8 Conclusion

This paper explores a methodological approach that is novel in that the detection of political media bias is done entirely in a sterile and objective environment. The main coverage outcome measure of accuracy is objective, and the quantification of further information from the textual data is machine annotated and automated, which while not free of bias or errors, does not require any subjective human judgements.

In the discussion, I place the findings in an institutional-specific media-political machinery. The notion of accuracy developed in this paper can, at the very least, detect differences in media engagement strategy even if media slant cannot be directly detected.

On a final note, the approach in this paper does not suggest that the extant literature that uses coverage intensity to measure slant is lacking. Instead, using intensity in the media of Singapore is insufficient to detect partisan differences in coverage. Using coverage accuracy can be

replicated for any media, and might be useful in contexts where slant is less overt. Other potential future applications of coverage accuracy include coverage of financial reports, and in science journalism to assess how accurately the media represents scientific findings.

Appendices

Appendix A

Appendix

A.0.1 Background to Parliamentary Speeches

This section synthesises 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 comes mostly 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 day's 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 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 on 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. Here, there is room for unscripted responses, to the extent that the supplementary questions have not been pre-empted and are not made known to colleagues in advance. Deviations from scripted exchange within the incumbent-party members may also arise when Members want to further engage 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 the majority of the incumbent-party Members to prepare scripted responses to speeches from opposition Members.

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

A.0.2 OLS Estimate of the Opposition Dummy is a Lower Bound on the True *Magnitude*

The following shows that treating ministerial rank as a bad control suggest 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 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. Hence even conditioning on rank does not recover the true causal effect because of the composition of the politicians and their rank.

As a simplification, the following abstracts away from other regressors (X), or more specifically assuming that opposition status does not determine X other than rank. In reality, there is also a whole range of ministerial rank/type, but these are all collapsed 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

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

and let the higher rank indicator for politician i be

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

Following Angrist and Pischke (2009)'s treatment of bad controls in the context of this paper, opposition status (D_i) determines both media coverage (y_i) and the rank (r_i):

$$\begin{aligned} y_i &= y_{0i} + (y_{1i} - y_{0i})D_i \\ r_i &= r_{0i} + (r_{1i} - r_{0i})D_i, \end{aligned}$$

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 in turn 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:

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

The observed difference in political media coverage between oppositions and ruling-party politicians can be decomposed into two two parts. The first term in the last line of the equation above 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 given assignment to 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 plausible set of assumptions suggests that the bias is likely positive. First, the average of the 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 constitute a lower bound on the true magnitude of the opposition effect.

A.0.3 Additional Tables and Figures

TABLE A.0.1
DAILY NEWSPAPER SUBSCRIPTION

Daily Newspaper	Language	Unique Digital Subscription	Print Subscription	Digital Subscription	Total Subscription
1 Berita Harian (Berita Minggu)	Malay	911	44'600	2'500	47'100
2 The Business Times	English	6'658	29'200	18'500	47'700
3 Lianhe Zaobao	Chinese	13'727	148'600	39'300	187'900
4 Lianhe Wanbao	Chinese	1'137	82'500	9'100	91'600
5 The New Paper (The New Paper Sunday)	English	757	70'200	40'400	110'600
6 Shin Min Daily News	Chinese				120'200
7 The Straits Times (The Sunday Times)	English	60'871	304'300	177'400	481'700
8 Tamil Murasu (Tamil Murasu Sunday)	Tamil				12'800

Source: *SPH 2015 Annual Report*.

a) Subscriptions of the eight daily newspaper wholly-owned by Singapore Press Holdings.

b) Parenthesis indicates the Sunday edition of the daily. For instance, *The Straits Times* is published as *The Sunday Times* on Sundays.

c) In addition to total subscriptions, unique digital subscriptions, print subscriptions, and digital subscriptions are shown where available.

TABLE A.0.2
DIFFERENCES IN OBJECTIVITY AND POLARITY

	Dependent variable is the difference in objectivity/polarity scores from					
	Quote to			Sentence containing quote to		
	Full speech	Speech paragraph	Speech sentence	Full speech	Speech paragraph	Speech sentence
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Differences in objectivity						
Opposition	0.003 (0.012)	0.010 (0.012)	−0.011 (0.011)	0.014 (0.010)	0.020* (0.011)	0.000 (0.010)
Panel B. Differences in polarity						
Opposition	−0.003 (0.009)	−0.005 (0.009)	0.002 (0.007)	0.000 (0.009)	−0.002 (0.009)	0.005 (0.008)
Time fixed-effects	Yes	Yes	Yes	Yes	Yes	Yes
Length controls	Yes	Yes	Yes	Yes	Yes	Yes
Individual controls	Yes	Yes	Yes	Yes	Yes	Yes
Article controls	Yes	Yes	Yes	Yes	Yes	Yes
Topic controls	Yes	Yes	Yes	Yes	Yes	Yes
Ministerial controls	Yes	Yes	Yes	Yes	Yes	Yes
<i>N</i>	14,887	14,885	14,885	14,886	14,885	14,885

Notes—Observations are at the quote level. The first set compares the quote to each of three speech components: (i) the full speech, (ii) the speech paragraph containing the quote, and (iii) the speech sentence(s) containing the quote; the second set compares the sentence containing the quote to three speech components enumerated above. In panel A, the dependent variable is the difference in objectivity from speech to quote; in panel B, the dependent variable is the difference in polarity from speech to quote. The polarity and objectivity measures are generated using the open-source *TextBlob* Pattern Analyzer (<https://textblob.readthedocs.io>). Robust standard errors adjusted for clusters by newspaper article in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE A.0.3
POLITICAL COVERAGE DURING PRE-ELECTION PERIODS

Election periods	Variable	Log of quote length by word count		Substring quote accuracy measure		Bag-of-words quote accuracy measure	
		Full	Subsample	Full	Subsample	Full	Subsample
		(1)	(2)	(3)	(4)	(5)	(6)
<i>Panel A. 6 & 3-months before a general election</i>							
$E_t = 1$ if 6 months before a general election $N_{E_t} = 1630$ $N_{opp} = 103$	Opposition	-0.15*** (0.05)	-0.06 (0.22)	-1.19* (0.69)	-2.99 (3.34)	-2.09*** (0.69)	-5.33 (3.28)
	Opposition $\times E_t$	0.05 (0.12)	. (0.12)	-3.10 (2.08)	. (2.08)	-4.08* (2.16)	. (2.16)
$E_t = 1$ if 3 months before a general election $N_{E_t} = 1207$ $N_{opp} = 80$	Opposition	-0.14*** (0.05)	-0.35 (0.29)	-1.38** (0.70)	3.58 (4.48)	-2.22*** (0.69)	-1.70 (3.66)
	Opposition $\times E_t$	0.04 (0.15)	. (0.15)	-0.61 (2.20)	. (2.20)	-3.05 (2.42)	. (2.42)
<i>Panel B. 3-months before a by-election</i>							
$E_t = 1$ if 3 months before a by-election $N_{E_t} = 1636$ $N_{opp} = 92$	Opposition	-0.13*** (0.05)	-0.04 (0.19)	-1.47** (0.73)	-3.23 (2.37)	-2.60*** (0.74)	-1.80 (2.29)
	Opposition $\times E_t$	-0.07 (0.12)	. (0.12)	0.55 (1.27)	. (1.27)	2.60*** (0.98)	. (0.98)
<i>Panel C. 3-months & 1-month before any election</i>							
$E_t = 1$ if 3 months before any election $N_{E_t} = 2843$ $N_{opp} = 172$	Opposition	-0.14*** (0.05)	-0.01 (0.14)	-1.43** (0.73)	-0.34 (2.05)	-2.40*** (0.73)	-2.30 (1.82)
	Opposition $\times E_t$	-0.01 (0.10)	. (0.10)	0.03 (1.29)	. (1.29)	0.06 (1.32)	. (1.32)
$E_t = 1$ if 1 month before any election $N_{E_t} = 974$ $N_{opp} = 54$	Opposition	-0.15*** (0.05)	0.15 (0.43)	-1.35* (0.71)	1.45 (5.40)	-2.43*** (0.71)	-1.33 (4.33)
	Opposition $\times E_t$	0.06 (0.16)	. (0.16)	-1.92 (1.60)	. (1.60)	0.89 (1.15)	. (1.15)

Notes—Observations are at the quote-level. The dependent variable in columns (1)–(2) is log of quote length; in columns (3)–(4) it is quote accuracy using measure (1); and in columns (5)–(6) it is quote accuracy using measure (2). Odd-numbered columns uses the full sample, while even-numbered columns uses the subsample observations for the relevant periods just before the elections. N_{E_t} indicates the number of observations in the relevant pre-election period; N_{opp} indicates the number of opposition quote observations in the same period. The models are the same as in the baseline specifications in Table IV, but with an additional Opposition $\times E_t$ term in the full sample model, which estimates indicates whether there is an additional effect during the pre-election periods. In the subsample models, the Opposition term estimates whether there is a difference between the opposition and ruling party politicians in the pre-election periods. 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.

TABLE A.O.4
MINISTERIAL RANK BY PARTY

Ministerial Rank		Non-opposition	Opposition	Total
Rank	Description			
PM	Prime Minister	491	0	491
DPM	Deputy PM	1185	0	1185
Minister	MR1-4 rank	6497	0	6497
SMS	Senior Minister of State	1012	0	1012
MOS	Minister of State	610	0	610
Mayor	Mayor (of 1 of 5 districts)	222	0	222
Sps	Senior Parliamentary Secretary	98	0	98
Parl Sec	Parliamentary Secretary	87	0	87
Speaker	Speaker of Parliament	38	0	38
MP	Member of Parliament (base rank)	3082	800	3882
NCMP	Non-constituency MP	0	306	306
NMP	Nominated MP	475	0	475
Total		13,797	1106	14,903

TABLE A.O.5
NEWSPAPER SECTION BY PARTY

Newspaper section	Non-opposition	Opposition	Total
Home	961	83	1044
Insight	50	7	57
Money	21	0	21
News	28	0	28
Opinion	11	0	11
Others	22	3	25
Prime News	1929	152	2081
Review - Insight	14	0	14
Singapore	8210	626	8836
Sports	14	0	14
ST	645	34	679
Think	12	3	15
Top of the News	1871	198	2069
World	9	0	9
Total	13,797	1106	14,903

(A) GUI example without text

(B) GUI example with text
FIGURE A.o.1

Graphics user interface with quote matching and extraction

TABLE A.o.6
SPECIFICATION CHECKS FOR QUOTE FRAGMENTS, ARTICLE-SPEECH LEVEL

	Baseline results (1)	Negative Binomial (2)	Journa- list FE (3)	No Ministerial controls (4)	No Translated quotes (5)	Clustering		Topic distributions of textual content					
						Cluster by speech (6)	Cluster by journalist (7)	Speech $K = 50$ (8)	Speech $K = 100$ (9)	Article $K = 30$ (10)	Article $K = 50$ (11)	Sentence topics (12)	Parsimonious topics (13)
Opposition	0.328*** (0.089)	0.167*** (0.045)	0.300*** (0.101)	-0.046 (0.067)	0.300*** (0.090)	0.328*** (0.083)	0.303** (0.119)	0.318*** (0.087)	0.348*** (0.091)	0.406*** (0.089)	0.341*** (0.089)	0.317*** (0.089)	0.395*** (0.088)
<i>N</i>	7,087	7,087	6,210	7,087	6,980	7,087	6,210	7,087	7,087	7,087	7,087	7,087	7,087

Notes—This Table reports specification checks for the baseline result in Table III where the coverage of opposition politicians are made up of more quote fragments than those of the ruling-party politicians. Column (2) models the count of quote fragments using the negative binomial regression. Column (3) includes journalist fixed-effects and beat dummies. Column (4) excludes ministerial controls on the grounds that ministerial type is a bad control. Columns (5) excludes observations that are recorded as translations (from vernacular to English). Columns (6) and (7) adjusts standard errors for clusters by speech and journalist instead of newspaper articles. Columns (8)–(13) tests various specifications of the topic distributions. Column (12) uses the topical distribution of the sentence containing the quote instead of the quote itself. Column (13) uses the most parsimonious well-performing topic distributions— $K=30$ for the news articles, and $K=50$ for the speeches. Robust standard errors in parentheses are clustered at news articles except in columns (6)–(7).

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE A.O.7
SPECIFICATION CHECKS FOR QUOTE ACCURACY, ARTICLE-SPEECH LEVEL

	Baseline results (1)	Journa- list FE (2)	No Ministerial controls (3)	No Translated quotes (4)	Clustering		Topic distributions for textual content					
					Cluster at speech (5)	Cluster at journalist (6)	Speech $K = 50$ (7)	Speech $K = 100$ (8)	Article $K = 30$ (9)	Article $K = 50$ (10)	Sentence topics (11)	Parsimonious topics (12)
Panel A. Dependent variable is bag-of-words quote accuracy measure												
Opposition	-2.5*** (0.785)	-2.75*** (0.899)	-1.95*** (0.594)	-2.33*** (0.778)	-2.5*** (0.779)	-2.71*** (0.944)	-2.44*** (0.784)	-2.59*** (0.799)	-2.62*** (0.784)	-2.43*** (0.765)	-2.35*** (0.773)	-2.61*** (0.789)
N	7,087	6,210	7,087	6,980	7,087	6,210	7,087	7,087	7,087	7,087	7,087	7,087
Panel B. Dependent variable is bag-of-words quote accuracy measure (No stopwords)												
Opposition	-3.3*** (0.947)	-3.62*** (1.07)	-2.59*** (0.721)	-3.09*** (0.931)	-3.3*** (0.937)	-3.58*** (1.21)	-3.07*** (0.936)	-3.46*** (0.966)	-3.51*** (0.945)	-3.19*** (0.92)	-3.12*** (0.938)	-3.36*** (0.937)
N	7,087	6,210	7,087	6,980	7,087	6,210	7,087	7,087	7,087	7,087	7,087	7,087

Notes—This Table reports specification checks for the baseline result in Table III where the coverage of opposition politicians are less accurate than those of the ruling-party politicians. Column (2) includes journalist fixed-effects and beat dummies. Column (3) excludes ministerial controls on the grounds that ministerial type is a bad control. Columns (4) excludes observations that are recorded as translations (from vernacular to English). Columns (5) and (6) adjusts standard errors for clusters by speech and journalist instead of newspaper articles. Columns (7)–(12) tests various specifications of the topic distributions. Column (11) uses the topical distribution of the sentence containing the quote instead of the quote itself. Column (12) uses the most parsimonious well-performing topic distributions— $K=30$ for the news articles, and $K=50$ for the speeches. Robust standard errors in parentheses are clustered at news articles except in columns (5)–(6).

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE A.o.8
SPECIFICATION CHECKS FOR QUOTE LENGTH, QUOTE-LEVEL

					Subsamples		Clustering		Topic distribution of textual content					
	Baseline results (1)	Negative Binomial (2)	Journa- list FE (3)	No ministerial controls (4)	No translations (5)	No low similarity (6)	Cluster at speech (7)	Cluster at journalist (8)	Speech $K = 50$ (9)	Speech $K = 100$ (10)	Article $K = 30$ (11)	Article $K = 50$ (12)	Sentence topics (13)	Parsimonious topics (14)
Panel A. Dependent variable is log of quote length by <i>word</i> count														
Opposition	-0.138*** (0.047)	-0.120*** (0.036)	-0.086* (0.048)	-0.163*** (0.039)	-0.128*** (0.049)	-0.138*** (0.048)	-0.138*** (0.048)	-0.084** (0.042)	-0.154*** (0.047)	-0.127*** (0.047)	-0.131*** (0.047)	-0.128*** (0.048)	-0.136*** (0.048)	-0.146*** (0.047)
<i>N</i>	14,887	14,887	13,513	14,887	14,682	14,423	14,887	13,513	14,887	14,887	14,887	14,887	14,887	14,887
Panel B. Dependent variable is log of quote length by <i>character</i> count														
Opposition	-0.123*** (0.044)	-0.112*** (0.035)	-0.074* (0.045)	-0.147*** (0.036)	-0.117** (0.045)	-0.121*** (0.044)	-0.123*** (0.045)	-0.072* (0.040)	-0.139*** (0.043)	-0.113** (0.044)	-0.118*** (0.044)	-0.113** (0.045)	-0.120*** (0.044)	-0.132*** (0.044)
<i>N</i>	14,885	14,887	13,511	14,885	14,680	14,421	14,885	13,511	14,885	14,885	14,885	14,885	14,885	14,885
Panel C. Dependent variable is log of quote length by <i>word</i> count (No stopwords)														
Opposition	-0.089** (0.039)	-0.092*** (0.035)	-0.049 (0.039)	-0.100*** (0.032)	-0.084** (0.041)	-0.091** (0.040)	-0.089** (0.040)	-0.047 (0.034)	-0.101*** (0.039)	-0.080** (0.040)	-0.084** (0.039)	-0.079** (0.040)	-0.090** (0.040)	-0.095** (0.039)
<i>N</i>	14,887	14,887	13,513	14,887	14,682	14,423	14,887	13,513	14,887	14,887	14,887	14,887	14,887	14,887
Panel D. Dependent variable is log of quote length by <i>character</i> count (No stopwords)														
Opposition	-0.081** (0.041)	-0.084** (0.035)	-0.038 (0.041)	-0.108*** (0.034)	-0.078* (0.042)	-0.081** (0.041)	-0.081* (0.042)	-0.035 (0.037)	-0.095** (0.040)	-0.073* (0.041)	-0.077* (0.041)	-0.070* (0.042)	-0.082** (0.041)	-0.089** (0.040)
<i>N</i>	14,872	14,887	13,499	14,872	14,667	14,409	14,872	13,499	14,872	14,872	14,872	14,872	14,872	14,872

Notes—This Table reports specification checks for the baseline result in Table IV where the quotes of opposition politicians are shorter than those of the ruling-party politicians. In panels C & D, stopwords (e.g. "the", "or", "a", "we" "be") are removed before the relevant measures are computed. Column (1) presents the baseline result. Column (2) models quote length as a count variable of words and characters using the negative binomial regression. Column (3) includes journalist fixed-effects and beat dummies. Column (4) excludes ministerial controls on the grounds that ministerial type is a bad control. Columns (5) and (6) excludes observations that are recorded as translations (from vernacular to English) and observations which have low quote accuracy (observations with both similarity measures below 75 are excluded). Columns (7) and (8) adjusts standard errors for clusters by speech and journalist instead of newspaper articles. Columns (9)–(14) tests various specifications of the topic distributions. Column (13) uses the topical distribution of the sentence containing the quote instead of the quote itself. Column (14) uses the most parsimonious well-performing topic distributions— $K=30$ for the news articles, and $K=50$ for the speeches. Robust standard errors in parentheses are clustered at news articles except in columns (7)–(8).

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE A.o.9
SPECIFICATION CHECKS FOR QUOTE ACCURACY, QUOTE-LEVEL

	Baseline Regression (1)	Journa- list FE (2)	No ministerial controls (3)	No translations (4)	Clustering		Topic distribution of textual content					
					Cluster by speech (5)	Cluster by journalist (6)	Speech $K = 50$ (7)	Speech $K = 100$ (8)	Article $K = 30$ (9)	Article $K = 50$ (10)	Sentence topics (11)	Parsimonious topics (12)
Panel A. Dependent variable is substring quote accuracy measure												
Opposition	-1.455** (0.701)	-1.953*** (0.742)	-1.570*** (0.562)	-1.302* (0.699)	-1.455** (0.739)	-1.928** (0.753)	-1.223* (0.698)	-1.516** (0.712)	-1.736** (0.700)	-1.515** (0.701)	-1.569** (0.712)	-1.488** (0.695)
Observations	14,887	13,513	14,887	14,682	14,887	13,513	14,887	14,887	14,887	14,887	14,887	14,887
Panel B. Dependent variable is substring quote accuracy measure (No stopwords)												
Opposition	-1.662** (0.723)	-2.226*** (0.770)	-1.587*** (0.571)	-1.495** (0.718)	-1.662** (0.759)	-2.220*** (0.779)	-1.333* (0.717)	-1.746** (0.738)	-1.990*** (0.725)	-1.699** (0.725)	-1.742** (0.732)	-1.667** (0.717)
Observations	14,887	13,513	14,887	14,682	14,887	13,513	14,887	14,887	14,887	14,887	14,887	14,887
Panel C. Dependent variable is bag-of-words quote accuracy measure												
Opposition	-2.434*** (0.707)	-2.594*** (0.750)	-1.802*** (0.566)	-2.321*** (0.718)	-2.434*** (0.726)	-2.573*** (0.839)	-2.200*** (0.710)	-2.284*** (0.707)	-2.660*** (0.706)	-2.349*** (0.689)	-2.301*** (0.701)	-2.446*** (0.705)
Observations	14,887	13,513	14,887	14,682	14,887	13,513	14,887	14,887	14,887	14,887	14,887	14,887
Panel D. Dependent variable is bag-of-words quote accuracy measure (No stopwords)												
Opposition	-3.244*** (0.902)	-3.454*** (0.967)	-2.437*** (0.713)	-3.131*** (0.914)	-3.244*** (0.927)	-3.440*** (1.174)	-2.817*** (0.893)	-3.077*** (0.900)	-3.590*** (0.905)	-3.108*** (0.881)	-3.065*** (0.901)	-3.181*** (0.892)
Observations	14,887	13,513	14,887	14,682	14,887	13,513	14,887	14,887	14,887	14,887	14,887	14,887

Notes—This Table reports specification checks for the baseline result in Table III where the quotes of opposition politicians are less accurate than those of the ruling-party politicians. In panels C & D, stopwords (e.g. "the", "or", "a", "we" "be") are removed before the relevant measures are computed. Column (1) presents the baseline result. Column (2) includes journalist fixed-effects and beat dummies. Column (3) excludes ministerial controls on the grounds that ministerial type is a bad control. Columns (4) excludes observations that are recorded as translations (from vernacular to English). Columns (5) and (6) adjusts standard errors for clusters by speech and journalist instead of newspaper articles. Columns (7)–(12) tests various specifications of the topic distributions. Column (11) uses the topical distribution of the sentence containing the quote instead of the quote itself. Column (12) uses the most parsimonious well-performing topic distributions— $K=30$ for the news articles, and $K=50$ for the speeches. Robust standard errors in parentheses are clustered at news articles except in columns (5)–(6).

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

FIGURE A.0.2
Effect Sizes of Opposition Status on Bag-of-words Accuracy

FIGURE A.O.3
Count of quotes over the years by partisanship

FIGURE A.O.4
Count of quotes over parliaments by partisanship

FIGURE A.o.5
Quote length over time

FIGURE A.o.6
Speech length over time

FIGURE A.o.7
Article length over time

FIGURE A.o.8
Bag-of-words quote accuracy measure over time

FIGURE A.0.9
Distribution of quote length

FIGURE A.0.10
Distribution of quote length at article-speech level

FIGURE A.O.11
Distribution of accuracy measures

FIGURE A.O.12
Distribution of speech and quote objectivity

FIGURE A.0.13
Distribution of speech and quote polarity

Chapter 2

WOMEN ON BOARDS AND FINANCIAL OUTCOMES: EVIDENCE FROM THE SINGAPORE EXCHANGE

2.1 Introduction

Board gender diversity has been linked to better financial performance in the popular press (Carter and Wagner 2011; Dawson et al. 2012; Devillard et al. 2013; Dawson et al. 2016). The report by Dawson et al. (2012), in particular, sounds promising—firms with more women on their boards have higher returns on equity (ROE), lower leverage, have higher market premium (price-to-book ratio), and higher growth in net income. Such reports, however, do not address endogeneity concerns.

In this chapter, I focus on three financial outcomes: (i) ROE, (ii) firm value, and (iii) firm risk. To identify the causal benefit of female board representation on the above three measures, I use a sample of Singapore firms listed on the Singapore Exchange for the years 2000–17 together with an institution-specific variation for identification. Specifically, the main identification strategy uses two sources of variation: (i) the gradual increase of women representation in politics and (ii) firm-government linkage. Board gender composition data at the firm-year level are aggregated from the individual director-level data from *BoardEx* while

data on firm characteristics come from the *Bloomberg* database.

In the first stage, the difference-in-differences estimates suggest that Tier 1 government-linked companies (GLCs)—those companies with the closest linkage to the government—respond more to the increase in women representation in parliament by appointing more women on their boards. This finding suggests the government can exert some influence on private firms, with this influence stronger in the GLCs, and can do so without introducing mandatory gender quotas that are potentially detrimental (e.g. Ahern and Dittmar 2012).

Contrary to reports in the popular press, the second-stage analyses do not reveal any effect of female board representation on firm value, profitability, or firm risk. The diagnostics however indicate that the 2SLS identification suffers from a weak instrument problem, which potentially explains the lack of power in the second stage. I deal with this problem in three ways: (i) by considering test statistics robust to weak IVs, (ii) by using a dynamic panel system GMM model, and (iii) by considering an alternative and stronger first stage. The conclusion remains unchanged. Overall, the results in this paper suggest that the apparent relationship between having women on boards and firm performance can be attributed to permanent firm characteristics, such as corporate culture.

In the discussion, I note three main explanations for the findings. The first has to do with the presumptions about gender differences. Croson and Gneezy (2009) for example, document differences in preferences of individuals by gender in lab settings, but these differences may not carry over to less general settings (e.g. when payoffs are high, Holt and Laury 2002), or in lab settings specifically simulated to reflect financial trading activities (Deaves et al. 2009). Another interpretation is that the sample lacks sufficient variation to capture effects that arise only with a critical mass of three women directors (Kanter 1977; Carter and Wagner 2011; Konrad et al. 2008; Schwartz-Ziv 2017). A third reason is that benefits from women representation need not be financial (as in Adams and Ferreira, 2009).

This paper contributes to the literature that investigates the effect of

female board representation on firm value, where Adams and Ferreira (2009) finds no effect and Ahern and Dittmar (2012) find a negative causal relationship but only under a dramatic mandatory quota. The findings also complement those in Sila et al. (2016) who find no effect of women representation on firm risk. In particular, the findings on the source of endogeneity in this paper are similar to Sila et al. (2016) where the spurious relationship can be attributed to unobserved heterogeneity across firms, such as corporate culture. More generally, this paper relates to the literature that links women on boards and financial performance (Scholtz and Kieviet 2018), Post and Byron (2014) and Pletzer et al. (2015) in particular provide wider surveys, although they are mostly limited to descriptive studies.¹

The outline of the paper is as follows. Section 2.2 provides some institutional background and documents how I build the panel using individual board data. Section 2.3 lays out the identification assumptions. Sections 2.4 and 2.5 present the first and second-stage results. Section 2.6 places limits to claims and presents results from an alternative first stage. Section 3.6 discusses the findings. Section 3.7 concludes.

2.2 INSTITUTIONAL BACKGROUND AND DATA

2.2.1 Data on Board Directors

Using the Bloomberg terminal, the key criterion to select firms are those in the annual top 100 firms by market capitalisation. This yields 191 firms.² The Bloomberg database also provides the annual financial characteristics of the firms in the sample.

¹This paper also relates to the small set of literature on the "glass cliff hypothesis" suggesting that women are overrepresented in precarious leadership positions (Haslam and Ryan 2008). The OLS estimates suggest that the firms with greater risk tend to be those firms with more women board directors, though the link is not causal.

²After filtering out duplicate listings (identified by firm name or ticker) and firms based outside Singapore, those firms that appear at least twice in the yearly top 100 firms by market cap are retained in the sample, yielding 212 listings. Of these, 21 are eventually dropped from the sample, mostly because they are double listings, or because the board director data cannot be found in BoardEx for the sample period 2000–17. Table B.0.1 enumerates the dropped firms and the reasons.

Data on individual board directors come from BoardEx, which contains data on director names, year of birth, start and end date of roles, whether the roles are supervisory or executive,³ gender, and size of their network. The individual director data are then parsed into a long-panel format where the *effective* count of directors is based on the sitting duration over the year. For example, for a two-person board in a year where the first director sits on the board for the full year but the second sits on the board for only half the year. Their effective durations are 1 and 0.5, respectively. If the second director is a woman, then this firm has 0.5 women for that year (or that 1/3 of the board is women).

In cases where the start or end date of a director role is incomplete (e.g. a recorded date of 0ct 2002 instead of 15 0ct 2002, I use three simple imputation methods: i) use longest-possible duration, ii) use shortest-possible duration, and iii) drop the individual data. Under the longest-possible duration, an incomplete date of 0ct 2002 will be imputed as 1 0ct 2002 (the earliest) if it is a start date. If 0ct 2002 is the end date, then it will be imputed as 31 0ct 2002 (the latest). Imputation using the shortest possible duration is the opposite. Under the third method, I do a listwise deletion for directors with both start and end dates missing. The main analyses use the longest-possible imputation method, but the benchmark first-stage results are robust to the two alternative imputation methods.^{4 5}

2.2.2 Temasek Holdings and Government-Linked Companies

Defining Government-Linked Companies (GLCs). Most of the paper distinguishes between GLCs that are *Tier 1* GLCs—those firms with

³Singapore's Company Act recognises a few types of directors: executive, non-executive, independent, nominee, de factor, and shadow. The BoardEx recognises executive vs supervisory (non-executive), so it is unclear whether supervisory directors under the BoardEx classification captures some or all of the non-executive director roles in the context of Singapore.

⁴Gender is also occasionally missing in the BoardEx data, but in those rare occasions, it turns out that a simple internet search for the individual reveals the gender either through pictures or the use of gender pronouns in press releases.

⁵Figures B.0.1 and B.0.2 summarises the distribution of director counts.

direct ownership of at least 20% by Temasek Holdings (one of the state's two Sovereign Wealth Funds), and the *Tier 2* GLCs—firms with an *effective* ownership of 15%. Direct ownership shares of the Tier 1 GLCs come from the annual reports of Temasek Holdings which started in 2004. Tier 2 GLCs shares are inferred using the latest available data of shares ownership from the Bloomberg database and corroborated with previous studies on GLCs in Singapore (Ramírez and Tan 2004; Ang and Ding 2006).⁶

In practice, the effective ownership varies from year to year, and some effective ownership shares of Tier 2 GLCs cannot be unambiguously inferred from the Bloomberg database. For those cases, I use an indicator for whether a Tier 2 GLC is in doubt. Coding them as non-GLCs or excluding them from the sample does not change the results.

For the Tier 1 firms, there is little ambiguity. Most of them have a consistently high percentage of their shares held directly by Temasek Holdings, and even though the 20% cutoff seems arbitrary, in practice most of the Tier 1 GLCs are far away from the cutoff with ownership typically around 40–60%.⁷

Table I tabulates the breakdown of the GLCs by year and by sector. Notable sectors in which GLCs operate are financials, industrials (e.g. marine, transportation, manufacturing), and real estate.⁸ The sample period 2000–17 includes twelve Tier 1 GLCs, twenty-five Tier 2 GLCs (the status of six Tier 2 GLCs are questionable), and 154 non-GLC firms. Since the GLCs in the sample are listed in the stock exchange, they are subject to the same market competition, regulations, and profit

⁶The primary purpose at the time of Temasek Holdings's inception in 1974 was to hold and manage assets previously held by the state, assets which were themselves used to drive the state's development in the mid to late 1900s. And, because Temasek Holdings was incorporated under the Singapore Companies Act, it is subject to the same regulations and requirements that typical private companies face. The sole owner of Temasek Holdings is the government of Singapore (or more precisely, the Ministry of Finance), and under this arrangement distributes dividends to the government.

⁷Table B.0.2 lists all firms in the sample, their GLC status, the average ownership percentage of Tier 1 GLCs over the years as reported in the annual reports, and indicates which of the Tier 2 GLCs are questionable.

⁸Table B.0.3 provides the summary by the more refined GICS industry group classification

TABLE I
BREAKDOWN OF GLCS, BY YEAR AND SECTOR

Year	By Year			
	Non-GLCs	Tier 2	Tier 1	Total
2000	86	9	11	106
2001	90	9	11	110
2002	97	11	11	119
2003	103	12	11	126
2004	108	14	12	134
2005	114	15	12	141
2006	121	17	12	150
2007	130	18	12	160
2008	132	19	12	163
2009	132	21	12	165
2010	137	22	12	171
2011	141	23	12	176
2012	142	23	12	177
2013	147	24	12	183
2014	150	24	12	186
2015	151	25	12	188
2016	151	25	12	188
2017	152	25	12	189
Total	2,284	336	212	2,832

Sector (GICS)	By Sector (firm-year counts)			
	Non-GLCs	Tier 2	Tier 1	Total
Consumer Discretionary	215			215
Consumer Staples	197		14	211
Energy	73			73
Financials	137	18	18	173
Health Care	49			49
Industrials	275	78	126	479
Information Technology	170		18	188
Materials	92			92
Real Estate	516	119	18	653
Telecommunication Services	1	30	18	49
Utilities	30	10		40
Missing	529	81		610
Total	2,284	336	212	2,832

maximisation incentives as the non-GLCs listed firms.⁹

GLCs in Other Studies. Two studies have looked at the link between GLCs and firm value in Singapore, finding a positive relationship. Ramírez and Tan (2004) use a 1994–98 sample and find that GLCs have 20% more market value relative to their non-GLC competitors ($p < 0.05$). Replicating their findings using the current data using a similar set of controls for the period 2000–2017 produces similar estimates (Table B.0.4)—GLCs have 16% higher firm value ($p < 0.01$). Ang and Ding (2006) also look at firms' GLC status and firm value using

⁹The GLC firms are *government-linked*, and not *state-owned* because there exist at least 2 degrees of separation: the Government of Singapore wholly owns Temasek Holdings—a private holdings company—which in turn holds shares private firms. Further separation exists if the GLC directly owned by Temasek Holdings holds shares of yet other firms. These are the Tier 2 GLCs in this paper.

FIGURE I
INCREASES IN THE REPRESENTATION OF WOMEN IN PARLIAMENT

percentage ownership by Temasek Holdings as a continuous measure of government linkage from a 1990–2000 sample. They find that a 1 percentage-point increase in Temasek Holdings ownership of a firm is associated with a 0.009 increase in Tobin's Q ($p < 0.05$). A partial replication of that result (Table B.0.5) also produces similar findings. Replicating their results suggests that the GLCs in the current sample are comparable to ones used in other studies.¹⁰

2.2.3 Share of Women in Parliament

The primary data source for the share of women in parliament is the records of the individual parliamentarians, which include information on their gender, ministerial portfolio, and ministerial type. I aggregate these individual records up to the individual-year, and then to the year level for the share of women in parliament. A secondary source comes from the IPU (Inter-Parliamentary Union).

The primary data source offers two alternative measures of women representation in parliament. The first is the *pay* share of the women in

¹⁰Using the current sample, the above results are not robust to including firm fixed effects, more refined industry group classification, and clustering of observations at the firm level.

TABLE II
SUMMARY STATISTICS

	Summary by GLC status				Differences by GLC status		
	Non-GLCs	Tier 2 GLCs	Tier 1 GLCs	All	Non-GLC–Tier 2	Non-GLC–Tier 1	Tier 1–Tier 2
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Board characteristics							
Board size	5.895 (3.015)	7.232 (3.456)	9.072 (1.705)	6.317 (3.116)	−0.869*** (−4.60)	−4.406*** (−17.54)	3.537*** (11.86)
Number of supervisors	4.455 (2.493)	6.397 (3.425)	8.179 (1.774)	4.986 (2.795)	−1.288*** (−8.16)	−4.477*** (−21.71)	3.189*** (11.45)
Number of women	0.434 (0.655)	0.581 (0.753)	0.762 (0.799)	0.478 (0.686)	−0.149*** (−4.76)	−0.403*** (−9.33)	0.253*** (4.05)
At least 1 woman	0.374 (0.484)	0.500 (0.501)	0.574 (0.496)	0.405 (0.491)	−0.107*** (−4.68)	−0.274*** (−8.62)	0.166*** (4.10)
At least 2 women	0.326 (0.469)	0.410 (0.493)	0.511 (0.501)	0.351 (0.478)	−0.0780*** (−3.60)	−0.253*** (−8.39)	0.175*** (4.45)
At least 3 women	0.0665 (0.249)	0.0843 (0.279)	0.148 (0.356)	0.0758 (0.265)	−0.0309** (−2.76)	−0.0966*** (−6.09)	0.0657** (2.67)
Proportion of women	7.681 (13.88)	6.817 (8.549)	8.562 (9.262)	7.682 (13.09)	−0.648 (−0.78)	−0.580 (−0.58)	−0.0676 (−0.07)
Firm characteristics							
Market capitalisation	99.00 (1307.1)	2.578 (2.188)	11.02 (14.30)	81.70 (1178.6)	65.02 (1.11)	55.14 (0.75)	9.886*** (11.85)
Appearances in sample	16.80 (2.489)	16.49 (2.675)	17.75 (0.971)	16.86 (2.425)	1.391*** (5.92)	−2.835*** (−8.96)	4.227*** (13.46)
Log(total assets)	7.063 (1.730)	7.517 (0.902)	9.059 (1.117)	7.292 (1.721)	−0.500*** (−4.98)	−2.222*** (−16.87)	1.722*** (16.65)
Revenue	92.62 (1072.3)	1.265 (1.262)	7.294 (5.765)	76.04 (967.1)	57.54 (1.31)	51.25 (0.89)	6.290*** (21.28)
Debt-to-equity	57.78 (110.3)	122.4 (297.0)	71.40 (75.71)	65.14 (138.0)	−34.98*** (−5.04)	−23.89*** (−3.36)	−11.09 (−0.70)
Return on assets	5.017 (15.28)	10.88 (11.99)	5.302 (5.366)	5.597 (14.45)	−30.02 (−0.63)	−25.88 (−0.42)	−4.134*** (−5.13)
Cash-to-assets	14.81 (13.89)	14.85 (14.07)	12.66 (7.091)	14.61 (13.43)	1.926* (2.52)	1.648 (1.74)	0.278 (0.29)
Log(Tobin's q)	0.174 (0.512)	0.484 (0.513)	0.293 (0.312)	0.214 (0.506)	−0.151*** (−4.74)	−0.118** (−3.03)	−0.0337 (−0.91)
R&D expenditure	262.2 (3196.3)	0 (0)	11.34 (26.80)	214.1 (2882.5)	198.2 (1.08)	187.4 (0.94)	10.73*** (6.28)
Systematic risk (beta)	1.241 (10.38)	1.256 (9.834)	0.395 (5.347)	1.163 (9.969)	−0.807 (−1.20)	0.0932 (0.11)	−0.900 (−1.43)

This Table reports summary statistics for the full sample and the subsamples by firms' GLC status. The sample comprises 3'438 firm-year observations from 2000–17. Board data are aggregated from individual director board data from the *BoardEx* database. Firm financial data come from the *Bloomberg* database. Non-ratio and non-log figures are reported in millions. ***, **, * denotes significance at the 1%, 5%, and 10% level for the *t*-test of differences in columns (5)–(7). Standard deviations in parentheses for columns (1)–(4); *t*-statistic in parentheses in columns (5)–(7). Observations unweighted.

Parliament—the pay of all women divided by the pay of all parliamentarians, where pay is increasing in ministerial rank. This measure captures the notion that not all members in Parliament are equal. The second alternative measure is the *elected* share of women in Parliament—the number of elected women divided by the total number of elected parliamentarians.¹¹ Figure I provides some insight into the variation of women representation in parliament using the three measures, with the timing of the election into parliament indicated by the four vertical lines.

Summary Statistics. Table I presents summary statistics for the board characteristics aggregated from the 48,072 individual director data from BoardEx and firm financial characteristics from Bloomberg, by both the full sample and also by firms' GLC status. The firm-year panel over the period 2000–17 comprises 2,832 firm-year observations from 191 firms and 18 years.¹²

2.3 Identification Strategy

The main hypothesis of interest is that more representation of women on boards can benefit firm financials. One source for this expectation is fundamental gender differences. Experimental studies, for instance, find women more generous and cooperative (Cadsby and Maynes 1998; Eckel and Grossman 1998; Ortmann and Tichy 1999). Another explanation is that a board functions as a monitoring mechanism ensuring that the firm exercises proper management (Fama and Jensen 1983). This mechanism is weakened in the presence of networks or old boys' clubs, and introducing more women on boards disrupts the networks and reinforces the monitoring. Adams and Ferreira (2009) for example, find that men in boards with more women have higher attendance, and the women themselves had higher attendance than the men. They also find greater sensitivity between CEO turnover and stock performance, and

¹¹For cases where a parliamentarian had a change in salary (switching between political portfolios), the earliest salary is used.

¹²Table B.0.6 lists the financial variables from Bloomberg, the descriptions, and the relevant field names.

a greater proportion of CEO compensation coming from equity-based pay.

To test if having more women on boards improve firm financial performance, I estimate:

$$y_{st} = \varphi_0 + \varphi_1(\text{women on boards})_{st} + \delta_s + \tau_t + \Phi' X_{st} + u_{st}, \quad (2.1)$$

where y_{st} is the financial outcome of firm s in year t ; which is in turn, (i) profitability (ROE), (ii) firm value (Tobins' q), and (iii) firm risk (beta from the CAPM). X_{st} is a vector of firm characteristics which are potentially correlated with both board composition and financial performance. The full specification includes firm fixed effects (δ_s), which removes permanent differences between firms such as corporate culture. Also included are year fixed effects (τ_t), which takes out idiosyncratic annual market shocks affecting all firms so that φ_1 , the main coefficient of interest, identifies the effect of female board representation on firm financial outcomes at a given firm-year observation. The year fixed effects also account for non-binding changes in the code of corporate governance (around 2005). The main challenge is in obtaining consistent estimates of φ_1 .

A preliminary look at the data in Figure II suggests that firm value, profitability, but also firm risk, increase with the number of women on boards. While firm fixed effects eliminate the unobserved permanent heterogeneity among firms, other observed and unobserved firm characteristics may drive firms' board composition, while affecting financial performance. Furthermore, it may well be the case that a reverse relationship exists, in that firms with better financial performance are better placed to source for and appoint women candidates onto their boards. It is therefore unclear what interpretations can be made about estimates from an OLS regression under equation (2.1).

To overcome these concerns I exploit the institutional setting of Singapore using firms from the Singapore Exchange. Specifically, in the first stage of a two-stage least squares (2SLS) regression analysis, I instrument the number of women on boards using the joint effect of firm linkage to the government (*GLC*) and the variation of the share of

FIGURE II
WOMEN DIRECTORS AND FIRM PERFORMANCE (BINNED MEANS)

women in the Parliament of Singapore (*share of women in parl.*):

$$(women\ on\ boards)_{st} = \gamma_0 + \gamma_1(GLC_s \times share\ of\ women\ in\ parl_t) + \mathbf{\Gamma}'\mathbf{X}_{st} + \delta_s + \tau_t + \varepsilon_{st}, \quad (2.2)$$

where GLC is the dummy indicator for whether firm s is a GLC. The above equation (2.2) is a general form of the difference-in-differences (DD) strategy where the two dimensions of variation are firms' GLC-status and the representation of women in parliament. Figure III gives some insight into the variation of women representation in both parliament and on boards and also provides some evidence that the representation of women on boards is similar for all firms before the sharp increase in the share of women in parliament.

As discussed in the previous section, there is considerable correlation between GLC status and firm financial performance, specifically firm value. One concern with the instrument in equation (2.2) is that the joint effect is simply picking up on variation between firms by government linkage. If this is true, then the exclusion restriction is violated. To mitigate such concerns, all main regressions include firm fixed effects, which removes time-invariant heterogeneity across firms, including

GLC status. This ensures that γ_1 is not merely picking up the effects of government linkage alone on board composition. While only suggestive, in the following section I show that government linkage (using the limited data on percentage shares owned by Temasek Holdings) alone cannot predict the number of women on boards.¹³

A second concern is that the variation in the share of women in parliament is picking up yearly fluctuations in the market, or that better market conditions make it more likely that women candidates put themselves forward as political candidates. To mitigate such concerns, all regressions include year fixed effects. This prevents the estimates from picking up market trends. Moreover, better market conditions, if anything, make it less likely that potential candidates give up their private careers in pursuit of a public one.

In the sample, GLC status varies only at the firm level, and the share of women in parliament varies only at the year level. Another concern might be that the joint effect of the two variables potentially captures only firm or industry-specific time-trends since performance is also industry dependent. To mitigate such concerns, in further checks I include sector-specific linear trends, and the results are unaffected.¹⁴

The key remaining assumption is on the exogeneity of the share of women in parliament. Since board composition and firm performance are likely to be highly correlated, the direction of change in female representation should go from parliament to corporate boards of directors, and never the other way round. Hence, yet another concern is that the rise of the share of women in parliament is influenced by or anticipating the rise of women representation on boards. This would constitute another violation of the instrument restriction criterion. I argue anecdotally that this is not the case—that it is corporate firms that follow the lead of the government, as implied directly by the DD analysis in equation (2.2), rather than the reverse.

¹³Moreover, a replication of the results in Ang and Ding (2006) using the sample in this paper suggests that once firm fixed effects are included, government linkage is no longer predicts firm value (Table B.0.5).

¹⁴Industry of all firms in the sample are classified using the GICS (Global Industry Classification Standard), obtained from Bloomberg.

FIGURE III

INCREASES IN THE REPRESENTATION OF WOMEN ON CORPORATE BOARDS

For example, Tan (2016) suggests that the rise of women politicians in Singapore has to do with the establishment of "safe seats" within group representations because of the political dominance of the ruling party, and where most seats in the Parliament of Singapore are contested in group-represented constituencies. In these constituencies, parties put forth a group of four to six candidates with each assigned to a ward in that constituency, and the winning party is determined at the group level with the most consolidated votes. These safe seats refer to the political dominance of the ruling party (with supermajority since independence) and how its dominance allows it to form groups that are typically anchored by a senior politician with a cabinet position, winged by three to five other spots, with one strategically allocated to a woman. Anecdotally, the performance of the group hinges mostly on the senior anchor, and less on the other candidates, whether held by women or men.

To illustrate this strategic allocation of political candidates by the ruling party, at the time of this writing 16 constituencies operate under group representation (with 13 other constituencies under single-member constituency), 15 of which are held by the ruling party. Of these 15 ruling-party constituencies, only 2 have more than 1 seat held by a

cabinet member, another 2 have more than 1 seat held by a woman, but by far most constituencies (10 of them) have exactly 1 seat held by a woman. This strategic placement of women candidates in group representation, together with the ruling party's shift away from patriarchy with an eye on the younger and more progressive electorate (Tan, 2015, 2016), suggests that the primary, and possibly only, motivation of women representation in politics is political.¹⁵

The above also cements the idea that the strategic allocation of political candidates is a top-down process. The choice on which women (or men) to stand at which constituencies is a centralised decision made by the inner circle of the ruling party, with candidates undergoing meetings and shadowing a mentor before the decision is made by the selection committee (Tan 2015). No primary election is held to select the candidates. Hence, it is highly unlikely that the timing of women political representation is capturing sentiments for women representation among the public since the allocation of candidates is not determined through electoral competition.

Finally, to account for dependence within observations, all regressions cluster the standard errors at the firm level. This imposes no restriction on the error correlations within firms, and therefore allows for the errors to be serially correlated, as is usually the case with financial variables.¹⁶¹⁷

¹⁵There are also suggestions that the group constituencies with larger district magnitudes tend to stand more women since they represent more of the electoral and constitutes a form of ticket balancing (Studlar and Welch 1991; Salmond 2006).

¹⁶Statistical significance is unaffected when using clustering of standard errors by industry level, and two-way clustering by industry *and* firm.

¹⁷It would have also been possible to use an identification strategy that comes from politically salient events. For instance, at least three events occurred in the year 2004: 1) legislation change allowing citizenship to be passed down to a child by descent from both parents (previously only possible through the father), 2) the appointment of a woman as the first full minister with her own portfolio, and 3) Temasek Holdings appointed its a woman as its first CEO. These events could have potentially also played a signalling role to private firms about the representation of women, but it turns out that using 2004 as a post-event dummy in a difference-in-differences approach does not satisfy the pre-trends assumption (Figure B.0.3). Going ahead with this first stage does not change the 2SLS conclusion anyway.

TABLE III
THE EFFECT OF WOMEN PARLIAMENT REPRESENTATION ON WOMAN BOARD
REPRESENTATION

	Dep. var. is (effective) count of women directors				
	Both supervisory & executive			Supervisory	Executive
	(1)	(2)	(3)	(4)	(5)
GLC $\times$ share of women in parl.	0.025** (0.011)	0.024** (0.011)	0.023* (0.014)	0.020* (0.012)	0.003 (0.005)
Tier 1 $\times$ share of women in parl.	0.045** (0.020)	0.047** (0.021)	0.049** (0.021)	0.046*** (0.017)	0.003 (0.008)
Tier 2 $\times$ share of women in parl.	0.015 (0.011)	0.013 (0.011)	0.001 (0.015)	-0.002 (0.015)	0.003 (0.007)
Year fixed effects	X	X	X	X	X
Firm fixed effects		X	X	X	X
Finance controls (lagged and first differenced)			X	X	X
Mean of Dep. Var.	0.374	0.374	0.435	0.324	0.112
N	3247	3247	2714	2714	2714

Notes—Dependent variables are the (effective) count of women directors in a firm-year, computed using individual director seat occupation data from BoardEx. The dependent variable in columns (1)–(3) is the count of all women directors; in column (4) it is the count of women directors in supervisory positions; and in column (5) the count of women directors in executive positions. First panel uses a naive GLC indicator that takes on value 1 for all firms with at least 15% ownership by Temasek Holdings. Second panel uses two GLC indicators, making a distinction between Tier 1 GLCs and Tier 2 GLCs. Tier 1 GLCs are those with at least 20% direct ownership by Temasek Holdings. Lagged and first-differenced finance controls are the log of total assets, debt to equity ratio, ROA, and cash holdings to total assets. All models control for the total count of directors on the right-hand-side, which is significant and has the expected (positive) sign. Robust standard errors clustered at the firm level are in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

2.4 First-Stage Results

2.4.1 Results

Table III presents the first-stage results, where the joint effect of GLC status and the share of women in parliament predict the number of women on boards. In the first panel are the results from equation (2.2) with a single and broad GLC indicator (=1 for both Tier 1 and Tier 2 GLCs). Conditional only on the time trend and board size in column (1), the estimate for β is positive and significant ($p < 0.05$). Column (2) adds firm fixed effects and the estimate barely changes.

The models in columns (1) and (2) are parsimonious. Adding time-varying firm financial covariates is possible but risks including bad controls (Angrist and Pischke 2009)—those controls that are also determined by the main variable of interest. Nonetheless, column (3)

adds firm financial covariates which allow the size of a firm, leverage, profitability, and liquidity to determine changes in board composition (Kaplan 1995; Ferris et al. 2003). These covariates enter as lagged and first-differenced terms to mitigate issues of bad control. The estimate remains unchanged, but the standard errors increase slightly, and significance drops to the 10% level.

The second panel distinguishes between Tier 1 and Tier 2 GLCs (replacing the GLC dummy in equation (2.2) to two dummies, one for each tier), and the overall finding is that the response of the GLCs come mainly from the Tier 1 GLCs. In the sample period, the largest increase of women in parliament is approximately 20 pp. Using the estimate of 0.49 ($p < 0.05$) from column (3), a 20 pp increase in the share of women in parliament leads to Tier 1 GLCs having a full headcount more of women directors relative to the non-GLCs. This is 1.6 times of the standard deviation, and more than twice the mean, of the effective count of women directors.

The results from columns (4)–(5) indicate that virtually all the increase in women directors comes from those appointed into supervisory seats. I interpret the DD finding as one where the Tier 1 GLCs are more likely to pick up cues from the government on higher women representation. The results in Table III also suggest that the instrument is unlikely to be confounded by the financial covariates. For this reason, the second-stage regressions below omit the financial covariates for parsimony.

2.4.2 ADDITIONAL RESULTS AND ROBUSTNESS

At Least 1, 2, and 3 Women Directors. Using dummy indicators for having at least 1, 2, and 3 women directors, it turns out that the higher number of women on Tier 1 GLC firms arises because of a higher probability of having at least 1 and 2 more women directors than the Tier 2 GLC and non-GLC firms. Table B.0.7 reports the results, and the estimates from columns (2) and (4) suggest that for a 20 pp increase of women in parliament, Tier 1 GLCs are 60–70% more likely to have at least 1 to 2 more women directors than non-GLC firms ($p < 0.01$). There is however no significant difference in the probability of having at least 3 women

FIGURE IV
SUMMARY OF FIRST-STAGE RESULTS

directors. I interpret this DD finding as one where Tier 1 GLCs pick up on government cues on women representation only if they have a low number of women directors (< 2) to begin with.

As a falsification check, I show that the additional increase of women directors on the boards of Tier 1 GLCs does not arise from appointing more directors (Table B.0.8). Figure IV summarises the results from the first-stage analysis, reporting the estimated DD coefficient for Tier 1 and Tier 2 GLCs (non-GLCs are the base category). Tier 1 GLCs have a higher increase in women representation, and this increase comes from women in supervisory directorships instead of executive ones. The increase can be attributed to a higher probability of at least 1 and 2 more women directors.

Robustness. I examine the robustness of the first-stage results to alternative definitions of the variables and specifications. First, the results are robust to different measures of women representation on boards. As discussed above, Tier 1 GLCs are more likely to have at least 1 to 2 women directors. I also consider the percentage of board directors who are women as the outcome measure (Table B.0.9), with observations weighted by board size. The results are qualitatively similar, though less precise. Second, the results are robust to the different imputation

methods for cases where the start or end date of individual directorship duration is missing.

Third, I consider four alternative measures of women representation in parliament: (1) share of women in parliament using the secondary IPU data, (2) *pay* share, (3) *elected* share, and (4) *elected pay* share of women in parliament (Table B.0.12). The estimates are all nearly identical to the benchmark results. I also interpret this as indicating that the headcount representation of women in parliament is a sufficiently salient form of representation.

Fourth, I use the data at the individual board director level. I define a dummy for when a woman holds a director seat in a firm-year observation for at least half the year. I then regress this dummy on the same set of variables as in the benchmark first-stage result in Table III, with additional individual board director controls—quadratics for director age and director network size from the BoardEx database. The estimates from the 38'402 director level observations in Table B.0.13 suggest that a director seat in Tier 1 GLCs is more likely to be filled by a woman for at least half the year by about 0.13% relative to the Tier 2 GLCs and non-GLC firms ($p < 0.01$). This is more than 4% of the probability ($\approx 3\%$) that a woman holds the directorship for at least half the year.

Table B.0.14 contains additional robustness checks, including coding the questionable Tier 2 GLCs as non-GLCs or dropping them from the sample, including time trends for a firm's sector and industry, allowing standard errors to be clustered at the industry level and both firm and industry level. I also weight observations by board size, market capitalisation, asset size, and the number of appearances in the yearly top 100 firms of the exchange by market capitalisation to give greater weight to larger firms with potentially lower errors in the records and data. The results are unchanged, or the effect even larger in magnitude.

As an additional falsification test, I check and show that the joint effect of government linkage and the timing of women representation in parliament cannot be replicated using just government linkage alone. I use voluntarily disclosed data on the percentage of shares ownership

by Temasek Holdings in their annual reports for Tier 1 GLCs starting in 2004. The continuous measure of ownership allows me to retain the firm fixed effects in the panel regressions. As it turns out, the estimates are never statistically significant, even when using Heckman correction because percentage ownership is only disclosed for Tier 1 GLCs in the sample (Table B.0.15). Finally, I also replicate the main result using the lagged values of women parliament representation, and the estimates remain largely unchanged (Table B.0.16).¹⁸

2.5 Second-Stage Results

2.5.1 The Effect of Female Board Representation

Return on Equity. The first outcome I consider is ROE, an accounting-based measure of market performance that measures how well a firm's senior management can generate income from its equity. The OLS estimates in column (1) of Table IV provides a simple comparison of means. Consistent with suggestions that female board representation improves ROE, the OLS estimate is positive and significant, suggesting that an additional woman in the boardroom increases ROE per year by 2.1 pp ($p < 0.05$). This is 19% of the average. As discussed above, since this OLS estimate lacks a causal interpretation, I turn to other methods to address endogeneity.

Column (2) includes firm fixed effects. The estimated coefficient for women directors becomes negative and is no longer significant. This suggests that the observed relationship between female board representation and ROE reported in past literature potentially stems from time-invariant firm characteristics that are correlated with both firm outcomes and board composition.¹⁹

¹⁸The share ownership by Temasek Holdings is imputed as zero for all non-GLC firms even though in practice there may be a small percentage from indirect ownership. Using GLC tiers dummies to measure linkage, while dropping the firm fixed effects and using industry fixed effects to control for heterogeneity across industry gives the same qualitative results (untabulated).

¹⁹As a further check, I run a Sargan-Hansen test where the null is that both the random effects and fixed effects panel specification are consistent, and the alternative hypothesis is that only the fixed effects estimator is consistent, with the test robust

TABLE IV
THE EFFECT OF FEMALE BOARD REPRESENTATION ON ROE

	Dependent variables are				
	ROE		No. of Women Directors	ROE	
	(1)	(2)	(3)	(4)	(5)
No. of women directors	2.100** (0.944)	−1.073 (2.342)		1.079 (7.463)	0.801 (1.241)
Tier 1 × share of women in parl.			0.040*** (0.015)		
<i>N</i>	2903	2903	2903	2903	2712
Year fixed effects	X	X	X	X	X
Firm fixed effects		X	X	X	X
First-stage <i>F</i>				6.64	
Anderson-Rubin Wald				.02	
Hansen					109.69
Diff-in-Hansen					0.34
AR(1)					−1.32
AR(2)					−.13
Model	OLS	Fixed Effects	2SLS First-stage	2SLS	Arellano-Bond Two-step

The dependent variable in the structural equations is ROE (return on equity) at the firm-year level. In column (4), the effective (heteroskedastic-robust) *F*-statistic is for the null that the instruments are jointly equal to 0, and is appropriate for non-i.i.d. errors. The Anderson-Rubin Wald χ^2 statistic is for the null that endogenous regressors in the structural equation are jointly equal to zero; this is a weak-IV robust test. Hansen's *J* is for the joint null that the instruments are valid—uncorrelated with the error term in the structural equation (2.1). For the system GMM results in column (5), the 1st and 2nd lags are included in the level and difference equations; Hansen's *J* is reported for all 115 instruments including the GMM-type instruments; Diff-in-Hansen reports the statistics for the same subset of strictly exogenous variables used in column (4) (interaction of GLC tier and share of women in parliament); and the AR tests show that the system is dynamically complete. Robust standard errors in parentheses. Columns (2)–(5) cluster standard errors at the firm level.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

The fixed-effects panel regression, however, does not properly account for a simultaneity bias. Columns (3)–(4) present estimates from the 2SLS specification that does. I drop the Tier 2 GLC interaction with the share of women in parliament in the second-stage analyses since it has no explanatory power in the first stage. In column (4), the 2SLS estimate suggests that the effect of female board representation is positive at 1.08, but is not statistically significant. The standard errors are large, a sign of a weak instrument, which is also indicated by the low first-stage *F*-statistic of 6.6 (heteroskedasticity-robust).

I address the weak instrument problem in three ways. First, I apply weak-IV robust inference using the Anderson-Rubin statistic for the

under heteroskedasticity. The χ^2_{17} -statistic is 79.7 ($p < 0.01$), rejecting the null and confirming that the results in column (1) are at least partly driven by time-invariant factors such as corporate culture, which influences both firm performance and board composition.

null hypothesis that the endogenous variable in the structural equation (2.1)—(effective) count of women directors—is equal to zero. Consistent with the 2SLS estimate, the insignificant Anderson-Rubin χ^2 of 0.02 from column (4) of Table IV suggests the null hypothesis that the estimate of the women on boards is zero, cannot be rejected.

The second way I address the weak-IV problem is to turn to the dynamic panel system GMM specification (Arellano and Bover 1995; Blundell and Bond 1998), estimated in column (5). This specification uses past realisations of the outcome and endogenous variables as additional GMM-type instruments. In addition, the estimate in column (5) enters the interaction of GLC status and the share of women in parliament as strictly exogenous instruments. The GMM estimate in column (5) is in line with the 2SLS estimate, with the estimate remaining statistically insignificant, having a smaller magnitude, and much smaller standard errors. The Hansen statistic for all the GMM-type and 2SLS-type instruments is 109.7, and the difference-in-Hansen statistic for the strictly exogenous instruments used in the 2SLS is 0.34, both statistically insignificant and consistent with the instruments satisfying the exogeneity assumption. The AR tests for autocorrelation show that the model satisfies the requirement that the system is dynamically complete and the differenced residuals do not exhibit AR(2) behaviour. As a third solution, I consider an alternative first-stage below in Section 2.6.2.²⁰

Firm value using Tobin's q. Table V reruns the same set of regressions with firm value as the dependent variable, measured by Tobin's q (in logs) which is less susceptible to accounting manipulation than the ROE measure. Findings from the existing literature is mixed, with evidence of a positive relationship (Carter et al. 2003), negative relationship (Ahern and Dittmar 2012; Darmadi 2013), and no relationship (Adams and Ferreira 2009; Carter et al. 2010). As it turns out, the findings for firm value in this paper are entirely consistent with the above findings with

²⁰All system GMM estimators in this paper use first differences to remove time-invariant firm fixed effects, with two-step estimation to compute the optimal weighting matrix, and reports the Windmeijer-corrected cluster-robust errors. The GMM model with ROE as the outcome uses the 1st and 2nd lags in the level and difference equations. The IV-type instrument is the interaction of GLC status and share of women in parliament.

TABLE V
THE EFFECT OF FEMALE BOARD REPRESENTATION ON FIRM VALUE

	Firm Value (ln of Tobin's q)			
	(1)	(2)	(3)	(4)
No. of women directors	0.056*** (0.014)	0.019 (0.030)	-0.035 (0.131)	0.013 (0.018)
<i>N</i>	2858	2858	2858	2667
Year fixed effects	X	X	X	X
Firm fixed effects		X	X	X
First-stage <i>F</i>			7.09	
Anderson-Rubin Wald			.07	
Hansen				53.63
Diff-in-Hansen				2.58
AR(1)				-3.76***
AR(2)				-2.05**
Model	OLS	Fixed Effects	2SLS	Arellano-Bond Two-step

The dependent variable in the structural equations is ln(Tobin's q) at the firm-year level. For the system GMM results in column (4), the 4th lag is included in the level and difference equations; Hansen's J is reported for all 73 instruments including the GMM-type instruments; Diff-in-Hansen reports the statistics for the same subset of strictly exogenous variables used in column (3) (interaction of GLC tier and share of women in parliament). Robust standard errors in parentheses. Columns (2)–(4) cluster standard errors at the firm level.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

ROE. The positive correlation disappears once endogeneity is accounted for.

The results for firm value in Table V complement the findings in Ahern and Dittmar (2012), who use a sample of Norwegian firms and exploit a new law in 2003 requiring 40% of board directors to be women. Using the pre-quota women representation as an instrument, they find that firms with larger increases in female board representation, as mandated by the law, experienced larger declines in firm value. In the context of Singapore however, no such quota exists, and the results here suggest that when firms are free to choose their directors, there is no difference between having more men or women directors in terms of firm value.²¹

The findings of no relationship between female board representation and firm value also complement the findings in Adams and Ferreira (2009). Their OLS results show a positive correlation between women

²¹The system GMM with ln(Tobin's q) as the outcome variable uses the 4th lag of the outcome and endogenous variable in the level and difference equations, since this specification has the most satisfactory diagnostics while being parsimonious to avoid the instrument proliferation problem (Roodman 2009).

TABLE VI
THE EFFECT OF FEMALE BOARD REPRESENTATION ON SYSTEMATIC RISK

	Systematic Risk (Firm beta)			
	(1)	(2)	(3)	(4)
No. of women directors	0.886*** (0.255)	0.432 (0.408)	6.028* (3.111)	0.519 (0.418)
<i>N</i>	2578	2578	2578	2390
Year fixed effects	X	X	X	X
Firm fixed effects		X	X	X
First-stage <i>F</i>			4.78	
Anderson-Rubin Wald			6.96**	
Hansen				59.27
Diff-in-Hansen				0.17
AR(1)				−4.58***
AR(2)				−.88
Model	OLS	Fixed Effects	2SLS	Arellano-Bond Two-step

The dependent variable in the structural equations is systematic risk (measured using beta in the CAPM) at the firm-year level. For the system GMM results in column (4), only the 2nd lag is included in the level and difference equations; Hansen's *J* is reported for all 78 instruments including the GMM-type instruments; Diff-in-Hansen reports the statistics for the same subset of strictly exogeneous variables used in column (3) (interaction of GLC tier and share of women in parliament). Robust standard errors in parentheses. Columns (2)–(4) cluster standard errors at the firm level.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

board representation and firm value but the sign flips to negative when including firm fixed effects, and the estimates are never statistically significant at conventional levels.²²

Systematic firm risk. The third and last outcome I consider is a firm's systematic risk, as measured by beta from the CAPM (capital asset pricing model). Beta is a measure of how sensitive a firm's stock is to non-diversifiable market risk. The conclusion is the same as before, where a positive correlation from OLS disappears with the fixed effects, 2SLS, and the GMM specifications.²³

²²When using the connections between the men and women on board as an instrument for female board representation, Adams and Ferreira (2009) find a stronger negative correlation, though this result is not robust and doubt has since been cast on the validity of board connection as an instrument for female board representation (Sila et al. 2016).

²³The system GMM model with firm risk as the outcome variable uses the 2nd lags in the level and difference equations. Using the first lag does not pass the Hansen test that the instruments (both GMM and the IV) are exogenous. The diagnostics suggest that the interaction of GLC status and share of women in parliament satisfy the exogeneity assumption, and the autocorrelation tests confirm that the model is dynamically complete.

The finding from column (1) is somewhat consistent with the "glass cliff" hypothesis, referring to how leadership positions occupied by women tend to be more precarious than those of the men, with anecdotal observations in both financial and political contexts (Ryan and Haslam 2005). This can at least be partially explained by firm fixed effects, though it still suggests that firms more affected by market conditions are also the firms with more women in the boardroom. There is however little evidence from the 2SLS and GMM estimates that the relationship runs the other way—where more women on boards increase firm risk and create precariousness in the first place.

2.6 Additional Results

2.6.1 Limits to Identification and External Validity

Here, I lay bare the institutional-specific threats to identification and follow that up with a discussion on the extent of external validity, given identification. The identification strategy taken in this study is to instrument the number of women on boards using the interaction of GLC status and the representation of women in parliament. While this identification strategy is institution-specific, so are plausible threats to identification.

First, the exclusion restriction requires that there is no corporate-related performance advantage tied specifically to GLCs. In other words, it may be inconsistent to argue at once that firms are influenced by the government on board composition through their government linkage yet receive no advantage through corporate advantage through the same channel. More specifically, one way in which the government can influence board composition is through a network of board director candidates who have government connections or experience. Just like appointing senior defence officials as advisors to firms in the military-industrial complex, appointing directors with government links may confer GLCs an advantage if GLCs engage more in business practices, products, or services that are heavily subjected to government purview.

TABLE VII
SECOND STAGE RESULTS—ALTERNATIVE FIRST-STAGE

	First-stage	Second-stage		
	Increase in women dir.	ROE	Ln of Tobin's q	Systematic Risk
	(1)	(2)	(3)	(4)
Increase in no. of women directors		53.834 (34.152)	−0.289 (0.794)	3.682 (12.590)
% Men on board × share of women in parl.	0.045*** (0.014)			
Year fixed effects	X	X	X	X
Firm fixed effects	X	X	X	X
First-stage <i>F</i>		10.41	8.6	8.1
Anderson-Rubin Wald		2.96*	.12	.08
<i>N</i>	2535	2472	2451	2322

This Table documents the two-stage least squares results using specification (2.3) as the alternative first-stage. The main estimand of interest in the structural equation is the *increase* in the number of women directors on the board for that firm-year observation. The instrument is the interaction of the number of men on the board of directors in the previous year, interacted with the share of women in parliament. Robust standard errors clustered at the firm level are in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

However, the above flaw suggests that if GLCs establish their boards to gain a corporate advantage, then the 2SLS estimates from Tables IV–VI should suffer from an upward bias but this does not appear to be a concern given the null results. Another concern is that the two-way fixed effects approach is insufficient to rule out performance trends influencing the estimates. Table B.0.18 repeats the 2SLS analyses with sector-by-year trends, and the qualitative conclusion does not change.

A related but entirely different problem also arises from using GLC status as the excluded instrument, since it limits the extent of external validity, endogeneity concerns notwithstanding. Specifically, since the first stage relies on variation in board composition that comes through GLC status, the approach says very little about how the average firm in general, which is not a GLC, will benefit from having more women on their boards.

2.6.2 Alternative First-Stage Identification

The main 2SLS estimates from Tables IV–VI do not say much about how firm financials would change with more women on boards for the

average firm, since most of them are not GLCs, with identification notwithstanding.

To address this concern in part, and at the same time to address the weak IV problem, I turn to an alternative excluded instrument in the first stage. The approach is analogous to the one in Stevenson (2010), and does not hinge on a firm's connection to the government. Specifically, to instrument for an *increase* in women board representation, I use the lagged percentage of men in the boardroom, interacted with the share of women representation in parliament:²⁴

$$\begin{aligned} & (\text{Increase in women on boards})_{st} \\ &= \pi_0 + \pi_1 (\% \text{men on boards} \times \text{share of women in parl.})_{st} + \delta_s + \tau_t + \Pi' \mathbf{X}_{s,t} + \eta_{st}. \end{aligned} \quad (2.3)$$

This avoids the 2SLS estimates capturing an effect local to the GLCs. Moreover, the channel through which any potential causal effect is identified is more general, to the extent that firms that increase women board representation by more are those with more men to begin with.

Table VII reports the second-stage results using the alternative first-stage specification from equation (2.3). As it turns out, the results are qualitatively similar to the one using the original first-stage specification, in that an increase in the number of women directors has no identified effect on the three firm financial outcomes.

Column (1) reports the first-stage results, where the estimated coefficient of the instrument—% men on boards $\times$ share of women in parl.—suggest that if all directors on board in the previous year is a 100%, then a 20 pp increase in the share of women in parliament is associated with a 0.9 increase in the number of women in the boardroom ($p < 0.01$). The first-stage F -statistic indicates that a weak IV is less of an issue, where the $F = 10.4$ is closer to a bias of 15% of the OLS estimator. The estimates from columns (2)–(4) lead to the same conclusion as

²⁴To estimate the causal implication of high school athletic participation of girls on their college attendance and labour force participation, Stevenson (2010) instruments participation level of girls in sports after Title IX using the pre-Title IX participation level of the boys, interacted with cohort.

before—that an increase in women on boards has no impact on the firm financials. Including sector-by-year trends do not meaningfully change the conclusions (Table B.0.18).

2.7 Discussion

2.7.1 Inherent Limits to Causal Inference

Here, I point out a fundamental problem with the goal of causal inference that focuses on the impact of board gender composition on firm outcomes. This can be generalised to diversity concerns not related to gender or to the board.

Most studies on the subject of board and management diversity on firm outcomes, to my knowledge, are descriptive in that they compare firms with more women directors to firms with fewer along dimensions relating to financial performance and social responsibility type of measures (Post and Byron 2014; Pletzer et al. 2015). These studies provide insights on what kind of equilibrium conditions are related to board composition, and in the context of this study at least, one such equilibrium condition appears to be firm connection to the government.

To study the causal impact of board gender composition, on the other hand, requires variation in the firm environment that is unrelated to any corporate dimension. One such study is the one by Ahern and Dittmar (2012), where they use the exogenous and drastic policy reform in Norway requiring that 40% of the board be women. However, a direct causal interpretation is still difficult in that setting if one views the pre-policy composition as optimal and that the post-policy composition increases cost to the firm because of the cost in adjustment and compliance. In this regard, even an exogenous policy reform that changes board composition does not shed very clean insight into the causal effect of board diversity. If anything, it highlights the 'shadow price' of increasing diversity, since firms that are further away from the quota pre-reform are the ones that suffer bigger losses in firm value. Furthermore, even in post-reform,

there will likely be heterogeneity in firm access to women director candidates, since firms still retain some discretion in who to hire, as opposed to a truly random assignment of candidates to boards.

In this sense, a change in board composition post-reform can still be considered endogenous. Unfortunately, this chapter will not be able to contribute anything further to resolve this conundrum. Furthermore, with regards to causal inference, the channel through which diversity can affect firm performance is still an open question, with many studies having essentially a reduced-form approach. Several hypotheses for potential channels include augmenting the monitoring role of the board (Adams and Ferreira 2009; Fama and Jensen 1983) and reducing litigation (Arnaboldi et al. 2021).

2.7.2 A nudge is enough?

A related issue is that diversity does not matter until a critical mass is reached (Arnaboldi et al. 2021). For the interest of the context in this chapter, one interpretation that cannot be ruled out with the data, is that there is insufficient variation in women representation in the sample frame to be able to capture any meaningful and potential impact.

This, unfortunately, turns on the identification mechanism that assumes that the government can subtly influence an increase in women representation without overt quotas. If a critical mass in women representation is essential (Arnaboldi et al. 2021; Carter and Wagner 2011; Konrad et al. 2008; Schwartz-Ziv 2017), then perhaps firms in Singapore cannot attain it with government nudges alone (only 7.6% of the firm-year observations have at least 3 women in the boardroom from Table I). Moreover, as evident in the results from Table III, most of the increase in women directors comes from supervisory directors rather than executive directors. The above suggests a glass ceiling, and one should therefore expect the impact of an increase from supervisory directors to be marginal.

2.8 FINAL REMARKS

This paper tests whether having more women on boards improves firm financials, by using variation of women representation in politics to instrument for the variation of women representation on boards. There is little evidence of a causal impact, which can be explained in at least three ways.

First, various studies using real-world observations find that differences in gender can bring about changes (e.g. Dollar et al., 2001; Edlund and Pande, 2002; Oswald and Powdthavee, 2010). In a seminal paper on political mandates for gender quota, Chattopadhyay and Duflo (2004) use India's 73rd constitutional amendment as the exogenous change to find that villages where councils have seats reserved for women experienced increases in public goods of greater concern to women, suggesting differential preferences by gender. These findings however may not translate into a financial setting, and the evidence from India is likely influenced by developmental factors that further differentiate the preferences of men and women.

In financial settings specifically, Holt and Laury (2002) find that gender differences in risk-aversion disappear with high payoffs, and Deaves et al. (2009) find women just as (over-) confident as the men. Adams and Funk (2011) use a large survey of board directors and find their gender-based differences do not reflect those of the general population.

A second explanation is the lack of variation in the sample, if a critical mass effect of three women is needed before changes can be pushed through (Carter and Wagner 2011; Konrad et al. 2008; Schwartz-Ziv 2017).²⁵ A related issue is that most of the increase in women directors comes through supervisory rather than executive directors with more latitude in implementing changes. A third explanation is that changes from higher women representation need not be financial (as in Adams

²⁵In the appendix, Table B.0.17 supports the critical mass rhetoric, insofar as the point estimates for at least 1, at least 2, and at least 3 women directors, respectively increase in magnitude. A dramatic jump in magnitude, in particular, is seen in the increase from two to three women directors. The standard errors however are large, and the point estimates themselves are never statistically significant at conventional levels.

and Ferreira 2009; Gender Diversity Taskforce 2014), and that there could be more holistic and non-financial impact of gender diversity not identified in this study.

Finally, the policy question remains open on whether subtle nudges in women representation are sufficient if diversity only works with at some critical mass.

Appendices

Appendix B

Data Appendix

TABLE B.O.1
LIST OF DROPPED FIRMS

Ticker	Firm Name	Reason
SPH100	Singapore Press Holdings	2nd listing for SPH
SBIF	ABF SINGAPORE BOND INDX FUND	index
KPLD2	Keppel Land ltd	2nd listing for KPLD
MLTA	Mapletree Logistics Trust-A	2nd listing for MLT
YHS1	Yeo Hiap Seng ltd-1	2nd listing for YHS
STTF	SPDR STRAITS TIMES INDEX ETF	index
WINGT1	Wing Tai Holdings ltd-100	2nd listing for WINGT
M110	Mobileone ltd-10	2nd listing for M1
GENS	Genting Singapore Plc	Accidental omission
CIH	CIH Ltd	no recorded board role in sample period
LBRY	DP Marine Pte Ltd	no recorded board role in sample period
COMF	Comfort Group Ltd	no recorded board role in sample period
GES	GES International Ltd	no recorded board role in sample period
GPK	Goodwood Park Hotel Ltd	no recorded board role in sample period
JTL	Jurong Technologies Industrial	no recorded board role in sample period
MAGIC	Mapletree North Asia Commercial Trust	Name changed from Mapletree Greater China Commercial Trust in 2017
CP	Frasers Centrepoint Ltd/old	Name changed to Frasers Centrepoint Ltd (FCL)
INDIA	ISHARES MSCI India Index ETF	index
DCL	Delgro Corp Ltd	cannot find in BoardEx database
KTF	KTF limited	cannot find in BoardEx database
OUT	UOB Trust ltd	cannot find in BoardEx database

Notes—Firm ticker and names are as reported in the Bloomberg terminal. No recorded board role for a firm in the sample period happens either because the recorded roles fall outside the sample period, or that there were roles falling inside the sample period, but either the start or end dates were missing.

TABLE B.0.2
LIST OF FIRMS IN SAMPLE PERIOD 2000–17

Firm Name	Ticker	GLC	Tier	Average share held directly by Temasek Holdings (2004–17)	Questionable GLCs
Neptune Orient Lines Ltd	NOL	Yes	1	66.1	—
Stats Chippac Pte Ltd	STAT	Yes	1	60.2	—
Smrt Corp Ltd	MRT	Yes	1	56.6	—
Singapore Airlines Ltd	SIA	Yes	1	55.7	—
Singapore Telecommunications	ST	Yes	1	54.7	—
Singapore Tech Engineering	STE	Yes	1	51.6	—
Sembcorp Industries Ltd	SCI	Yes	1	49.3	—
Sats Ltd	SATS	Yes	1	42.7	—
Capitaland Ltd	CAPL	Yes	1	40.9	—
Dbs Group Holdings Ltd	DBS	Yes	1	28.5	—
Olam International Ltd	OLAM	Yes	1	22.8	—
Keppel Corp Ltd	KEP	Yes	1	22.3	—
Tiger Airways Holdings Ltd	TGR	Yes	2	—	—
Capitaland Retail China Trus	CRCT	Yes	2	—	—
Singapore Food Industries	SFI	Yes	2	—	—
Sia Engineering Co Ltd	SIE	Yes	2	—	—
Keppel Infrastructure Trust	KIT	Yes	2	—	Yes
K1 Ventures Ltd	KONE	Yes	2	—	—
Capitaland Commercial Trust	CCT	Yes	2	—	—
Mapletree Commercial Trust	MCT	Yes	2	—	—
Sembcorp Logistics Ltd	1025435D	Yes	2	—	—
Mapletree North Asia Commere	MAGIC	Yes	2	—	—
Keppel Dc Reit	KDCREIT	Yes	2	—	Yes
Mapletree Logistics Trust	MLT	Yes	2	—	—
Starhub Ltd	STH	Yes	2	—	—
Capitaland Mall Trust	CT	Yes	2	—	—
M1 Ltd	M1	Yes	2	—	Yes
Keppel Shipyard Ltd	KPHZ	Yes	2	—	Yes
Raffles Holdings Ltd	RHL	Yes	2	—	Yes
Singapore Post Ltd	SPOST	Yes	2	—	—
Keppel Land Ltd	KPLD	Yes	2	—	—
Capitamalls Asia Ltd	CMA	Yes	2	—	Yes
Ascott Residence Trust	ART	Yes	2	—	—
Mapletree Industrial Trust	MINT	Yes	2	—	—
Keppel Telecom & Transport	KPTT	Yes	2	—	—
Keppel Reit	KREIT	Yes	2	—	—
Sembcorp Marine Ltd	SMM	Yes	2	—	—
Sakari Resources Ltd	SAR	—	0	—	—
Hotel Grand Central Ltd	GRAN	—	0	—	—
Cosco Shipping International	COS	—	0	—	—
Csa Holdings Ltd	CSA	—	0	—	—
Golden Agri-Resources Ltd	GGR	—	0	—	—
Want Want Holdings Ltd	WANT	—	0	—	—
Times Publishing Ltd	TIMES	—	0	—	—
Amtek Engineering Ltd	AMTEK	—	0	—	—
Mfs Technology Ltd	MFS	—	0	—	—
Singapore Exchange Ltd	SGX	—	0	—	—
Huan Hsin Holdings Ltd	HUAN	—	0	—	—
Perennial Real Estate Holdin	PREH	—	0	—	—
Ho Bee Land Ltd	HOBEE	—	0	—	—
Fortune Reit	FRT	—	0	—	—
Fragrance Group Ltd	FRAG	—	0	—	—
Silverlake Axis Ltd	SILV	—	0	—	—
Far East Hospitality Trust	FEHT	—	0	—	—
Brook Crompton Holdings Ltd	BC	—	0	—	—
Elec & Eltek Int Co Ltd	ELEC	—	0	—	—
United Overseas Bank Ltd	UOB	—	0	—	—
China Energy Ltd	CEGY	—	0	—	—
Oversea-Chinese Banking Corp	OCBC	—	0	—	—
Krisenergy Ltd	KRIS	—	0	—	—
Wilmar International Ltd	WIL	—	0	—	—
Kim Eng Holdings Ltd	KEH	—	0	—	—
Vertex Venture Holdings Ltd	VERTX	—	0	—	—
Chuan Hup Holdings Ltd	CH	—	0	—	—
Sbs Transit Ltd	SBUS	—	0	—	—
Guocoland Ltd	GUOL	—	0	—	—

Continued on next page

Table B.0.2 – continued from previous page

Firm Name	Ticker	GLC	Tier	Average share held directly by Temasek Holdings (2004–17)	Questionable GLCs
Hong Leong Finance Ltd-Forei	SFINF	—	0	—	—
Sheng Siong Group Ltd	SSG	—	0	—	—
City Developments Ltd	CIT	—	0	—	—
Parkway Holdings Ltd	PWAY	—	0	—	—
Fu Yu Corp Ltd	FUYU	—	0	—	—
Oxley Holdings Ltd	OHL	—	0	—	—
Frasers Hospitality Trust	FHT	—	0	—	—
Gallant Venture Ltd	GALV	—	0	—	—
Oceanus Group Ltd	OCNUS	—	0	—	—
Suntec Reit	SUN	—	0	—	—
Citic Envirotech Ltd	CEL	—	0	—	—
Banyan Tree Holdings Ltd	BTH	—	0	—	—
United Industrial Corp Ltd	UIC	—	0	—	—
Goodpack Ltd	GPACK	—	0	—	—
Jardine Cycle & Carriage Ltd	JCNC	—	0	—	—
Cerebos Pacific Ltd	CER	—	0	—	—
Hong Leong Asia Ltd	HLA	—	0	—	—
Boustead Singapore Ltd	BOCS	—	0	—	—
Sph Reit	SPHREIT	—	0	—	—
Haw Par Corp Ltd	HPAR	—	0	—	—
Straits Trading Co Ltd	STRTR	—	0	—	—
Otto Marine Ltd	OTML	—	0	—	—
Starhill Global Reit	SGREIT	—	0	—	—
Great Eastern Holdings Ltd	GE	—	0	—	—
Cdl Hospitality Trusts	CDREIT	—	0	—	—
Heineken Asia Mtn Pte Ltd	APB	—	0	—	—
Robinson & Co Pte Ltd	ROB	—	0	—	—
Singapore Press Holdings Ltd	SPH	—	0	—	—
Osim International Pte Ltd	OSIM	—	0	—	—
Cefc International Ltd	CEFC	—	0	—	—
Sinarماس Land Ltd	SML	—	0	—	—
Bukit Sembawang Estates Ltd	BS	—	0	—	—
Yoma Strategic Hldgs Ltd	YOMA	—	0	—	—
Raffles Education Corp Ltd	RLS	—	0	—	—
Halcyon Agri Corp Ltd	HACL	—	0	—	—
Gp Batteries Intl Ltd	GP	—	0	—	—
Natsteel Broadway Ltd	NBKY	—	0	—	—
Cache Logistics Trust	CACHE	—	0	—	—
Allgreen Properties Ltd	AG	—	0	—	—
Frasers Centrepont Ltd/Old	CP	—	0	—	—
Gmg Global Ltd	GGL	—	0	—	—
Asian Pay Television Trust	APT	—	0	—	—
Netlink Nbn Trust	NETLINK	—	0	—	—
Hwa Hong Corp Ltd	HWAH	—	0	—	—
Cwt Ltd	CWT	—	0	—	—
Ascendas Real Estate Inv Trt	AREIT	—	0	—	—
Wing Tai Holdings Ltd	WINGT	—	0	—	—
Super Group Ltd	SUPER	—	0	—	—
Ying Li International Real E	YINGLI	—	0	—	—
Singapore Computer Systems	SCS	—	0	—	—
Idt Holdings Singapore Ltd	IDT	—	0	—	—
Jaya Holdings Ltd	JAYA	—	0	—	—
Venture Corp Ltd	VMS	—	0	—	—
Ezra Holdings Ltd	EZRA	—	0	—	—
Pan Pacific Hotels Group Ltd	PPAC	—	0	—	—
United Engineers Ltd	UEM	—	0	—	—
Sc Global Developments Ltd	SCGD	—	0	—	—
Raffles Medical Group Ltd	RFMD	—	0	—	—
Dimension Data Asia Pacific	DAT	—	0	—	—
Ascendas Hospitality Trust	ASCHT	—	0	—	—
Hi-P International Ltd	HIP	—	0	—	—
Delfi Ltd	DELFI	—	0	—	—
Hotel Properties Ltd	HPL	—	0	—	—
Indofood Agri Resources Ltd	IFAR	—	0	—	—
Gsh Corp Ltd	GSH	—	0	—	—
Oue Commercial Real Estate I	OUECT	—	0	—	—
Yanlord Land Group Ltd	YLLG	—	0	—	—
Globalfoundries Singapore Pt	CSM	—	0	—	—

Continued on next page

Table B.0.2 – continued from previous page

Firm Name	Ticker	GLC	Tier	Average share held directly by Temasek Holdings (2004–17)	Questionable GLCs
Japfa Ltd	JAP	—	0	—	—
United Test And Assembly Cen	UTAC	—	0	—	—
Ezion Holdings Ltd	EZI	—	0	—	—
Fraser's Centrepont Trust	FCT	—	0	—	—
Aims Amp Capital Industrial	AAREIT	—	0	—	—
Ascendas India Trust	AIT	—	0	—	—
Liongold Corp Ltd	LIGO	—	0	—	—
Oue Ltd	OUE	—	0	—	—
Oue Hospitality Trust	OUEHT	—	0	—	—
Pan-United Corp Ltd	PAN	—	0	—	—
Mewah International Inc	MII	—	0	—	—
Hong Leong Finance Ltd	HLF	—	0	—	—
Hyflux Ltd	HYF	—	0	—	—
Ascott Ltd/The	SCOT	—	0	—	—
Republic Hotels & Resorts Ltd	REPUB	—	0	—	—
China Merchants Hldgs Pac Lt	CMH	—	0	—	—
Yeo Hiap Seng Ltd	YHS	—	0	—	—
Industrial & Commercial Bank	ICB	—	0	—	—
GI Ltd	GLL	—	0	—	—
Midas Holdings Ltd	MIDAS	—	0	—	—
Informatics Education Ltd	INFO	—	0	—	—
Nsl Ltd	NSL	—	0	—	—
Parkwaylife Real Estate	PREIT	—	0	—	—
Global Logistic Properties L	GLP	—	0	—	—
Crystal Trust	CRYS	—	0	—	—
Metro Holdings Ltd	METRO	—	0	—	—
Htl International Hldgs Ltd	HWA	—	0	—	—
Lafe Corp Ltd	LAPE	—	0	—	—
Gk Goh Holdings Ltd	GKG	—	0	—	—
Wheelock Properties (S) Ltd	WP	—	0	—	—
Singapore Land Ltd	SL	—	0	—	—
Tat Hong Holdings Ltd	TAT	—	0	—	—
Siic Environment Holdings Lt	SIIC	—	0	—	—
Uob-Kay Hian Holdings Ltd	UOBK	—	0	—	—
Wbl Corp Ltd	WBL	—	0	—	—
Golden Energy & Resources Lt	GER	—	0	—	—
First Real Estate Invt Trust	FIRT	—	0	—	—
Esr-Reit	EREIT	—	0	—	—
First Resources Ltd	FR	—	0	—	—
Parkson Retail Asia Ltd	PRA	—	0	—	—
Biosensors International Gro	BIG	—	0	—	—
Singapore Petroleum Co Ltd	SPC	—	0	—	—
Hutchison Port Holdings Tr-U	HPHT	—	0	—	—
Uol Group Ltd	UOL	—	0	—	—
Fraser And Neave Ltd	FNN	—	0	—	—
Unisteel Technology Ltd	USTL	—	0	—	—
China Aviation Oil Singapore	CAO	—	0	—	—
Pacific Century Region Devel	PAC	—	0	—	—
Fraser's Commercial Trust	FCOT	—	0	—	—
Mdr Ltd	MDR	—	0	—	—
Gp Industries Ltd	GPI	—	0	—	—
Comfortdelgro Corp Ltd	CD	—	0	—	—
Creative Technology Ltd	CREAF	—	0	—	—
Ara Asset Management	ARA	—	0	—	—
Lippo Malls Indonesia Retail	LMRT	—	0	—	—
Mcl Land Ltd	MCL	—	0	—	—
Beyonics Technology Ltd	BT	—	0	—	—

Notes—Firm ticker are those used under the Bloomberg Database. The GLC column indicates firms that have a close linkage to the Singapore government, defined as a 15% effective ownership by the sovereign wealth fund Temasek Holdings. Average shares of Tier 1 GLCs is a (harmonic) mean of the shares of direct Temasek Holdings ownership over the period 2004–17, which comes from the annual reports. Questionable GLC column indicates those firms where it is unclear whether they should be categorised as a Tier 2 GLC.

TABLE B.0.3
SECTOR-INDUSTRY (GICS) BREAKDOWN

Sector	Industry	Non-GLCs	Tier 2	Tier 1	Total
Consumer Discretionary	Distributors	18			18
	Diversified Consumer Services	36			36
	Hotels, Restaurants & Leisure	72			72
	Household Durables	18			18
	Media	36			36
	Multiline Retail	36			36
	Specialty Retail	36			36
	Total	252			252
Consumer Staples	Food & Staples Retailing	18		18	36
	Food Products	216			216
	Household Products	18			18
	Total	252		18	270
Energy	Energy Equipment & Services	36			36
	Oil, Gas & Consumable Fuels	72			72
	Total	108			108
Financials	Banks	36		18	54
	Capital Markets	54			54
	Consumer Finance	36			36
	Diversified Financial Services	18	18		36
	Insurance	18			18
	Total	162	18	18	198
Health Care	Health Care Equipment & Supplies	18			18
	Health Care Providers & Services	18			18
	Pharmaceuticals	18			18
	Total	54			54
Industrials	Aerospace & Defense			18	18
	Air Freight & Logistics		36		36
	Airlines		18	18	36
	Commercial Services & Supplies	54			54
	Construction & Engineering	18			18
	Industrial Conglomerates			36	36
	Machinery	90	18		108
	Marine			18	18
	Professional Services	18			18
	Road & Rail	36		18	54
	Trading Companies & Distributors	72			72
	Transportation Infrastructure	36	18	18	72
	Total	324	90	126	540
Information Technology	Electronic Equipment, Instruments & Comp	144			144
	Semiconductors & Semiconductor Equipment			18	18
	Software	18			18
	Technology Hardware, Storage & Periphera	18			18
	Total	180		18	198
Materials	Chemicals	18			18
	Construction Materials	18			18
	Containers & Packaging	18			18
	Metals & Mining	54			54
	Total	108			108
Real Estate	Equity Real Estate Investment Trusts (RE	342	180		522
	Real Estate Management & Development	396	18	18	432
	Total	738	198	18	954
Telecommunication Services	Diversified Telecommunication Services	18		18	36
	Wireless Telecommunication Services		36		36
	Total	18	36	18	72
Utilities	Multi-Utilities		18		18
	Water Utilities	36			36
	Total	36	18		54
	Missing	540	90		630

TABLE B.0.4
THE EFFECT OF GLC STATUS ON LOG OF FIRM VALUE

Partial replication of the result in Ramírez and Tan (2004) where GLCs are associated with higher firm value (log of Tobin's q, with GLCs are defined as firms with at least 20% ownership by Temasek Holdings.				
	(1)	(2)	(3)	(4)
GLC (both Tier 1 and 2)	0.162 (0.020)*** [0.066]**		0.137 (0.021)*** [0.071]*	
Tier 1 GLC		0.215 (0.026)*** [0.078]***		0.177 (0.031)*** [0.099]*
Tier 2 GLC		0.135 (0.025)*** [0.083]		0.121 (0.024)*** [0.080]
Year fixed effects	X	X	X	X
Firm finance controls	X	X	X	X
Sector fixed effects	X	X		
Industry fixed effects			X	X
F-test: Finance controls = 0	$F = 42.55^{***}$	$F = 37.94^{***}$	$F = 41.36^{***}$	$F = 37.05^{***}$
R^2	0.275	0.276	0.349	0.350
N	2533	2533	2533	2533

Notes on partial replication—Dependent variable is the log of q. GLC in this (partial) replication is a dummy for firms where Temasek Holdings holds at least 5% of voting shares. The regression results with GLC defined by Tiers make a distinction for Tier 1 GLCs with at least 20% direct ownership by Temasek Holdings. Sample period for Ramírez and Tan (2004) is from 1994–98, and include non-SGX-listed firms owned by statutory boards. The results in this table uses only SGX-listed firms for the period 2000–17. The five financial covariates in this table are EPS (earnings-to-price), log of total assets, debt-to-equity ratio, ROA (return-on-assets), and beta—the exact same covariates used in Ramírez and Tan (2004). Robust standard errors reported in parentheses; robust standard errors adjusted for clusters by 186 firms reported in brackets.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE B.0.5
THE EFFECT GLC STATUS (BY OWNERSHIP %) ON FIRM VALUE

Partial replication of the result in Ang and Ding (2006) where linkage to government, measured by % ownership, is linked to higher firm value (approx. of q).				
	(1)	(2)	(3)	(4)
GLC share	0.188 (0.095)** [0.187]	0.237 (0.107)** [0.116]	0.211 (0.121)* [0.116]	0.044 (0.267) [0.116]
Year fixed effects	X	X	X	X
Firm finance controls	X	X	X	X
Sector fixed effects		X		
Industry fixed effects			X	
Firm fixed effects				X
<i>F</i> -test: Finance controls = 0	<i>F</i> = 35.53***	<i>F</i> = 30.86***	<i>F</i> = 34.8***	<i>F</i> = 86.68***
<i>R</i> ²	0.054	0.078	0.083	0.284
<i>N</i>	2078	2078	2078	2078

Notes on partial replication—Dependent variable is an approximation of Tobin's *q*. Sample period in Ang and Ding (2006) is from 1990–2000, with percentage of government for both direct owned SGX-listed firms and the indirect ones—firms with shares directly held by Temasek Holdings. This table uses data from 2004–17, and includes only the directly owned GLCs since this is the data available from the Temasek Holdings annual reports released since 2004. Finance covariates in this table includes the log of total assets, total debt to total assets, and ROE. Ang and Ding (2006) also includes a dummy for non-duality (of CEO and board chairman) and percentage of foreign ownership, but the former is sporadic and the second is unavailable in this dataset. Robust standard errors reported in parentheses; robust standard errors adjusted for clusters by 166 firms reported in brackets.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE B.o.6
DESCRIPTION OF FIRM FINANCIAL CHARACTERISTICS FROM BLOOMBERG

Variable	Field Name	Field Mnemonic	Description
Market capitalisation	Historical Market Cap	HISTORICAL_MARKET_CAP	Total market value of all a company's outstanding shares at period-end date, calculated as Shares Outstanding $\times$ Last Closing Price.
Total assets	Total Assets	BS_TOT_ASSET	Total of all short and long-term assets as reported on firm's Balance Sheet.
Revenue	Revenue	SALES_REV_TURN	Amount of sales generated by a company after the deduction of sales returns, allowances, discounts, and sales based taxes.
Debt-to-equity	Total Debt to Total Equity	TOT_DEBT_TO_TOT_EQY	Total debt divided by total shareholders' equity, computed as (Short and Long Term Debt) / (Shareholders' Equity) $\times$ 100
Return on assets	Return on Assets	RETURN_ON_ASSET	Indicator of how profitable a company is relative to its total assets, in percentage. Computed as (Trailing 12M Net Income / Average Total Assets) $\times$ 100
Cash-to-assets	Cash to Total Assets	CASH_TO_TOT_ASSET	Ratio to measure the percentage of cash and near cash over total assets, computed as (Cash and Near Cash) / Total Assets $\times$ 100
Tobin's q	Tobins Q Ratio	TOBIN_Q_RATIO	Ratio of the market value of a firm to the replacement cost of the firm's assets, computed as (Market Cap + Total Liabilities + Preferred Equity + Minority Interest) / Total Assets
R&D Expenditure	R&D Expense	IS_RD_EXPEND	Total research & development expenditures incurred which includes R&D in profit and loss account and capitalised R&D .
Systematic risk (beta) Raw Beta		EQY_RAW_BETA	Volatility measure of the percentage price change of the security given a one percent change in a representative market index. The beta value is determined by comparing the price movements of the security and the representative market index for the past two years of weekly data.

FIGURE B.O.1
Count of directors

FIGURE B.O.2
At least 2 women directors

FIGURE B.o.3
No parallel trends pre-2004

TABLE B.o.7
TIER 1 GLCs MORE LIKELY TO HAVE AT LEAST 1 AND 2 WOMEN DIRECTORS

	At least 1 woman director		At least 2 woman directors		At least 3 woman directors	
	(1)	(2)	(3)	(4)	(5)	(6)
Tier 1 $\times$ share of women in parl.	0.029*** (0.011)	0.031*** (0.011)	0.032*** (0.012)	0.035*** (0.012)	0.012 (0.008)	0.013 (0.009)
Tier 2 $\times$ share of women in parl.	0.006 (0.008)	−0.010 (0.010)	0.003 (0.007)	−0.009 (0.011)	0.004 (0.004)	0.003 (0.006)
Year fixed effects	X	X	X	X	X	X
Firm fixed effects	X	X	X	X	X	X
Finance controls (lagged and first differenced)		X		X		X
F -test: Finance controls = 0		$F = 4.09^{***}$		$F = 4.02^{***}$		$F = 1.69^*$
R^2	0.261	0.165	0.226	0.149	0.095	0.083
N	3247	2614	3247	2614	3247	2614

Dependent variables are indicator variables for having at least 1, 2, and 3 (effective) count of women directors in a firm-year observation. Estimates come from linear probability models. Odd-numbered columns include year and firm F.E.; even-numbered columns allow common firm financial characteristics to predict the number of women on boards. Controls are the exact same set as in the benchmark results in Table III. Robust standard errors clustered at the firm level are in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE B.o.8
FALSIFICATION TEST—EFFECT OF WOMEN PARLIAMENT REPRESENTATION ON BOARD SIZE

	All directors		Supervisory directors		Men directors	
	(1)	(2)	(3)	(4)	(5)	(6)
GLC $\times$ share of women in parl.	−0.022 (0.050)		0.031*** (0.010)		−0.023* (0.014)	
Tier 1 $\times$ share of women in parl.		−0.058 (0.055)		0.039*** (0.010)		−0.049** (0.021)
Tier 2 $\times$ share of women in parl.		0.010 (0.075)		0.024 (0.015)		−0.001 (0.015)
Year fixed effects	X	X	X	X	X	X
Firm fixed effects	X	X	X	X	X	X
Finance controls (lagged, and first-differenced)	X	X	X	X	X	X
F -test: Finance controls = 0	$F = 264.93^{***}$	$F = 265.27^{***}$	$F = 26.15^{***}$	$F = 26.01^{***}$	$F = 135.55^{***}$	$F = 133.84^{***}$
Mean of Dep. Var.	5.750	5.750	4.594	4.594	5.314	5.314
R^2	0.233	0.234	0.900	0.900	0.948	0.949
N	2714	2714	2714	2714	2714	2714

Notes—Columns(1)–(2) perform a basic falsification test to show that increases in women representation cannot be easily explained by increases in total board size. The dependent variable in columns (3)–(4) is the count of directors in supervisory seats; and in columns (5)–(6) it is the count of directors in executive seats. Odd-numbered columns use the naive GLC indicator; even-numbered columns makes a distinction between Tier 1 GLCs and Tier 2 GLCs. Definitions of GLCs and controls are the exact same as in the benchmark results in Table III. Robust standard errors clustered at the 191 firms are in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE B.0.9
TIER 1 GLCS HAVE HIGHER PERCENTAGE OF WOMEN DIRECTORS

	% Women directors		% Women supervisory directors	
	(1)	(2)	(3)	(4)
Tier 1 $\times$ share of women in parl.	0.595*** (0.179)	0.436* (0.242)	0.518*** (0.146)	0.413** (0.204)
Tier 2 $\times$ share of women in parl.	0.217 (0.200)	-0.079 (0.237)	0.089 (0.196)	-0.149 (0.227)
Year fixed effects	X	X	X	X
Firm fixed effects	X	X	X	X
Finance controls (lagged and first differenced)		X		X
χ^2 -test: Finance controls = 0		$F = 45.4^{***}$		$F = 34.91^{**}$
R^2	0.027	0.112	0.025	0.097
N	2696	2351	2696	2351

Dependent variables are percentages of women directors in a firm-year observation. Observations are weighted by the size of the board. Odd-numbered columns include year and firm F.E.; even-numbered columns allow common firm financial characteristics to predict the percentage of women on boards. Controls are the exact same set as in the benchmark results in Table III. Observations weighted by board size. Robust standard errors clustered at the firm level are in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE B.0.10
FIRST-STAGE MAIN RESULTS, USING SHORTEST BOARD DURATION

	Dep. var. is (effective) count of women directors				
	Both supervisory & executive			Supervisory	Executive
	(1)	(2)	(3)	(4)	(5)
GLC $\times$ share of women in parl.	0.025** (0.010)	0.024** (0.011)	0.023* (0.014)	0.020* (0.012)	0.003 (0.005)
Tier 1 $\times$ share of women in parl.	0.046** (0.020)	0.047** (0.020)	0.049** (0.021)	0.046*** (0.016)	0.003 (0.008)
Tier 2 $\times$ share of women in parl.	0.015 (0.011)	0.013 (0.011)	0.000 (0.015)	-0.002 (0.014)	0.003 (0.007)
Year fixed effects	X	X	X	X	X
Firm fixed effects		X	X	X	X
Finance controls (lagged and first differenced)			X	X	X
Mean of Dep. Var.	0.370	0.370	0.432	0.321	0.111
N	3247	3247	2714	2714	2714

Notes— Dependent variables are the (effective) count of women directors in a firm-year, computed using individual director seat occupation data from BoardEx. This table is a replication of Table III, but where missing board data at the individual level is imputed using the shortest possible duration instead of the longest used in the main analyses. Definition of GLCs and controls are the exact same set as in the benchmark results in Table III. Robust standard errors clustered at the firm level are in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE B.O.11
FIRST-STAGE MAIN RESULTS, DROP MISSING INDIVIDUAL BOARD DATA

	Dep. var. is (effective) count of women directors				
	Both supervisory & executive			Supervisory	Executive
	(1)	(2)	(3)	(4)	(5)
GLC $\times$ share of women in parl.	0.026** (0.010)	0.024** (0.011)	0.021 (0.014)	0.019 (0.012)	0.002 (0.006)
Tier 1 $\times$ share of women in parl.	0.048** (0.021)	0.046** (0.021)	0.046** (0.022)	0.045** (0.018)	0.002 (0.008)
Tier 2 $\times$ share of women in parl.	0.016 (0.010)	0.013 (0.011)	-0.002 (0.015)	-0.005 (0.013)	0.002 (0.007)
Year fixed effects	X	X	X	X	X
Firm fixed effects		X	X	X	X
Finance controls (lagged and first differenced)			X	X	X
Mean of Dep. Var.	0.340	0.340	0.396	0.290	0.106
N	3247	3247	2714	2714	2714

Notes— Dependent variables are the (effective) count of women directors in a firm-year, computed using individual director seat occupation data from BoardEx. This table is a replication of Table III, but where missing board data at the individual level is dropped listwise instead of the using the longest possible duration imputation used in the main analyses. Definition of GLCs and controls are the exact same set as in the benchmark results in Table III. Robust standard errors clustered at the firm level are in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE B.O.12
DIFFERENT MEASURES OF WOMEN PARLIAMENT REPRESENTATION

Representation of women in parliament using	Dep. var. is the (effective) count of supervisory women directors				
	Share	Pay share	Elected share	Elected pay share	IPU
	(1)	(2)	(3)	(4)	(5)
Tier 1 $\times$ share of women in parl.	0.046*** (0.017)				
Tier 2 $\times$ share of women in parl.	-0.002 (0.015)				
Tier 1 $\times$ pay share of women in parl.		0.044** (0.018)			
Tier 2 $\times$ pay share of women in parl.		0.007 (0.015)			
Tier 1 $\times$ elected share of women in parl.			0.048*** (0.017)		
Tier 2 $\times$ elected share of women in parl.			0.004 (0.016)		
Tier 1 $\times$ elected pay share of women in parl.				0.044** (0.018)	
Tier 2 $\times$ elected pay share of women in parl.				0.008 (0.015)	
Tier 1 $\times$ IPU share of women in parl.					0.048*** (0.016)
Tier 2 $\times$ IPU share of women in parl.					0.000 (0.016)
Year fixed effects	X	X	X	X	X
Firm fixed effects	X	X	X	X	X
Finance controls (lagged and first differenced)	X	X	X	X	X
F-test: Finance controls = 0	$F = 5.48^{***}$	$F = 5.58^{***}$	$F = 5.47^{***}$	$F = 5.57^{***}$	$F = 5.35^{***}$
R^2	0.177	0.179	0.181	0.179	0.179
N	2714	2714	2714	2714	2714

This table explores the robustness of the baseline result in column (1) using alternative measures of women representation in parliament. Share in column (1) is the percentage headcount of women members in parliament. Pay share in column (2) is the percentage of women members' pay in parliament based on their ministerial type. Elected share in column (3) is the percentage headcount of elected (as opposed to nominated) women members in parliament. Elected pay share in column (4) is the pay share of elected women members. IPU in column (5) is the secondary data source of percentage headcount of women members in parliament from the Inter-Parliamentary Union. Definitions of GLCs and controls are the exact same as in the benchmark results in Table III. Robust standard errors clustered at the 191 firms are in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE B.O.13
TIER 1 BOARD SEATS MORE LIKELY TO BE A WOMAN

	Dependent variable is $B_i = 1$ if a woman held the director seat for at least half the year		Dependent variable is $1(\text{gender}_i = \text{woman}) \times \text{duration of seat}$ held for the year	
	(1)	(2)	(3)	(4)
GLC $\times$ share of women in parl.	0.0008** (0.0004)		0.0008** (0.0004)	
Tier 1 $\times$ share of women in parl.		0.0013*** (0.0004)		0.0013*** (0.0004)
Tier 2 $\times$ share of women in parl.		0.0001 (0.0005)		-0.0000 (0.0005)
Year fixed effects	X	X	X	X
Firm fixed effects	X	X	X	X
Individual controls	X	X	X	X
Finance controls (lagged, and first-differenced)	X	X	X	X
Mean of Dep. Var.	0.0297	0.0297	0.0299	0.0299
R^2	0.1039	0.1040	0.1060	0.1061
N	38402	38402	38402	38402

Dependent variables here are the indicators for whether a board director seat is taken by a woman. Observations are at the individual-firm-year level. In addition to the controls used in the benchmark results in Table III, individual controls include quadratics for a director's age and network size, both recorded from individual data from the BoardEx database. Robust standard errors in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE B.O.14
ADDITIONAL ROBUSTNESS CHECKS FOR FIRST-STAGE

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Tier 1 $\times$ share of women in parl.	0.046*** (0.017)	0.046*** (0.017)	0.045*** (0.016)	0.051*** (0.017)	0.062*** (0.018)	0.046*** (0.016)	0.046*** (0.016)	0.047*** (0.017)	0.088*** (0.028)	0.069*** (0.019)	0.049** (0.021)
Tier 2 $\times$ share of women in parl.	-0.002 (0.015)	-0.002 (0.019)	-0.002 (0.019)	0.000 (0.015)	0.003 (0.017)	-0.002 (0.013)	-0.002 (0.013)	-0.001 (0.019)	-0.028 (0.023)	0.014 (0.031)	0.000 (0.016)
Year fixed effects	X	X	X	X	X	X	X	X	X	X	X
Firm fixed effects	X	X	X	X	X	X	X	X	X	X	X
Finance controls (lagged and first differenced)	X	X	X	X	X	X	X	X	X	X	X
R^2	0.177	0.177	0.171	0.257	0.358	0.177	0.177	0.196	0.362	0.319	0.214
N	2714	2714	2640	2714	2714	2714	2714	2440	2651	2714	2685

This Table performs additional robustness tests for the results in Table III. Column (1) is the benchmark result from column (4) of Table III. Column (2) codes the questionable GLCs as non-GLC firms. Column (3) drops the questionable GLCs entirely. Column (4) and (5) include linear trends for a firm's sector and industry. Column (6) cluster standard errors at the industry level. Column (7) allows for two-way clustering by both firm and industry. Columns (8)–(11) weights observations by board size, market capitalisation, asset size, and no. of appearance in top 100 in terms of market capitalisation. Otherwise, definition of GLCs and controls are as in Table III, and robust standard errors clustered at the firm level are in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE B.O.15
FALSIFICATION TEST—GOVERNMENT LINKAGE

	All women directors		Women supervisory directors		Women executive directors	
	(1)	(2)	(3)	(4)	(5)	(6)
% shares owned by TH	0.001 (0.004)	−0.006 (0.005)	0.002 (0.003)	−0.008 (0.005)	−0.001 (0.001)	0.002* (0.001)
Year fixed effects	X	X	X	X	X	X
Firm fixed effects	X	X	X	X	X	X
Finance controls (lagged and first differenced)	X	X	X	X	X	X
Heckman correction		X		X		X
Inverse mills ratio		.83*		.74*		.09
Mean of Dep. Var.	0.443	0.443	0.325	0.325	0.119	0.119
R ²	0.172		0.124		0.093	
N	2074	2074	2074	2074	2074	2074

This Table documents the falsification test, showing that firms' linkage to the government alone does not predict the number of women on boards. The main measure of government linkage in this Table is measured by the percentage of Tier 1 GLC shares owned by Temasek Holdings, as reported in their annual reports beginning in 2004. The % shares ownership however is only observed for Tier 1 GLCs in the sample, so shares for the non-GLCs are imputed as zero on the assumption that % shares, even if non-zero, is trivial. Tier 2 GLCs have % shares ownership that is by definition positive and non-trivial, so these observations are dropped. Odd-numbered columns are results from fixed effects panel regression, retaining the firm fixed effects for unobserved time-invariant characteristics. Even-numbered columns are results using the Heckman model to correct for the selection bias of observing Tier 1 GLC shares owned by Temasek Holdings (TH). Robust standard errors clustered at the firm level are in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE B.O.16
THE (LAGGED) EFFECT OF WOMEN PARLIAMENT REPRESENTATION ON WOMAN BOARD REPRESENTATION

	Dep. var. is (effective) count of women directors				
	Both supervisory & executive			Supervisory	Executive
	(1)	(2)	(3)	(4)	(5)
GLC × share of women in parl.	0.022** (0.009)	0.021** (0.009)	0.023** (0.011)	0.020** (0.010)	0.003 (0.004)
Tier 1 × share of women in parl.	0.037** (0.016)	0.038** (0.016)	0.040** (0.016)	0.039*** (0.012)	0.001 (0.006)
Tier 2 × share of women in parl.	0.015 (0.010)	0.013 (0.009)	0.007 (0.013)	0.003 (0.013)	0.004 (0.006)
Year fixed effects	X	X	X	X	X
Firm fixed effects		X	X	X	X
Finance controls (lagged and first differenced)			X	X	X
Mean of Dep. Var.	0.374	0.374	0.435	0.324	0.112
N	3247	3247	2714	2714	2714

Notes—Replication of main first-stage results in Table III, but where the share of women in parliament is lagged by a year. Robust standard errors clustered at the firm level are in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE B.O.17
SECOND STAGE RESULTS—AT LEAST 1, 2, AND 3 WOMEN DIRECTORS

	Dependent variable is								
	ROE			Log of q			Firm beta		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
At Least 1 Woman	0.002 (11.885)			−0.110 (0.202)			5.299 (3.573)		
At Least 2 Women		1.068 (11.415)			−0.095 (0.198)			5.458 (3.389)	
At Least 3 Women			54.312 (49.924)			0.130 (0.524)			24.820 (17.894)
Year fixed effects	X	X	X	X	X	X	X	X	X
Firm fixed effects	X	X	X	X	X	X	X	X	X
Kleibergen-Paap F	5.112	4.477	1.267	5.463	4.880	1.416	5.195	4.425	1.008
Anderson-Rubin Wald	2.65	2.65	2.65	3.01	3.01	3.01	7.6**	7.6**	7.6**
Hansen (ρ -value)	0.173	0.173	0.189	0.162	0.155	0.118	0.100	0.107	0.375
N	2903	2903	2903	2858	2858	2858	2578	2578	2578

This Table uses dummy indicators for at least 1, 2, and 3 women directors as the main estimand of interest in the structural equation (2.1). Robust standard errors clustered at the firm level are in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE B.O.18
SECOND STAGE RESULTS—SECTOR-BY-YEAR TRENDS

Dep. var. is	Excluded instrument in first stage is Tier 1 $\times$ share of women in parl.% Men on boards $\times$ share of women in parl.					
	ROE (1)	Log of q (2)	Firm beta (3)	ROE (4)	Log of q (5)	Firm beta (6)
No. of women directors	-19.200 (23.324)	-0.146 (0.120)	10.771* (5.996)			
Increase in no. of women directors				32.803 (24.300)	-0.727 (0.685)	14.039 (12.732)
Year fixed effects	X	X	X	X	X	X
Firm fixed effects	X	X	X	X	X	X
Sector-by-year fixed effects	X	X	X	X	X	X
First-stage F	7.65	7.36	5	9.42	8.4	7.24
Anderson-Rubin Wald	.61	1.65	6.3**	1.82	.95	1.35
N	2903	2858	2578	2472	2451	2322

This Table replicates the 2SLS results from Tables IV, V, VI, and VII, with sector-by-year fixed effects. Sectors of firms are as classified under the GICS (Global Industry Classification Standard), obtained from Bloomberg. Robust standard errors clustered at the firm level are in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

Chapter 3

WOMEN'S WAVE OR THE BLUE WAVE? RESULTS FROM THE 2018 U.S. HOUSE ELECTIONS

3.1 Introduction

The 2018 U.S. midterm elections—which took place during the peak of the MeToo movement—saw women candidates achieve historic gains. These elections took place halfway through the first term of Republican President Donald Trump when Republicans held a majority in both the House of Representatives and the Senate. In the House, all 435 seats were up for election and the Republicans lost 40 seats—the most since the 1974 midterm elections.¹ Overall, the 2018 midterms have the highest number of women candidates voted into Congress. The House, in particular, had 235 women candidates, with 102 of them winning, and most (89) running under the Democratic banner (Center for American Women and Politics 2018). In this paper, I test the assertion that the 2018 election was a "MeToo election".

To test whether the house Republican candidates incurred a backlash from the MeToo movement, I first download all tweets containing the MeToo hashtag in the year 2018, leading right up to the general elections

¹The Democrats gained a net of 49 seats in the 1974 post-Watergate House elections.

on Nov 6. I then match the tweets to U.S. counties using the Twitter user geolocation. This county-level variation in tweets is the measure of the MeToo movement. As validation, I use microdata from a new survey to confirm that the tweets are highly correlated with the pro-women and anti-Republican sentiment of their county residents.

The empirical strategy is a difference-in-differences approach, comparing the electoral performance of U.S. house candidates across counties and across parties (and gender). This approach mitigates concerns that candidates select into districts, such as Democratic or women candidates competing only in or campaigning harder in districts with high support of the MeToo movement. The findings are nuanced. It turns out that the expected advantage of Democratic women candidates and the disadvantage for Republican men occurs only in Republican strongholds, where there is a high Republican vote share in the 2016 presidential elections. Given a standard deviation increase in the republican presidential vote share, a standard deviation increase in the tweet density measure is associated with a 0.96 percentage point advantage for Democratic women candidates, while the Republican men incur a 0.45 percentage point disadvantage. This effect, however, can also be found in the prior 2016 House elections, suggesting that the findings reflect a temporal trend in voting.

In light of this, I turn to tests for potential political economy mechanisms instead that are related to the elections and the grassroots movement. Like the main results, I find subtle effects in candidacy selection. The tweet density measure increases the probability of having a Democratic woman challenger in seats with a Republican incumbent, and also in seats held by a Democratic man. I also test and find that the tweets measure captures an increase in turnout as well as a decrease in the Republican vote share at the county level. This suggests that rather than changing vote patterns, a major channel is turnout. Another puzzle is in why the MeToo effect is detected only in Republican strongholds. I test and find evidence that a substantial part of the results is driven by counties from districts that were marginally won by the Republicans in the previous 2016 House elections. However, I do not find evidence that the results are driven by straight-party voting, where the results are

driven by counties in states with split delegation senate elections and the House candidates benefited from a straight ticket.

In the discussion, I argue and show evidence that the MeToo tweet measure is unlikely to be capturing inauthentic political activity masquerading as grassroots. Moreover, the tweet measure correlates well with individual-level attitudes. One explanation for why the tweets capture an increase in turnout and a decrease in Republican vote share lies in changing demographics, where places that have high Republican support are also those places that are shedding voters over time. Another issue is that the county-level variation exploited implies a limit on generalising of the results, to the extent that the within-district county averages largely capture areas that are more Republican, more White, and more rural, suggesting that the MeToo movement was a complex one.

This study relates to the literature on the political economy of the mass media, specifically those that look at how varying access to media outlets, and the varying political coverage by the media can affect electoral outcomes (Adena et al. 2015; Boas and Hidalgo 2011; DellaVigna and Kaplan 2007; Enikolopov et al. 2011; Ferraz and Finan 2008; Gentzkow 2006; Larreguy et al. 2014; Lim et al. 2015; Miner 2015; Oberholzer-Gee and Waldfogel 2009). In these studies, the variation from the media comes mostly from changes that are already politically motivated, such as changes in radio broadcasting from the reign of the Weimar government to the Nazi party (Adena et al. 2015), and how the media coverage of malfeasant incumbents affected their vote share (Ferraz and Finan 2008; Larreguy et al. 2014).

A novel feature of this study is how it documents social media influence (as in Fujiwara et al. 2020), arising from what is essentially a grassroots movement, on the House elections. This, as opposed to the influence of more traditional media outlets such as print (Lim et al. 2015), radio (Adena et al. 2015; Boas and Hidalgo 2011; Ferraz and Finan 2008; Larreguy et al. 2014), and broadcast (DellaVigna and Kaplan 2007; Oberholzer-Gee and Waldfogel 2009). Further, the grassroots aspects of the MeToo movement mean that the measures are not directly influenced

by political candidates.² The study by Fujiwara et al. 2020 in particular, exploits early exposure to twitter in the nascency and provides one of the early pieces of evidence that social media can affect electoral outcomes—in favour of the Democratic party.

Another contribution is on how an independent media platform influences electorate behaviour against the incumbent (Enikolopov et al. 2011; Fujiwara et al. 2020; Miner 2015). Miner (2015) finds that the rise of internet access in Malaysia accounted for a large drop in points for the 40-year incumbent party during the 2008 elections. Enikolopov et al. (2011) find that differential access to the only independent national TV channel decreased the Russian government party's vote share during the 1999 parliamentary elections. Neither the internet nor the independent TV channel are centrally controlled nor have a formal political allegiance. The setting in this paper is similar, where the MeToo movement started as an independent grassroots movement. This paper also contributes to a growing literature on the effects of protest movements and rallies (Acemoglu et al. 2018; Campante et al. 2017; Feinberg et al. 2019; Larreboure and Gonzalez 2021; Lilley and Wheaton 2019).

Finally, the paper contributes to the literature on how access to media sources can influence voter turnout (Campante et al. 2017; DellaVigna and Kaplan 2007; Enikolopov et al. 2011; Gentzkow 2006; Oberholzer-Gee and Waldfogel 2009), and the literature on media bias (Gentzkow and Shapiro 2010; Groseclose and Milyo 2005; Larcinese et al. 2011), with the implicit assumption that a politically biased media sways voter decisions. In the context of this paper, the MeToo movement in social media became overtly pro-Democratic, and one interpretation is that the movement persuaded voters to turn out to vote for the women and Democratic candidates.³

²Unlike in Boas and Hidalgo (2011) for example, where incumbents have an advantage in gaining access to community radio before elections, which increases their vote share.

³Another contribution is that this paper bears some evidence on ethnic and gender-based voting (Abrajano and Alvarez 2005; Flanagan 2018; Holli and Wass 2010; Matsubayashi and Ueda 2011). In this paper, the interactions of county ethnic percentage makeup and candidate (predicted via name) ethnicity are systematically correlated to vote share, though the variation it explains in vote share is small. This paper also bears some indirect non-experimental evidence of expressive voting, a behaviour that does not originate from the belief that the vote is instrumental in the election

The rest of the paper is as follows. The next Section 3.2 describes the matching of twitter data to U.S. counties. Section 3.3 discusses legal and political implications of the MeToo movement. Section 3.4 discusses the empirical strategy and presents the results at the candidate level. Section 3.5 tests for political mechanisms. Section 3.6 discusses the findings. Section 3.7 concludes.

3.2 Data

Getting historical tweets data. I use a third-party custom-written Python library *GetOldTweets-python* to download all tweets containing the "MeToo" hashtag in 2018, leading up to the elections on November 6. The total number of tweets found in this period with the MeToo hashtag is 1'915'322.⁴

Geolocation of Twitter users. The tweets metadata include usernames, which I use to query the Official Twitter API for the users' geolocation. The 1'915'322 tweets come from 700'891 usernames. Of the 700'891 usernames, 158'857 cannot be found through the Twitter API at the time of query, and another 21'521 users can be found but the geolocation string is left empty. For the remaining 520'513, I obtain the Twitter users' geolocation tagged to the account. This disclosure of geolocation by users is tagged to their user account and is completely voluntary and without standardised formatting. To parse the geolocation strings I write a series of hard-coded rules to identify U.S. city-state, if applicable. This allows me to successfully parse 130'433 (25% of 520'513) geolocations into a standard city-state (e.g. Grand Rapids, Michigan). Finally,

outcome (Fischer 1996; Tyran 2004; Hillman 2010). Another related paper is Stephens-Davidowitz (2014) who uses racial animus proxied using Google search data and finds that Obama did comparably worse than his Democratic peers by vote share in places with higher racial animus.

⁴The *GetOldTweets-python* library is written by Jefferson Henrique and is hosted at <https://github.com/Jefferson-Henrique/GetOldTweets-python>. Two lines of code are changed to handle current changes in the underlying browser's HTML formatting so that the Twitter username can be retrieved. The Official Twitter API has a 7-day limit on past tweets at the time of this writing. Under the hood, the third-party API scrapes the Twitter Search which allows users to find historical tweets containing certain keywords. Search results appear in a scroll loader which loads more tweets through calls to a JSON provider as a user continues scrolling down, without a definite limit.

I match the city-states to their primary counties using the *United States Cities Database*,⁵ where primary counties are the centroids of the city as defined by the U.S. Geological Survey. Appendix C.0.1 provides more details.

2018 House election data. The primary source for the 2018 House of Representatives election returns comes from the individual states' Secretary of State. I hand-collect the returns of individual candidates at the county level using their Election Department's report. From this, I collect data on 40 states. I supplement data on 4 more states (Arkansas, Michigan, Nevada, and New Mexico) using the *MIT Election Data and Science Lab's* unofficial results.^{6 7} For the individual political candidates, I also record gender and incumbency. I infer candidate race (or ethnic) using their names (both first and last) through the *NamePrism* API (Ye et al. 2017). From this, each of the 1'022 candidates has an indicator for whether their (predicted) ethnic is White, Black, Hispanic, or Others.⁸

County-level covariates. County demographics come from the ACS (American Community Survey) 5-year estimates for 2012–16. The 14 variables include population and voting population sizes, demographic composition by ethnicity, gender, age, foreign-born, and education, income and unemployment data, and the rural-urban distribution. County-level density measure of high-speed internet connection—computed as the ratio of the number of residential units with at least 200 kbps in at least one direction to the total number of households—for June 2017 come from the FCC (Federal Communications Commission)¹⁰

⁵<https://simplemaps.com/data/us-cities>

⁶<https://github.com/MEDSL/2018-elections-unofficial>.

⁷At the time of collection, Alaska's Secretary of State (SOS) page on election results cannot be found, Connecticut, Kansas, Mississippi, and Missouri SOS page lacks voting results at the county or precinct level, and Minnesota has yet to publish their election results.

⁸The NamePrism is a supervised classifier developed using 74 million names from an email company. Their Naive Bayes classifier infers nationality/ethnicity using both first and last names (to mitigate migration and marriage), with the likelihood estimated using the homophily principle in communication pattern—people of the same type communicate more frequently and recently.

¹⁰<https://www.fcc.gov/general/form-477-county-data-internet-access-services>.

FIGURE I
INTENSITY OF TWEETS WITH MeToo HASHTAG IN 2018⁹

For past elections, both the county-level 2016 House elections and Presidential elections data come from the *MIT Election Data and Science Lab*.

Summary. Figure I shows the intensity of the MeToo tweets throughout the year 2018 right up to the election on the 6th of November. I make two observations here. The first is that the intensity of the tweets is

¹⁰January 7 Golden Globe awards; January 20 a million people took part in the second annual Women's March on the anniversary of President Donald Trump's oath of office, voicing disapproval of his administration and encouraging people to vote; January 28 Actor Jeremy Piven accused of sexual assault by three more women; February 25 Monica Lewinsky writes an essay about her experience with Bill Clinton; March 4 Oscars; April 16 *The New York Times* and *The New Yorker* won the Pulitzer Prize gold medal for public service for their work on the Harvey Weinstein scandal and sexual assault in general; April 26 Bill Cosby finally found guilty of sexual assault; May 10 Spotify no longer plays R. Kelly; May 25 Harvey Weinstein is taken into police custody; June 5 17 states have their primary elections; July 6 Canada PM Justin Trudeau denies need to conduct investigation of sexual misconduct against him; July 27 a *New Yorker* article reports that CBS will investigate allegations of sexual misconduct by its CEO Leslie Moonves; August 20 accuser Asia Argento herself accused of sexual misconduct; and September 16 a *Washington Post* article revealed Christine Blasey Ford was a victim of sexual assault by then Supreme Court nominee Brett Kavanaugh. See for example <https://www.chicagotribune.com/lifestyles/ct-me-too-timeline-20171208-htmlstory.html> for a curation of MeToo events.

FIGURE II
GEOGRAPHICAL DISTRIBUTION OF MeToo TWEETS IN 2018

relatively consistent throughout, without a single salient spike. In fact, a few spikes occur which can be traced to a number of identifiable events such as the Cosby hearing and the (second) Kavanaugh confirmation hearing. The second observation is that even though the geolocations of Twitter users cannot all be parsed into identifiable U.S. counties—some because they are unambiguously outside the U.S.—the plot shows that the time trend of the global tweets and the identifiable U.S. counties tweets are similar, suggesting that there is no systematic difference in the tweets that can and cannot be matched to U.S. counties. Figures II and III provide insight into the geographical variation in the MeToo tweets intensity (logs) and vote share of the political candidates. There is substantial geographical variation in the tweets and the vote share of candidates, by both state and county.

The final sample is for 44 U.S. states, with 388 House congressional districts, 2'652 counties, and 1'022 House election candidates, of which 767 are from the two main parties. This gives 8'653 candidate-county-level observations. Districts, where a single candidate wins by default, are not included in the sample.

FIGURE III
GEOGRAPHICAL DISTRIBUTION OF WOMEN AND DEMOCRATIC VOTE SHARE

3.3 Context and Implications of MeToo

The MeToo Movement. The phrase "Me Too" began more than a decade ago in 2006 on the myspace social network, when Tarana Burke used it in her local community to encourage Black and Hispanic girls, as well as other women to come forth with their accounts of sexual misconduct (Gibson et al. 2019). Social media became the place where these accounts can be made available to the mass public, and the MeToo movement picked up massive momentum in 2017 when celebrities lent their voices and experiences, notably on the microblogging and social networking service Twitter. Tweets of this nature use the MeToo hashtags.

Legal Implications of the Movement. The MeToo movement, with widespread attention in social media, is more than just window dressing. First, the

attention on sexual harassment issues has gained traction in Congress, with Democrats sponsoring the BE HEARD Act (Bringing an End to Harassment by Enhancing Accountability and Rejecting Discrimination Act) with bipartisan support to extend harassment protections to workers at small businesses and independent contractors (North 2019).¹¹ Second, courts tended to apply the *Faragher* defense—when employers can show they took reasonable measures to prevent or redress harassment—in favour of employers, and the MeToo movement may pressure courts to be more narrow on what they consider reasonable (Tippett 2018). Third, some states (including California, New York, and Pennsylvania at the time of writing) are considering or have already passed bills to limit the extent of non-disclosure agreements, including its use in cases of sexual misconduct (Tippett 2018).

Fourth, at least two judges—Judge Aaron Persky in California and Judge Michael Corey in Alaska—at the time of writing have been recalled as a reaction to their lenient sentencing of specific sexual assault cases in 2018, in spite of favourable judicial performance evaluation. The recall campaigns ride on the MeToo movement and the contemporaneous controversy surrounding the confirmation hearings for Supreme Court nominee Brett Kavanaugh, which was itself tangled with the movement. Before this, the most recent recall of a state judge dates back to 1977 (Singer 2019).¹² ¹³

Electoral Implications of the Movement? The media in particular, has framed the 2018 midterm elections as a "#MeToo election", asserting

¹¹Under Title VII of the Civil Rights Act, there is no explicit reference to harassment, and courts generally treat issues of sexual harassment as a form of discrimination (Tippett 2018).

¹²In states with the retention election system, nonpartisan commissions nominate qualified judicial candidates to the governor, who then appoint a nominee to an open seat. Appointed judges then face periodic retention elections without another challenger, the only decision voters have to make is whether to retain or recall the judge. Some states have judicial performance evaluations in place for these elections so that the electorate can make informed decisions (Singer 2019).

¹³In a similar turn of events, former Connecticut U.S. house representative and Democrat Elizabeth Esty was publicly pressured to resign, after it became known that she attempted to cover up sexual misconduct by her chief of staff. She retired and the vacated seat was later won by Democrat Jahana Hayes, the first Black woman to represent Connecticut in Congress. See for example <https://edition.cnn.com/2018/03/30/politics/elizabeth-esty-staffer-abuse/index.html>.

FIGURE IV
WOMEN IN CONGRESS

that women candidates will benefit from the movement.¹⁴ Figure IV shows the jump in both women candidates running and voted into Congress in the 2018 elections, affirming the fact that the 2018 elections are historic for the representation of women in Congress (Center for American Women and Politics 2018).

Another possibility is that Democratic candidates benefit since the MeToo movement became tied to partisan attitudes. The nomination and confirmation hearings of Kavanaugh was a particularly salient politically charged episode.¹⁵ In fact, from the timeline of the MeToo movement from Figure I, the peak as reflected on Twitter came right after the second hearing. The incumbent Republican President Trump himself, accused of sexual harassment, is a subject of the movement. There were women marches shortly after the 2016 Presidential election as an objection to Trump's election (note 10). Figure V suggests this (negative) correlation between the MeToo movement and the house Republican vote share.

¹⁴Deckman (2018) drew links months before the elections, and Peaker (2018) after.

¹⁵See for example the media piece by Walsh (2018).

FIGURE V
REPUBLICAN VOTE SHARE AND MeToo MOVEMENT

3.3.1 Determinants of Tweet Density

I first check if past and existing trends can determine the intensity of the MeToo tweets in 2018. The full model I estimate is:

$$\tau_c = \alpha + \beta_1 \nu_{c, 2016}^{\text{Rep., House}} + \beta_2 \nu_{c, 2012-16}^{\text{Rep., Pres.}} + \Gamma X_c + \varepsilon_c, \quad (3.1)$$

where τ_c is log county-level tweet density—the number of (identified) county-level MeToo tweets in 2018 (before the elections) divided by county population. $\nu_{c, 2016}^{\text{Rep., House}}$ are the 2016 house Republican vote share and turnout, $\nu_{c, 2012-16}^{\text{Rep., Pres.}}$ are the 2012–16 equivalent, and X_c are the county census variables. Standard errors are clustered at the congressional districts.

Column (1) of Table II includes only the full interaction of county high-speed internet connection density and percentage females as a control, which is positively correlated with the MeToo movement. Column (2) includes controls for the 2016 House election and presidential election outcome. The previous link between internet connection and percentage disappears, but the Republican vote shares on the other hand are statistically significant. I interpret this as an indication that the movement

is predominantly political rather than gender-based.

Column (3) adds the county census demographics, which are highly correlated with tweet density, as anticipated and indicated by the joint F -statistic. This is likely because urban areas and education are highly correlated with the MeToo movement. I show below, however, that accounting for these demographics does not change the main results.

Column (4) adds the congressional district fixed effects. With this, the estimates capture within-district determinants of the MeToo tweet density in the year 2018, leading up to the elections. Turnout in the 2016 presidential election is now positively associated with the tweets measure ($\rho < 0.05$). The Republican vote share in the 2016 House and presidential elections, however, is no longer significant, indicating that the county-level MeToo tweets are not correlated with the past election results within the House congressional districts themselves. Column (5) uses the two-party Republican vote share measures (votes received by Republican candidates divided by votes received by both Republican and Democratic candidates), and the results are similar.

3.3.2 Selection of Women Candidates

In Table C.0.2, I also check what covariates are linked to the presence of women candidates for the 388 U.S. congressional districts in the sample. Specifically, the model I estimate is:

$$I_{ds} = \alpha + \beta\tau_d + \Gamma\mathbf{X}_d + \Delta\mathbf{Z}_d + \text{state}_s + \varepsilon_{ds}, \quad (3.2)$$

where I is the dummy for the presence of women candidates at the districts; τ_d is the district-level log tweet density; $\mathbf{Z}_d$ are dummies for whether the seat is open, has a woman incumbent, or has a Republican incumbent; $\mathbf{X}_d$ are all other district-level controls including the aggregated county census controls and past electoral trends. All regressions include state fixed effects, with standard errors clustered at the 44 states in the sample.

Assuringly, the selection of women candidates is orthogonal to the occurrences of the MeToo tweets and past electoral trends in both the House

and the Presidential elections. Strong predictors (both economically and statistically) of women challenging incumbents come from the political seat characteristics. Women are more likely to challenge when the seat is open, and when the incumbent is Republican ($\rho < 0.01$).

3.4 Results

3.4.1 Empirical Strategy

To identify the effect of the MeToo movement on the 2018 House elections, the baseline empirical strategy I use is the difference-in-differences (DD) strategy, comparing the vote share of individual candidates across counties, which vary in their intensity and density of the MeToo tweets. Specifically, I regress the vote share of individual candidates at the district-county level,¹⁶ on the interaction of candidate party, gender, and the density of the MeToo tweets at the county levels:

$$\begin{aligned} \nu_{icd} = & \alpha + \beta^{RW} RW_i \tau_c + \beta^{DW} DW_i \tau_c + \beta^{RM} RM_i \tau_c + \beta^{DM} DM_i \tau_c \\ & + \text{Candidate}_i + \Delta_1 \nu_{c,2016}^{\text{Rep., House}} + \Delta_2 \nu_{c,2012-2016}^{\text{Rep., Pres.}} + \Gamma X_{ic} + \varepsilon_{icd}, \end{aligned} \quad (3.3)$$

where ν_{icd} is the vote share of 2018 house candidate i in district-county cd ; τ_c is the county-level log of tweet density (county MeToo tweets divided by population); where R (or D) indicates candidate from the Republican (Democratic) party, and W (or M) indicates a woman (man) candidate, so that RW for instance, indicates a Republican woman candidate. So if there is indeed an advantage for the Democratic women candidates in counties with high incidences of the MeToo tweets, then $\beta^{DW} > 0$.

The full specification includes the interaction of candidate party and gender with past electoral outcomes. $\nu_{c,2016}^{\text{Rep., House}}$ is the full interaction of the 2016 house Republican vote share and candidate party; and $\nu_{c,2012-2016}^{\text{Rep., Pres.}}$ is the full interaction of the 2016 presidential Republican candidate vote share and candidate party. This prevents the DD estimates from picking

¹⁶Some districts have boundaries that run across counties.

up existing political support for the parties. The full sample regressions also include the dummy interaction for all third-party candidates.¹⁷

The vector X_{ic} are the county census demographics that enter as full interactions with party status. This prevents the DD estimates from capturing how votes differ by the basic demographics (Edlund and Pande 2002; Herron and Sekhon 2005; Oswald and Powdthavee 2010). X_{ic} also includes the interaction of candidate ethnic (African American, Hispanic, Others, and White) with the percentage composition of the corresponding ethnic at the county level, and similarly with gender. This allows for voting heuristics, where voters cast their ballot based on the ethnicity or gender of the candidates (as in Abrajano and Alvarez 2005; Holli and Wass 2010; Stephens-Davidowitz 2014; Flanagan 2018).

The baseline specification (3.3) includes candidate fixed effects, which removes county-invariant candidate characteristics, including party, incumbency, and open-seat contests. The candidate fixed effects also prevent the DD estimates from picking up past private and public office credentials, seniority in committees, as well as campaigning and overall support in a district. The standard errors are clustered by candidates.¹⁸

¹⁹

In what scenarios would the DD estimates be biased? An important identifying assumption in the DD specification (3.3) is that candidate campaigning across counties of a district is uniform. And, if there are heterogeneities in campaigning across counties, then they must be orthogonal to candidate party (gender) or to the prevalence of the MeToo movement at the county level. That is, Democratic or women candidates are not just campaigning harder in geographical areas with a higher level of interest in the MeToo movement, as proxied by the MeToo tweets in 2018.

¹⁷Coefficients for third-party candidates not reported to conserve on space.

¹⁸In an appendix robustness check in Table C.0.4, using non-nested two-way clustering of standard errors for the house candidates and county does not change the results.

¹⁹A well-known pattern during the midterms is a depression of the vote share for candidates who are from the same party as the sitting president. The party status dummies for candidates, nested within the candidate fixed effects, account for any of such potential swing so that the reported coefficients are interpreted as changes beyond the regular midterm swings.

The results will also be biased if the MeToo tweet density captures the intent to vote for women candidates, and that women candidates only run in districts with high occurrences of the tweets. I show in Table C.0.2 however, that the tweets are orthogonal to the presence of women candidates in districts. Moreover, the DD specification identifies within rather than cross-district variations. The remaining assumption is that women (Democratic) candidates are not selecting into districts with high variation of the MeToo movement, while the men (Republican) candidates are simply selecting into districts with low variation, but where the aggregated district measure of tweet density for both the women and men (Democratic and Republican) are statistically identical. I find this selection behaviour unlikely.²⁰

Another form of bias comes from a few layers of measurement errors. First, tweets containing a MeToo hashtag in 2018 are only a proxy for how engaged county citizens are in the MeToo movement. Further, the engagement can go in either direction—pro-feminist or anti-feminist—though I show below that the tweets do proxy for the expected pro-feminist direction. Second, the MeToo tweet measure is itself measured with error, since only a subset of the global tweets (Twitter users) can be successfully matched to US counties, and some days have missing records (Figure I). Finally, the Twitter user geolocation record might itself be inaccurate, since a user may no longer (or have never) reside in the reported area. All these work against the results, reducing the precision of the estimates.

3.4.2 Average Effect on Candidate Vote Share

Columns (1)–(2) of Table III reports the results from estimating equation (3.3). All reported coefficients are in absolute terms.²¹ In column (1), only the DD estimates for candidate party gender and candidate fixed

²⁰In an appendix robustness check in Table C.0.4, I show that excluding districts where the within district variation in the MeToo tweets is less than the 90th percentile does not change the results.

²¹So that the coefficients can be interpreted without requiring back-of-envelope differencing/addition The full report of the three-way interaction between candidate party-gender, log tweet density, and the 2016 presidential Republican 2016 vote share is in Table C.0.3.

effects are included ($\Delta_1 = \Delta_2 = \Gamma = 0$ in equation (3.3)), and the estimates are as anticipated, suggesting that the movement had an effect by both candidate party and gender lines. Democratic candidates have an advantage in counties with high MeToo tweet density, while both Republican candidates face a disadvantage. The DD estimate for Democratic women, in particular, implies they get a 2.4 percentage point advantage with a standard deviation increase in the tweet density measure (1.164), while the Republican men get a 2.6 percentage point disadvantage.

Column (2) enters past electoral controls and county demographics as full interaction with candidate party, together with a set of controls for ethnic and gender-based voting. In this most demanding specification, the average effect of the MeToo tweets on candidate vote share by party and gender disappears. In addition, the F-tests for ethnic and gender voting are highly significant, suggesting that there is a strong statistical tendency for voters to vote along their own gender and party line, even if the overall variation it explains in candidate vote share is small.²²

The average effects in columns (1)–(2) of Table III however, might hide heterogeneous effects. Since the MeToo movement is linked to partisan attitudes, and in particular that the MeToo movement is highly linked to the disapproval of the Republican party (in its handling of sexual harassment issues), I test below whether there is a backlash of Republican candidates in Republican strongholds.

3.4.3 Heterogeneous Effect/Backlash, by Existing Republican Support

The conditional plot in Figure VI provides visual indication of a heterogeneous effect of the MeToo movement, where the three binned scatters are for Republican vote share between 0%–50%, 40%–60%, and 50%–100%. The anticipated disadvantage of Republican candidates in the 2018 House elections comes only in counties with a high Republican vote share ($> 50\%$) in the 2016 Presidential election.

²²The result does not change with the main-party sample.

FIGURE VI
CONDITIONAL EFFECT OF MeToo MOVEMENT²³

To test this formally, I add the 2016 presidential Republican vote share to the above interaction of party-gender and log tweet density. The full specification, with absolute effects for ease of interpretation, is:

$$\begin{aligned}
 \nu_{icd} = & \alpha + \beta^{RW}(RW_i)\tau_c + \gamma^{RW}(RW_i)\tau_c\nu_{c,2016}^{\text{Rep., Pres.}} + \beta^{DW}(DW_i)\tau_c + \gamma^{DW}(DW_i)\tau_c\nu_{c,2016}^{\text{Rep., Pres.}} \\
 & + \beta^{RM}(RM_i)\tau_c + \gamma^{RM}(RM_i)\tau_c\nu_{c,2016}^{\text{Rep., Pres.}} + \beta^{DM}(DM_i)\tau_c + \gamma^{DM}(DM_i)\tau_c\nu_{c,2016}^{\text{Rep., Pres.}} \\
 & + \text{Candidate}_i + \Delta_1\nu_{c,2016}^{\text{Rep., House}} + \Delta_2\nu_{c,2016}^{\text{Rep., Pres.}} + \Gamma X_{ic} + \varepsilon_{icd},
 \end{aligned} \tag{3.4}$$

where the main coefficients of interest are the γ^j 's. For example, $\gamma^{DW} > 0$ implies a positive effect of the MeToo movement on Democratic women candidates in counties with high existing Republican support. Similarly, $\gamma^{RM} < 0$ implies that Republican men candidates do worse in these same counties. The Republican presidential vote share in specification (3.4) is centered at 50%, so that the DD estimates β 's can be easily interpreted as the effect of the MeToo tweets on candidate vote share when the 2016 presidential Republican vote share is split right down the middle.

Columns (3)–(5) of Table III confirms the above hypotheses. In column (3), the anticipated advantage for Democratic women and Republican men candidates are present in counties with high existing Republican support. The estimated coefficients of β^{DW} and γ^{DW} suggest that for

²³Binned scatter plots are with past electoral trends already partialled out.

counties with high MeToo tweets, Democratic women face a disadvantage when there is a 50-50 split, and this effect reverses in counties with high Republican support. The estimate of γ^{RM} is negative, implying that the Republican men face a disadvantage in the same counties where the Democratic women get an advantage. Column (4) uses only the two-party vote shares on both sides of the equation, and the results are similar.

The estimate of γ^{DW} from column (4) implies that in counties with a standard deviation increase in the presidential Republican vote share above the 50-50 split (67.5% Republican vote share), a standard deviation increase in the county log MeToo tweet density (1.164) gives Democratic women candidates a 0.96 percentage point vote share advantage ($\rho < 0.01$) relative to their peers in counties with a 50–50 split in the Republican vote share.²⁴ Republican men on the other hand, incur a 0.45 percentage point disadvantage ($\rho < 0.05$).²⁵ The estimate of γ^{DM} is only marginally significant, suggesting that the advantage of Democrats came mainly through their women candidates. To further place the estimates in perspective, counties with a 50-50 split or with lower than 50% existing support for the Republican party are a minority in the two-party sample (approximately 22%). The estimates β^{DW} and γ^{DW} suggest that an absolute advantage begins after the 69% Republican vote share mark (approximately 46% of the two-party observations).²⁶

Column (5) uses district fixed effects instead of candidate fixed effects, and the results barely change, indicating that both observed and unobserved characteristics of the candidates, including experience, grassroots campaigning and support, and funding, are unlikely to be driving the results.

²⁴Or, $0.047 \times (67.5 - 50) \times 1.164$.

²⁵Or, $-0.023 \times (67.5 - 50) \times 1.164$

²⁶I do not fully understand this pattern. One explanation might be a systematic difference between the kind of MeToo tweets which appears in Democratic vs. Republican counties, where those appearing in Democratic counties are tweets opposing the MeToo movement, while the MeToo tweets in Republican counties are genuinely representing the spirit of the MeToo movement and having an anti-Republican sentiment. Another explanation might be that the MeToo tweets capture mostly capture the expected pro-women and anti-Republican sentiment in a county, but the tweets mobilised Republican voters in retaliation in the Democratic counties with a woman candidate.

3.4.4 Robustness

Table IV examines robustness of the results in column (4) of Table III. The estimated joint effect of the MeToo movement and the Republican vote share for the Democratic women and Republican men is consistent throughout, while the effect for the Democratic men is not. In column (1), the tweets measure is computed using only tweets that occur before June, the earliest month in which a substantial number of states (17) have their primary elections. This mitigates the astroturfing concern (which I discuss below in Section 3.6.1).

Column (2) excludes districts with only one or two counties. This ensures that the DD specification picks up the intended within-district effect of the tweets, and is not driven by districts with a small number of counties. In columns (3) and (4), counties with low turnout ($< 2,000$) and counties with low voting-aged population ($< 2,000$) are dropped to mitigate concerns that the MeToo effect is present only in small areas. In column (5), districts with low geographical variation in the tweets measure (low standard deviation across counties of a district) are dropped to mitigate another astroturfing concern (Section 3.6.1). In column (6), counties on extreme tails of high-speed internet connectivity are excluded so that the results are representative of the average geographical area by internet use.²⁷

Finally, column (7) excludes districts where incumbents have vacated their seats. In the 2018 House elections, 36 Republicans and 18 Democrats did not seek re-election. In the two-party sample, open seats account for approximately 15% of the observations. Excluding these district observations with open seats does not change the results.²⁸

²⁷Outliers in high-speed broadband connection might include those municipalities where internet access is either partially or fully provided by the local governments, and these areas are arguably more left-leaning with the public provision of what is otherwise a private good. Omitting these places suggests that the results are not simply driven by these pro-Democratic areas.

²⁸In Table C.0.4 of the appendix, I also check that the results hold with non-nested two-way clustering of the house candidates and counties; with a general specification of the log tweet density measure where the coefficients of log tweets and log population are allowed to differ; with the incontiguous Hawaii state observations dropped from the sample, and with the tweets measure computed using only the tweets without other hashtags present to prevent it from picking up other grassroots sentiments and tweets with overt political angles (e.g. "#bluewave").

3.4.5 Back to the 2016 House Elections

The above suggests a robust correlation between the 2018 House election returns and the MeToo measure. This cross-sectional result does not negate a temporal trend (as in Lilley and Wheaton 2019). To test if the results from Table III are capturing existing trends, I use the 2016 House election. Table C.0.7 reports the result and uses the same regression specification aside from the temporal switch in the dependent variable. If the MeToo movement provides an advantage to Democratic candidates, its effect should not be present in the 2016 House elections. The estimates suggest otherwise, being very similar to the baseline. Areas, where Democratic candidates in the 2018 House elections appear to benefit from the grassroots MeToo movement, are also areas where Democratic candidates have an advantage in the 2016 House elections before the MeToo movement reached its peak. This effect may be attributable to underlying shifts in demographics (more on this in Section 3.6.2). Given the failure of the 2016 placebo test, in the following section, I turn to tests of potential channels through which political agents leverage the grassroots movement.

3.5 Exploring Channels and Interpretations

3.5.1 Candidacy as Strategic Reaction

Here, I focus on whether there was a strategic reaction to the MeToo movement through candidacy. One potential channel through which the movement works is by increasing the likelihood that non-traditional political candidates who fit the MeToo zeitgeist stand for elections. Section 3.3.2 tests and shows that the measure of log tweet density does not predict the selection of women candidates at the district level, condition on the past election returns, demographics, and the state fixed effects. A more authentic test, however, given the subtle results above, is to examine whether the MeToo tweets affected candidacy only in particular classes of political seats.

FIGURE VII
Candidacy as Strategic Reaction

Specifically, I augment the tests from equation (3.2) and Table C.0.2 and test whether the probability of having a Democratic woman challenger in: (i) districts with Republican incumbents and (ii) districts with Democratic man incumbents, is moderated by the MeToo measure. Figure VII reports the marginal effects over the range of the MeToo measure, controlling for electoral trends, demographics, and the state fixed effects. The black solid lines inside the 95% confidence intervals (constructed from standard errors clustered by states) report the marginal effects at different levels of the MeToo tweets measure, and the stacked histograms indicate the distribution of house seats for the particular class of interest. The confidence intervals are, as expected, increasingly wide at levels of the log tweet density with fewer observations.

As shown in the left panel of Figure VII, and consistent with the story that candidacy was a strategic reaction to the MeToo movement, the marginal effect, while non-linear, increases with the log tweet density (left panel of Figure VII). This suggests that the probability of a district having a Democratic woman challenger to a Republican incumbent increases in districts with higher MeToo movement. While the single coefficient estimate is not significant (as reported in the figure), the kernel estimates suggest that the effect is in fact positive and significant at moderately high levels and high density of the MeToo tweet measure. From the left panel of Figure VII, what may perhaps be more surprising is that a similar effect exists, but for districts where the incumbent is a Democratic man. When a district has an incumbent from the Democratic Party, its effect on the probability of having a women candidate from

the same party is negative, as expected. This negative effect, however, starts reversing in districts with high MeToo movement. The plot of the estimates non-trivially implies that at moderately high levels and density of the MeToo tweet measure, having a Democratic woman candidate in a district is conditionally independent of whether the district has an incumbent from the same party who is a man. Moreover, going further to the right of the MeToo measure distribution means that having a district with a Democratic man incumbent increases the probability that a woman candidate from the same party stands for election during the 2018 house elections, although the estimate is imprecise because of the lack of observations.

These results during the general elections necessarily imply similar machinery in the primaries, suggesting that candidacy as a strategic reaction to the MeToo movement started earlier during the primary elections. Figure I is consistent with this interpretation, which shows a local peak in the MeToo tweet intensity on June 5 exactly, when 17 states have their primary elections.

Overall, the results above suggest that a key political economy mechanism is through candidate selection, where opportunistic political actors catch the winds of change in the MeToo movement. Woman candidates from the Democratic Party are more likely to stand in districts with a Republican incumbent, and also in districts where the seat is held by a Democratic man. This effect is not present for Republican woman challengers when the parties are reversed (Figure C.0.3). This suggests that the movement shifted incentives in candidacy, where women who are traditionally more outsider and less centrist in the established Democratic party machinery are more likely to be frontrunners.

3.5.2 Turnout as a Channel

The 2018 midterm elections set a record high in turnout.²⁹ A natural question is whether the MeToo movement had a part to play in turning

²⁹[washingtonpost.com/news/monkey-cage/wp/2018/11/20/](https://www.washingtonpost.com/news/monkey-cage/wp/2018/11/20/)

[americans-just-set-a-turnout-record-for-the-midterms-voting-at-the-highest-rate-since-1914-tl](#)

out voters. To test this, I estimate the model:

$$t_c^{\text{House 2018}} - t_c^{\text{House 2016}} = \alpha + \beta_1 \tau_c + \beta_2 \nu_{c, 2016}^{\text{Rep., Pres.}} + \gamma(\tau \cdot \nu^{\text{Rep., Pres.}})_c + \Gamma \mathbf{X}_c + \varepsilon_c, \quad (3.5)$$

where t_c is log total votes cast in the House elections in county c , so that the dependent variable is the log change in total votes cast from 2016 to 2018, which is interpreted as a percentage change. The demographics controls now include both levels and trends (using the ACS 5-year estimates from 2012–16 and 2015–19), including the percentage of citizen voting-age population.

Table V presents the results, which is consistent with the heterogeneous effect in Table III. First, columns (1)–(2) show that the log tweet density measure does not predict change in turnout. In columns (3)–(4), I replace the tweets measure with the log tweet *intensity* measure (without dividing by county population), and the interaction term (γ in equation (3.5)) is now positive. The estimate in column (4) implies that for a standard deviation increase in log tweets intensity (1.946), every 10 percentage point increase in the Republican vote share increases turnout by 1.17% ($\rho < 0.05$).

As a falsification test, I repeat the regressions in Table V, but with the increase in 2012–16 presidential turnouts as the dependent variable, and the results confirm that no such trend exists before 2016 (Table C.0.5). Figure VIII show that the MeToo movement on Twitter begins in full force only from 16 Oct 2017.³⁰

The results from Table V suggest that the intensity of the MeToo movement is what matters for turnout. The finding on turnout connects with a set of existing literature. For example, DellaVigna and Kaplan (2007) find that the Republican-leaning Fox News increased turnout (and also the Republican vote share) in the 2000 presidential elections. Campante et al. (2017) in particular, provides some insight into the MeToo movement as a grassroots protest. In the context of Italy, they find that the internet facilitated local online grassroots protest movements and that the new Italian political party in 2009 (M5S) grew out of those

³⁰The equivalent is to test 2014–16 House turnout, but county-level House returns are available only from 2016 onwards.

online protest groups and is overrepresented by supporters who did not vote in the previous elections.

3.5.3 County-level Vote Changes

If turnout is a channel, then the anti-incumbency effect of the MeToo movement should also be observed through changes in the district-county level vote shares. To test this, I regress the district-county level change in the Republican House vote share from 2016 to 2018.

Table VI documents the results, which suggest that in places with a Republican stronghold, there is a fall in the house Republican vote share from 2016 to 2018. In columns (1) and (3), the estimates imply that for a standard deviation increase in the log tweet density measure (1.17) and the Republican vote share (17.7), the all-party Republican vote share drops by 0.59 percentage points ($\rho < 0.01$), and for the two-party vote share, it is a 0.28 percentage point drop ($\rho < 0.05$). The drop in the Republican two-party vote share is about half the magnitude of the all-party vote share, consistent with a shift of votes mostly from the independent (rather than Republican) to Democratic.³¹

Columns (2) and (4) include the change in log turnout between 2016 and 2018 and the estimates, while only marginally significant, have a negative sign which is consistent turnout as a channel of the MeToo effect. The estimate from column (4) implies that a standard deviation percentage increase in turnout (0.44) decreases the Republican two-party vote share by 1.4 percentage points ($\rho < 0.1$).

As a falsification test, column (5) checks that the estimates are not capturing existing downward trends in Republican support by geography—that counties with a high Republican vote share in 2016 are not those with a drop in the 2012–16 Republican presidential vote share. Column (6) checks that the estimates are not capturing existing downward trends by anti-Republican sentiment—that counties with a high Republican

³¹In Table C.0.8 of the appendix, I repeat the regressions using log tweets as an intensity measure instead of the log tweets density measure (normalised by county population), and the results are more significant overall, both in terms of economic and statistical significance.

FIGURE VIII
EXTENDED TIMELINE OF *MeToo* TWEETS³³

vote share in 2012 are those with a drop in the 2012–16 presidential Republican vote share.³²

3.5.4 Marginal Districts

A straightforward explanation for the heterogeneous effect from Section 3.4.3, is that many of the strongly Republican counties are in districts that are marginal. To test this, I repeat the analysis from Table III, but only for counties of Republican districts with a low margin of victory in the 2016 house elections. If the results are driven by Republican counties situated in districts that are marginal, then the estimates should be substantially different.

Column (2) of Table VII reports the results. Only the main-party candidates are included and district observations without a main-party challenger are dropped for simplicity. The estimates are indeed larger in

³²The falsification tests have the same conclusion with the all-party presidential Republican vote share.

³³The MeToo movement blew up on Twitter on October 16, 2017 when Alyssa Milano started using the MeToo hashtag to encourage people to share their stories. This day is the peak of the movement so far, as indicated by the peak in the figure. Figure C.0.2 plots the timeline in level terms.

magnitude, by two to three times compared to the baseline full sample in column (1), suggesting that the increases in voter turnout are orchestrated by political agents for a district-level advantage. The F -test using the county-level census demographics tests for comparability, with the value of 3.04 suggesting that the included subsample is systematically different compared with the excluded sample, but not by a large magnitude.

Alternatively, the results may be driven by straight-ticket voting, where voters primarily turn out to vote in the senate elections and the house election candidates benefit from the straight-tickets.³⁴ Column (3) tests this for states with a senate election and where there is a split delegation—with one Republican and one Democratic senator. If straight-ticket voting is driving the results, then the estimates should be larger in the counties of those states. The estimates, however, are not substantially different from the baseline in column (1).

A third alternative is that political parties turn out voters to build up a support base for the next election cycle in 2020 which includes the presidential elections. As is well known, past turnout strongly predicts future turnout, and political parties thus have incentives to turn out partisans even if it is immaterial in the current elections. If this is the case, one should observe substantially larger effects in Republican states where there was a low margin of victory in the 2016 presidential elections. The results in column (4) do not provide support for this.³⁵

³⁴Straight-ticket or straight-party voting here refers to the behaviour of voting for candidates that are from the same party in a ballot, as a strategic choice, and not to the logistics where voters are given an option to check a box on the ballot for a straight-party vote.

³⁵Table C.0.9 reports the results here as interactions in the full sample and the conclusions drawn are qualitatively similar, with the only detected effect in the triple interactions coming from the subsample of Republican districts with a low margin of victory in the 2016 house elections.

3.6 Discussion

3.6.1 What do the Tweets Measure?

3.6.1.1 Astroturfing?

A basic sanity check concerns an implicit assumption in this paper. Are the tweets a proxy of the grassroots movement or are they from astroturfing? Astroturfing generally refers to the practice of coordinated but inauthentic political activity masquerading as true grassroots activity. Notably, astroturfing is easy/cheap to directly implement on online platforms, including social media. If the MeToo tweet measure originates from astroturfing, it will change the fundamental interpretation and theme of how grassroots movement can affect society (Acemoglu et al. 2018; Campante et al. 2017; Feinberg et al. 2019; Larreboure and Gonzalez 2021; Lilley and Wheaton 2019). While definitive evidence is difficult, arguments can be made against it.

First, the correlation between the MeToo tweet measures and county demographics provides evidence against astroturfing. The tweets measure is highly correlated with percentage females, Hispanic, foreign-born, aged 29 and under, and college education or higher, in the expected positive direction, while being negatively correlated with the percentage of residents living in rural areas (Figure C.0.6). This does not square with a broad-based generation of fake grassroots.

Second, as indicated in Figure I, both the global MeToo tweets and those that are successfully matched to U.S. counties have spikes in intensity that coincide with high-profile MeToo events. This supports the assumption that the tweets are capturing grassroots sentiments (Figure I and note 10).

A check against astroturfing is to cut off aggregation of the 2018 tweets measure before June when most primary elections occur.³⁶ Astroturfing might begin early in the year, but candidates are not yet finalised and funds, if any, diverted to astroturfing will likely yield higher benefits

³⁶17 states have their primaries in June. See <http://www.ncsl.org/research/elections-and-campaigns/2018-state-primary-election-dates.aspx>.

much closer to the elections in November. Another potential sign of astroturfing is when the MeToo tweets are highly uniform in a district. Table IV shows that the conclusion is unaffected when using only pre-primary (pre-June) tweets and when dropping districts with low within-district variation in the tweets measure.

3.6.1.2 Sexual harassment and disapproval of the Republican party (VOTER survey)

The above discussion suggests that the MeToo tweets measure is able to capture sentiments on the ground, and do not originate from inauthentic political activity. Another check that MeToo tweets measure pro-women and anti-Republican sentiment comes from the 2018 *VOTER* (Views of the Electorate Research) survey (Democracy Fund Voter Study Group 2018), which tracks about 8,000 individuals from 2012–18, though some individuals drop out of the study from 2016–18. I use the reported ZIP code and match them to (primary) counties using the FIPS from the U.S. Cities Database. Out of the 2,649 counties in the sample, 1,352 counties can be successfully matched to the VOTER microdata. 7,491 individuals are ultimately matched to the county-level tweet density data.³⁷

Table VIII presents the results using the microlevel VOTER data on attitudes towards sexual harassment. All regressions control for individual characteristics, their political interest and knowledge, and their voting history.³⁸ Overall, the tweets capture individual pro-feminist (anti-sexism) and anti-Republican attitudes. First, I regress an aggregated "sexism" score based on six questions that proxy for attitudes toward gender roles and sexual harassment, which is increasing in

³⁷Change variables are computed for individuals that have been tracked throughout.

³⁸The set of individual characteristics include gender, race, education, employment, birth cohort (by decade), income, marital status, and the number of children. The set of controls for voting history and tendency include whom the respondent would have voted for in a presidential election and for congress when asked in 2012 ((1) Democratic, (2) Republican, (3) Other/not sure/would not vote), plus the two indicators who whether the respondent always vote the same party. 1,809 respondents (23.4%) indicate that they always vote Republican, 2,287 (29.6%) indicate that they always vote Democratic, and the remaining 3641 (47.1%) indicate they vote for both. The regressions also control for interest and knowledge in current affairs and politics on a four-point scale.

sexism.³⁹ The tweets measure is negatively correlated with the sexism measure in both the 2016 and 2018 waves in columns (1)–(2), as anticipated. The tweets measure, however, does not predict any change in sexism (column (3)). Column (4) indicates the tweets measure is not just picking up concerns about broader "problems in society".

In column (5), the tweets measure implies that respondents from counties with higher MeToo tweet incidences are less approving of the Republican party. The tweets measure, however, does not predict approval of the Democratic party in column (6). As expected, whether an individual always votes Democratic or Republican is also highly correlated with the party's approval. Overall, the results from Table VIII indicate that the county-level MeToo tweets in 2018 are indeed correlated with attitudes of the electorate towards harassment and the political parties' handling of it.⁴⁰

3.6.2 Changing Demographics

A limitation on a causal interpretation is in changing demographics. As a falsification test, I repeat the main results from Table III, but with the 2016 House returns as the dependent variable. Since the MeToo movement in 2018 cannot travel back in time, there should be no detected correlation between the MeToo measure and the 2016 outcomes. Table C.0.7 reports the results and suggest that the county-level trends detected in Table III already exist in 2016.

One possible explanation is an underlying time trend in Republican counties towards a higher Democratic vote share. Evidence of this trend is also hinted at in Table V to the extent that counties with a higher Republican vote share in the 2016 presidential elections have

³⁹For example, one question gets respondents to respond to the statement "*Women who complain about harassment often cause more problems than they solve*". Responses go from a scale of 1–4 (strongly agree to strongly disagree).

⁴⁰Table C.0.6 of the appendix uses the VOTER microdata with further evidence that the tweets capture an anti-Republican sentiment. The results imply that the log tweet density measure of an individual's county is negatively and statistically associated with the probability of voting Republican in the presidential and congressional elections, conditional on the same controls in Table VIII. Further, the tweet density measure also increases the probability that the individual switch their vote to the Democratic party from 2016 to 2018.

systematically lower turnout in the 2018 house elections, and in Table VI to the extent that change in turnout is negatively correlated with the Republican presidential vote share (columns (5)–(6)). Part of this may be driven by cross-county migration or the coming of age of young voters. The results for changes include changes in demographics, using the 5-year estimates from 2012–16 to 2015–19, but these may not fully address underlying trends.

3.6.3 County-level Variation

Using county-level variation of the MeToo measure and vote shares, in principle, addresses concerns with unobserved confounders at the congressional district level. Moreover, county borders are infrequently adjusted and are therefore not directly affected by gerrymandering present at other geographical delineations. The county-level variation, however, poses two broad problems.

First, the use of the difference-in-differences at the county level relies on MeToo high variation within counties of a district. Districts with low cross-county variation in the MeToo movement and with a low number of counties do not contribute much to the estimation. Trivially, any county that is coterminous with its district does not contribute to identification. The extent to which districts with many counties drive the results implies limitations to generalising the results to the average US county.

A second institution-specific issue is the difference in county vs district border changes. Districts are apportioned by population size every ten years, according to the US Constitution. County borders, however, are more idiosyncratic, infrequently adjusted, and are more determined by historic episodes such as the colonial land grant era than by contemporary population size. The growing urbanisation of the US, the decennial apportionment of the districts, combined with the lack of mandate for county border adjustments imply that suburban and rural districts are geographically large, and tend to contain many counties, while the urban districts are geographically small and contain fewer counties.

Furthermore, the above implies that most of the counties are more rural, more White, and more Republican, which suggests that the detected effect of the MeToo measure at the county level are concentrated in geographically large areas with a small share of the US voting population. To this extent, results involving the additional interaction with the Republican vote share say more about the MeToo movement and the Republican rural and suburban counties rather than the average US county.

3.7 Conclusion

This paper investigates the connection between the MeToo movement and the 2018 midterm House elections. The candidate-level results indeed suggest an advantage for Democratic women candidates, and a disadvantage for Republican men candidates, but only in counties with a high MeToo tweet measure and high Republican support in the 2016 presidential elections. However, this pattern also exists in the 2016 elections, implying that the effect of the MeToo movement comes through channels other than changing voting patterns.

Exploring additional channels suggests two main explanation for the observed advantage of Democratic women candidates in the 2018 House elections. The first is that the movement changed established political party machinery, where districts with a higher measured MeToo movement have increased probability of having a Democratic women challenger when the seat is held by a man from the same party. The second is that the MeToo intensity measure captures an increase in turnout in Republican strongholds, and this is consistent with a drop in Republican vote share in the same counties.

This study adds to the literature on the political economy of the mass media and its effect on electoral outcomes. The focus is on the House elections instead of the Senate because only a third of Senate seats are up for election. The House is an important part of the national legislation and, keeping with the theme of the grassroots is the legislation that is more responsive to what their constituencies need. Furthermore, the

House has the power to initiate impeachment, as is the case with the Republican president at the time of this writing.⁴¹

Potential avenues of research include whether the MeToo effect persists into the next round of elections, which include the Republican president, a subject of the movement himself. In terms of minority representation, one may be interested in whether the 2018 congressional composition substantially changed public goods provision (as in Chattopadhyay and Duflo 2004; Pande 2003). Given how state judges were recalled for being lenient in sexual assault cases, another potential study relating to legal realism is on whether the movement induces harsher sentencing in sexual assault crimes.⁴²

3.8 Tables

⁴¹On allegations that President Trump leveraged his position in the White House to pressure foreign leaders into investigating his political opponents in the coming elections.

⁴²At the time of writing, Twitter CEO Jack Dorsey announced a ban on all paid political advertising, stating that political messages "should be earned, not bought" (Rajan 2019). While Facebook is the dominant platform for political advertising, Twitter's policy acknowledges the potential influence of political campaigns on social media.

TABLE I
SUMMARY STATISTICS

	Mean	Std. Dev.	Min	Max	Obs.
Log no. of MeToo tweets	2.396	(2.361)	0.000	10.601	8654
<i>Candidate characteristics</i>					
Challenger (%)	68.847	(46.315)	0.000	100.000	8654
Woman (%)	20.857	(40.631)	0.000	100.000	8654
Black (%)	0.404	(6.347)	0.000	100.000	8654
Hispanic (%)	3.547	(18.499)	0.000	100.000	8654
White (%)	94.754	(22.297)	0.000	100.000	8654
<i>District Seat characteristics</i>					
Republican incumbent (%)	65.091	(47.671)	0.000	100.000	8654
Democratic incumbent (%)	17.056	(37.614)	0.000	100.000	8654
Open seat (%)	17.587	(38.073)	0.000	100.000	8654
No main challenger (%)	1.502	(12.165)	0.000	100.000	8654
<i>Electoral variables</i>					
2016 House Rep. vote share (%)	63.859	(21.597)	0.000	100.000	8482
2016 House turnout ('000)	96.860	(282.419)	0.000	3129.539	8654
2012 Pres. Rep. vote share (%)	57.538	(15.554)	5.978	95.862	8646
2012 Pres. turnout ('000)	96.032	(281.350)	0.000	3181.067	8654
2016 Pres. Rep. vote share (%)	60.323	(16.936)	8.296	96.033	8646
2016 Pres. turnout ('000)	102.954	(304.243)	0.000	3434.308	8654
<i>Census variables 2012–16 ACS average</i>					
Population ('000)	259.121	(851.295)	0.076	10'057.155	8642
Black (%)	8.737	(13.066)	0.000	81.533	8642
Hispanic (%)	10.804	(14.750)	0.000	98.959	8642
White (%)	74.841	(20.668)	0.760	100.000	8642
Foreign born (%)	6.110	(7.443)	0.000	52.230	8642
Female (%)	50.060	(2.163)	21.513	56.418	8642
Age 29 and under (%)	37.575	(5.399)	11.842	70.981	8642
Age 65 and over (%)	17.150	(4.576)	3.855	53.106	8642
Median HH income ('000)	50.089	(13.765)	18.972	125.672	8642
Unemployment (%)	7.077	(3.000)	0.000	29.927	8642
HS or less (%)	13.928	(6.317)	1.279	51.479	8642
College or more (%)	22.512	(10.198)	2.985	80.210	8642
Rural population (%)	51.696	(33.614)	0.000	100.000	8646

Notes—Observations are at the county level. Ethnic of house candidates are inferred using the Name Prism API (Ye et al., 2017). Republican vote share is computed as total number of vote cast for the Republican party divided by the total number of votes cast. House vote shares reported in this Table is the all-party vote share. Presidential vote shares are always two-party vote shares. County census variables come from the ACS (American Community Survey) 5-year estimates for 2012–16. Observations unweighted.

TABLE II
SELECTION OF TWEET DENSITY IN COUNTIES

	<i>ln</i> (tweets density) in 2018 with meToo hashtag				
	(1)	(2)	(3)	(4)	(5)
House Rep. vote share in 2016		0.009*** (0.002)	0.007*** (0.002)	0.003 (0.004)	0.003 (0.004)
House elections turnout in 2016		0.017** (0.007)	0.007 (0.005)	−0.005 (0.008)	−0.004 (0.008)
Pres. Rep. vote share in 2016		−0.020*** (0.004)	−0.009** (0.004)	−0.005 (0.005)	−0.007 (0.005)
Pres. Rep. vote share change (2012–16)		−0.018*** (0.006)	0.011** (0.006)	0.014 (0.012)	0.006 (0.012)
Pres. election turnout in 2016		0.002 (0.010)	0.001 (0.006)	0.023** (0.010)	0.021** (0.010)
Pres. election turnout change (2012–16)		0.005 (0.010)	0.003 (0.007)	−0.014 (0.010)	−0.015 (0.010)
%Female × (High-speed connection)	0.105*** (0.036)	0.030 (0.032)	−0.044 (0.032)	−0.002 (0.033)	−0.004 (0.033)
Constant	−7.416*** (1.219)	−8.131*** (1.151)	−8.444*** (1.286)	−7.824*** (1.488)	−7.685*** (1.488)
<i>Control variables</i>					
2016 House elections		X	X	X	X
2012–16 Presidential election		X	X	X	X
Census Control			X	X	X
U.S. House District F.E.				X	X
F-test: County census controls = 0			$F = 23.55^{***}$	$F = 14.33^{***}$	$F = 14.82^{***}$
R^2	0.066	0.128	0.273	0.384	0.384
N	2466	2427	2427	2427	2427

Notes—Observations are at the county level. The dependent variable is tweet density—the (natural) log of the ratio of the number of tweets in 2018 which contains the MeToo hashtag, to population size. High-speed connection is the ratio of residential households in a county with high-speed internet connections from the FCC. Republican vote share is the votes received by the Republican candidate (party) in the Presidential (House) election, divided by the total number of votes cast. Turnout is the number of votes cast divided by the number of voting-aged population. County census controls for demographics come from the ACS (American Community Survey) 5-year estimates for 2012–16—they include 14 demographic variables of ethnic, gender, age, education, and foreign-born composition, income and employment rate, and rural-urban composition data. Column (5) uses the two-party Republican vote share—number of votes received by the Republican candidate divided by votes received by both the Republican and Democratic candidates. Robust standard errors in parentheses clustered at the 388 U.S. House congressional districts.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE III
THE EFFECT OF THE MeToo MOVEMENT ON CANDIDATE VOTE SHARE

	Heterogeneous effect, by presidential Republican vote share in 2016				
	All-party vote share			Two-party vote share	
	(1)	(2)	(3)	(4)	(5)
Log tweet density $\times$ (Rep. woman)	-3.557*** (0.993)	-0.674 (0.574)	-0.707 (0.724)	-0.472 (0.713)	-0.193 (0.699)
Log tweet density $\times$ (Dem. woman)	2.073*** (0.492)	0.074 (0.188)	-0.655*** (0.253)	-0.930*** (0.287)	-0.636** (0.312)
Log tweet density $\times$ (Rep. man)	-2.218*** (0.410)	0.008 (0.157)	0.300 (0.235)	0.408 (0.255)	0.302 (0.272)
Log tweet density $\times$ (Dem. man)	2.316*** (0.515)	0.279 (0.232)	-0.029 (0.354)	-0.439 (0.431)	-0.592 (0.436)
Log tweet density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Rep. woman)			-0.023 (0.030)	-0.027 (0.029)	-0.027 (0.028)
Log tweet density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Dem. woman)			0.037*** (0.014)	0.047*** (0.013)	0.043*** (0.015)
Log tweet density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Rep. man)			-0.021** (0.009)	-0.022** (0.009)	-0.024** (0.010)
Log tweet density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Dem. man)			0.014 (0.014)	0.027* (0.015)	0.035** (0.015)
<i>Control variables</i>					
Candidate fixed effects	X	X	X	X	
District fixed effects					X
2016 House & 2012–16 Pres. election		X	X	X	X
County census demographics		X	X	X	X
Racial & gender voting		X	X	X	X
F-test: House & 2012–16 Pres. election = 0		$F = 297.71^{***}$	$F = 13.07^{***}$	$F = 15.63^{***}$	$F = 12.17^{***}$
F-test: Census controls = 0		$F = 3.82^{***}$	$F = 3.55^{***}$	$F = 4.14^{***}$	$F = 2.63^{***}$
F-test: Racial & gender voting = 0		$F = 3.81^{***}$	$F = 4.53^{***}$	$F = 9.16^{***}$	$F = 2.94^{***}$
Main-party candidates only				X	X
R^2	0.907	0.975	0.977	0.952	0.886
N	8634	8470	8470	6234	6234

Notes—The dependent variable is the candidate vote share at the district-county level. Tweet density is the (natural) log of MeToo tweets in 2018 divided by county population. Past electoral controls include: (1) 2016 house Republican vote share, (2) 2016 presidential Republican vote share, and (3) 2012–16 presidential Republican vote share change, fully interacted with party. County census controls for demographics come from the ACS 5-year estimates for 2012–16—they include 14 demographic variables of ethnic, gender, age, education, and foreign-born composition, income and employment rate, and rural-urban composition data. Controls for voting by racial and gender lines include interacting politician gender and ethnic (White, Black, Hispanic, and Others) with the corresponding county ethnic percentage. Ethnic of a politician is inferred using their names through the NamePrism API (Ye et al., 2017). Columns (4)–(5) includes only main-party candidates and uses two-party vote shares on both sides of the equation. Robust standard errors in parentheses are clustered by candidates.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE IV
ROBUSTNESS

	Robustness check for Column (4) of Table III						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log tweet density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Rep. woman)	-0.039 (0.032)	-0.039 (0.036)	-0.028 (0.030)	-0.028 (0.030)	-0.027 (0.029)	-0.029 (0.038)	-0.012 (0.025)
Log tweet density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Dem. woman)	0.047*** (0.014)	0.049*** (0.014)	0.051*** (0.014)	0.051*** (0.014)	0.047*** (0.013)	0.043*** (0.013)	0.044*** (0.014)
Log tweet density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Rep. man)	-0.021** (0.009)	-0.025*** (0.008)	-0.028*** (0.009)	-0.025*** (0.009)	-0.022** (0.009)	-0.030** (0.012)	-0.024*** (0.009)
Log tweet density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Dem. man)	0.022* (0.013)	0.039*** (0.013)	0.033** (0.016)	0.028* (0.015)	0.027* (0.015)	0.014 (0.016)	0.032** (0.016)
<i>Control variables</i>							
Candidate fixed effects	X	X	X	X	X	X	X
2016 House & 2012–16 Pres. election	X	X	X	X	X	X	X
County census demographics	X	X	X	X	X	X	X
Racial & gender voting	X	X	X	X	X	X	X
Pre-primary MeToo tweets	X						
Counties > 2		X					
Turnout > 2000			X				
Voting population > 2000				X			
District std. dev. tweets > 5th percentile					X		
High-speed internet bw. 5th & 95th percentile						X	
Non-vacated seats							X
R^2	0.952	0.953	0.946	0.949	0.950	0.957	0.953
N	6234	5872	5821	5983	6065	5289	5311

Notes—This Table presents a set of robustness checks for column (4) of Table III. In column (1), the tweets measure is cut off before June, when most (17 states) of the primary elections took place. In the column (2) sample "Counties > 2", districts with 1 or 2 counties are dropped. In column (3), the sample "Turnout > 2000" excludes counties with fewer than 2,000 votes cast in the 2018 House elections. In column (4), the sample "Voting population > 2000" excludes counties with an estimated ACS voting-aged population of fewer than 2,000. In column (5), the sample "Std. dev. tweets > 5th percentile" excludes districts where the geographical variation in the MeToo tweets is below the 5th percentile. In column (6), the sample "High-speed internet bw. 5th & 95th percentile" includes only counties where the high-speed internet measure from the FCC is between the 5th & 95th percentile. In column (7), the sample "Non-vacated seat" drops open-seat districts where the incumbent has retired. All controls are otherwise the same. Robust standard errors in parentheses are clustered by candidates.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE V
THE EFFECT OF THE MeToo MOVEMENT ON TURNOUT

	Measure of county-level MeToo movement (τ) is			
	ln(No. of tweets divided by population)		ln(No. of tweets)	
	(1)	(2)	(3)	(4)
τ	0.0221 (0.0144)	-0.0162 (0.0347)	0.0145 (0.0121)	-0.0221 (0.0180)
Pres. 2016 Rep. vote share	-0.0128*** (0.0025)	-0.0081* (0.0046)	-0.0127*** (0.0025)	-0.0136*** (0.0025)
$\tau \times$ (Pres. 2016 Rep. vote share)		0.0006 (0.0004)		0.0006*** (0.0002)
District fixed effects	X	X	X	X
Census Control	X	X	X	X
F-test: Electoral controls = 0	$F = 19.19^{***}$	$F = 4.32^{***}$	$F = 19.26^{***}$	$F = 18.1^{***}$
F-test: County census = 0	$F = 2.72^{***}$	$F = 2.74^{***}$	$F = 2.47^{***}$	$F = 2.53^{***}$
R^2	0.6551	0.6557	0.6543	0.6556
N	3102	3102	3102	3102

Notes—Observations are at the county level. The dependent variable is the log of total county votes cast in the 2018 House elections minus the same variable for the 2016 House elections. In columns (1)–(2), the measure of the MeToo movement is the log of county-level MeToo tweets divided by county population; in columns (3)–(4) the measure is the log of county-level MeToo tweets. Pres. 2016 Rep. vote share is the two-party county-level vote share of the Republican candidate in the 2016 presidential election. County census controls include the 14 demographic variables and additionally the percentage of citizen voting-age population; these are entered as both levels and changes from the ACS 5-year estimates for 2012–16 and the ACS 5-year estimates for 2015–19, except for the percentage rural population available only from the decennial census. Controls also include the turnout for both the 2016 Presidential and House elections, and the 2012 presidential Republican vote share. Robust standard errors in parentheses are clustered by districts.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE VI
CHANGE IN REPUBLICAN VOTE SHARE, DISTRICT-COUNTY LEVEL

	Change in <i>house</i> Republican vote share between 2016 and 2018				Falsification	
	Change in all-party vote share		Change in two-party vote share		Change in <i>presidential</i> Republican vote share between 2012 and 2016	
	(1)	(2)	(3)	(4)	(5)	(6)
Log tweet density $\times$ (Pres. 2016 Rep. vote share)	−0.0302*** (0.0090)	−0.0288*** (0.0088)	−0.0149** (0.0074)	−0.0137* (0.0072)	−0.0033 (0.0027)	
Log tweet density $\times$ (Pres. 2012 Rep. vote share)						0.0017 (0.0028)
Log tweet density	0.4280* (0.2332)	0.4699** (0.1988)	0.1150 (0.2573)	0.1536 (0.2175)	0.0239 (0.0553)	−0.1444 (0.1644)
Change in log(total House votes) 2016–18		−3.0346* (1.6873)		−3.0140* (1.6980)	0.0094 (0.1008)	−0.6947** (0.2975)
Change in log(total Pres. votes) 2012–16					−2.8255 (1.9617)	−1.5249 (1.6380)
<i>Control variables</i>						
District fixed effects	X	X	X	X	X	X
Past electoral controls	X	X	X	X	X	X
County census demographics	X	X	X	X	X	X
R^2	0.8905	0.8953	0.9057	0.9103	0.8766	0.9118
N	3102	3102	3102	3102	3102	3102

Notes—Observations are at the district-county level. The dependent variable is the change in Republican vote share. In columns (1)–(2), it is the all-party change in Republican vote share in the House elections from 2016–18. In columns (3)–(4) the dependent variable is the same variable for the two-party vote share. In columns (5)–(6), the dependent variable is the change in the presidential Republican (two-party) vote share from 2012–16. County census controls include the 14 demographic variables and additionally the percentage of citizen voting-age population; these are entered as both levels and changes from the ACS 5-year estimates for 2012–16 and the ACS 5-year estimates for 2015–19, except for the percentage rural population available only from the decennial census. The electoral control variables in columns (1)–(4) include the house Republican vote share in 2016, and the change in presidential Republican vote share from 2012–16; in columns (5)–(6) the electoral controls are the house Republican vote share in 2016, and the change in presidential Republican vote share from 2008–12. Robust standard errors in parentheses are clustered by the 388 U.S. House congressional districts.

- *** Significant at the 1 per cent level.
- ** Significant at the 5 per cent level.
- * Significant at the 10 per cent level.

TABLE VII
SELECTED DISTRICTS AND STATES

	(1)	Subsamples		
		(2)	(3)	(4)
Log tweet density $\times$ (Rep. woman)	-0.463 (0.667)	1.406 (0.974)	-2.278 (1.813)	-2.701* (1.444)
Log tweet density $\times$ (Dem. woman)	-0.909*** (0.286)	-2.067*** (0.604)	-1.064** (0.490)	-1.063*** (0.319)
Log tweet density $\times$ (Rep. man)	0.420* (0.237)	1.458*** (0.491)	1.012*** (0.354)	0.813*** (0.299)
Log tweet density $\times$ (Dem. man)	-0.208 (0.414)	-2.188*** (0.650)	-0.484 (0.521)	-0.341 (0.367)
Log tweet density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Rep. woman)	-0.027 (0.027)	-0.065 (0.042)	0.027 (0.042)	-0.098* (0.055)
Log tweet density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Dem. woman)	0.046*** (0.014)	0.123*** (0.035)	0.051** (0.023)	0.049*** (0.014)
Log tweet density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Rep. man)	-0.023*** (0.008)	-0.074*** (0.028)	-0.036** (0.018)	-0.035*** (0.012)
Log tweet density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Dem. man)	0.016 (0.014)	0.115** (0.045)	0.011 (0.016)	0.025* (0.014)
<i>Control variables</i>				
Candidate fixed effects	X	X	X	X
2016 House & 2012–16 Pres. election	X	X	X	X
County census demographics	X	X	X	X
Racial & gender voting	X	X	X	X
<i>F</i> -test: House & 2012–16 Pres. election = 0	<i>F</i> = 18.19***	<i>F</i> = 5.35***	<i>F</i> = 2.47**	<i>F</i> = 8.91***
<i>F</i> -test: Census controls = 0	<i>F</i> = 4.74***	<i>F</i> = 11.93***	<i>F</i> = 8.48***	<i>F</i> = 5.08***
<i>F</i> -test: Racial & gender voting = 0	<i>F</i> = 11.17***	<i>F</i> = 14.22***	<i>F</i> = 13.39***	<i>F</i> = 34.25***
<i>F</i> -test for balance	—	<i>F</i> = 3.04**	<i>F</i> = 1.45	<i>F</i> = 2.62**
Main-party candidates only	X	X	X	X
<i>R</i> ²	0.950	0.937	0.957	0.955
<i>Sample</i>				
Full sample	X			
Rep. districts with low margin (< 10%) in 2016		X		
States with senate elections & split delegation			X	
Rep. states in 2016 Pres. elections with low margin (< 10%)				X
Candidate observations	645	72	142	258
Candidate-county observations	6079	640	1296	2388

Notes—The dependent variable is the candidate two-party vote share at the district-county level. Observations are subsamples of Table III. Only main-party candidates are included. Districts without a main-party challenger are dropped. Column (1) includes the full sample with the previous conditions. Column (2) is the subsample of candidate-county observations in districts with a low margin of victory in the 2018 house elections. Column (3) includes only observations in states with a senate election and where the senate has a split delegation (one Democratic and one Republican senator). Column (4) includes only observations in states where the Republicans won in the 2016 presidential elections and where the margin of victory is small. All columns include an *F*-test which tests for how similar the included observations are to the excluded observations, using the 13 county-level demographics. Robust standard errors in parentheses are clustered by candidates.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE VIII
CORRELATION BETWEEN THE MeToo MOVEMENT AND INDIVIDUAL ATTITUDES (VOTER DATA)

	Sexism 2016 (Range 1 to 24)	Sexism 2018 (Range 1 to 24)	Change in sexism (Range -23 to 23)	1(Allegations indicative of wider problems)	Approval of Rep. party in handling harassment (Range 1 to 4)	Approval of Dem. party in handling harassment (Range 1 to 4)
	(1)	(2)	(3)	(4)	(5)	(6)
Log of tweet density	-0.096*** (0.036)	-0.134*** (0.050)	0.009 (0.034)	0.011* (0.006)	-0.035*** (0.012)	0.013 (0.013)
1(Always vote for Democrats)	-0.289*** (0.112)	-0.204 (0.138)	0.016 (0.113)	0.035* (0.018)	-0.103*** (0.036)	0.116*** (0.035)
1(Always vote for Republicans)	0.823*** (0.117)	1.032*** (0.171)	0.057 (0.128)	-0.015 (0.025)	0.242*** (0.040)	-0.107*** (0.040)
<i>Control variables</i>						
Individual characteristics	X	X	X	X	X	X
Voting history & tendency	X	X	X	X	X	X
Political interest & knowledge	X	X	X	X	X	X
F-test: Individual characteristics = 0	F = 12.78***	F = 9.34***	F = 1.27	F = 3.84***	F = 1.33*	F = 3.02***
F-test: Voting tendency = 0	F = 546.05***	F = 242.62***	F = .66	F = 101.44***	F = 315.92***	F = 216.49***
F-test: Political interest & knowledge = 0	F = 4.45***	F = .63	F = 1.06	F = .11	F = 3.03**	F = 1.32
R ²	0.393	0.393	0.015	0.187	0.351	0.307
N	6625	3908	3816	3972	3931	3934

Notes—Observations are individual respondents in the Democracy Fund VOTER (Views of the Electorate Research) survey. All regressions control for individual characteristics including gender, race, education, employment, birth cohort (by decade), income, marital status, and number of children. Voting history & tendency controls include which party the individual would have for congress and president in 2012, and an indicator for whether the individual always for for the same party. Political interest and knowledge controls for the level of interest and knowledge the individual has in politics and current affairs. The dependent variable in columns (1)–(2) is an aggregated score from *sexism1*–*sexism6* in the VOTER survey, which is increasing in "sexism". The dependent variable in column (3) is the change in this score for the same individual from 2016–18. The dependent variable in column (4) is a dummy for whether the respondent thinks that recent allegations of sexual harassment and assault reflect widespread problems in society. The dependent variable in column (5) and (6) is the approval rating of the Republican and Democratic party in the handling of harassment and assault in politics. Robust standard errors clustered by counties reported in parentheses.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

Appendices

Appendix C

Data Appendix

C.0.1 Data Details

To download and map the tweets to counties, I proceed as follows:

1. I use the *GetOldTweets-python* pseudo-API by Jefferson Henrique (<https://github.com/Jefferson-Henrique/GetOldTweets-python>) which scrapes the Twitter Search browser for tweets containing the MeToo hashtag. At the time of use, I need to make changes to two lines of the code to retrieve the author's username as noted in the *issues* of the repository. With the usernames, I query the Official Twitter API which returns their user geolocation strings.
2. I use a series of hard-coded rules to parse the various user-input geolocations into a standardised U.S. city-state format (e.g. Philadelphia, Pennsylvania). I first retain only characters in a string that are ASCII characters, so that non-English and symbols are removed. After retaining only ASCII characters, 87'123 geolocation strings have a character length of 7 or less, indicating a sizeable number of Twitter users key-in non-ASCII.
3. I then check whether the geolocation string can be unambiguously identified as a non-U.S. country. If so, these are filtered out immediately. Using the ISO-3166 country names and codes, 89'115 (16% of 520'513) of the geolocation strings are immediately identified as Twitter users who list a non-U.S. country as their location.

4. For the remaining geolocation strings, I check if they can be identified as a U.S. city-state by searching for both state names and postal codes as well as city names within the string. Pseudo-code listing 1 provides the specific hard-coded rules used. The set of rules allows me to successfully parse 130'433 (25% of 520'513) geolocations into a standard U.S. city-state. A relatively small percentage of geolocation strings, 19'590 or 3.8%, is stated as the United States, but omits information about the state, the city, or both.
5. Finally, I match the tweets by U.S. city-state to their primary counties using the *United States Cities Database*. The primary counties are defined by the U.S. Geological Survey, which takes the centroid of a city and then recording the county in which the centroid lies.

Pseudocode 1: Parsing Geolocation

```

foreach geoloc(ation string) do
  if comma not in geoloc then
    /* [Step 1] Check if unambiguously a non-U.S. country */
    if len(geoloc)==2 then
      | check if geoloc matches a non-U.S. country using ISO alpha-2 code
    else if len(geoloc)==3 then
      | check if geoloc matches a non-U.S. country using ISO alpha-3 code
    else
      | check if geoloc matches a non-U.S. country using ISO country name name
    if not unambiguously non-U.S. country in [Step 1] then
      /* [Step 2] Try geoloc string as U.S. city w/o state info. */
      | try geoloc as a city named after state (e.g. 'utah' as Utah City, Utah)
      if that fails then
        | try geoloc as a uniquely named U.S. city (e.g. 'chicago' as Chicago City in Illinois)
    if still not identified as a U.S. city-state in [Step 2] then
      /* [Step 3] Try geoloc string as U.S. city with state info. */
      | check if a comma is implied in either order (e.g. 'philadelphia pa' and 'pa
      | philadelphia' as Philadelphia City, Pennsylvania)

  if comma in geoloc then
    /* [Step 4] Check if unambiguously a non-U.S. country */
    | check if one side of comma is unambiguously a non-U.S. country as in [Step 1] (e.g.
    | 'beunos aires, argentina' should be filtered out)
    if not unambiguously a non-U.S. country then
      /* [Step 5] Try as a U.S. city-state */
      if one side of comma in geoloc has len==2 then
        | use as State postal code and the other side as city (e.g. 'avon, al' as Avon City,
        | Alabama)
      else
        | try one side as a State name and the other as a city name (e.g. 'avon, alabama'
        | and 'alabama, avon' as Avon City, Alabama)
    if still not identified as U.S. city-state then
      /* [Step 6] Try one side as city-state and the other an indicator of the
      | States (e.g. 'us', 'usa', 'united states', 'united states of america'
      | */
      if one side is indicator of the U.S. then
        | try other side as city named after state (e.g. 'utah, usa' as Utah City, Utah)
        if that fails then
          | try geoloc as a uniquely named U.S. city (e.g. 'chicago, united states' as
          | Chicago City in Illinois)

```

C.0.2 Extra Figure and Tables

TABLE C.O.1
Examples of Parsing Twitter User Geolocation

User Geolocation	State	(Primary) County
nomadic	—	—
sandy oaks, tx	Texas	Bexar
los angeles, ca	California	Los Angeles
calcinato, lombardia	—	—
pensacola, fl	Florida	Escambia
london, england	—	—
victoria, bc, canada	—	—
virginia	—	—
washington, dc	District of columbia	District Of Columbia
dallas, tx	Texas	Dallas
south	—	—
ca	—	—
united states	—	—
michigan, usa	—	—
bordeaux, aquitaine	—	—
oxford, ms	Mississippi	Lafayette
chicago	Illinois	Cook
port townsend, wa	Washington	Jefferson
ut ,	—	—
namak haram in pakistan	—	—
lagos, nigeria	—	—
boston, ma	Massachusetts	Suffolk
grittydelphia via la,nyc,gb	—	—
pakistan	—	—
oakland, ca	California	Alameda
united states	—	—
st louis, mo	Missouri	St. Louis (City)
kitchener, ontario	—	—
san francisco, ca	California	San Francisco
stanford, ca	California	Santa Clara
probably on the floor somewhere	—	—
chicago, il	Illinois	Cook
houston, tx	Texas	Harris
micromsmemumbaiwala	—	—
mother earth	—	—
houston, tx	Texas	Harris
cleveland, tn	Tennessee	Bradley
oregon, usa	—	—
tuscaloosa	Alabama	Tuscaloosa
new york	New york	New York
provo, ut	Utah	Utah
united states	—	—
grand rapids, mi	Michigan	Kent
the village	Oklahoma	Oklahoma
san francisco	California	San Francisco
murcia, espana	—	—
mount greenwood, chicago	—	—
morgantown, wv	West virginia	Monongalia
las vegas, nv	Nevada	Clark
new jersey, usa	—	—
whalley, bc	—	—

Notes—This Table provides 50 examples of parsing twitter users' geolocation. *User Geolocation* column is the self-declared geolocation of users. *State* column is the identified State in the U.S., and the *(Primary) County* column is the identified U.S. county based on the city-state. Primary Counties are identified using the *United States Cities Database* from <https://simplemaps.com/data/us-cities> where primary counties of cities are identified by the U.S. Geological Survey and U.S. Census Bureau by taking the centroid of a city and then recording the county in which the centroid lies.

FIGURE C.O.3
Candidacy as Strategic Reaction (Republican Women Challengers)

FIGURE C.0.4
CORRELATIONS OF TWEETS AND REPUBLICAN VOTE SHARE, BY STATE (PART 1)

FIGURE C.0.5
CORRELATIONS OF TWEETS AND REPUBLICAN VOTE SHARE, BY STATE (PART 2)

FIGURE C.o.6

CORRELATION (BINNED SCATTERS) BETWEEN COUNTY DEMOGRAPHICS, TWEETS, AND 2018 HOUSE ELECTIONS

TABLE C.0.2
SELECTION OF WOMEN CANDIDATES INTO DISTRICTS

	Dep. var. are indicators for				
	At least 1 woman candidate	Exactly 1 woman candidate	H2H man-woman main party	At least 1 woman challenger	At least 1 main-party woman challenger
	(1)	(2)	(3)	(4)	(5)
Log tweets density	−0.009 (0.067)	−0.029* (0.016)	−0.058 (0.057)	0.005 (0.072)	−0.017 (0.064)
<i>Past Electoral controls</i>					
House Rep. vote share in 2016	−0.001 (0.004)	−0.002 (0.002)	0.002 (0.005)	−0.000 (0.005)	0.001 (0.005)
House turnout 2016	0.006 (0.006)	−0.002 (0.003)	0.007 (0.006)	0.007 (0.007)	0.007 (0.006)
Pres. Rep. vote share in 2016	0.008 (0.007)	0.006*** (0.002)	0.004 (0.009)	0.006 (0.008)	0.005 (0.009)
Pres. Rep. vote share change (2012–16)	−0.013 (0.016)	−0.003 (0.008)	−0.024 (0.015)	−0.007 (0.016)	−0.014 (0.013)
Pres. turnout 2016	−0.001 (0.014)	−0.003 (0.006)	−0.005 (0.015)	−0.001 (0.015)	−0.000 (0.015)
<i>Political Seat controls</i>					
Open seat	0.337*** (0.109)	0.010 (0.069)	0.252*** (0.089)	0.445*** (0.110)	0.477*** (0.102)
Incumbent is woman	0.576*** (0.080)	−0.137*** (0.037)	0.551*** (0.147)	−0.146 (0.140)	−0.040 (0.077)
Incumbent is Republican	0.256*** (0.079)	−0.026 (0.044)	0.175* (0.094)	0.361*** (0.067)	0.341*** (0.064)
Incumbent is Rep. woman	−0.072 (0.115)	0.149*** (0.055)	−0.735*** (0.209)	0.369* (0.185)	0.301** (0.133)
State fixed effects	X	X	X	X	X
Census Control	X	X	X	X	X
F-test: County census = 0	6***	9.8***	4.54***	2.65**	1.69*
R ²	0.280	0.255	0.221	0.235	0.243
Probability (Unconditional)	0.541	0.067	0.405	0.430	0.376
N	388	388	388	388	388

Notes—Observations are House congressional districts. Results are estimated using the linear probability model. Dependent variable in column (1) is the dummy for at least one woman candidate in the district; in column (2) it is the dummy for exactly one woman candidate; in column (3) it is a dummy for when there is a head-to-head between a man and woman candidate from the major party; in column (4) it is a dummy for at least one woman candidate who is a challenger; and in column (5) it is a dummy for at least one woman candidate who is a challenger from one of the two major parties. Census controls are aggregated from the county to the district level. Observations weighted by the total votes cast in the 2016 Presidential election. Robust standard errors in parentheses are clustered at states.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE C.0.3
FULL REPORT OF INTERACTED COEFFICIENTS, FOR PARTY AND GENDER

	All-party	Two-party
	(1)	(2)
Rep. woman $\times$ (Log tweet density)	−0.707 (0.724)	−0.472 (0.713)
Dem. woman $\times$ (Log tweet density)	−0.655*** (0.253)	−0.930*** (0.287)
Rep. man $\times$ (Log tweet density)	0.300 (0.235)	0.408 (0.255)
Dem. man $\times$ (Log tweet density)	−0.029 (0.354)	−0.439 (0.431)
Rep. woman $\times$ (Pres. 2016 Rep. vote share)	0.330 (0.298)	0.543* (0.281)
Dem. woman $\times$ (Pres. 2016 Rep. vote share)	−0.342** (0.149)	−0.483*** (0.137)
Rep. man $\times$ (Pres. 2016 Rep. vote share)	0.401*** (0.130)	0.623*** (0.099)
Dem. man $\times$ (Pres. 2016 Rep. vote share)	−0.446*** (0.140)	−0.572*** (0.138)
Rep. woman $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	−0.023 (0.030)	−0.027 (0.029)
Dem. woman $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	0.037*** (0.014)	0.047*** (0.013)
Rep. man $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	−0.021** (0.009)	−0.022** (0.009)
Dem man $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	0.014 (0.014)	0.027* (0.015)
<i>Control variables</i>		
Candidate fixed effects	X	X
2016 House & 2012–16 Pres. election	X	X
County census demographics	X	X
Racial & gender voting	X	X
Main-party candidates only		X
R^2	0.977	0.952
N	8470	6234

Notes—This Table reports the full coefficients of the interaction between party, gender, log tweet density, and the 2016 presidential Republican vote share. The coefficients here corresponds column (4) of Table III.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE C.o.4
ADDITIONAL ROBUSTNESS CHECKS

	Additional robustness check for Column (4) of Table III					
	(1)	(2)	(3)	(4)	(5)	(6)
Rep. woman $\times \tau \times$ (Pres. 2016 Rep. vote share)	−0.016 (0.032)	−0.027 (0.029)	−0.027 (0.027)	−0.027 (0.027)	−0.024 (0.030)	−0.058 (0.056)
Dem. woman $\times \tau \times$ (Pres. 2016 Rep. vote share)	0.037*** (0.012)	0.047*** (0.014)	0.047*** (0.014)	0.047*** (0.014)	0.049*** (0.015)	0.048** (0.023)
Rep. man $\times \tau \times$ (Pres. 2016 Rep. vote share)	−0.016* (0.008)	−0.022** (0.009)	−0.022*** (0.008)	−0.022*** (0.008)	−0.023** (0.009)	−0.018* (0.009)
Dem. man $\times \tau \times$ (Pres. 2016 Rep. vote share)	0.020 (0.015)	0.027* (0.015)	0.027 (0.016)	0.027* (0.014)	0.023 (0.014)	0.039** (0.018)
<i>Control variables</i>						
Candidate fixed effects	X	X	X	X	X	X
2016 House & 2012–16 Pres. election	X	X	X	X	X	X
County census demographics	X	X	X	X	X	X
Racial & gender voting	X	X	X	X	X	X
Main-party candidates only						
General log(tweets) - log(population)	X					
Drop Hawaii		X				
Two-way cluster Candidate and county			X			
Two-way cluster Candidate and District-county				X		
MeToo tweets without other hashtags					X	
Std. dev. tweets < 90th percentile						X
R^2	0.953	0.952	0.951	0.951	0.952	0.960
N	6234	6224	6122	6122	6234	5592

Notes—This Table presents additional robustness checks for column (4) Table III. In column (1), the specification is more general, with log(tweets) and log(population) entering the model separately so that their coefficients are allowed to differ. In column (2), observations from Hawaii are dropped. Columns (3) and (4) adjust standard errors by two-way non-nested clustering of the house candidates and county. In column (5), the tweets measure is computed using only tweets with a single (the MeToo) hashtag. In column (6), only districts where the standard deviation in the MeToo tweets is lower than the 90th percentile are included. In column (1) the reported coefficient is for log(tweets), in columns (2)–(6) the tweets measure is the log tweet density measure—log(tweets/population). All controls are otherwise the same, and robust standard errors in parentheses are otherwise clustered by candidates.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE C.O.5
THE EFFECT OF THE MeToo MOVEMENT ON TURNOUT
(FALSIFICATION—CHANGE IN TURNOUT PRESIDENTIAL ELECTION 2012–16)

	Measure of county-level MeToo movement (τ) is			
	ln(No. of tweets divided by population)		ln(No. of tweets)	
	(1)	(2)	(3)	(4)
τ	0.0006 (0.0008)	−0.0022 (0.0034)	0.0012 (0.0008)	0.0030 (0.0021)
Pres. 2016 Rep. vote share		−0.0019*** (0.0007)		−0.0021*** (0.0007)
$\tau \times$ (Pres. 2016 Rep. vote share)		0.0000 (0.0001)		−0.0000 (0.0000)
District fixed effects	X	X	X	X
Census Control	X	X	X	X
F-test: County census = 0	$F = 30.01^{***}$	$F = 24.12^{***}$	$F = 27.66^{***}$	$F = 25.26^{***}$
R^2	0.6985	0.7025	0.6988	0.7028
N	3158	3158	3158	3158

Notes—Observations are at the county level. The dependent variable is the log of total county votes cast in the 2016 Presidential elections minus the same variable for the 2012 Presidential elections. In columns (1)–(2), the measure of the MeToo movement is the log of county-level MeToo tweets divided by county population; in columns (3)–(4) the measure is the log of county-level MeToo tweets. Pres. 2016 Rep. vote share is the two-party county-level vote share of the Republican candidate in the 2016 presidential election. County census controls include the 14 demographic variables and additionally the percentage of citizen voting-age population; these are entered as both levels and changes from the ACS 5-year estimates for 2012–16 and the ACS 5-year estimates for 2015–19, except for the percentage rural population available only from the decennial census. Controls also include the 2008–2012 presidential elections turnout and Republican vote share. Robust standard errors in parentheses are clustered by districts.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE C.o.6
THE EFFECT OF MeTOO ON INDIVIDUAL VOTING (VOTER DATA)

	1(Voted Republican) in 2016	in 2018	Change in vote from Dem. to Rep.
	(1)	(2)	(3)
Log tweet density	-0.012*** (0.003)	-0.009*** (0.003)	-0.005** (0.003)
1(Always vote for Democrats)	-0.077*** (0.012)	-0.077*** (0.012)	0.003 (0.011)
1(Always vote for Republicans)	0.066*** (0.009)	0.092*** (0.013)	-0.022*** (0.007)
<i>Control variables</i>			
Individual characteristics	X	X	X
Voting history & tendency	X	X	X
Political interest & knowledge	X	X	X
<i>F</i> -test: Individual characteristics = 0	<i>F</i> = 4.66***	<i>F</i> = 1.78**	<i>F</i> = 1.38*
<i>F</i> -test: Voting history & tendency = 0	<i>F</i> = 2736.44***	<i>F</i> = 2405.58***	<i>F</i> = 5.04***
<i>F</i> -test: Political interest & knowledge = 0	<i>F</i> = .5	<i>F</i> = .91	<i>F</i> = .78
<i>R</i> ²	0.723	0.764	0.033
<i>N</i>	6020	3466	3204

Notes—Observations are individual respondents in the Democracy Fund VOTER (Views of the Electorate Research) survey. The dependent variable in column (1) is a dummy for whether the respondent voted Republican in the 2016 Presidential. The dependent variable in column (2) is a dummy for whether the respondent *would* have voted Republican for Congress in 2018 (recorded in April). Base category is to vote Democrat. The dependent variable in column (3) captures whether the respondent changes vote from 2016–18: 1 if vote changes from Democratic to Republican, 0 if no change, -1 if from Republican to Democratic party. All regressions control for individual characteristics including gender, race, education, employment, birth cohort (by decade), income, marital status and number of children. Voting history & tendency controls include which party the individual would have for congress and president in 2012, and an indicator for whether the individual always for the same party. Political interest and knowledge controls for the level of interest and knowledge the individual has in politics and current affairs. Robust standard errors clustered by counties.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE C.0.7
*THE MEDIATED EFFECT OF THE MeToo MOVEMENT, 2016 HOUSE
 ELECTIONS*

	All-party	Two-party	
	(1)	(2)	(3)
A. Gender dimension only			
Woman $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	0.015 (0.014)	0.018 (0.015)	−0.011 (0.017)
Man $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	0.008 (0.006)	0.009 (0.007)	0.016** (0.008)
B. Party dimension only			
Rep. $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	−0.015* (0.008)	−0.018** (0.008)	−0.012 (0.010)
Dem. $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	0.032*** (0.010)	0.040*** (0.009)	0.040*** (0.011)
C. Party & Gender			
Rep. woman $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	−0.039 (0.040)	−0.044 (0.039)	−0.057 (0.038)
Dem. woman $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	0.040** (0.017)	0.041** (0.016)	0.009 (0.020)
Rep. man $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	−0.014 (0.009)	−0.016* (0.009)	−0.011 (0.010)
Dem. man $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	0.031*** (0.011)	0.040*** (0.011)	0.049*** (0.013)
<i>Control variables</i>			
Candidate fixed effects	X	X	
District fixed effects			X
2008–12 Pres. election	X	X	X
County census demographics	X	X	X
Racial & gender voting	X	X	X
Main-party candidates only		X	X
N	7822	6055	6055

Notes—The dependent variable is the 2016 house candidate vote share at the district-county level. Tweet density is the (natural) log of MeToo tweets in 2018 divided by county population. Column (1) reports the results for the full all-party sample; column (2) reports the results for the main-party sample and uses two-party vote shares on both sides of the equation. Past electoral results include the 2008 and 2012 presidential Republican vote share. All controls are otherwise the same as in Table III. Robust standard errors in parentheses are clustered by candidates.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE C.0.8
CHANGE IN REPUBLICAN VOTE SHARE, DISTRICT-COUNTY LEVEL (LOG TWEETS)

	Change in <i>house</i> Republican vote share between 2016 and 2018				Falsification	
	Change in Republican all-party vote share		Change in Republican two-party vote share		Change in <i>presidential</i> Republican vote share between 2012 and 2016	
	(1)	(2)	(3)	(4)	(5)	(6)
Log tweet density $\times$ (Pres. 2016 Rep. vote share)	−0.0435*** (0.0042)	−0.0422*** (0.0041)	−0.0195*** (0.0044)	−0.0182*** (0.0043)	0.0002 (0.0018)	
Log tweet density $\times$ (Pres. 2012 Rep. vote share)						0.0002 (0.0019)
Log tweet density	0.2439 (0.1803)	0.2826* (0.1541)	0.1328 (0.1988)	0.1648 (0.1698)	−0.1029** (0.0498)	−0.1288 (0.1116)
Change in log(total House votes) 2016–18		−2.8934* (1.5798)		−2.9671* (1.6629)	0.0168 (0.1009)	−0.6764** (0.2895)
Change in log(total Pres. votes) 2012–16					−2.6034 (1.9449)	−1.3049 (1.6145)
<i>Control variables</i>						
District fixed effects	X	X	X	X	X	X
Past electoral controls	X	X	X	X	X	X
County census demographics	X	X	X	X	X	X
R^2	0.8970	0.9014	0.9070	0.9115	0.8768	0.9121
N	3102	3102	3102	3102	3102	3102

Notes—This Table replicates the regressions in Table VI, except that log tweets are used instead of log tweet density (log of tweets divided by county population). All specifications are otherwise the same. Robust standard errors in parentheses are clustered by the 388 U.S. House congressional districts.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE C.o.9
ADDITIONAL EFFECTS BY STATE AND DISTRICT

	Differences by State I == 1 if State has						Differences by Districts I == 1 if District has	
	Two Rep. senators	Split delegation	No senate elections	Battleground states	Rep. & Battleground states	Senate elections (split delegation)	Head-to-head bw. man & woman	Rep. districts & low margin
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Rep. woman $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	-0.031 (0.033)	-0.031 (0.033)	-0.019 (0.032)	-0.015 (0.036)	-0.016 (0.032)	-0.032 (0.032)	-0.083* (0.050)	-0.025 (0.032)
Dem. woman $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	0.050*** (0.017)	0.050*** (0.017)	0.055*** (0.015)	0.054** (0.023)	0.043** (0.018)	0.047*** (0.016)	0.115*** (0.036)	0.029** (0.014)
Rep. man $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	-0.026*** (0.010)	-0.026*** (0.010)	-0.021* (0.011)	-0.020 (0.013)	-0.019 (0.011)	-0.024** (0.009)	-0.020* (0.011)	-0.015* (0.009)
Dem. man $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	0.019 (0.016)	0.019 (0.016)	-0.004 (0.018)	0.012 (0.022)	0.013 (0.020)	0.018 (0.016)	0.014 (0.016)	0.005 (0.014)
<i>Additional differences by State/District</i>								
I $\times$ Rep. woman $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	0.026 (0.048)	0.026 (0.048)	-0.021 (0.040)	-0.090 (0.062)	-0.097 (0.067)	0.062 (0.056)	0.078 (0.058)	-0.033 (0.049)
I $\times$ Dem. woman $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	-0.012 (0.026)	-0.012 (0.026)	-0.021 (0.034)	-0.019 (0.025)	0.006 (0.025)	0.003 (0.027)	-0.087** (0.037)	0.077* (0.039)
I $\times$ Rep. man $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	0.005 (0.018)	0.005 (0.018)	-0.003 (0.018)	-0.005 (0.016)	-0.012 (0.016)	-0.014 (0.019)	-0.006 (0.016)	-0.052* (0.028)
I $\times$ Dem. man $\times$ (Log tweet density) $\times$ (Pres. 2016 Rep. vote share)	-0.013 (0.022)	-0.013 (0.022)	0.043 (0.027)	0.007 (0.026)	0.003 (0.025)	-0.008 (0.024)	-0.005 (0.034)	0.113** (0.050)
I $\times$ Rep. woman $\times$ (Log tweet density)	-0.099 (1.680)	-0.099 (1.680)	1.472 (1.174)	-2.844 (1.775)	-3.070 (1.870)	-2.380 (2.259)	-0.160 (1.510)	2.320* (1.223)
I $\times$ Dem. woman $\times$ (Log tweet density)	0.081 (0.561)	0.081 (0.561)	0.657 (0.672)	0.416 (0.556)	-0.046 (0.580)	-0.245 (0.552)	1.282* (0.662)	-1.340* (0.727)
I $\times$ Rep. man $\times$ (Log tweet density)	0.405 (0.477)	0.405 (0.477)	0.245 (0.496)	-0.416 (0.463)	0.304 (0.432)	0.850* (0.460)	0.043 (0.441)	0.932 (0.613)
I $\times$ Dem. man $\times$ (Log tweet density)	-0.246 (0.745)	-0.246 (0.745)	-1.287* (0.756)	0.598 (0.763)	0.349 (0.727)	-0.335 (0.822)	0.861 (1.060)	-2.016** (0.807)
<i>Control variables</i>								
Candidate fixed effects	X	X	X	X	X	X	X	X
2016 House & 2012–16 Pres. election	X	X	X	X	X	X	X	X
County census demographics	X	X	X	X	X	X	X	X
Racial & gender voting	X	X	X	X	X	X	X	X
Main-party candidates only	X	X	X	X	X	X	X	X
R ²	0.952	0.952	0.952	0.952	0.952	0.951	0.953	0.952
N	6185	6185	6185	6185	6185	6185	6185	6185

Notes—This Table replicates columns (3)–(5) of Table III, except that an additional interaction is entered into the model to capturing any differences of the MeToo effect by state or district. In column (1), the additional interaction is a dummy for states where both senators are Republican; in column (2), it is for states where the senate is split; in column (3), it is in states where there were no senate elections in 2018; in column (4), it is for battleground states defined as states with less than a 10% margin in the 2016 presidential elections; in column (5), it is for battleground states defined as states with less than a 10% margin in the 2016 presidential elections, and where the Republican candidate won; in column (6), it is for states with senate elections and where there is split delegation (one Democratic and one Republican senator); in column (7), it is for districts with a head-to-head between a woman and man candidate from the main parties in the 2018 House elections; and in column (8), it is for Republican districts where the winning margin is less than 10% in the 2016 House elections. All other controls are the same. Robust standard errors in parentheses are clustered by candidates.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

TABLE C.O.10
THE INTENSIVE AND EXTENSIVE MARGINS OF MeToo MOVEMENT

	This Table breaks down the tweet density into MeToo tweets per MeToo author and MeToo authors per population		
	All-party	Two-party	
	(1)	(2)	(3)
<i>Log MeToo tweets per MeToo twitter</i>			
Log tweets per twitter $\times$ (Pres. 2016 Rep. vote share) $\times$ (Rep. woman)	-0.142* (0.081)	-0.092 (0.077)	-0.135* (0.079)
Log tweets per twitter $\times$ (Pres. 2016 Rep. vote share) $\times$ (Dem. woman)	0.134*** (0.048)	0.102** (0.040)	0.050 (0.042)
Log tweets per twitter $\times$ (Pres. 2016 Rep. vote share) $\times$ (Rep. man)	-0.119*** (0.028)	-0.076*** (0.025)	-0.035 (0.026)
Log tweets per twitter $\times$ (Pres. 2016 Rep. vote share) $\times$ (Dem. man)	0.100*** (0.032)	0.064** (0.032)	0.037 (0.033)
<i>Log MeToo twitter density</i>			
Log twitter density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Rep. woman)	0.018 (0.031)	-0.006 (0.024)	0.010 (0.027)
Log twitter density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Dem. woman)	0.009 (0.016)	0.031** (0.015)	0.040** (0.018)
Log twitter density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Rep. man)	0.003 (0.011)	-0.006 (0.011)	-0.021 (0.014)
Log twitter density $\times$ (Pres. 2016 Rep. vote share) $\times$ (Dem. man)	-0.004 (0.018)	0.016 (0.019)	0.034* (0.020)
<i>Control variables</i>			
Candidate fixed effects	X	X	
District fixed effects			X
2016 House & 2012–16 Pres. election	X	X	X
County census demographics	X	X	X
Racial & gender voting	X	X	X
Main-party candidates only		X	X
R^2	0.975	0.952	0.886
N	8634	6234	6234

Notes—This Table replicates columns (3)–(5) of Table III, except that in this Table the log tweet density measure is decomposed into a log MeToo tweets per MeToo author and a log MeToo author density (log MeToo author at the county level divided by county population). All other controls are the same. Robust standard errors in parentheses are clustered by candidates.

*** Significant at the 1 per cent level.

** Significant at the 5 per cent level.

* Significant at the 10 per cent level.

Bibliography

- Abrajano, M. A. and R. M. Alvarez (2005). A Natural Experiment of Race-Based and Issue Voting: The 2001 City of Los Angeles Elections. *Political Research Quarterly* 58(2), 203–218. 118, 129
- Acemoglu, D., T. A. Hassan, and A. Tahoun (2018). The power of the street: Evidence from Egypt’s arab spring. *Review of Financial Studies* 31(1), 1–42. 118, 142
- Adams, R. B. and D. Ferreira (2009). Women in the Boardroom and Their Impact on Governance and Performance. *Journal of Financial Economics* 94(2), 291–309. 65, 66, 72, 84, 85, 86, 91, 92
- Adams, R. B. and P. Funk (2011). Beyond the Glass Ceiling: Does Gender Matter? *Management Science* 58(2), 219–235. 92
- Adena, M., R. Enikolopov, M. Petrova, V. Santarosa, and E. Zhuravskaya (2015). Radio and the Rise of The Nazis in Prewar Germany. *The Quarterly Journal of Economics* 130(4), 1885–1939. 5, 117
- Ahern, K. R. and A. K. Dittmar (2012). The Changing of the Boards: The Impact on Firm Valuation of Mandated Female Board Representation. *The Quarterly Journal of Economics* 127(1), 137–197. 65, 66, 84, 85, 90
- Altonji, J., T. Elder, and C. Taber (2005). Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools. *Journal of Political Economy* 113(1), 151–184. 3, 22
- Anderson, S. P. and J. McLaren (2012). Media Mergers and Media Bias With Rational Consumers. *Journal of the European Economic Association* 10(4), 831–859. 8
- Ang, J. S. and D. K. Ding (2006). Government ownership and the performance of government-linked companies: The case of Singapore. *Journal of Multinational Financial Management* 16(1), 64–88. 68, 69, 75, 102
- Angrist, J. D. and J.-S. Pischke (2009). *Mostly Harmless Econometrics: An Empiricist’s Companion*. Princeton : Princeton University Press, c2009. 3, 21, 25, 46, 78

- Arellano, M. and O. Bover (1995). Another look at the instrumental variable estimation of error-components models. *Journal of Econometrics* 68(1), 29–51. 84
- Arnaboldi, F., B. Casu, A. Gallo, E. Kalotychou, and A. Sarkisyan (2021). Gender diversity and bank misconduct. *Journal of Corporate Finance*, 101834. 91
- BBC News (2013). Singapore profile - Media. 8
- Bernhardt, D., S. Krasa, and M. Polborn (2008). Political Polarization and the Electoral Effects of Media Bias. *Journal of Public Economics* 92(5), 1092–1104. 4, 37
- Berry, F. C. (1967). A Study of Accuracy in Local News Stories of Three Dailies. *Journalism Quarterly* 44(3), 482–490. 12, 37
- Besley, T. and R. Burgess (2002). The Political Economy of Government Responsiveness: Theory and Evidence from India. *The Quarterly Journal of Economics* 117(4), 1415–1451. 5
- Besley, T. and A. Prat (2006). Handcuffs for the Grabbing Hand? Media Capture and Government Accountability. *American Economic Review* 96(3), 720–736. 4, 36
- Blei, D., A. Y. Ng, and M. I. Jordan (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research* 3, 993–1022. 13
- Blundell, R. and S. Bond (1998). Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics* 87(1), 115–143. 84
- Boas, T. C. and F. D. Hidalgo (2011). Controlling the Airwaves: Incumbency Advantage and Community Radio in Brazil. *American Journal of Political Science* 55(4), 869–885. 5, 117, 118
- Cadsby, C. B. and E. Maynes (1998). Gender and Free Riding in a Threshold Public Goods Game: Experimental Evidence. *Journal of Economic Behavior & Organization* 34(4), 603–620. 72
- Campante, F., R. Durante, and F. Sobbrío (2017). Politics 2.0: The Multifaceted Effect of Broadband Internet on Political Participation. *Journal of the European Economic Association* 16(4), 1094–1136. 118, 138, 142
- Carter, D. A., F. D’Souza, B. J. Simkins, and W. G. Simpson (2010, sep). The Gender and Ethnic Diversity of US Boards and Board Committees and Firm Financial Performance. *Corporate Governance: An International Review* 18(5), 396–414. 84
- Carter, D. A., B. J. Simkins, and W. G. Simpson (2003, feb). Corporate Governance, Board Diversity, and Firm Value. *Financial Review* 38(1), 33–53. 84

- Carter, N. M. and H. M. Wagner (2011). Women Board Directors: A Comparison of Economic Results. pp. 1–3. 64, 65, 91, 92
- Center for American Women and Politics (2018). Women Candidates for Congress 1974-2018. Technical report. 115, 125
- Chang, J., S. Gerrish, C. Wang, and D. M. Blei (2009). Reading Tea Leaves: How Humans Interpret Topic Models. *Advances in Neural Information Processing Systems* 22, 288—296. 14
- Chattopadhyay, R. and E. Duflo (2004). Women as Policy Makers: Evidence from a Randomized Policy Experiment in India. *Econometrica* 72(5), 1409–1443. 92, 147
- Chiang, C.-F. and B. Knight (2011). Media Bias and Influence: Evidence from Newspaper Endorsements. *The Review of Economic Studies* 78(3), 795–820. 5
- Corneo, G. (2006). Media Capture in a Democracy: The Role of Wealth Concentration. *Journal of Public Economics* 90(1), 37–58. 5, 36
- Croson, R. and U. Gneezy (2009). Gender Differences in Preferences. *Journal of Economic Literature* 47(2), 448–474. 65
- D'Alessio, D. (2003). An Experimental Examination of Readers' Perceptions of Media Bias. *Journalism & Mass Communication Quarterly* 80(2), 282–294. 35
- Darmadi, S. (2013, jun). Do women in top management affect firm performance? Evidence from Indonesia. *Corporate Governance: The international journal of business in society* 13(3), 288–304. 84
- Dawson, J., R. Kersley, and S. Natella (2012). Gender diversity and corporate leadership. *Credit Suisse* (August). 64
- Dawson, J., R. Kersley, and S. Natella (2016). The CS Gender 3000: The Reward for Change. *Credit Suisse AG* (September), 52. 64
- Deaves, R., E. Luders, and G. Y. Luo (2009). An Experimental Test of the Impact of Overconfidence and Gender on Trading Activity. *Review of Finance* 13(3), 555–575. 65, 92
- Deckman, M. (2018). #MeToo and the Midterms. *Gender Watch 2018*. 125
- DellaVigna, S. and E. Kaplan (2007). The Fox News Effect: Media Bias and Voting. *Quarterly Journal of Economics* 122(3), 1187–1234. 5, 117, 118, 138
- Democracy Fund Voter Study Group (2018). Views of the Electorate Research Survey. 143
- Devillard, S., S. Sancier, C. Werner, I. Maller, and C. Kossoff (2013). Gender Diversity in Top Management: Moving Corporate Culture, Moving Boundaries. 64

- Dollar, D., R. Fisman, and R. Gatti (2001). Are women really the “fairer” sex? Corruption and women in government. *Journal of Economic Behavior & Organization* 46(4), 423–429. 92
- Durante, R. and B. Knight (2012). Partisan Control, Media Bias, and Viewer Responses: Evidence from Berlusconi’s Italy. *Journal of the European Economic Association* 10(3), 451–481. 4, 8
- Eckel, C. C. and P. J. Grossman (1998). Are Women Less Selfish Than Men?: Evidence From Dictator Experiments. *The Economic Journal* 108(448), 726–735. 72
- Edlund, L. and R. Pande (2002). Why Have Women Become Left-Wing? The Political Gender Gap and the Decline in Marriage. *The Quarterly Journal of Economics* 117(3), 917–961. 92, 129
- Eisensee, T. and D. Strömberg (2007). News Droughts, News Floods, and U. S. Disaster Relief. *The Quarterly Journal of Economics* 122(2), 693–728. 5
- Enikolopov, R., M. Petrova, and E. Zhuravskaya (2011). Media and Political Persuasion: Evidence from Russia. *American Economic Review* 101(7), 3253–3285. 4, 5, 117, 118
- Fama, E. F. and M. C. Jensen (1983, jun). Separation of Ownership and Control. *The Journal of Law and Economics* 26(2), 301–325. 72, 91
- Feinberg, A., R. Branton, and V. Martinez-Ebers (2019). The Trump Effect: How 2016 Campaign Rallies Explain Spikes in Hate. 118, 142
- Ferraz, C. and F. Finan (2008). Exposing Corrupt Politicians: The Effects of Brazil’s Publicly Released Audits on Electoral Outcomes. *The Quarterly Journal of Economics* 123(2), 703–745. 4, 117
- Ferris, S. P., M. Jagannathan, and A. C. Pritchard (2003, jun). Too Busy to Mind the Business? Monitoring by Directors with Multiple Board Appointments. *The Journal of Finance* 58(3), 1087–1111. 79
- Fischer, A. J. (1996). A further experimental study of expressive voting. *Public Choice* 88(1), 171–184. 119
- Flanagan, F. X. (2018). Race, Gender, and Juries: Evidence from North Carolina. *The Journal of Law and Economics* 61(2), 189–214. 118, 129
- Fujiwara, T., K. Müller, and C. Schwarz (2020). The Effect of Social Media on Elections: Evidence from the United States. *SSRN Electronic Journal*. 117, 118
- Gender Diversity Taskforce (2014). Gender diversity on boards: A Business Imperative. 93
- Gentzkow, M. (2006). Television and Voter Turnout. *The Quarterly Journal of Economics* (3), 931. 5, 117, 118

- Gentzkow, M. and J. M. Shapiro (2006). Media Bias and Reputation. *Journal of Political Economy* 114(2), 280–316. 37
- Gentzkow, M. and J. M. Shapiro (2010). What Drives Media Slant? Evidence From U.S. Daily Newspapers. *Econometrica* 78(1), 35–71. 1, 4, 37, 118
- George, C. (2012). *Freedom From the Press: Journalism and State Power in Singapore*. Singapore : NUS Press, c2012. 4, 8, 36
- Gerber, A. S., J. G. Gimpel, D. P. Green, and D. 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), 135–150. 5
- Gibson, C., S. Davenport, T. Fowler, C. B. Harris, M. Prudhomme, S. Whiting, and S. Simmons-Horton (2019). Understanding the 2017 “Me Too” Movement’s Timing. *Humanity & Society* 43(2), 217–224. 123
- Gibson, R. and D. Zillmann (1993). The Impact of Quotation in News Reports on Issue Perception. *Journalism Quarterly* 70(4), 793–800. 35
- Groeling, T. and S. Kernell (1998). Is Network News Coverage of the President Biased? *The Journal of Politics* 60(4), 1063–1087. 37
- Groseclose, T. and J. Milyo (2005). A Measure of Media Bias. *Quarterly Journal of Economics* 120(4), 1191–1237. 1, 4, 118
- Haslam, S. A. and M. K. Ryan (2008). The road to the glass cliff: Differences in the perceived suitability of men and women for leadership positions in succeeding and failing organizations. *The Leadership Quarterly* 19(5), 530–546. 66
- Herron, M. C. and J. S. Sekhon (2005). Black Candidates and Black Voters: Assessing the Impact of Candidate Race on Uncounted Vote Rates. *The Journal of Politics* 67(1), 154–177. 129
- Hillman, A. L. (2010). Expressive behavior in economics and politics. *European Journal of Political Economy* 26(4), 403–418. 119
- Ho, D. E. and K. M. Quinn (2008). Measuring Explicit Political Positions of Media. *Quarterly Journal of Political Science* 3(4), 353–377. 4
- Holli, A. M. and H. Wass (2010). Gender-based voting in the parliamentary elections of 2007 in Finland. *European Journal of Political Research* 49(5), 598–630. 118, 129
- Holt, C. A. and S. K. Laury (2002). Risk Aversion and Incentive Effects. *American Economic Review* 92(5), 1644–1655. 65, 92
- Hotelling, H. (1929). Stability in Competition. *The Economic Journal* 39(153), 41. 37

- Jurafsky, D. and J. H. Martin (2000). *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition*. Prentice Hall series in artificial intelligence. Upper Saddle River, N.J. : Prentice Hall, 2000. 9
- Kanter, R. M. (1977, mar). Some Effects of Proportions on Group Life: Skewed Sex Ratios and Responses to Token Women. *American Journal of Sociology* 82(5), 965–990. 65
- Kaplan, S. N. (1995, mar). Corporate governance and incentives in German companies: Evidence from top executive turnover and firm performance. *European Financial Management* 1(1), 23–36. 79
- Konrad, A. M., V. Kramer, and S. Erkut (2008). Critical mass: The impact of three or more women on corporate boards. *Organizational Dynamics* 37(2), 145–164. 65, 91, 92
- Larcinese, V., R. Puglisi, and J. M. Snyder (2011). Partisan Bias in Economic News: Evidence on the Agenda-setting Behavior of U.S. Newspapers. *Journal of Public Economics* 95(9), 1178–1189. 1, 4, 37, 118
- Larreboure, M. and F. Gonzalez (2021). The Impact of the Women’s March on the U.S. House Election. 118, 142
- Larreguy, H. A., J. Marshall, and J. 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*, 1–57. 5, 117
- Lehrer, A. (1989). Between Quotation Marks. *Journalism Quarterly* 66(4), 902–941. 13
- Lilley, M. and B. Wheaton (2019). Trump Rallies and Hate Crimes: A Comment on Feinberg et al. (2019). 118, 135, 142
- Lim, C. S. H., J. M. Snyder, and D. Strömberg (2015). The Judge, the Politician, and the Press: Newspaper Coverage and Criminal Sentencing across Electoral Systems. *American Economic Journal: Applied Economics* 7(4), 103–135. 5, 117
- Lott, J. R. and K. A. Hassett (2014). Is Newspaper Coverage of Economic Events Politically Biased? *Public Choice* 160(1), 65–108. 4
- Marshall, H. (1977, mar). Newspaper Accuracy in Tucson. *Journalism Quarterly* 54(1), 165–169. 12
- Matsubayashi, T. and M. Ueda (2011). Political Knowledge and the Use of Candidate Race as a Voting Cue. *American Politics Research* 39(2), 380–413. 118
- McCarthy, P. M. and S. Jarvis (2010). MTLD, vocd-D, and HD-D: A validation study of sophisticated approaches to lexical diversity assessment. *Behavior Research Methods* 42(2), 381–392. 32

- McMillan, J. and P. Zoido (2004). How to Subvert Democracy: Montesinos in Peru. *Journal of Economic Perspectives* 18(4), 69–92. 36
- Miner, L. (2015). The Unintended Consequences of Internet Diffusion: Evidence from Malaysia. *Journal of Public Economics* 132, 66–78. 4, 5, 117, 118
- MIT Election Data and Science Lab (2018a). County Presidential Election Returns 2000-2016.
- MIT Election Data and Science Lab (2018b). U.S. House of Representatives Precinct-Level Returns 2016.
- Mullainathan, S. and A. Shleifer (2005). The Market for News. *The American Economic Review* 95(4), 1031. 37
- Mullainathan, S. and J. Spiess (2017). Machine Learning: An Applied Econometric Approach. *Journal of Economic Perspectives* 31(2), 87–106. 4
- North, A. (2019). Democrats’ sweeping new anti-harassment bill, explained. 124
- Oberholzer-Gee, F. and J. Waldfogel (2009). Media Markets and Localism: Does Local News en Español Boost Hispanic Voter Turnout? *American Economic Review* 99(5), 2120–2128. 117, 118
- Ortmann, A. and L. K. Tichy (1999). Gender Differences in the Laboratory: Evidence From Prisoner’s Dilemma Games. *Journal of Economic Behavior & Organization* 39(3), 327–339. 72
- Oster, E. (2017). Unobservable Selection and Coefficient Stability: Theory and Evidence. *Journal of Business & Economic Statistics*, 1–18. 3, 22
- Oswald, A. J. and N. Powdthavee (2010). Daughters and Left-Wing Voting. *The Review of Economics and Statistics* 92(2), 213–227. 92, 129
- Pande, R. (2003). Can Mandated Political Representation Increase Policy Influence for Disadvantaged Minorities? Theory and Evidence from India. *American Economic Review* 93(4), 1132–1151. 147
- Peaker, H. (2018). This was a #Metoo election and those who say differently ignore the ‘Pink Wave’ at their peril. *The Telegraph*. 125
- Petrocik, J. R. (1996). Issue Ownership in Presidential Elections, with a 1980 Case Study. *American Journal of Political Science* 40(3), 825–850. 19
- Petrova, M. (2008). Inequality and Media Capture. *Journal of Public Economics* 92(1), 183–212. 36
- Pletzer, J. L., R. Nikolova, K. K. Kedzior, and S. C. Voelpel (2015). Does Gender Matter? Female Representation on Corporate Boards and

- Firm Financial Performance - A Meta-Analysis. *PLOS ONE* 10(6), e0130005. 66, 90
- Post, C. and K. Byron (2014, nov). Women on Boards and Firm Financial Performance: A Meta-Analysis. *Academy of Management Journal* 58(5), 1546–1571. 66, 90
- Puglisi, R. (2011). Being The New York Times: The Political Behaviour of a Newspaper. *B.E. Journal of Economic Analysis and Policy* 11(1). 1, 4, 19
- Puglisi, R. and J. M. Snyder (2011). Newspaper Coverage of Political Scandals. *The Journal of Politics* 73(3), 931–950. 1, 4
- Qian, N. and D. 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), 463–499. 1, 4
- Qin, B., D. Strömberg, and Y. Wu (2017). Why Does China Allow Freer Social Media? Protests versus Surveillance and Propaganda. *Journal of Economic Perspectives* 31(1), 117–140. 4
- Qin, B., D. Strömberg, and Y. Wu (2018). Media Bias in China. *American Economic Review* 108(9), 2442–2476. 1, 4
- Rajan, A. (2019). Twitter to ban all political advertising. *BBC*. 147
- Ramírez, C. D. and L. H. Tan (2004). Singapore Inc. Versus the Private Sector: Are Government-Linked Companies Different? *IMF Staff Papers* 51(3), 510–528. 68, 69, 101
- Roodman, D. (2009, feb). A Note on the Theme of Too Many Instruments. *Oxford Bulletin of Economics and Statistics* 71(1), 135–158. 85
- Ryan, M. K. and S. A. Haslam (2005, jun). The Glass Cliff: Evidence that Women are Over-Represented in Precarious Leadership Positions. *British Journal of Management* 16(2), 81–90. 87
- Salmond, R. (2006). Proportional Representation and Female Parliamentarians. *Legislative Studies Quarterly* 31(2), 175–204. 77
- Scholtz, H. and S. Kieviet (2018). The Influence Of Board Diversity On Company Performance Of South African Companies. *Journal of African Business* 19(1), 105–123. 66
- Schwartz-Ziv, M. (2017, apr). Gender and Board Activeness: The Role of a Critical Mass. *Journal of Financial and Quantitative Analysis* 52(2), 751–780. 65, 91, 92
- Sila, V., A. Gonzalez, and J. Hagendorff (2016). Women on board: Does boardroom gender diversity affect firm risk? *Journal of Corporate Finance* 36, 26–53. 66, 86
- Singapore Statutes Online (1974). Newspaper and Printing Presses Act.

- Singer, J. M. (2019). Judicial Recall and Retention in the #MeToo Era. *Court Review* 55(1), 36–43. 124
- Snyder, J. M. and D. Strömberg (2010). Press Coverage and Political Accountability. *Journal of Political Economy* 118(2), 355–408. 5
- Stephens-Davidowitz, S. (2014). The cost of racial animus on a black candidate: Evidence using Google search data. *Journal of Public Economics* 118, 26–40. 119, 129
- Stevenson, B. (2010). Beyond the Classroom: Using Title IX to Measure the Return to High School Sports. *The Review of Economics and Statistics* 92(2), 284–301. 89
- Strömberg, D. (2004). Radio's Impact on Public Spending. *Quarterly Journal of Economics* 119(1), 189–221. 5
- Studlar, D. T. and S. Welch (1991, jun). Does District Magnitude Matter? Women Candidates in London Local Elections. *Western Political Quarterly* 44(2), 457–466. 77
- Sutter, D. (2012). Is the Media Liberal? An Indirect Test Using News Magazine Circulation. *Applied Economics* 44(27), 3521–3532. 4
- Tan, N. (2014). Institutionalized Succession and Hegemonic Party Cohesion in Singapore. In A. Hicken and E. M. Kuhonta (Eds.), *Party System Institutionalization in Asia: Democracies, Autocracies, and the Shadows of the Past*, pp. 49–73. Cambridge: Cambridge University Press. 22
- Tan, N. (2015). Party Quotas and rising women politicians in Singapore. *Politics and Gender* 11(1), 196–207. 77
- Tan, N. (2016). Why Are Gender Reforms Adopted in Singapore? Party Pragmatism and Electoral Incentives. *Pacific Affairs* 89(2), 369–393. 76, 77
- Tippett, E. C. (2018). The Legal Implications of the MeToo Movement. *Minnesota Law Review* 103(1), 229–302. 124
- Torruella, J. and R. Capsada (2013). Lexical Statistics and Tipological Structures: A Measure of Lexical Richness. *Procedia - Social and Behavioral Sciences* 95, 447–454. 32
- Tyran, J.-R. (2004). Voting when money and morals conflict: an experimental test of expressive voting. *Journal of Public Economics* 88(7), 1645–1664. 119
- Walsh, J. (2018). ‘Kavanaugh’s Revenge’: Every Democratic senator in a competitive midterm race who voted against Brett Kavanaugh lost. 125
- Ye, J., S. Han, Y. Hu, B. Coskun, M. Liu, H. Qin, and S. Skiena (2017). Nationality Classification using Name Embeddings. *CIKM*. 120, 148, 150