



A University of Sussex PhD thesis

Available online via Sussex Research Online:

<http://sro.sussex.ac.uk/>

This thesis is protected by copyright which belongs to the author.

This thesis cannot be reproduced or quoted extensively from without first obtaining permission in writing from the Author

The content must not be changed in any way or sold commercially in any format or medium without the formal permission of the Author

When referring to this work, full bibliographic details including the author, title, awarding institution and date of the thesis must be given

Please visit Sussex Research Online for more information and further details

Data Science in Political Economy

Daniele Guariso

Submitted for the degree of
Doctor of Philosophy in Economics



University of Sussex Business School

University of Sussex

January 2021

Statement

I hereby declare that this thesis has not been and will not be, submitted in whole or in part to another University for the award of any other degree.

I also hereby declare that chapter 2 is single-authored. The research was carried out under the guidance of my supervisors Professor Richard Dickens and Dr Vikram Pathania. I am also grateful to Professor David Weir and Simon Wibberley that introduced me to *Method52*, the software developed by the Text Analysis Group at the University of Sussex that I used for the data collection. Chapter 3 is joint work together with Dr Omar A. Guerrero (The Alan Turing Institute, UCL) and Dr Ulrich Matter (University of St. Gallen). Through fruitful discussions with them, I elaborated on the initial research question, conducted the data collection and preparation, the statistical analysis and wrote the manuscript. They conceived the original methodology that I further developed and helped me with framing the narrative through edits to the paper. Chapter 4 is a joint work together with Dr Omar A. Guerrero (The Alan Turing Institute, UCL). He granted me access to the raw corpus of Mexican newspapers and provided me with valuable insights into the institutional context. He also helped me building the narrative by editing the chapter. Thanks to constructive exchanges with him, I developed the research idea and empirical strategy, preprocessed the data and conducted the statistical analysis, together with taking the lead in writing the paper.

Signature:

Daniele Guariso

UNIVERSITY OF SUSSEX

DANIELE GUARISO

DOCTOR OF PHILOSOPHY IN ECONOMICS

DATA SCIENCE IN POLITICAL ECONOMYSUMMARY

This thesis provides three applications of Data Science methodologies in Political Economy, combining them with established techniques in the literature. Its aim is to show how important is to keep an open dialogue with different disciplines to broaden the standard toolkit of an empirical economist and be able to tackle research questions in a novel way.

In chapter 2, I study the effects of terrorist attacks on British politicians' immigration rhetoric on social media. I scrape the Members of the Parliament's Twitter accounts and identify the immigration-related Tweets, which are then leveraged to frame a natural experiment. Looking at the 2017 Manchester bombing as my main event study, I find a substantial decrease in the expected number of immigration-related Tweets after the incident. I hypothesise that this "muting effect" results from the risk-averse attitude of politicians during the election campaign.

Chapter 3 explores strategic voting in the United Nations General Assembly (UNGA). We first predict the expected behaviour of country representatives in the UNGA. Next, we construct a network that describes the structure of the deviations underlying the observed votes. The graph is used to compute a *Reciprocity Index*. Through this statistic, we find that deviations from the expected votes are systematically *not* reciprocated. The conclusions are consistent with a narrative of vote buying and question the unweighted voting system of the Assembly.

In chapter 4, we investigate government's information processing and its implications for policy responses. To study how the Mexican government processes a specific signal (opinion pieces from newspapers), we devise a *News Index* that creates a link between informational inputs and policy outputs. We find that changes in the index are associated with policy overreactions. The findings are further assessed through a natural experiment. Overall, the results are consistent with the dominant theory of disproportionate information processing in government's decisions.

Acknowledgements

I guess, dear reader, that most of the time, in the acknowledgments of a PhD thesis, you will probably find something like “this thesis comes at the end of a wonderful journey” or “this PhD has been such a gratifying and enriching experience” and other blah, blah, blah.

That is, without doubt, all true also in my case.

However, on second thought, I think that I will opt for something a bit different. The truth is I would have really appreciated someone telling me at the very beginning: “the PhD has nothing to offer but blood, toil, tears and sweat”. So, instead of the pretty common “very rewarding journey” I will go for “incredibly tough fight”. The PhD has taught me far more than the stuff you can read in the textbooks, and I like to think that it did it the hard way.

I mean, this thesis went through an uncountable number of mental breakdowns, multiple (and self-diagnosed) forms of impostor syndrome, hundreds of hours running meaningless codes and realising at 3am that the whole problem was a missing comma at line 231, the Brexit negotiations, four different UK governments (if I counted them right), the US Presidency of Donald Trump and several other catastrophes, including a global pandemic.

Yes, dear reader, that is right, even a global pandemic.

Of course, I am just kidding, it has not been that bad, but first, I always like a bit of drama, then the term “fight” gives that epic touch proper of every great narration.

Anyway, if I managed to survive to this fight, it is just thanks to the people that I have met and that supported me during these four years.

First, I am extremely grateful to Omar Guerrero, who has been a real mentor to me, especially in the second half of my PhD. He has been always ready to share his knowledge, discuss and provide insightful feedback, in a thorough and supportive way. He taught me the importance of keeping a multidisciplinary and heterodox approach as a social scientist. His genuine enthusiasm for scientific research has been a true source of inspiration for me.

Needless to say, I am indebted to my supervisors, Richard Dickens and Vikram Pathania. Their advice and positive attitude saved me from despair more than once. They managed to be always very approachable and patiently guide me during these years, despite my poor command of this language

and the fact that I am an incredibly (and I mean incredibly) slow learner. Most importantly, they gave me the opportunity to find my own way in that twisted labyrinth that is the PhD process. I also want to thank Barry Reilly, not just because he managed to teach me Econometrics (which means that he is a truly amazing teacher), but also because he was the first one who really encouraged me to apply for the PhD programme.

I am grateful for the financial support provided by the University of Sussex Business School, that allowed me to pursue this degree. Then, of course, my appreciation and gratitude go to the Sussex community, the Brighton people and my PhD colleagues: Amrita, Marco, Mattia, Maika, Kai, Inma, Olive, Guillermo, Stella, Fjolla, Erendira, Juan Manuel, Margherita, Anuja, Cecilia, Eva-Maria, Michael, Sweta, Nick, Marta, Edgar, Antonia, Shilan, Farai, Mohamed, Diego, Lucy, Jonathan, Eugenia, Egidio, Hector, Jorge, Rashaad, Nihar, Manuel, Monika, Wiktoria, Chiara, Maggie, Bridget, Pavel, Melina and Irene. A special mention goes to the members of the Intelligencija: Filippo, Andrea, Pier, Nicolò, Bernardo and Emanuele, with whom I shared the thrilling excitement and profound disappointment that constitute the daily life of a researcher, together with deep philosophical and existential questions fuelled by more than a pint. I am also very grateful to the vibrant community of The Alan Turing Institute, that hosted me for almost a year. There I had the opportunity to widen my perspective within a truly interdisciplinary and intellectually stimulating environment, which significantly contributed to my development as a researcher.

Then, I would like to thank all my friends in Italy, especially the *5+1* crew: Sergio, Teo, Fede, Pakke and Ciro. We might have chosen different paths guys, but we are still there for each other.

However, nothing of this would have been possible without the support and love of my family. My grandmother Fausta, who did not have the chance to see the end of it, but I am sure that she is celebrating with me right now. My brother Gabriele and my sister Claudia with her family, William, Emanuele and Irene.

Finally, I would like to dedicate this PhD thesis to my parents, Tiziana and Giorgio.

Vi ringrazio non solo per avermi dato la possibilità di studiare, ma per il vostro incondizionato affetto e supporto durante questi anni. Per aver trovato il tempo di ascoltare le mie lamentele e frustrazioni. Per aver sempre creduto nelle mie capacità. Per avermi trasmesso la gioia della scoperta, la brama di sapere e l'importanza di vivere intensamente le proprie passioni.

Siete e rimarrete per sempre i miei più grandi insegnanti.

Contents

List of Tables	ix
List of Figures	xiii
1 Introduction	1
2 Terrorist Attacks and Immigration Rhetoric: A Natural Experiment on British MPs	5
2.1 Introduction	5
2.2 Literature Review	8
2.3 Methodology	11
2.3.1 Construction of the Dataset	12
2.3.2 Statistical Model	14
2.3.3 Natural Experiment Setting	15
2.4 Data Description	17
2.4.1 Twitter Data	17
2.4.2 Data at the Constituency Level	17
2.4.3 Tweeting Trends	18

2.5	Analysis	21
2.5.1	Baseline Results	21
2.5.2	Heterogeneity	25
2.5.3	Robustness Checks	34
2.6	Discussion and Conclusion	37
3	Strategic Voting in the United Nations General Assembly	42
3.1	Introduction	42
3.2	Literature Review	45
3.3	Institutional Setting and Data Description	48
3.3.1	The United Nations General Assembly	48
3.3.2	Data Description	49
3.4	Methodology	51
3.4.1	Empirical Strategy	51
3.4.2	Estimation of the <i>Reciprocity Index</i>	53
3.5	Predicting Voting Behaviour	57
3.5.1	Estimation of Latent Topics	57
3.5.2	Topic-based Model vs. Traditional Approaches	62
3.5.3	Alternative Classifiers	64
3.6	Analysis of Strategic Voting in the UNGA	66
3.7	Implications of Anti-reciprocal Deviations	67
3.7.1	The Role of UNSC Permanent Members	68
3.7.2	Generalised Inequality in Resources	70

3.8	Conclusion	79
4	Information Processing and Policy Overreactions: Evidence from Mexico	82
4.1	Introduction	82
4.2	Literature Review	84
4.3	Institutional Setting and Data Description	87
4.3.1	Institutional Framework	87
4.3.2	The Importance of Opinion Journalism	88
4.3.3	Data on Public Expenditure	89
4.3.4	Data on Opinion Columns	90
4.4	Government's Information Processing and Budget Allocation	91
4.4.1	Parallel and Serial Processing	91
4.4.2	Building a <i>News Index</i> : Methodology	93
4.4.3	<i>News Index</i> : Analysis and Results	95
4.5	News and Budget Allocation: A Natural Experiment	101
4.5.1	The Setting	101
4.5.2	Analysis and Results	105
4.5.3	Testing Alternative Hypotheses	107
4.6	Conclusion	110
5	Conclusion	113
	Bibliography	117
	Appendix A Appendix to Chapter 2	129

A.1	List of Words for Boolean Search	129
A.2	Examples of Tweets	130
A.3	Additional Descriptive Statistics	133
Appendix B Appendix to Chapter 3		135
B.1	Web Pages in the UN Digital Library	135
B.2	Example of UNGA Resolution	137
B.3	Kolmogorov-Smirnov Tests	138
B.4	Countries with Highest Degree by Session	139
B.5	Data Coverage over Time	141
B.6	Bootstrapping Procedure	142
B.7	Robustness Checks: <i>Reciprocity Index</i>	143
Appendix C Appendix to Chapter 4		148
C.1	Branches: Descriptive Statistics	148
C.2	Textual Data on Public Expenditures	150
C.3	Additional Robustness Checks: News Index	152
C.4	News Index: Descriptive Statistics	153
C.5	Additional Robustness Checks: Natural Experiment	154
C.6	Inferred Topics Used in the Analysis	155

List of Tables

2.1	Classifier Performance	13
2.2	Descriptive Statistics (MP/Constituency Level) 16/02/2017-01/06/2017	18
2.3	Randomness in Tweeting Behaviour	20
2.4	Effect on Total Number of Tweets (Westminster Attack)	22
2.5	Effect on Total Number of Tweets (Manchester Attack)	22
2.6	Effect on Total Number of Tweets (Manchester Attack, All Obs.)	22
2.7	Effect on Total Number of Immigration-related Tweets (Westminster Attack)	24
2.8	Effect on Total Number of Immigration-related Tweets (Manchester Attack)	24
2.9	Effect on Total Number of Immigration-related Tweets (Manchester Attack, All Obs.)	25
2.10	One-sided Z-test $H_0 : \beta_{Immigration} \leq \beta_{Total}$	25
2.11	Heterogeneity by Majority Share	26
2.12	One-sided Z-test $H_0 : \beta_{Subsample} \leq \beta_{FullSample}$	27
2.13	Heterogeneity by Betting Odds	27
2.14	One-sided Z-test $H_0 : \beta_{Subsample} \leq \beta_{FullSample}$	27
2.15	Heterogeneity by Share of UK-Born People	28
2.16	Wald Test for Joint Significance (UK-Born)	28

2.17	Heterogeneity by Share of White British People	29
2.18	Wald Test for Joint Significance (White British)	29
2.19	Heterogeneity by Level of Unemployment	30
2.20	Wald Test for Joint Significance (Unemployment)	30
2.21	Heterogeneity by Incumbency	31
2.22	Wald Test for Joint Significance (Incumbency)	31
2.23	Heterogeneity by EU Referendum Results	32
2.24	Wald Test for Joint Significance (EU Referendum)	32
2.25	Leave Share and Immigration-related Tweets During Election	33
2.26	Robustness Check 1: NB2 with Total Tweets as Covariate	34
2.27	Robustness Check 2: Poisson RE with Total Tweets as Covariate	35
2.28	Robustness Check 3: London's Constituencies Excluded	35
2.29	Robustness Check 4: Probability of Immigration-related Tweets	35
2.30	Robustness Check 5: MPs not Standing in 2017 Elections Excluded	36
2.31	Robustness Check 6: Days of Suspended Campaigning Activity Excluded	37
2.32	Attention to Immigration-related Tweets	38
3.1	Descriptive Statistics	49
3.2	Descriptive Statistics for Labels	58
3.3	Labels Frequency (Top 12)	59
3.4	Model Comparison (I) with Five Fold Cross-validation	64
3.5	Model Comparison (II) with Five Fold Cross-validation	65
3.6	Deviations' Covariates	78

4.1	Absolute Difference between Expenditure Approved and Paid for Selected Branches . . .	90
4.2	Distribution of Articles across Newspapers and their Average Circulation	91
4.3	Coverage across the Newspapers	92
4.4	Parallel and Serial Information Processing	97
4.5	Parallel and Serial Information Processing Subsample of Opinion Pieces	98
4.6	Parallel and Serial Information Processing Non-negative Changes	99
4.7	Parallel and Serial Information Processing (Only <i>Reforma</i>)	100
4.8	Parallel and Serial Information Processing Non-negative Changes (Only <i>Reforma</i>) . . .	101
4.9	Difference-in-Differences Design	105
4.10	Natural Experiment	106
4.11	Natural Experiment Alternative Hypothesis	110
A.1	MP Distribution by Political Party	133
A.2	Comparison Demographic Characteristics	134
A.3	Reasons for Exclusion	134
B.1	Kolmogorov-Smirnov Tests for Degree Distributions	138
B.2	Countries with Highest Outdegree and Indegree by Session (31 st to 50 th)	139
B.3	Countries with Highest Outdegree and Indegree by Session (51 st to 72 nd)	140
C.1	Descriptive Statistics for All Branches (I)	148
C.2	Descriptive Statistics for All Branches (II)	149
C.3	Textual Data on the Budget Programmes (I)	150
C.4	Textual Data on the Budget Programmes (II)	151

C.5	Serial Information Processing with Alternative Topic Model	152
C.6	Descriptive Statistics for the <i>News Index</i>	153
C.7	Natural Experiment with Alternative Topic Model	154

List of Figures

2.1	Timeline of the Events	16
2.2	Total Volume of Tweets by Day	19
2.3	Total Volume of Immigration-Related Tweets by Day	19
2.4	Average of Daily Tweets by Political Party	20
2.5	Leave Share and Immigration-related Tweets During Election Top 20 Constituencies . .	33
2.6	Example of MP Reaction to Anti-immigration Rhetoric Following the Westminster Attack	40
3.1	Evolution of Sponsorship, <i>Yes</i> Votes, Resolutions and Membership over Time Sessions 31 st to 72 nd (1976-2018)	50
3.2	The Methodology	55
3.3	Word Frequencies and Most Representative Documents for Selected Topics	61
3.4	Prevalence of Selected Topics over Time (2000-2015)	61
3.5	Evolution of the <i>Reciprocity Index</i> over Time	67
3.6	Reciprocity Estimator for the Permanent Members of the UNSC	70
3.7	Rate of Sponsorship of the Permanent Members of the UNSC	71
3.8	Evolution of the <i>Reciprocity Index</i> over Time Permanent Members of the UNSC Omitted	72

3.9	Evolution of the Correlation between Population and Deviations	73
3.10	Evolution of the Correlation between Income and Deviations	74
3.11	Evolution of the Correlation between Experience and Deviations	75
3.12	Evolution of the Correlation between SC Membership and Incoming Deviations	77
4.1	Word Clouds for Selected Topics	95
4.2	Earthquakes ShakeMaps	102
4.3	Proportion of Earthquake-related Opinion Pieces (September 2017)	103
A.1	Examples of Irrelevant Tweets Picked Up with the Boolean Search	130
A.2	Examples of False Positives for the <i>Relevant</i> Category	131
A.3	Examples of False Positives for the <i>Irrelevant</i> Category	131
A.4	Examples of Tweets with Different Polarities	132
B.1	Examples of Web Pages in the United Nations Digital Library (1)	135
B.2	Examples of Web Pages in the United Nations Digital Library (2)	136
B.3	Sample of a UNGA's Resolution Text	137
B.4	Coverage of Population and Income Data over Time Sessions 31 st to 72 nd (1976-2018)	141
B.5	Evolution of the <i>Reciprocity Index</i> over Time Alternative Classifier	143
B.6	Evolution of the <i>Reciprocity Index</i> over Time Alternative Cutoff (μ)	144
B.7	Evolution of the <i>Reciprocity Index</i> over Time Alternative Time Frame (2 Sessions)	145
B.8	Evolution of the <i>Reciprocity Index</i> over Time Alternative Time Frame (3 Sessions)	146
B.9	Evolution of the <i>Reciprocity Index</i> over Time Narrow Margin Approach	147
C.1	Word Clouds for Selected Topics	155

Chapter 1

Introduction

In recent years, we have witnessed to the increasing availability of large-scale granular information on activities unmeasured before. From small-sample surveys, researchers can now rely on public sector datasets that cover entire populations, often interlinked through administrative records. In addition, not just Internet giants such as Facebook, Google, eBay and Amazon, but many other private companies have started to collect an impressive amount of real-time information on the behaviour of their customers. Such data can be accessed for scientific purposes either publicly or through suitable data-sharing agreements. This shift in the paradigm of data sources has enabled the development of new research designs and unveiled an array of previously unstudied research questions.

However, as pointed out by Einav and Levin (2014), to take full advantage of the opportunities offered by this data revolution, Economics must expand its set of traditional techniques with data mining methods (e.g., machine learning algorithms) that often can complement the standard toolkit employed by empirical researchers. Following the line of thought of these authors, my thesis proposes three empirical studies that combine methods borrowed by other disciplines, such as Computer Science and Computational Linguistics, with more conventional Econometrics techniques and research designs. The purpose is to fully exploit the potential of large and unstructured data.

While the broader theme that emerges from the three papers is the study of policymakers' behaviour, the salient features of this thesis are the data used and the combination of techniques employed, which allow tackling my research questions in a novel way.

More specifically, in chapter 2, I study the effects of exogenous shocks on the rhetoric of British

Members of Parliament (MPs) on social media. In particular, I focus on the impact of terrorist attacks on the issue of immigration. The existing literature has highlighted the dramatic effect that these events can have on the perception of immigrants (e.g., see Legewie 2013). However, while several studies have analysed the influence that media have on the process of attitudes' formation towards outgroups (Allen and Blinder 2013; Brader, Valentino, and Suhay 2008), little is known about the role that politicians might play. In fact, political leaders have strong incentives to exploit these tragic episodes to their own advantage, especially when the returns can be really high as during an election. As pivotal opinion leaders in the society, they have the power to shape public attitudes and the beliefs of their electorate.

To conduct my analysis, I collect more than 1,500,000 Tweets from the active Twitter accounts of MPs using web scraping. Next, I rely on a Naïve Bayes classifier to identify those Tweets that are immigration-related. The panel structure of the dataset that I build is then leveraged to frame a sound natural experiment setting and capture the effect of terrorist attacks on the social media agenda of MPs. Looking at the 2017 Manchester bombing as my main event study, I detect a counterintuitive finding: a substantial decrease in the expected number of immigration-related Tweets occurred after the incident. I hypothesise that this “muting effect” results from the risk-averse strategic behaviour of politicians during the election campaign. However, the MPs' response shows remarkable heterogeneity according to the socio-economic characteristics of their constituencies.

The study provides new evidence on the social media behaviour of policymakers. A topic that has gained extreme relevance over the last years, both at the level of the political system and the society as a whole, as the Donald Trump's presidency in the US has recently shown.

In chapter 3 instead, we explore the issue of vote trading in the United Nations General Assembly (UNGA). The Assembly is the only one among the six major UN organs in which all member states have equal representation. Key principles in the institution are the equality and unweighted voting among member states. However, several testimonies about exchange of votes, coercion and even direct payments suggest that votes in the UNGA might not reflect the countries' true preferences (e.g., see Malone 2000; Eldar 2008, Carter and Stone 2015). In fact, economic and military power disparities between nations set strong incentives to not vote sincerely, but exploit the one country-one vote rule strategically. We focus on the fact that reciprocity between representatives characterise different forms of strategic voting: vote buying is a trade of one's vote for goods/money; coercion is a trade of a vote for non-retaliation; vote trading is a direct exchange of votes.

In order to empirically assess these cases, we formalise in a first step the expected behaviour of country representatives in the Assembly through a statistical model based on the latent topics of the UN resolutions, which we infer using Structural Topic Models (Margaret E. Roberts et al. 2014). This allows us in a second step to build a network that describes the pattern of the deviations from the predicted voting decisions. This graph can be leveraged to compute an aggregate statistic: the *Reciprocity Index*. The measure is devised to capture the non-random occurrence of reciprocated deviations and it can be used to assess the extent of different forms of strategic voting. We find that deviations from the expected votes in the UNGA are systematically *not* reciprocated. The conclusions are consistent with a narrative of vote buying (e.g., Dreher, Nunnenkamp, and Thiele 2008; Carter and Stone 2015) and state socialisation (Alderson 2001) and highlight the structural heterogeneity of the countries involved in the decision-making process.

Hence, our findings question the most distinctive features of the Assembly, namely the unweighted voting system and the equality principle. Shedding light on the actual functioning of one of the key UN organs that should embody multilateralism is crucial, given the increasing emergence of supranational issues, such as global warming, environmental sustainability or world-wide pandemics, that call for a coordinated response from the international community.

In chapter 4, we investigate government’s information processing and its implications for policy outputs. A prominent reference for this study is the model of choice for public policy proposed by B. D. Jones and Baumgartner (2005). Key feature of the framework is the way in which policymakers process the informational signals they attend to, that is then reflected in their policy decisions. However, how well alternative processing mechanisms fit real world data remains an open empirical question. We focus on Mexico and we look at short term adjustments to budget allocation, to assess how the federal government processes a specific type of informational signal: opinion pieces from newspapers, which have strong political relevance as cues for societal issues in the Mexican context.

For this purpose, we devise a *News Index* that creates a mapping between the main topics discussed in the opinion columns and the policy outputs (i.e., the budget programmes of the federal government). Our methodology is divided into two steps. In the first one, we employ Latent Dirichlet Allocation (Blei, Ng, and Jordan 2003) on a corpus of 35822 articles from 9 Mexican newspapers to infer the main topics covered in the media and their salience. In the second step, we link policy outcomes to these topics through the *News Index*, looking at how close the budget programmes are to the different issues identified. We then rely on the index to assess two different descriptive models of information

processing: parallel and serial. The findings provide robust evidence for the serial processing hypothesis and disproportionate policy reactions by the government, measured through the gap between approved and paid expenditure on a given budget programme.

We then take a step forward and try to analyse the relationship between informational shocks and government's disproportionate reactions within a causal setting. We exploit the different media coverage of two earthquakes that occurred in Mexico in 2017 to frame a Difference-in-Differences research design and uncover the relation of interest.

Overall, the results are consistent with a narrative of disproportionate information processing and selective attention in government's decisions, which is in line with the theoretical framework of B. D. Jones and Baumgartner (2005).

By stressing the role of limited attention in policymakers' reactions to external signals, our conclusions appear to have important implications in an information-rich world where everyone is forced to prioritise among a plethora of different cues, even the government.

Finally, to conclude the thesis, chapter 5 briefly summarises each study, discusses their limitations and provides possible avenues for future research.

Chapter 2

Terrorist Attacks and Immigration Rhetoric: A Natural Experiment on British MPs

2.1 Introduction

I analyse the consequences of two acts of terrorism occurred in the UK in 2017 on the rhetoric of the British Members of Parliament (MPs) on social media. In particular, I will focus on the issue of immigration, as several studies have shown how these dramatic events might shape public attitudes towards outgroups (e.g., see Legewie [2013](#)). If opinions and beliefs are indeed affected, it appears of interest to understand which information channels could mediate this effect.

The existing literature has studied in depth how media depict immigrants and how the frame they provide can influence people's opinions (Allen and Blinder [2013](#); Brader, Valentino, and Suhay [2008](#)). However, these studies seem to neglect the role that politicians might play in the process of attitudes' formation. In fact, political leaders could get leverage on these events to their own advantage, especially when they face high stakes as during an election campaign.

One of the obvious reasons for this gap in the literature is the lack of suitable data, but also the challenges of conceiving an appropriate research design. In this paper, I try to overcome these hurdles.

I focus on a specific information environment: the social media Twitter; given its increasing relevance as platform for news’ provision and political campaigning. Using web scraping, machine learning techniques and a natural experiment setting, I try to capture the change in the rhetoric of political elites caused by a terrorist attack. Politicians might exploit these dramatic events to foster the debate on immigration and divert attention from contextual problems. This would potentially create an implicit and dangerous link in the mind of the public between the threats posed by radical extremism and the presence of perceived outgroups in the community. Alternatively, they might seize the opportunity to signal their valence or ideological stance towards immigrants. In both scenarios, we would expect an increase in the amount of immigration-related information provided by politicians in the aftermath of a terrorist attack. This in turn would have an impact on public attitudes, as the high level of anxiety induced by these dreadful events can enhance information-seeking (Gadarian and Albertson 2014). However, what I observe instead is rather counterintuitive. The amount of relevant information, measured as Tweets related to immigration posted by a politician, actually decreases, on average, in the week after the terrorist attack. Moreover, when I focus on the event occurred during an election campaign, I find significant heterogeneity according to the characteristics of the MP or her constituency. The “muting effect” of the attack is more pronounced for politicians standing for marginal seats and elected in more restrictionist constituencies. In contrast, a smaller stock of foreign people and adverse economic conditions appear to lessen the “muting effect” on the expected number of immigration-related Tweets. Surprisingly, the political divide between MPs belonging to the incumbent government’s party and members of the main challenger does not seem to matter for the response to the event.

The first attack we consider took place on the 22nd of March 2017. The 52 years-old Briton Khalid Masood drove a grey Hyundai Tucson into pedestrians along the pavement in proximity of the Palace of Westminster in London, the seat of the Parliament. The perpetrator killed three civilians and injured more than 40 people of different nationalities; he then left the car and stabbed to death a police officer before being shot. Another wounded woman died in hospital two days after.¹ The last message sent by the attacker stated that he was waging jihad in revenge for the Western interventions in the Middle East.² The Islamic State claimed responsibility for the act but no evidence emerged that backed up the allegation.³ Prior to this attack, the last act of terrorism causing multiple casualties on the British mainland was the suicide bombing in London of the 7th of July 2005.

¹Source: *BBC News* (<http://www.bbc.co.uk/news/live/uk-39355505>, retrieved on the 5th of January 2018).

²Source: *The Independent* (<https://www.independent.co.uk/news/uk/crime/last-message-left-by-westminster-attacker-khalid-masood-uncovered-by-security-agencies-a7706561.html>, retrieved on the 5th of January 2018).

³Source: *BBC News* (<https://www.bbc.com/news/uk-39408786>, retrieved on the 5th of January 2018).

The second event that we look at is the Manchester bombing occurred on the 22nd of May 2017. After the concert of the singer Ariana Grande, the 22 years-old British born Salman Ramadan Abedi detonated an explosive device in the foyer area of the Manchester Arena, causing the death of 22 people (10 of them aged under 20) and injuring more than 500.⁴ The ISIS claimed again responsibility, stating in a post on the social media that “[...] *one of the soldiers of the caliphate was able to place an explosive device within a gathering of the crusaders in the city of Manchester*”.⁵ This second attack occurred after the announcement on the 18th of April of a snap election by the British Prime Minister Theresa May, whose stated purpose was to gain a large majority to strengthen her position in the upcoming Brexit negotiations.⁶ The election took place on the 8th of June. The majority Conservative government lost 13 seats (shifting to 317) and was forced to secure a confidence and supply deal with the Democratic Unionist Party.⁷ The main challenger instead, the Labour party, won 262 seats, with a net gain of 30 seats from the previous election.

There are several reasons for the choice of these two particular events. First, even if not strictly identical, they represent the same type of shock and share common characteristics, as the nature of the attack (religious radicalism) and origin of the offender (English). Second, the two acts of terrorism occurred within a relatively short period of time and the subjects of my treatments all belong to the 56th UK Parliament, so they experienced the same institutional context. Finally, the two incidents embody two distinctive treatment conditions, where the incentives faced by the MPs were substantially different. However, given the intrinsic interest in the high stakes faced by political elites during an election campaign, my main object of study will be the Manchester attack.

In order to assess the effects of such shocks on the behaviour of politicians, I will first revise the literature related to the determinants and correlates of public attitudes towards immigration and the emotional and behavioural responses to terrorist attacks. I will also briefly mention the increasing applications of machine learning techniques in Economics, and how this study contributes to this growing body of research. In section 2.3, I will describe the methodology employed, how the data were gathered and the dataset constructed. Then, I will present the features of my classifier, the statistical model chosen and my identification strategy. Section 2.4 provides descriptive statistics on the data used, together with the time trends and general tweeting behaviour of politicians. Next, in section

⁴Source: *BBC News* (<http://www.bbc.co.uk/news/uk-england-manchester-41839277>, retrieved on the 5th of January 2018; <https://www.bbc.com/news/uk-40012738>, retrieved on the 20th of January 2021).

⁵Source: *The New York Times* (<https://www.nytimes.com/2017/05/23/world/europe/manchester-arena-attack-ariana-grande.html>, retrieved on the 5th of January 2018).

⁶Source: *The Guardian* (<https://www.theguardian.com/commentisfree/2017/may/02/europe-landslide-victory-theresa-may-brexit>, retrieved on the 5th of January 2018).

⁷Source: *BBC News* (<http://www.bbc.co.uk/news/uk-politics-40403434>, retrieved on the 5th of January 2018).

2.5, I will move to the core of the paper with the analysis of the impact of terrorist attacks on the social media agenda of British MPs. In addition, I will explore several channels that might mediate heterogeneous effects among the MPs, together with performing robustness checks on the baseline estimations. Finally, in section 2.6, I will complete the study by discussing my results and framing the direction for future research.

2.2 Literature Review

This study is placed at the crossroads of different bodies of research. International migration represents in our times one of the most challenging issues from a social, economic and political perspective. During the last 20 years, there has been a growing interest in the determinants and trends of public attitudes towards immigration (Ceobanu and Escandell 2010) and how the rising inflows of outgroups might be correlated with the upsurge of populist and xenophobic movements in Europe (e.g., Whitaker and Lynch 2011). In a recent review of the literature, Hainmueller and Hopkins (2014) underline how perceived threats to intangible social constructs, like the national economy or identity, are among the main shaping factors of attitudes' formation. In particular, the authors put emphasis on the perceptions of sociotropic threats, especially cultural, as opposed to concerns related to material self-interest.

At the same time, research has focused on exceptional circumstances that might endanger the integration of migrant people. This can be due to the emotional impact of such events or because they are perceived as signals of assimilation's failure. In this strand of the literature, we observe a growing number of studies that analyse the social and psychological effects of terrorist attacks. These dreadful incidents can have substantial consequences on natives' attitudes and perceptions of ethnic minorities and foreigners (Cohu, Maisonneuve, and Testé 2016; Legewie 2013; Schüller 2016). In addition, they have major psychological repercussions. They affect risk perception and increase the willingness to trade off civil liberties for increased public security (Bozzoli and Müller 2011), they negatively impact expectations (Coupe 2017) and lead to high levels of anxiety and anger (Huddy et al. 2005; Vasilopoulos 2018). Besides, these emotional reactions can pervade actual behaviour. Hanes and Machin (2014) document an exacerbation of hate crimes against Asians and Arabs following the London bombings in 2005. Moreover, terrorist attacks, as dramatic and dreadful events, might question the effectiveness of the government and the political system, leading to an impact on electoral outcomes (Montalvo 2010). If these events have such grievous consequences on the social structure, it seems sensible to try

to understand through which channels they might affect public opinions. In fact, if sociotropic threats are more influential than the egotropic ones, acts of terrorism could be exploited by political elites to shape mass attitudes towards immigration, appealing to social constructs such as national identity or local economic conditions. Alternatively, they can seek the opportunity to signal their position on the political spectrum about the issue. This could be even more relevant during elections; since, as shown by Kendall, Nannicini, and Trebbi (2015), even in the short run voters do update their beliefs when receiving new information on the ideology or stance of a candidate.

As a matter of fact, opinion leaders, and so political leaders, are likely to represent the main hubs in an information acquisition's network (Galeotti and Goyal 2010). In this structure, they constitute core nodes with high indegree: the central pillars in the information environment of the voters. Their role becomes even more relevant in the aftermath of a terrorist attack. It has been shown through experimental evidence that the anxiety induced by unfamiliar threatening conditions triggers political response and information seeking, with a bias on distressing news (Brader, Valentino, and Suhay 2008; Gadarian and Albertson 2014). Hence, voters in the wake of such events are likely to be much more sensitive to any information provided by opinion leaders, which in turn can flow through a variety of communication channels.

In the last decade, one of these channels, social media, has arisen for its rapid diffusion and development. At the same time, we have witnessed to its growing impact on news' provision. As an example, surveys from the Pew Research Center show that the share of US adults getting news on social media has increased from 49% in 2012 to 62% in 2016 (Gottfried and Shearer 2016). In addition, this trend has been matched with a widespread uptake of these platforms by political leaders, with a consequential effect on their electoral performance. Recent studies find a positive association between social media-based campaigning (specifically, the activity on Twitter) and voting outcomes in the UK (Bright et al. 2020). Thus, it appears relevant to assess the role of these emerging information channels and their strategic use by politicians.

Research on the use of social media, and in particular Twitter, as a communication and electoral tool by political elites is still in early stages, but with a growing number of findings (see Jungherr (2016) for a survey of the literature). Interestingly, small sample studies suggest that the personal use of Twitter by politicians might diverge from what we would expect in a communication environment strategically coordinated, where members collectively advocate party policies (Adi, Erickson, and D. G. Lilleker 2014). At the same time, tweeting behaviour seems to transcend partisanship, and common

patterns emerge, at least among major political parties (Evans, Cordova, and Sipole 2014). In fact, this microblog becomes a channel for expressing individual lines on policy, due to the personalisation that this particular hybrid platform allows. The construction of a blurred private/public personality by politicians is meant to induce empathy from voters and could in turn reflect strategic behaviour aimed at earning personal support (Jackson and D. Lilleker 2011). This personalisation of their professional figure is in line with Impression Management Theory (E. E. Jones and Pittman 1982), which provides a taxonomy of attitudes through which individuals try to actively manage the public perception of themselves. Nonetheless, it appears that the focus of their messages has predominantly a political theme, especially during election campaigns (Evans, Cordova, and Sipole 2014).

It has been observed that political elites use their Twitter account for constituency service: it is a convenient channel for reaching crucial audiences quickly and effectively. However, even if it seems mainly a unidirectional channel of communication, politicians do interact with the Twitter community in order to attack an opponent, debate, or taking a position on a specific issue (Graham et al. 2013). Said that, the evidence shows that the microblog represents a powerful way of self-promotion, leveraged to maximise the impact on the electorate (Jackson and D. Lilleker 2011).

At the same time, the influence of politicians on the information environment of the public is indirectly amplified. This peculiar channel offers the possibility of manipulating the flow of the national dialogue through its impact on the agenda of traditional mass media (Kreiss 2016). Qualitative research shows how professional journalists do use Tweets from political leaders to shape their coverage in terms of issues and events. They also obtain from them background information, polling data and quotes that subsequently include in their articles (Parmelee 2014). However, it is important to underline that the role that Twitter might play in traditional media's agenda-building is very context-dependent, and it is likely to change according to the institutional setting under analysis.

In any case, the freedom of expressing personal beliefs and opinions offered by this social media might turn out to be a double-edged sword. Since journalists can rely on this microblog as a way to monitor politicians' view and inform their agenda, the exposure of political leaders to criticism and attacks is magnified, leading to a careful use of the platform. In fact, it is not uncommon in the UK context that hasty Tweets led to subsequent public condemnation, requiring formal excuses. As an example, in 2013 the Prime Minister David Cameron had to face open criticism after that a member of his staff endorsed by error an offensive Twitter account (Adi, Erickson, and D. G. Lilleker 2014). It is clear that in such a setting unexpected events might engender strategic responses, especially when stakes

are high, as during an election campaign.

While the role of traditional mass media in depicting immigrants has been analysed extensively (e.g., Allen and Blinder 2013), to the best of my knowledge it appears that poor quantitative research has been conducted on the role and behaviour of political elites and their strategic use of social media. This paper thus tries to partially fill this gap by proposing one of the first empirical studies on the effects of terrorist attacks on the immigration rhetoric of politicians on social media.

From a methodological perspective, this work adds to the growing literature on applications of machine learning techniques in economic and social research. Nowadays, such algorithms are spreading in different fields of Economics and Political Science, often with the aim of selecting the relevant co-variates in an empirical model (Belloni, Chernozhukov, and Hansen 2014) or capturing heterogeneous treatment effects (Wager and Athey 2018). Applications range from predicting consumer demand (Bajari et al. 2015) to test theories of risky and ambiguous behaviour (Peysakhovich and Naecker 2017), with an increasing emphasis on estimating causal effects (Athey and Imbens 2015). This paper is thus an attempt to combine what Leo Breiman called the two cultures of statistical modelling (Breiman 2001b). The first one, based on stochastic data models, aims at capturing causal relationships between variables. The second one employs learning algorithms to maximise the accuracy of out-of-sample predictions. In my analysis, I leverage the latter to improve the quality of the data used. In addition, the granularity of the information retrieved is exploited to frame a natural experiment design that is likely to allow the interpretation of the parameter of interest as a causal effect. However, this relationship is estimated through a standard statistical model, which is meant to describe the underlying data-generating process.

2.3 Methodology

The methodology employed in this paper can be divided into three parts. The first illustrates the process of collecting the relevant information and building the dataset, whereas the second one presents the statistical model preferred to carry out the empirical analysis. We then conclude by describing the natural experiment setting of the study.

2.3.1 Construction of the Dataset

The construction of the dataset can be further broken down into two steps. In the first one I collected all the most recent Tweets of British MPs with an active Twitter account through the Twitter API. One limitation is that only the latest 3,200 Tweets can be collected per account.⁸ However, as explained later on in this section, only the Members of the Parliament for whom I have complete information for the different time periods were considered in the analyses. The collection was executed on the 27th of September 2017, bringing 1,504,088 Tweets. I then started to select all the *relevant* Tweets through a Boolean search. *Relevant* Tweets are defined as all the Tweets containing one or more of the words listed in Appendix A.1.⁹ The choice was mainly informed by the report of Allen and Blinder (2013) that documents all the major words correlated with the terms immigrants, migrants, asylum seekers and refugees (or variations) in the 20 main British newspapers between 2010 and 2012. These terms plus synonyms of the dominant correlates and other current relevant words (e.g., free movement) form the final list. The amount of data was thus reduced to around 20,600 Tweets. However, the dataset to this point still contained lots of Tweets irrelevant for my analysis, due to the variability in the semantics of the chosen words in different contexts. Figure A.1 in Appendix A.2 provides some examples of these problematic Tweets. Hence, I further improved the quality of my dataset by relying on machine learning (ML) techniques.

In the subsequent step, I trained a classifier that was able to effectively reproduce the decision-making process and distinguish between Tweets that were relevant to my research from those that were not. The underlying predictive model is a semi-supervised multinomial Naïve Bayes coupled with the feature marginals (FM) algorithm, as proposed by Lucas and Downey (2013).¹⁰ A Naïve Bayes model was preferred as it is relatively fast to train, it has been proven effective for text analysis and it suits well semi-supervised learning algorithms (Kober and Weir 2015). FM was chosen as it has been shown to outperform other standard algorithms both in text topic classification and sentiment analysis (Lucas and Downey 2013). Its main feature being that it does not have to iteratively compute multiple passes over the unlabelled data for each new task (contrary to the expectation-maximisation algorithm, for example). It instead precomputes a set of statistics (i.e., the marginal probability of each word) over the unlabelled data in advance. These statistics are then used as constraints in the optimisation

⁸Native Retweets are counted in this total (https://developer.twitter.com/en/docs/tweets/timelines/api-reference/get-statuses-user_timeline.html).

⁹The Tweets containing words for which those chosen are substrings were retained as well (e.g., all the Tweets with the word immigrant were preserved, since they contain the word migrant, that is present in our list).

¹⁰Semi-supervised learning is an approach that tries to leverage the information on both the unlabelled and labelled data in order to learn the target function (Lucas and Downey 2013).

problem, in order to improve the estimates of the class-conditional probability of each word. This procedure is particularly suitable for improving the estimates of words that have not been seen in the labelled data. The subset of human-coded (i.e., labelled) Tweets used during the training of the classifier was made up by 600 items. Data were preprocessed by normalising URLs, punctuation was filtered and all tokens were made lowercase before extracting the features. Bigrams and trigrams were used in addition to unigrams as features for the classification, to capture more complex grammatical structures.

According to the ML literature the quality of an algorithm is assessed through its out-of-sample prediction performance (Varian 2014). Hence, I manually labelled a subset of 900 Tweets (630 were *relevant*, 270 were not) in order to form a gold standard dataset against which the classifier was tested. Table 2.1 reports the performance of the classifier. The overall accuracy (i.e., the proportion of Tweets which were assigned to the correct category) is 0.878, well above the value of 0.7 recommended by Van Rijsbergen (1979) for scientific research. The precision value states the proportion of all documents the classifier believed belonging to a given class that were truly belonging to that category. Using standard hypothesis testing notation, this can be thought of as $(1 - \alpha)$, that is one minus the likelihood of a Type I error. The recall value is the proportion of all documents belonging to a particular category, which the classifier labelled as belonging to that class. It can be thought of as the power of the test $(1 - \beta)$. The F-score represents the harmonic average of precision and recall. However, for the purpose of my analysis, the most significant statistic is 1-precision for the *relevant* class, as it captures the proportion of false positives for that category in the classification exercise. The proportion of Tweets erroneously labelled as *relevant* by my classifier was less than 10%, a rather small value (examples are provided in figure A.2 in Appendix A.2).¹¹ The final dataset was made up by 14,817 immigration-related Tweets, spanning a period of over 9 years.

Table 2.1: Classifier Performance

Label	Precision	Recall	F-score	Accuracy
<i>Relevant</i>	0.908	0.919	0.913	
<i>Irrelevant</i>	0.805	0.781	0.793	
				0.878

It must be noted that the *relevant* category captures all texts generally related to the issue of immigration posted by an English MP over the time period considered in the study. Thus, I do not perform

¹¹However, also the proportion of false positives for the *irrelevant* class is meaningful, as it points out that my analysis can only be a lower-bound estimation. Examples of this kind of Tweets are shown in figure A.3 in Appendix A.2.

a sentiment analysis, as I do not discriminate between Tweets that have positive, negative or neutral polarity with respect to this topic.¹² The choice was driven by the potential pitfalls of this type of analysis in my particular setting. The performance of a sentiment classifier is crucially dependent on the domain’s consistency of the data used (Barbieri, Ronzano, and Saggion 2015; Deriu et al. 2017) and the context of the words in a given textual corpus (Saif et al. 2016; Teng, Vo, and Zhang 2016). Even if the former is well-defined through my two-step procedure, the latter is very likely to change in the aftermath of a terrorist attack, due to the emotional reactions that such events engender. Thus, a sentiment analysis could mistakenly interpret a shift in the context and choice of words used to convey feelings as a change in the amount of Tweets with a given polarity.

2.3.2 Statistical Model

To model the data-generating process I opt for the zero-inflated negative binomial (ZINB), as it is suitable for over-dispersed data (i.e., the conditional mean is not equal to the conditional variance) that present excess of zeros (Cameron and Trivedi 2013). The main idea underlying this model is to include a separate component (π) that inflates the likelihood of observing a zero. Thus, the ZINB assumes that the zero observations arise from two different sources, a structural one (given by π) and a sampling one (given by the base count density $f_2(y)$) (Hu, Pavlicova, and Nunes 2011). Equation 2.1 presents the generalisation of the model. In my application the base count density $f_2(y)$ is a NB2 (Hilbe 2011).

$$Pr(y = j) = \begin{cases} \pi + (1 - \pi) f_2(0) & \text{if } j = 0 \\ (1 - \pi) f_2(j) & \text{if } j > 0 \end{cases} \quad (2.1)$$

The inflation factor π might be a constant or depend on a set of regressors in a binary outcome model. In my case, the inflation factor is a (logistic) function of the total number of Tweets posted by the MP in a given day. The insight is simple: the likelihood of observing a non-zero for the dependent variable of interest (the total number of immigration-related Tweets) is correlated to the daily Twitter activity of the politician. The more she tweets, the more likely she is to talk soon or later about immigration.

Moreover, the ZINB has already been used in other scientific fields to model Twitter data (e.g., Williams and Burnap 2016). In section 2.5.3, as a robustness check, I estimate my baseline model

¹²See figure A.4 in Appendix A.2 for examples of Tweets with different polarities.

using a standard negative binomial and adding the total number of Tweets at the MP level as a covariate.

2.3.3 Natural Experiment Setting

I study the effects of terrorist attacks on the number of immigration-related Tweets posted by an English MP on a given day. In order to accomplish this purpose, I exploit the panel structure of my dataset in an event study framework. Due to the exogeneity in the timing of these acts of terrorism, the estimates are likely to provide the average effect of the treatment (the attack) on the treated (the MPs). However, the time windows chosen are crucial for my identification strategy.

It has been noted that Twitter data is particularly volatile, and messages are generally triggered by specific events related to the topic under study (Wibberley, Weir, and Reffin 2014). Thus, I eliminate from the analysis two main events directly related to the issue of immigration that caused a peak in the frequency of politicians' Tweets about this topic.¹³ The first one took place on the 7th of March, when the amendment proposed by the Conservative MP Heidi Allen to properly audit local council capacity to house unaccompanied child refugees was defeated in the Parliament. The amendment was grounded on the Home Office's sudden abandonment of the Dubs Scheme on refugees occurred in February and caused a significant contention on the issue among MPs.¹⁴ The other triggering event was on the 29th of May: The Battle for Number 10, a live TV debate between the incumbent Prime Minister Theresa May and the leader of the Labour Party Jeremy Corbyn. During the discussion, immigration was a major theme of confrontation, prompting all kinds of remarks along the political spectrum.¹⁵ I exclude the day of the first event from the analysis of the Westminster attack, whereas The Battle for Number 10 will constitute the upper temporal bound for the study of the Manchester bombing. The main purpose of these omissions is to capture the effect on the average tweeting behaviour of the MPs.

It is also worth mentioning that there is no direct reason why the two terrorist attacks (both committed by British citizens) should provoke a change in the immigration rhetoric of the MPs, apart from political gain. However, even if we were to assume an effect, we would not expect a long-lasting impact. Issue-Attention Cycle Theory posits that public attention to even major social problems suddenly peaks, but then rapidly fades away (Downs 1972). This hypothesis is consistent with the empirical results of


¹³See section 2.4 for an overview of the Tweets time trends.

¹⁴Source: *BBC News* (<http://www.bbc.com/news/uk-politics-39187290>, retrieved on the 5th of January 2018).

¹⁵Source: *The Guardian* (<https://www.theguardian.com/politics/blog/live/2017/may/29/paxman-interview-corbyn-may-sky-general-election-paxman-interviews-may-and-corbyn-politics-live>, retrieved on the 5th of January 2018).

Legewie (2013), and Williams and Burnap (2016), which document how the emotional and attitudinal effects of terrorist attacks are quite short-lived. Hence, I mainly expect a reaction from politicians only in the immediate aftermath of the incident. I will study a time interval that looks at the week after the event (including the day of the attack) and the week before. I will then expand it by also looking at two and three weeks prior to the incident. One main constraint of the analysis is that the further I extend the time interval the more likely I am to capture other, even if less known, triggering events, that might be systematically related to the response variable conditional on the attack (the temporal stability assumption of Legewie (2013)). Figure 2.1 shows a timeline of the relevant events considered in the analysis.

Figure 2.1: Timeline of the Events



Another important point to mention is the number of cross-sectional observations considered in the following analyses. As already noted above, I was only able to gather the 3,200 most recent Tweets for each active account. This constraint is reflected in the number of observations available for the two events. When comparing the two incidents, I will only consider the MPs for whom I have full information for both the attacks (519).¹⁶ When I will focus on Manchester, I will consider all the available active Twitter accounts for which I have information (548). In any case I look at a sizeable proportion of the members of The House of Commons.¹⁷ Appendix A.3 provides further descriptive statistics on the MPs included in the analyses and those excluded.

¹⁶Actually, in the comparison I also consider four more accounts that were created meanwhile when looking at Manchester.

¹⁷The Lower House of the British Parliament has 650 members. Thus, I analyse around 80% of them in the comparison and 84% when I just focus on Manchester.

2.4 Data Description

2.4.1 Twitter Data

The dataset employed in the analysis has a longitudinal structure. The cross-sectional unit of observation is a Member of the House of Commons. I record daily Twitter activity related to the issue of immigration as described in section 2.3. Thus, my main dependent variable is the number of immigration-related Tweets posted by an English MP in a given day. I also retain the daily number of Tweets posted by the MP and information on the account, as the number of followers, number of friends, number of statuses and its age.¹⁸ The characteristics of the account might be important correlates of the tweeting behaviour of the politicians, so I decide to keep them as controls. I also add demographics; as previous studies have shown how age and gender might affect the use of the microblog, in particular when considering issue-specific Tweets (Evans, Cordova, and Sipole 2014; Jackson and D. Lilleker 2011).

2.4.2 Data at the Constituency Level

In order to explore the heterogeneity of the effect across the MPs, I gather information on their constituencies. I consider the majority share of the incumbent MPs in 2015 general election and their betting odds for the 2017 election. I collect a proxy for the average unemployment level in 2016, measured as the proportion of economically active 16-64 years-old residents claiming Jobseeker's Allowance. This last information comes from the ONS Nomis database, and is meant to capture the perception of local competition over scarce resources. From the British Election Study (BES, 2017 results) I collect the share of UK-born and the share of people of white British ethnicity as measures of intergroup contact.¹⁹ This database also contains the estimates of the results for the 2016 EU Referendum at the constituency level, as computed by Hanretty (2017). This measure is meant to capture the salience of the issue of immigration in a given constituency. Unfortunately, the BES does not report information on Northern Ireland, so I have to systematically exclude its 18 constituencies (16 of them included in our dataset) when analysing the effect of these last variables. Table 2.2 presents descriptive statistics for the independent variables employed in the study.

¹⁸These variables change over time, but in my dataset are fixed, as they report the value on the day of the collection (27/09/2017). However, they are a good proxy for the type of node that the MP represents in the network structure of the Twitter community and her level of engagement with the platform.

¹⁹It is worth noting that these two proxies might not precisely capture the same concept. In fact, the proportion of UK-born also includes second-generation migrants, so it does not distinguish between multiple ethnic groups.

Table 2.2: Descriptive Statistics (MP/Constituency Level)
16/02/2017-01/06/2017

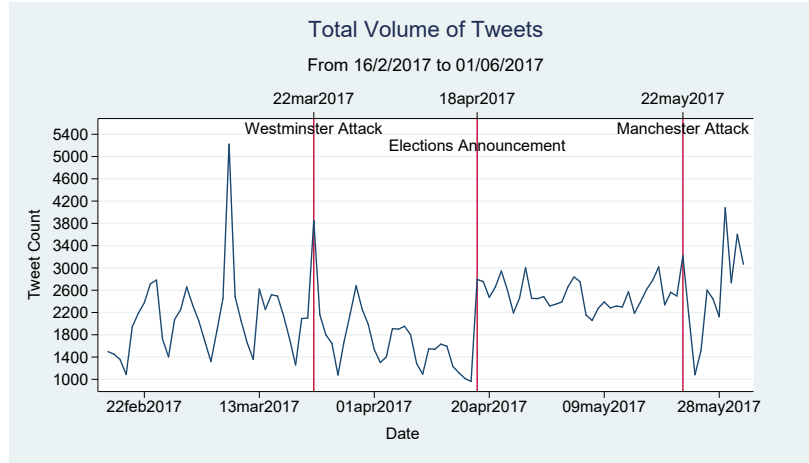
	Mean	Std. Dev.	Min.	Max.	N
<i>Followers Count</i>	25330.895	82282.903	329	1538565	58512
<i>Friends Count</i>	1792.217	2882.238	0	38861	58512
<i>Statuses Count</i>	8799.687	10891.670	16	82326	58512
<i>Male</i>	0.670	0.470	0	1	58512
<i>Age</i>	50.768	10.149	22	82	58512
<i>Age Account</i>	5.423	2.074	0	9	58206
<i>Majority Share (2015)</i>	23.536	14.079	0.100	72.300	58512
<i>Betting Odds (MP)</i>	0.278	0.757	0.002	9	55120
<i>Unemployment 2016 (Avg.)</i>	2.578	1.509	0.494	9.737	58512
<i>Leave Share</i>	51.603	11.673	20.481	75.650	56816
<i>White British Ethnicity (%)</i>	82.667	18.763	12.712	97.792	56816
<i>UK-Born (%)</i>	87.960	11.639	40.728	98.018	56816

2.4.3 Tweeting Trends

One first important question that we might want to ask is if Twitter is a meaningful way to capture politicians' rhetoric, and if these opinion leaders really use the platform to communicate with their electorate. Figure 2.2 shows the time trend of the total number of Tweets for the MPs on whom I have information for the whole time period considered. We can clearly see that, after the elections announcement, the average number of Tweets substantially increases and no longer displays that seasonal pattern observed before. It instead presents small fluctuations around a higher grand level until the day of the Manchester attack, when it drops dramatically. A similar decrease seems to occur after Westminster. It appears that, at least during the election campaign, the MPs did increase their use of Twitter, presumably to get more in touch with their voters and promote themselves.

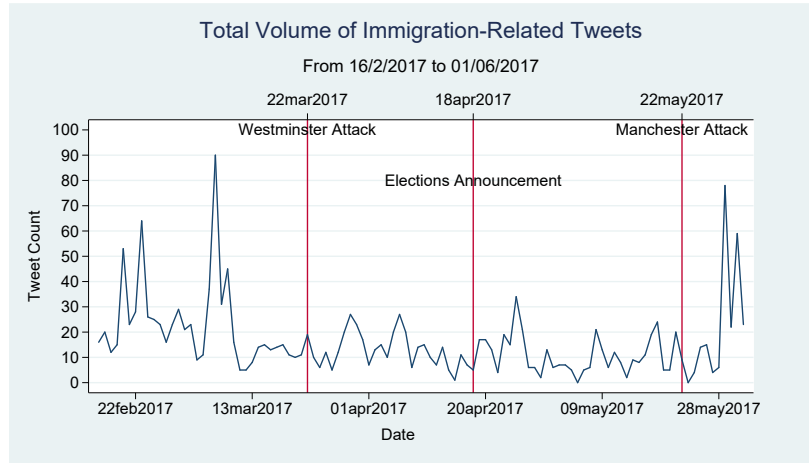
If we look at the trend of the immigration-related Tweets in figure 2.3, we do not observe, on average, a significant rise during the election. Hence, it does not seem that this was a topic particularly highlighted by the MPs in their political campaign strategies. This might be a consequence of immigration being a rather controversial and risky theme. As I have already mentioned in section 2.3, we have two major peaks: one in correspondence with the rejection of Heidi Allen's amendment, the other one on the day of The Battle for Number 10. We do notice a fall after the Manchester bombing, even if it does not seem as dramatic as for the total number of Tweets, but this might be due to the greater difference

Figure 2.2: Total Volume of Tweets by Day



in absolute values. The Westminster attack seems to cause a very short-lived drop, but the effect is not as clear as for the overall volume. However, these are just aggregate trends that do not take into account the individual characteristics of the MPs.

Figure 2.3: Total Volume of Immigration-Related Tweets by Day



Another important point to underline is the volatility of tweeting behaviour. As crude measure, table 2.3 reports the R^2 values of simple OLS regressions that capture the probability that a politician tweets. The dependent variable is a dummy that takes the value of 1 if the MP tweets in a given day and the independent variables are a full set of individual and day fixed effects. This exercise is carried out for both pre and post elections announcement periods. What we observe is a rather random behaviour in the likelihood of tweeting. The full set of covariates is able to explain less than 40% of the variation between groups in both cases.²⁰ Thus, I do not expect great predictive power from my

²⁰Here I try to provide a raw measure of the share of explained variance due to between-group variation, but I should highlight that one must be very cautious when interpreting the R^2 as a measure of goodness of fit with binary outcomes.

models and quite noisy estimates. However, given the number of unobservables at play, capturing a significant and robust effect would represent a rather neat result.

Table 2.3: Randomness in Tweeting Behaviour


	OLS (1) Pre-announcement	OLS (2) Election Campaign
R^2	.37	.39
N	31,408	22,495
$MP\ FE$	Yes	Yes
$Day\ FE$	Yes	Yes

Note: the dependent variable is a dummy for tweeting or not in a given day. Period before the announcement is 17/02/2017-17/04/2017. Period during the election is 18/04/2017-28/05/2017.


One more pattern that might be interesting to analyse is the difference in tweeting behaviour across political parties. Figure 2.4 presents the average daily tweeting activity of MPs by political affiliation. The values refer to 30 days before the elections announcement and 30 days after. As already noted with the time trends, we see a clear rise in the post-announcement period, and this is true for almost all political parties. We also observe that between the two major ones, Labour presents systematically higher values than the Conservatives. In addition, Plaid Cymru exhibits the most significant increase, with an average number of Tweets more than doubled after the announcement. These patterns seem to confirm the relevance of Twitter as a channel of communication and information exchange between the politicians and their electorate.

Figure 2.4: Average of Daily Tweets by Political Party

(a) Pre-announcement Period



(b) Post-announcement Period



2.5 Analysis

2.5.1 Baseline Results

I now present the baseline results on the impact of the two acts of terrorism on the total number of Tweets and immigration-related Tweets. Next, I will focus on the Manchester attack and I will explore the heterogeneity of the effect according to the different characteristics of the politician or her constituency.

Effect of the Terrorist Attacks on the Total Number of Tweets

In tables 2.4, 2.5 and 2.6 I explore the impact of the terrorist incidents of Westminster and Manchester on the daily tweeting activity of British MPs. I look at different time intervals: 3, 2 and 1 week before the attack, but I compare them only with the first week following the event, each time. The treatment is a dummy that takes the value of 1 on the day of the incident and the following six days. In each regression I add as covariates a dummy for being male, the age of the politician, number of followers, number of friends, number of statuses, and age of the account. Day-of-the-week fixed effects are included, in order to capture weekly seasonality in tweeting behaviour. The models are estimated through a negative binomial and errors are clustered at the MP level. Tables 2.4, 2.5 and 2.6 show the incident rate ratios (IRRs) for the treatment and p -values are reported in parentheses. The last table shows the estimates for the Manchester attack when I consider all the available MPs. As mentioned in section 2.3, the 7th of March is not considered in the estimation of the Westminster attack. The baseline regression model is presented in equation 2.2, where the vector \mathbf{x}_{it} contains the control variables.

$$E[TweetsCount_{it}|Treatment_t, \mathbf{x}_{it}] = \exp(\alpha + \beta Treatment_t + \mathbf{x}_{it}\gamma) \quad (2.2)$$

What we observe is a clear decrease in the number of Tweets after both the events. However, the effect is definitely more pronounced for the Manchester attack (a reduction between 11% and 18% in the expected number of Tweets) and it is always highly significant in every time interval. In addition, the magnitude of the effect fades away as I extend the time period. For Westminster, the pattern is less clear. The impact appears to be not significant in the proximity of the event, but gains relevance

as I enlarge the period analysed, with a magnitude that is less than 11%.²¹

Table 2.4: Effect on Total Number of Tweets (Westminster Attack)

	NB (1) 1 Week/1 Week	NB (2) 2 Weeks/1 Week	NB (3) 3 Weeks/1 Week
<i>Westminster Attack</i>	0.972 (0.386)	0.892 (0.000)	0.918 (0.002)
<i>N</i>	7,266	10,892	14,000

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported. 07/03/2017 not included.

Table 2.5: Effect on Total Number of Tweets (Manchester Attack)

	NB (1) 1 Week/1 Week	NB (2) 2 Weeks/1 Week	NB (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.818 (0.000)	0.876 (0.000)	0.885 (0.000)
<i>N</i>	7,322	10,983	14,643

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported.

Table 2.6: Effect on Total Number of Tweets (Manchester Attack, All Obs.)

	NB (1) 1 Week/1 Week	NB (2) 2 Weeks/1 Week	NB (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.822 (0.000)	0.875 (0.000)	0.882 (0.000)
<i>N</i>	7,672	11,508	15,343

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported. All the available MPs are used.

Effect of the Terrorist Attacks on the Number of Immigration-Related Tweets

Now I move to the core of the analysis; the dependent variable of interest is the daily number of immigration-related Tweets posted by a politician. My goal is to capture the impact of terrorist attacks on the social media agenda of British political elites. The features of the analysis are the same as in the previous subsection, but I now estimate my models through a zero-inflated negative binomial,

²¹As incident rate ratios are just the exponentiated coefficients, the proportional change in the expected count is simply given by $(\exp^{\beta} - 1) \times 100\%$, where β represents the coefficient of interest. Taking as an example the third column in table 2.4, the proportional change is computed as $(0.918 - 1) \times 100\% = -8.2\%$.

where the inflation factor π is a function of the total number of Tweets posted by a politician in a given day. The results are presented in tables 2.7, 2.8 and 2.9. It is important to reiterate that, as these dramatic incidents are not directly related with immigration, there is no reason why we should observe a distinct effect at all, apart from strategic behaviour. Moreover, if politicians wanted to seek the opportunity to signal their stance or ideology with respect to this issue, we should expect an increase in the number of immigration-related Tweets. We would also expect the same outcome if MPs were trying to shape public attitudes on the theme.

Considering the Westminster attack, the effect is rather imprecisely estimated, and it is only marginally significant when we look at the three weeks before the incident, but always implying a reduction. The results for the Manchester attack appear more accurate. The effect is quite substantial in magnitude: a reduction of around 30% in the expected number of immigration-related Tweets when comparing one week before to one week after, that decreases to around 27% when considering the two weeks prior to the event. Both the effects are statistically significant at the 5% level. The effect increases substantially when I consider all the available MPs for the Manchester attack: a decrease of around 38% compared to the week before, slightly less (around 36%) when considering the two weeks previous to the incident. Both the effects are statistically significant at the 1% level. The pattern appears to be quite similar to that of the total number of Tweets: the impact fades away as I extend the time interval. However, the magnitude of the shock is proportionally greater. Table 2.10 reports the results of a one-sided Z -test under the null hypothesis that the coefficient of the treatment effect on the immigration-related Tweets is less than or equal to the coefficient of the effect on the total number of Tweets. I cannot reject the null hypothesis for every time interval at any standard significance level: the impact of the attack appears to be more negative (i.e., greater in absolute value) for the immigration-related Tweets.²² Hence, what we observe is a rather counterintuitive “muting effect”: a substantial reduction in the number of immigration-related Tweets following an act of terrorism, and this seems to be particularly true during the election.

However, it is important to underline at this point that the two terrorist incidents are not strictly comparable. The Westminster attack was the first act of terrorism on British mainland after almost twelve years, whereas the Manchester one had definitely a greater death toll, and many of the individuals involved were young people, so the emotional reactions are likely to be different. Moreover, the second attack occurred during an election campaign.

²²Here I am comparing the results for the Manchester attack with all the available MPs.

A possible explanation for the observed behaviour is a risk-averse strategy adopted by the politicians. Being aware of the unpredictable reactions and emotional distress of their electorate, and knowing the potential link between terrorist attacks and attitudes towards immigration, they prefer not to expose themselves and being on the safe side by neglecting the topic in the aftermath of the event. Moreover, the difference in the estimated effect between the two episodes deserves further considerations. It appears that the “muting effect” on the immigration-related Tweets is clearly observed only for the second attack. This could be the result of a dynamic process, in which political leaders learn to avoid risky issues and tend to maximise this behaviour in high stakes situations, like an election. I will further examine this hypothesis in section 2.6.

In order to study the heterogeneity of the impact across different characteristics of the MPs or their constituencies, I will now focus on the Manchester attack, as it occurred during an election campaign and it is thus more suitable to analyse the different incentives that politicians might face.

Table 2.7: Effect on Total Number of Immigration-related Tweets
(Westminster Attack)

	ZINB (1) 1 Week/1 Week	ZINB (2) 2 Weeks/1 Week	ZINB (3) 3 Weeks/1 Week
<i>Westminster Attack</i>	0.966 (0.852)	0.895 (0.457)	0.785 (0.099)
<i>N</i>	7,266	10,892	14,000

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported. 07/03/2017 not included.

Table 2.8: Effect on Total Number of Immigration-related Tweets
(Manchester Attack)

	ZINB (1) 1 Week/1 Week	ZINB (2) 2 Weeks/1 Week	ZINB (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.704 (0.030)	0.734 (0.048)	0.913 (0.543)
<i>N</i>	7,322	10,983	14,643

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported.

Table 2.9: Effect on Total Number of Immigration-related Tweets (Manchester Attack, All Obs.)

	ZINB (1) 1 Week/1 Week	ZINB (2) 2 Weeks/1 Week	ZINB (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.623 (0.002)	0.638 (0.002)	0.797 (0.103)
<i>N</i>	7,672	11,508	15,343

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported. All the available MPs are used.

Table 2.10: One-sided Z-test $H_0 : \beta_{Immigration} \leq \beta_{Total}$

	Manchester Attack 1 Week/1 Week	Manchester Attack 2 Weeks/1 Week	Manchester Attack 3 Weeks/1 Week
<i>Z Statistic</i>	-1.749	-2.159	-0.715
<i>p-value</i>	0.960	0.984	0.763

Note: one-sided Z-test. Null hypothesis: the coefficient of the treatment effect on immigration-related Tweets is equal or less than the coefficient of the treatment effect on the total Tweets. Manchester attack, all the available MPs are used.

2.5.2 Heterogeneity

I now focus on the Manchester bombing and try to capture potential channels of heterogeneity in the effect among the MPs. I will use all the available MPs as I no longer consider the two events. I will explore the following factors that could mediate the effect: the “safety” of a politician’s seat, intergroup contact, competition over scarce resources, incumbency and the salience of the issue of immigration in the constituency.

The first channel that I am going to analyse is the relative strength of a MP’s position in her constituency. It should be underlined that the expected sign of the treatment effect is not straightforward. On one hand, we might think that the marginal utility coming from an additional Tweet for those MPs with a safe seat is lower, so they will tend to ignore the issue of immigration. On the other hand, these politicians might be willing to take a stance even on the riskier topics, due to the strength of their position. The same reasoning, but with opposite effects, applies to MPs standing in marginal seats. In order to test these contrasting hypotheses, I use two proxies for the relative risk of a politician’s position: the majority share in the 2015 general election and the last available betting odds at the constituency level for the 2017 election.²³ I perform the analysis on two different subsamples, focusing

²³Data on betting odds were retrieved on the 16th of January 2018 from BetOnPolitics.co.uk (now <https://www.bettingpro.com/>).

just on those MPs standing in marginal constituencies for the 2017 election. The first subsample is defined by all MPs standing for re-election in those constituencies where their marginal share of votes in 2015 was less than 10%. The second one is restricted to only those MPs standing in constituencies where their betting odds were greater than 0.1. Thus, I am only considering those politicians with a risky seat.

Table 2.11 and table 2.13 report the results of this exercise. What we observe is a substantial increase in the absolute magnitude of the effect, and this seems to hold for both proxies and in every time interval. For instance, if we look at the narrowest time window for the subsample defined by the majority share, the resulting reduction in the expected number of immigration-related Tweets is greater by around 20 percentage points compared to my baseline results. As in subsection 2.5.1, I perform a one-sided Z -test to compare the size of the effect in the subsample analysed with the full sample. The null hypothesis is that the coefficient of the treatment effect on the immigration-related Tweets in the subsample is less than or equal to the coefficient of the effect in the full sample. Table 2.12 reports the results for the subsample defined by the majority share and table 2.14 reports the results for the subsample defined by the betting odds. In both cases I cannot reject the null hypothesis for every time interval at any standard significance level. Hence, it appears that this risk-averse behaviour does depend on the relative strength of the politician and the “muting effect” of the terrorist attack is magnified for those leaders with a marginal seat. MPs tend to be even more cautious in their tweeting behaviour when their position is not safe.

Table 2.11: Heterogeneity by Majority Share

	ZINB (1)	ZINB (2)	ZINB (3)
	1 Week/1 Week	2 Weeks/1 Week	3 Weeks/1 Week
<i>Manchester Attack</i>	0.420 (0.013)	0.488 (0.033)	0.661 (0.135)
<i>N</i>	1,456	2,184	2,912

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and p -values are reported. The sample is restricted to MPs with a majority share in the 2015 election of less than 10 percentage points. Only MPs standing for the 2017 election are considered.

I now analyse a different channel through which the treatment might have a heterogeneous impact: the presence of a relevant stock of migrant people in the constituency. In order to explore this hypothesis, I will add an interaction of the treatment with the variable of interest, keeping the latter as a covariate to account for differences in levels. According to Intergroup Contact Theory, increased intergroup

Table 2.12: One-sided Z -test $H_0 : \beta_{Subsample} \leq \beta_{FullSample}$

	Manchester Attack 1 Week/1 Week	Manchester Attack 2 Weeks/1 Week	Manchester Attack 3 Weeks/1 Week
<i>Z Statistic</i>	-1.196	-0.835	-0.706
<i>p-value</i>	0.884	0.798	0.760

Note: one-sided Z -test. Null hypothesis: the coefficient of the treatment effect on immigration-related Tweets in the subsample is equal or less than the coefficient of the treatment effect in the full sample. Manchester attack, all the available MPs are used.

Table 2.13: Heterogeneity by Betting Odds

	ZINB (1) 1 Week/1 Week	ZINB (2) 2 Weeks/1 Week	ZINB (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.537 (0.007)	0.555 (0.011)	0.668 (0.088)
<i>N</i>	2,394	3,591	4,787

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and p -values are reported. The sample is restricted to MPs with betting odds greater than 0.1. Only MPs standing for the 2017 election are considered.

Table 2.14: One-sided Z -test $H_0 : \beta_{Subsample} \leq \beta_{FullSample}$

	Manchester Attack 1 Week/1 Week	Manchester Attack 2 Weeks/1 Week	Manchester Attack 3 Weeks/1 Week
<i>Z Statistic</i>	-0.733	-0.691	-0.860
<i>p-value</i>	0.768	0.755	0.805

Note: one-sided Z -test. Null hypothesis: the coefficient of the treatment effect on immigration-related Tweets in the subsample is equal or less than the coefficient of the treatment effect in the full sample. Manchester attack, all the available MPs are used.

relations reduce the conventional image of outgroups and enhance attitudes toward them (Legewie 2013). Hence, we should expect that, if the share of migrant population in a constituency is relatively small (i.e., the share of native people is large), voters might be more worried about the issue of immigration and the politician could exploit the event to signal her position on the political spectrum. Thus, we would expect those MPs to be more prone to expose themselves in the aftermath of the incident. My proxies for intergroup contact are the share of UK-born people and the share of people of white British ethnicity at the constituency level. However, these two measures present some drawbacks. First, they come from the 2011 Census, so they do not reflect the constituency's condition at the time of the event. Second, there might be concerns on how well these variables represent the same concept. In fact, the share of UK-born people also includes second-generation migrants, who might still be

considered outsiders from the other natives, so it does not discriminate between different ethnic groups. Moreover, as these variables come from the British Election Study 2017, I do not have information on the 16 constituencies of Northern Ireland present in my dataset.

Tables 2.15 and 2.17 show the results of the estimations. I report also a Wald test for the joint significance of the interaction term and the variable considered (tables 2.16 and 2.18). The terms are jointly statistically significant at less than 10% in every specification, and in all periods. The magnitudes do not differ substantially and the cumulated effects have positive sign. Yet, the estimated impact is quite small. For instance, considering table 2.15 and the closest time interval, if we raise the share of UK-born individuals in a constituency by 20 percentage points, the elected politician is predicted to increase its expected number of immigration-related Tweets by only 4.5% after the attack, compared to the others.²⁴

Table 2.15: Heterogeneity by Share of UK-Born People

	ZINB (1) 1 Week/1 Week	ZINB (2) 2 Weeks/1 Week	ZINB (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.130 (0.021)	0.135 (0.025)	0.194 (0.063)
<i>Interaction</i>	1.019 (0.079)	1.019 (0.074)	1.017 (0.098)
<i>UK-Born Share</i>	0.984 (0.036)	0.985 (0.011)	0.986 (0.026)
<i>N</i>	7,448	11,172	14,895

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported. All the available MPs are used.

Table 2.16: Wald Test for Joint Significance (UK-Born)

	Manchester Attack 1 Week/1 Week	Manchester Attack 2 Weeks/1 Week	Manchester Attack 3 Weeks/1 Week
<i>Wald Statistic</i>	5.557	7.679	6.414
<i>p-value</i>	0.062	0.021	0.040

Note: Wald test for the joint significance of *Interaction* and *UK-Born Share*.

The next assumption that I am going to test is related again to the contextual factors that might shape politicians' behaviour. Material concerns and perceived group deprivation could increase intergroup hostility. Adverse economic conditions might reduce collective resources and enhance outgroup threat,

²⁴The cumulated impact is computed as $(1.018526 \times .9839869)^{20}$, as the effect is multiplicative. The comparison group is represented by politicians affected by the treatment (i.e., tweeting after the attack), but belonging to constituencies with a share of UK-born people lower by 20 percentage points.

Table 2.17: Heterogeneity by Share of White British People

	ZINB (1) 1 Week/1 Week	ZINB (2) 2 Weeks/1 Week	ZINB (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.222 (0.004)	0.229 (0.006)	0.306 (0.027)
<i>Interaction</i>	1.013 (0.050)	1.013 (0.045)	1.013 (0.056)
<i>White British Share</i>	0.991 (0.073)	0.992 (0.038)	0.993 (0.068)
<i>N</i>	7,448	11,172	14,895

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported. All the available MPs are used.

Table 2.18: Wald Test for Joint Significance (White British)

	Manchester Attack 1 Week/1 Week	Manchester Attack 2 Weeks/1 Week	Manchester Attack 3 Weeks/1 Week
<i>Wald Statistic</i>	5.310	6.664	5.896
<i>p-value</i>	0.070	0.036	0.052

Note: Wald test for the joint significance of *Interaction* and *White British Share*.

as attitudes are likely to be shaped by the perceived impact of the outsiders at the community rather than at the individual level (Hainmueller and Hopkins 2014). Thus, we could expect that in constituencies facing downturns a politician would be more prone to exploit a terrorist event to highlight the issue of immigration and shift public attention toward this topic, using outgroups as scapegoats for the recession. Hence, we would anticipate a relatively higher number of immigration-related Tweets after the attack for those politicians elected in constituencies facing worse economic conditions. My proxy for competition over scarce resources is the average unemployment level in 2016, measured as the share of economically active residents aged between 16-64 years-old claiming Jobseeker's Allowance. Results are shown in table 2.19.

The estimated effect is quite large in magnitude, but it is only jointly statistically significant in the closest interval (see table 2.20). For instance, a politician elected in a constituency with an unemployment rate 2 percentage points higher, is predicted to increase the expect number of immigration-related Tweets by around 28% after the incident, compared to the others.²⁵ Thus, it seems that the impact differs, but just when taking into account the interval in the immediate proximity of the event.

²⁵The cumulated effect is computed as $(1.325405 \times .854322)^2$. The comparison group is represented by politicians affected by the treatment (i.e., tweeting after the attack), but belonging to constituencies with an unemployment rate lower by 2 percentage points.

Table 2.19: Heterogeneity by Level of Unemployment

	ZINB (1) 1 Week/1 Week	ZINB (2) 2 Weeks/1 Week	ZINB (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.277 (0.000)	0.378 (0.005)	0.499 (0.034)
<i>Interaction</i>	1.325 (0.007)	1.193 (0.122)	1.170 (0.138)
<i>Average Unemployment</i>	0.854 (0.040)	0.930 (0.363)	0.950 (0.437)
<i>N</i>	7,672	11,508	15,343

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported. All the available MPs are used.

Table 2.20: Wald Test for Joint Significance (Unemployment)

	Manchester Attack 1 Week/1 Week	Manchester Attack 2 Weeks/1 Week	Manchester Attack 3 Weeks/1 Week
<i>Wald Statistic</i>	7.831	2.398	2.200
<i>p-value</i>	0.020	0.301	0.333

Note: Wald test for the joint significance of *Interaction* and *Average Unemployment*.

Another important channel of heterogenous effects is that of incumbency. In particular, it is of interest to understand if politicians belonging to the party of the incumbent government act differently from the main challengers. If the “muting effect” is due to politicians strategically being cautious and avoiding a risky topic, we should expect a greater reduction for the members of the incumbent government’s party, as they might be deemed responsible for the current immigration policy. We thus select a subsample of the MPs: those only belonging to either the Conservatives (the incumbent) or the Labour Party (the main challenger). I re-estimate the model by adding a dummy variable for belonging to the Tories and an interaction with the treatment. Results are presented in table 2.21.

The effects are rather imprecisely estimated, and the two terms are jointly statistically significant (at the conventional levels) only when considering the largest time period (see table 2.22). However, the sign of the cumulated effect is as expected, but the magnitude is rather small. If we look at the third column, after the terrorist attack a Conservative MP is predicted to reduce her expected number of immigration-related Tweets by a further 5% compared to a Labour one.²⁶

²⁶The cumulated effect is computed as $(1.623777 \times .5876088)$. The impact is larger when considering the closest interval to the event (a reduction of around 14%), but it is only jointly significant at 11%. The comparison group is represented by Labour MPs affected by the treatment (i.e., tweeting after the attack).

Table 2.21: Heterogeneity by Incumbency

	ZINB (1) 1 Week/1 Week	ZINB (2) 2 Weeks/1 Week	ZINB (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.565 (0.018)	0.553 (0.009)	0.654 (0.047)
<i>Interaction</i>	1.695 (0.242)	1.566 (0.233)	1.624 (0.181)
<i>Conservative</i>	0.509 (0.037)	0.617 (0.045)	0.588 (0.014)
<i>N</i>	6,160	9,240	12,319

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported. Only observations for Conservatives and Labour are used.

Table 2.22: Wald Test for Joint Significance (Incumbency)

	Manchester Attack 1 Week/1 Week	Manchester Attack 2 Weeks/1 Week	Manchester Attack 3 Weeks/1 Week
<i>Wald Statistic</i>	4.375	4.204	6.519
<i>p-value</i>	0.112	0.122	0.038

Note: Wald test for the joint significance of *Interaction* and *Conservative*.

One last factor that might mediate politicians' strategic behaviour is the salience of the issue among their voters. However, the sign of the resulting effect is not straightforward. On one side, we could think that hostility towards immigrants in the electorate can be hazardously exploited by a politician to signal her ideology after the incident. On the other side, we might expect that, if voters are particularly sensitive to immigration and the politician adopts a risk-averse stance, she would avoid dealing with that issue in the aftermath of the attack, given the emotional impact that such events have on the public. My proxy for the salience of the topic is the share of votes for Leave in the 2016 EU Referendum. I assume that a higher proportion of Leave is suggestive of restrictionism in immigration policy, and so greater concerns about free movement of people. Unfortunately, the results of the Referendum are not available at the constituency level. Hence, I use the estimates computed by Hanretty (2017). As the data come from the British Election Study, I lose again information on the 16 constituencies of Northern Ireland present in my dataset. Results are shown in table 2.23.

The share of votes for Leave and its interaction with the treatment are jointly statistically significant in all time periods (see table 2.24) and their cumulated effect is rather substantial. Looking at the closest interval, if we increase in a constituency the share of votes for Leave by 10 percentage points,

the elected politician is predicted to reduce the expected number of immigration-related Tweets by an additional 23% after the attack, compared to the others.²⁷

Table 2.23: Heterogeneity by EU Referendum Results

	ZINB (1) 1 Week/1 Week	ZINB (2) 2 Weeks/1 Week	ZINB (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.173 (0.002)	0.324 (0.025)	0.485 (0.148)
<i>Interaction</i>	1.030 (0.019)	1.017 (0.135)	1.012 (0.261)
<i>Leave Share</i>	0.946 (0.000)	0.959 (0.000)	0.963 (0.000)
<i>N</i>	7,448	11,172	14,895

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported. All the available MPs are used.

Table 2.24: Wald Test for Joint Significance (EU Referendum)

	Manchester Attack 1 Week/1 Week	Manchester Attack 2 Weeks/1 Week	Manchester Attack 3 Weeks/1 Week
<i>Wald Statistic</i>	34.278	25.898	23.186
<i>p-value</i>	0.000	0.000	0.000

Note: Wald test for the joint significance of *Interaction* and *Leave Share*.

This last result motivates a closer look at the strategic behaviour of MPs belonging to those constituencies where the Leave vote scored high in the EU Referendum. Figure 2.5 displays the cumulated number of immigration-related Tweets during the election campaign for the twenty constituencies with the highest Leave share. From the chart it does not appear that politicians belonging to those areas were particularly keen on approaching the topic: 13 out of 20 did not touch upon the immigration issue at all during the election. In fact, when considering all constituencies, the Pearson’s correlation coefficient between the two variables is negative, even if not dramatically large (-0.2798). Thus, the pattern that seems to emerge is an inverse association: the greater the demand for restrictionism, the less the MP covers the issue of immigration in her electoral agenda. However, these findings are just suggestive, as MPs belonging to “Leave constituencies” might also tweet systematically less. I address this concern in table 2.25, which shows a regression of the daily number of immigration-related Tweets posted by a MP during the election (19/04/2017-07/06/2017) on the Leave share in her constituency.

²⁷The cumulated effect is computed as $(1.030285 \times .9459803)^{10}$. The comparison group is represented by politicians affected by the treatment (i.e., tweeting after the attack), but belonging to constituencies with a Leave share lower by 10 percentage points.

I estimate the model through a ZINB with the inflation factor π given by a logistic function of the daily number of Tweets, thus keeping into account the everyday use of the microblog by the politician. I cluster the errors at the MP level and use the same covariates as in my baseline. The result suggests again a negative relationship, and the coefficient is highly significant ($p < 0.01$). Taken together, these findings provide further evidence of politicians adopting a risk-averse attitude on immigration when their electorate is more sensitive to the issue.

Figure 2.5: Leave Share and Immigration-related Tweets During Election
Top 20 Constituencies

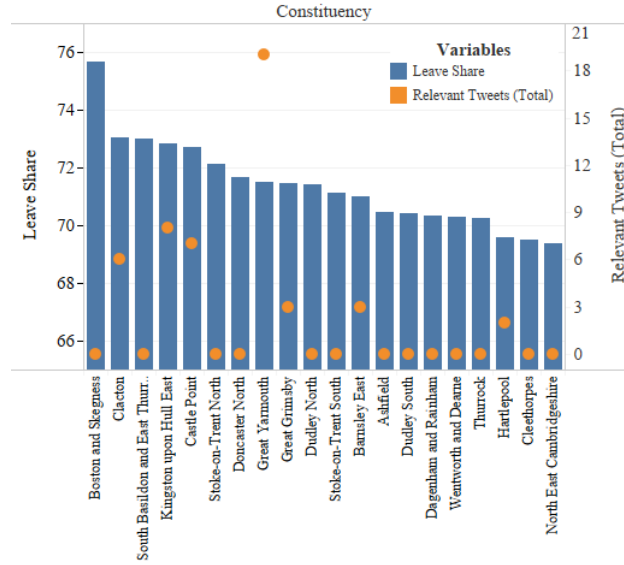


Table 2.25: Leave Share and
Immigration-related Tweets During
Election

ZINB (1) Election Campaign	
<i>Leave Share</i>	0.975 (0.001)
<i>N</i>	23,069

Note: day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRR and p -value are reported. Only MPs for whom I have complete information are used.

2.5.3 Robustness Checks

I now present a series of robustness checks for my baseline results on the Manchester attack, when I consider all the available MPs for whom I have complete information over the time periods considered. I start by performing different estimation strategies. First, I estimate the model with a standard negative binomial adding the total number of Tweets posted by the politician in a given day as a covariate. Results are reported in table 2.26. The effects are less precisely estimated, but the magnitudes do not change substantially: the difference in the IRRs remains between 2.3 and 3.3 percentage points from my baseline.

Table 2.26: Robustness Check 1: NB2 with Total Tweets as Covariate

	NB2 (1) 1 Week/1 Week	NB2 (2) 2 Weeks/1 Week	NB2 (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.646 (0.003)	0.671 (0.008)	0.820 (0.173)
<i>N</i>	7,672	11,508	15,343

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, total Tweets, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and p -values are reported. All the available MPs are used.

Next, I exploit the panel structure of the dataset to take into account unobserved individual heterogeneity. I estimate my baseline with a Poisson random effects (RE) model (Cameron and Trivedi 2013). This model is less demanding in terms of distributional assumptions than a negative binomial RE, but it is more efficient than a pooled Poisson when overdispersion is of the NB2 form (as I have assumed in my baseline). The choice of RE is justified by the fact that, given the exogeneity in the timing of the event, it is unlikely for the time-constant individual effect to be correlated with my treatment variable. I use the same covariates as in the baseline, adding the daily total number of Tweets posted by a MP as in the previous robustness check. Results are reported in table 2.27. The IRRs are very close to my baseline regressions and the effect is even more precisely estimated for the largest time interval.

The second exercise that I am going to perform is to exclude all MPs elected in London's constituencies, to control that my results are not driven by what is happening in the capital city.²⁸ Results are displayed in table 2.28. Again, my main conclusions are unaffected by this test: we still observe a substantial and significant decrease in the closest time intervals.

Subsequently, I slightly change the nature of my dependent variable in table 2.29: I construct a dummy

²⁸Out of the 73 constituencies of London, 65 are present in my dataset.

Table 2.27: Robustness Check 2: Poisson RE with Total Tweets as Covariate

	Poisson RE (1) 1 Week/1 Week	Poisson RE (2) 2 Weeks/1 Week	Poisson RE (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.615 (0.001)	0.632 (0.001)	0.777 (0.066)
<i>N</i>	7,672	11,508	15,343

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, total Tweets, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported. All the available MPs are used.

Table 2.28: Robustness Check 3: London's Constituencies Excluded

	ZINB (1) 1 Week/1 Week	ZINB (2) 2 Weeks/1 Week	ZINB (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.644 (0.008)	0.669 (0.011)	0.832 (0.225)
<i>N</i>	6,776	10,164	13,551

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported. 65 constituencies of London (out of 73) are excluded from the analysis.

at the MP level for posting an immigration-related Tweet in a given day. Hence, now I am looking at the likelihood of tweeting about immigration in the days following the attack. The results are consistent with my previous findings: focusing on the narrowest time interval, the probability of writing an immigration-related Tweet was around 38% less during the week after the incident, compared to the week before.²⁹

Table 2.29: Robustness Check 4: Probability of Immigration-related Tweets

	Logit (1) 1 Week/1 Week	Logit (2) 2 Weeks/1 Week	Logit (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.618 (0.003)	0.668 (0.007)	0.812 (0.152)
<i>N</i>	7,672	11,508	15,343

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, total Tweets, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. ORs and *p*-values are reported. All the available MPs are used.

Another interesting question that we might want to ask is if the effect captured for the Manchester attack is constrained to the MPs that were standing in the election or it is instead a more generalised

²⁹The model is estimated through a logit using the same covariates as in the baseline estimations, but keeping also the total number of Tweets by day per MP as independent variable. Errors are clustered at the MP level. Odds ratios (ORs) and *p*-values are reported.

result that applies to all politicians in charge. In fact, it should be pointed out that those politicians that were not trying to be re-elected, even if facing different incentives during the campaign, might adopt a strategic response as well, since their behaviour is likely to influence the odds of the candidate of the same party standing for their constituency. In table 2.30 I report my baseline estimations excluding the 23 MPs not standing in the 2017 election that are present in my dataset. If we compare the estimates with those in table 2.9 we notice that the IRRs are not affected by this exercise, the difference is less than 1 percentage point in every time interval. Hence, the observed behaviour seems to hold across all MPs. However, it should be highlighted that the politicians excluded are only a small portion of my sample, so their reaction to the terrorist attack should be substantially different to radically change the size and sign of the average effect estimated in my baseline.

Table 2.30: Robustness Check 5: MPs not Standing in 2017 Elections Excluded

	ZINB (1) 1 Week/1 Week	ZINB (2) 2 Weeks/1 Week	ZINB (3) 3 Weeks/1 Week
<i>Manchester Attack</i>	0.616 (0.002)	0.644 (0.003)	0.804 (0.124)
<i>N</i>	7,350	11,025	14,699

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported. Politicians not standing in the 2017 election are excluded from the analysis (23 in my dataset).

Finally, I address the suspension of the election campaign that occurred in the aftermath of the attack. After the Manchester bombing, to pay tribute to the victims, the leaders of all major parties agreed on suspending the campaigning activity, which was subsequently resumed at the local level after two days.³⁰ Hence, we might wonder if the “muting effect” that we observe it is not just the result of this political freeze. In order to explore this hypothesis I exclude the two days following the attack from my sample and re-estimate my models. However, this exercise implies a substantial reduction in the treatment group. In order to overcome the loss of efficiency given by a reduced sample size I narrow down the analysis to the two closest time intervals. This allows me to add 4 additional MPs for whom I have complete information over these periods. Table 2.31 reports the results for this last robustness check. The effect is less precisely estimated, but it is significant at conventional levels. We still observe a substantial proportional decrease in the expected number of immigration-related Tweets, around 30% when we compare the week before to the week following the attack. Thus, it seems that the

³⁰Source: *The Guardian* (<https://www.theguardian.com/politics/2017/may/23/general-election-campaigning-suspended-after-manchester-attack>, retrieved on the 13th of March 2018), *BBC News* (<http://www.bbc.co.uk/news/election-2017-40026416>, retrieved on the 13th of March 2018).

“muting effect” lasted even after this major campaigning shock.

Table 2.31: Robustness Check 6: Days of Suspended Campaigning Activity Excluded

	ZINB (1) 1 Week/1 Week	ZINB (2) 2 Weeks/1 Week
<i>Manchester Attack</i>	0.700 (0.048)	0.736 (0.058)
<i>N</i>	6,624	10,488

Note: the treatment takes the value of 1 on the day of the event and the subsequent 6 days, but 23/05 and 24/05 are excluded. Day-of-the-week dummies, sex, age, number of followers, number of friends, number of statuses, and age of the account are included. Errors are clustered at the MP level. IRRs and *p*-values are reported. All the available MPs are used.

2.6 Discussion and Conclusion

In this study I analyse the consequences of two acts of terrorism occurred in 2017 on the immigration rhetoric of British MPs. I focus on a specific information environment: the social media Twitter. My goal is to explore a potential channel through which these events might be exploited by political elites and in turn shape public opinion on immigrants. Natives’ attitudes towards outgroups are crucial for the social integration of minorities and the economic success of the community as a whole. Thus, it appears relevant to look at the role that political leaders might play in this process of perceptions’ formation. To answer this question, I scrape politicians’ Twitter accounts using text analysis and machine learning techniques in order to gather all their Tweets related to the issue of immigration. I then frame a natural experiment setting exploiting the exogeneity in the timing of the events and the granularity of the data gathered. I find a significant impact during the election campaign, but the direction of the effect is rather counterintuitive. In fact, political leaders might strategically exploit these dramatic episodes to foster the debate on immigration and divert attention from contextual problems. Alternatively, they might seek the opportunity to signal their ideology. In both cases, we would expect an increase in the amount of relevant information provided by politicians in the aftermath of a terrorist attack. In contrast, what we observe is a “muting effect”: a substantial decrease, on average, in the number of immigration-related Tweets in the week following the incident. My hypothesis is that, given the high stakes that they face during an election and the emotional distress caused by these dreadful events, MPs strategically prefer not to take a stance on a risky topic.

In order to further investigate this hypothesis, I construct a measure of attention to a Tweet and analyse if in the days following the attacks the public was more sensitive about the issue of immigration. Increased attention by their followers to this theme would justify a more risk-averse attitude, due to the open exposure to criticisms that the microblog implies. I conduct this exercise for the Westminster attack, in order to motivate a learning behaviour of politicians that would help to understand the observed results for the Manchester bombing. The attention variable is computed as $\ln(\text{Favourites} + \text{Retweets} + 1)$, for any immigration-related Tweet posted by a British MP in a given day.³¹ I include in the regression day-of-the-week dummies, sex, number of followers, number of friends, number of statuses, and another dummy for the message shared by the politician being a Retweet itself; as these are all factors that might affect the attention to a Tweet. The impact of the event is estimated through OLS and I compute robust standard errors. In the exercise, I compare the attention to immigration-related Tweets two weeks before and two weeks after the incident, so the treatment takes the value of 1 on the day of the attack and the subsequent 13 days. Results are reported in table 2.32. We can observe that in the aftermath of the event the public was definitely more sensitive about the issue. The attention to immigration-related Tweets posted by MPs increased by approximately 59%, compared to the two weeks before the attack, and the effect is statistically significant at 5%.³²

Table 2.32: Attention to
Immigration-related Tweets

	OLS (1) 2 Weeks/2 Weeks
<i>Westminster Attack</i>	0.466 (0.021)
<i>N</i>	428

Note: the dependent variable is $\ln(\text{Favourites} + \text{Retweets} + 1)$. The treatment takes the value of 1 on the day of the event and the subsequent 13 days. Day-of-the-week dummies, sex, number of followers, number of friends, number of statuses, and dummy for Retweet are included. Heteroscedasticity-consistent standard errors are computed. *P*-values are reported in parentheses.

These last findings appear to be consistent with the experimental results of Gadarian and Albertson (2014), who, building on Affective Intelligence Theory, show how anxious individuals exhibit increased and biased information seeking. Hence, it appears that the incident caused a greater attention to the

³¹I add 1 to the argument of the logarithm to account for Tweets that are neither favourite nor shared. The adjustment, however, should not be too problematic, as the percentage of zeroes in this restricted sample is around 2.3% (Wooldridge 2016). A log function is preferred to squeeze the distribution of Favourites and Retweets.

³²The effect is computed as $\exp^{0.466} - 1$.

issue of immigration and this attentiveness could have been exploited by political elites. However, for those who tried to do so after the Westminster attack, this turned out to be quite a risky strategy.

Donald Trump Jr., whose racist messages in the microblog already prompted widespread backlash,³³ was publicly denounced after his Tweet criticising London’s mayor Sadiq Khan in the aftermath of the Westminster attack. Among the critics, Wes Streeting, Labour MP for Ilford North, replied on the social media defining the US President’s son “*a disgrace*”, condemning his attempt to exploit the event for his own political gain.³⁴ Other European politicians, as the Front National leader Marine Le Pen or the Polish PM Beata Szydło, openly linked the attack to immigration policy and borders control.³⁵ This generated prompt reactions from different MPs and inflamed the debate in the Twitter community, especially among those users who blamed failed multiculturalism, as figure 2.6 shows. At the same time, Nigel Farage appeared on US television endorsing the hard-line immigration and anti-Muslim policies of President Trump. The former UKIP leader clearly connected the episode with British politics, blaming for the attack Tony Blair’s government which encouraged mass immigration and “*invited in terrorism*”.³⁶ He was then forced to draw back from his initial position and publicly admit no direct link between the event and the issue of immigration, once it was clear that the offender was actually British.³⁷

Thus, it seems that taking a stance was a dangerous strategy for both sides of the political spectrum, as it exposed the leaders to attacks and criticisms by opponents and the public. Hence, the observed “muting effect” for Manchester might be a consequence of politicians learning to avoid a risky topic when the electorate is more sensitive about the theme.

Digging deeper, I find significant heterogeneity in this “muting effect” according to the characteristics of the MPs or their constituencies, but also further evidence for a risk-averse attitude adopted by political elites in the aftermath of the attack.

A possible consequence of this reluctant behaviour is a potential mismatch between voters’ preferences and the actual type of politicians. Due to the increased information seeking and sensitivity after the

³³Source: *The Guardian* (<https://www.theguardian.com/us-news/2016/sep/20/donald-trump-jnr-compares-refugees-poisoned-skittles-twitter-reacted>, retrieved on the 6th of February 2018).

³⁴Source: *The Guardian* (<https://www.theguardian.com/uk-news/2017/mar/22/donald-trump-jr-tweet-london-mayor-sadiq-khan>, retrieved on the 6th of February 2018).

³⁵Source: *The Guardian* (https://www.theguardian.com/uk-news/2017/mar/23/anti-immigrant-politicians-link-london-attack-migrant-policy?utm_content=buffer9f72&utm_medium=social&utm_source=twitter.com&utm_campaign=buffer, retrieved on the 6th of February 2018).

³⁶Source: *The Independent* (<http://www.independent.co.uk/news/uk/home-news/nigel-farage-london-terror-attack-multiculturalism-blame-immigration-lbc-radio-ukip-mep-leader-a7645586.html>, retrieved on the 6th of February 2018).

³⁷Source: *The Independent* (<http://www.independent.co.uk/news/uk/politics/nigel-farage-admits-westminster-attack-immigration-fox-news-sean-hannity-a7650541.html>, retrieved on the 6th of February 2018).

Figure 2.6: Example of MP Reaction to Anti-immigration Rhetoric Following the Westminster Attack



event, the electorate might be more receptive to the few opinion leaders who are willing to expose themselves, irrespective of their quality. In particular, if the political leaders in charge adopt a risk-averse attitude and are less willing to take a stance, voters might become more sensitive to the rhetoric of anti-establishment parties and movements, which are not afraid to expose themselves given their firm position on such issues. However, a potential connection between the observed behaviour of politicians and the actual electoral outcomes is not pursued in this paper and it appears to be an interesting and unanswered question for future research.

A general concern with the analysis might be that the “muting effect” is a result of messages with extreme negative polarity being censored by Twitter itself. However, Twitter’s hateful conduct policy applies to rather extreme cases, such as *“promote violence against or directly attack or threaten other*

people”.³⁸ Hence, it seems very unlikely that the MPs will be so radical in their response to the event in such a critical juncture represented by the election campaign (given also the absence of far-right parties in the Parliament).

An interesting extension of my work would be trying to understand if the risk-averse behaviour that we observe is a decision of the single politician or a strategic response coordinated by the parties. The results presented in subsection 2.5.2 provide some suggestive evidence for the first hypothesis, as we detect greater variability at the individual level, with a stronger “muting effect” for MPs sitting on more marginal seats or belonging to constituencies where the issue of immigration is more salient. Instead, the response to the event among the two major parties does not seem to dramatically differ. However, a deeper study of the relationship between the MPs’ network on social media and the type of response to the attack among clusters of accounts could be carried out.

In addition, exploring other issues, more directly related to the nature of the attacks, as multiculturalism and Islamophobia, could provide more insights on the strategic reactions of politicians to these dreadful events.

One important limitation of the study is that it does not explore possible changes in the polarity of the Tweets. However, as explained in section 2.3, this is mainly due to my research design. In fact, one might wonder if the drop in the immigration-related Tweets is not capturing a reduction in the amount of messages, but a change in the wording around this theme instead. This is unlikely to be the case, as Twitter, with its 140 characters constraint,³⁹ does not allow for complex phrasing or involuted circumlocutions. Thus, my two-step classification should effectively cover the domain of interest (i.e., immigration). This conclusion will not be the same in a sentiment analysis framework. In such a setting, the choice of words characterising the polarity around the theme of interest is likely to change in response to the event and my classification exercise would not be able to capture this shift. Future work should try to overcome these constraints and capture if and how political elites shape their sentiment towards immigration following such shocks.

³⁸For Twitter’s hateful conduct policy, please visit <https://help.twitter.com/en/rules-and-policies/hateful-conduct-policy>.

³⁹The constraint refers to the one present during the period analysed.

Chapter 3

Strategic Voting in the United Nations General Assembly

Joint work with Omar A. Guerrero and Ulrich Matter

3.1 Introduction

Multilateralism is often justified by the democratic processes applied in international organisations where supranational decisions are reached. A democratic principle encountered in many international institutions is that member countries have the same weight in votes. From a democratic theory perspective this is rather peculiar, as it ignores the substantial variation in the number of citizens represented by each member country. From a politico-economic perspective, the real-world functioning of the equality principle is questionable, as it ignores the economic and military power disparities between the members. These disparities set strong incentives to not vote sincerely, but exploit the one country-one vote rule strategically in order to improve one's own position in international affairs.

A most prominent setting for such potential strategic interactions is the United Nations General Assembly (UNGA), the only one among the main six UN organs in which all member states have equal representation. While the principles of equality and unweighted voting among members have been at the core of the UNGA since its very foundation,¹ scholars in International Relations and

¹This is in stark contrast with other organs of the UN, as the Security Council (its executive body), where the perma-

Political Economy (e.g., Malone 2000; Eldar 2008, Carter and Stone 2015) have repeatedly questioned whether the UNGA indeed works under these conditions. Various testimonies about *exchange of votes*, *coercion* and even *direct payments*, suggest that votes in the UNGA do not necessarily reflect the countries' sincere preferences on international politics.²

It is challenging to disentangle different types of strategic voting from sincere votes because the countries involved might have incentives to conceal such behaviours. While one would expect some strategic interactions due to the heterogeneous composition of the UNGA and the different priorities of its member states, it is not clear whether these should be widespread and systemic. However, the literature on institutional design and voting behaviour suggests that one could infer certain types of potential strategic interactions from the institutional framework at hand (i.e., "the rules of the game") (Volden and Carrubba 2004; Aksoy 2012).

We build on the intuitive idea that the exchange of favours between countries in the form of support to proposed UN resolutions leaves patterns of reciprocity in the UNGA roll call data. For example, country A's representative is expected to vote against a particular resolution, given what we can observe about this country (location, form of government, economic development, etc.). Suppose now that country B sponsors this resolution and profits from its passage. In secrecy, B convinces A to support the draft in return for a favour. Thus, we observe that A deviates from the expected voting behaviour. If then later, we observe the same scenario vice versa, a pattern of reciprocity in favours (deviations) emerges. Following this idea, we model which type of strategic voting would be consistent with alternative patterns of reciprocity (or anti-reciprocity) in the data. Our conceptual point of departure is the UNGA's institutional setting (specifically, the equality principle) and the benchmark of sincere voting. We focus on the fact that reciprocity between representatives is a key aspect of prominent forms of strategic voting: vote buying is a trade of one's vote for goods/money; the other way around, coercion is (at the receiving end) a trade of a vote for non-retaliation; vote trading instead is a direct exchange of votes. Measuring such behaviours is challenging as these exchanges are generally not easy to capture (the parties involved in them have incentives to keep such deals hidden from the public). In terms of the simple example above, we can observe that A and B deviated, but we cannot

ment members (France, China, the United Kingdom, the Russian Federation, and the United States) hold veto power. However, as of March 2020, the use of this prerogative has varied a lot across the permanent members, ranging from the 16 vetoes cast by China to the 143 of the USSR/Russian Federation (<https://www.securitycouncilreport.org/un-security-council-working-methods/the-veto.php>).

²For example, Lockwood (2013) reports that in 2008 Iran bought the support of the Solomon Islands against Israel in future votes of the General Assembly for \$200,000 and technological aid. Another example is a United States law requiring the United States Agency for International Development (USAID) to make aid distribution dependent on the recipients voting in line with US interests in the UNGA (Carter and Stone 2015).

observe why. That is, we do not see whether or what they got in return for their potential favours. Thus, when applying statistical modelling to explain roll call data in the UNGA, we cannot necessarily rely on observables to capture such strategic interactions. Instead, we have to focus on the residuals: the part of voting decisions that statistical models cannot explain.³ By leveraging these residuals, our framework allows us to measure three distinct patterns of reciprocity in voting behaviour. First, deviations are random and no strategic voting (at least not related to reciprocity in our sense) took place. Second, there is a systematic pattern of reciprocity in vote favours, which is consistent with vote trading. Alternatively, there is a systematic pattern of anti-reciprocity, which is suggestive of coercion or vote buying.

In order to empirically study these cases, we formalise in a first step the expected behaviour of country representatives in the UNGA through statistical modelling. This allows us in a second stage to construct a network that describes the structure of the deviations underlying the observed voting decisions. We can then leverage the graph to compute an aggregate statistic: the *Reciprocity Index*. This measure is devised to capture the non-random occurrence of reciprocated deviations and it can be used to quantify the extent of different forms of strategic voting in the Assembly. In the second and third step, we specifically rely on and extend the framework suggested by Matter and Guerrero (2017).

While the absence of any type of strategic voting would be the “ideal” scenario, even a setting where vote trading were systematically prevalent could be interpreted as indicative of an effective equality principle and good functioning of the institution. In contrast, if vote trading is replaced by other forms of strategic interactions in the UNGA, it means that the equality principle is ineffective because votes are not considered a valuable “exchange currency” by all countries. In this case, the large asymmetries of power between the member states would play a much more preponderant role in the UNGA dynamics, manifesting themselves in the form of vote buying and coercion.

Through our methodology, we find that deviations from the expected votes in the UNGA are systematically *not* reciprocated. Our conclusions pass a series of robustness tests and are consistent with a narrative of vote buying and state socialisation (Alderson 2001). We also find weak evidence of vote trading *across* UN institutions. Our findings highlight the structural heterogeneity of the countries involved in the decision-making process and question the most distinctive features of the General Assembly: the unweighted voting system and the equality principle.

³Framing the study of strategic voting as an analysis of the residuals is an approach that stems from the seminal paper on vote trading of Stratmann (1992).

Our paper contributes to the literature in several ways. It provides a general framework to explicitly model and assess the strategic interactions that might occur in a deliberative organ. It does so by building a statistic that can effectively measure reciprocal or anti-reciprocal patterns of exchanges over time, providing deeper insights into the evolution of strategic voting across more than 40 sessions of the General Assembly.

The next section discusses the literature on voting dynamics in the UNGA and the phenomenon of strategic voting, including vote trading and vote buying, in several institutional settings. In section 3.3 we present the main characteristics of the UNGA and the data employed in the study. Section 3.4 introduces the methodological framework adopted, whereas section 3.5 focuses on assessing different predictive models of voting behaviour in the General Assembly. We then develop the main analysis in section 3.6. Section 3.7 discusses our baseline results and tests alternative hypotheses. Finally, we provide some reflections and conclusions in section 3.8.

3.2 Literature Review

The validity of formal equality among sovereign states is a controversial topic in international legal theory (Lockwood 2013). Nevertheless, it embodies the key feature of the General Assembly, the only organ in the UN system where all members are equally represented. Instead, the implications of unequal power relationships between countries have been mostly addressed in the context of the Security Council (O'Neill 1996; Voeten 2001). However, in this executive organ the disparity among players might be considered formally expressed in the institutional framework itself, given the veto power conferred to the permanent members.

With respect to the General Assembly, researchers in International Relations and Political Economy have mainly analysed the observed votes of countries to infer foreign policy preferences. The adopted approaches range from estimating dyadic similarity scores between states (Signorino and Ritter 1999) to ideal points on low-dimensional spaces through NOMINATE scaling (Voeten 2000). More recent developments employ Bayesian inference to discriminate between changes in the agenda of the General Assembly and shifts in countries' preferences over time (Bailey, Strezhnev, and Voeten 2017). This literature also focuses on the factors that affect such preferences. In particular, leadership turnover (Dreher and N. M. Jensen 2013, Smith 2016), regime type (e.g., see Brazys and Panke 2017b), the level of democracy and economic development have been shown to have a strong correlation with foreign

policy preferences (Kim and Russett 1996; Voeten 2004; Carter and Stone 2015; Smith 2016). Other studies have examined the voting behaviour of countries and the resulting dynamics of conflicts and alignments following major shocks to the international system, especially the structural changes caused by the end of the Cold War (Kim and Russett 1996; Voeten 2000).

However, there is an important caveat about these studies: the observed behaviour (e.g., a country's vote) is the outcome of the interplay between the formation of true preferences and the strategic interactions among political players. Seminal works on public choice have shown that, due to the mismatch between the intensity of preferences over the different decisions and the "rules of the game" generally adopted in a deliberative organ, actors with equal bargaining power can engage in negotiations to promote their own interests and enhance (theoretically) aggregate welfare (Tullock 1959; Buchanan and Tullock 1962; Tullock 1970). In these studies, logrolling (i.e., vote trading) is presented as a rational response to overcome the constraints imposed by the majority voting. That is, the equality principle (with the simple majority rule) sets incentives for strategic behaviour, given heterogeneous preference intensities.

Recent game-theoretical work instead has focused on characterising the dynamics of an actual market for votes (both uncoordinated or coordinated by leaders) under a simple majority rule. In such a setting, before making a decision, agents are free to buy and sell votes for a numeraire according to the intensity of their preferences (Philipson and Snyder 1996; Casella, Llorente-Saguer, and T. R. Palfrey 2012; Casella, T. Palfrey, and Turban 2014; Casella and Turban 2014). For instance, Philipson and Snyder (1996) show that, when the utility functions of the legislators are strictly concave and the distribution of most-preferred alternatives is skewed enough, a market directed by an intermediary can result in an equilibrium that represents a Pareto improvement compared to the outcome in the absence of trades (i.e., the median voter policy). These studies implicitly assume an effective equality principle because buyers cannot coerce sellers and neither the other way around. That is, there is no heterogeneity in the bargaining power of the agents.

Moving to the empirical literature, while there exist some works on vote trading at the national level, in particular on the US Congress (Stratmann 1992; Cohen and Malloy 2014), the issue has been addressed in international settings mainly from a normative perspective (Eldar 2008). In the context of the UN, the studies tend to report well-documented cases and the underlying dynamics between the states, especially with respect to the election of the non-permanent members of the Security Council (e.g., Malone 2000). However, the literature lacks of a quantitative assessment of this practice. One exception

is the work of Aksoy (2012), which explores how different decision rules and multidimensional issues affect the likelihood of vote trading within the European Union’s Council of Ministers. The author leverages the expert-based Decision Making in the EU (DEU) dataset to infer issues’ salience for the different members. Then, building on a spatial model of voting, she analyses the factors that influence legislators’ success in negotiating over the preferred drafts. The study exposes the practice of logrolling in a supranational environment, emphasising *within-legislation* logrolling (i.e., the bargaining at the level of the issues addressed within the same proposal).

In contrast to vote trading, vote buying has been widely studied in the International Relations and Political Economy literature. Despite the fact that some authors tend to analyse the two practices together (e.g., Eldar 2008), they present important theoretical differences. As pointed out by Lockwood (2013), vote trading is not inconsistent with political equality, since each voter is endowed with a single vote that acts as the only mean of exchange. On the other hand, vote buying appears to be much more in contrast with UN principles, especially if it systematically undercompensates poor countries or if it turns into direct coercion.⁴

Vote buying can take place through economic benefits for governing elites (Eldar 2008) or, alternatively, between donors and aid recipients by conditioning future financial flows. A substantial amount of research has documented a positive association between aid flows/loans and temporary membership to the Security Council (UNSC) (e.g., Dreher, Sturm, and Vreeland 2009a; Dreher, Sturm, and Vreeland 2009b). Special attention has been given to the role played by US aid and the salience of a UNSC’s seat (Kuziemko and Werker 2006). In the context of the General Assembly, researchers have mainly focused on the preponderant role of the US (Dreher, Nunnenkamp, and Thiele 2008; Carter and Stone 2015). Carter and Stone (2015) use the annual report of the State Department, *Voting Practices in the United Nations*, to identify votes that are relevant to American interests and explicitly model lobbying activity between the US and a potential recipient country as a two-stage game. The authors find that democracies that often tend to align with American interests do so because they are more likely to be subject to threats or promises of aid flows by the US.

In our study instead, we consider the role that all member states could play and the type of strategic interactions they might have. We propose a unified framework based on the concept of non-random

⁴Even if without a direct reference to it, the UN declaration on *Principles of International Law concerning Friendly Relations and Co-operation among States* condemns: “[...] the use of economic, political or any other type of measures to coerce another State in order to obtain from it the subordination of the exercise of its sovereign rights and to secure from it advantages of any kind.”. Source: A/RES/2625 (XXV) (<https://digitallibrary.un.org/record/202170?ln=en#record-files-collapse-header>).

reciprocity in directed networks (Garlaschelli and Loffredo 2004) to analyse the different forms of strategic voting that might occur within the UNGA and their implications for the underlying institutional setting (i.e, the equality principle).

3.3 Institutional Setting and Data Description

3.3.1 The United Nations General Assembly

The United Nations General Assembly (UNGA) represents the only forum of international policymaking within the UN where each of the 193 member states has equal representation and where a variety of global issues are debated. The UNGA meets annually in a regular session that usually runs from September to December, and it can resume in January until all items on the agenda are discussed. Proposals are generally drafted and initiated by a member state with the possible support of other co-sponsors. Drafts need to be submitted under a specific agenda item in order to be considered in a formal meeting of the UNGA. Every item in the agenda is assigned to one of the Main Committees or to the Plenary (even if it is possible to assign the same item to multiple bodies).⁵ Between the submission and the adoption of a draft (e.g., during the meeting of a Main Committee), other countries can decide to join the group of sponsors. The Main Committees always submit a report with their decisions and recommendations on the proposed drafts to the Plenary, which then takes final action.

Generally speaking, most of the decisions are made following the simple majority rule.⁶ Decisions involving elections (e.g., the Secretary General, non-permanent members of the Security Council, members of the Economic and Social Council) are held by secret ballot, but for every other resolution representatives may request a vote (either electronically or through a roll call). Countries then can cast a *yes* or *no* vote, and they can also abstain.

Even if the Assembly can take legally binding decisions with respect to its budget or internal procedures, most of its resolutions are non-binding for the members. However, they still carry significant political weight. In fact, they represent a clear statement of a country's stance regarding major international affairs.

⁵The six Main Committees of the UNGA deal with different issues: 1) disarmament and international security; 2) economic and financial matters; 3) social, humanitarian and cultural matters; 4) special political and decolonisation matters; 5) administrative and budgetary matters and 6) legal matters.

⁶However, for important questions, such as the admission of new members or issues related to peace and security, a majority of two-thirds is required.

3.3.2 Data Description

Our data contain information on votes and final sponsors of individual resolutions that were not adopted by consensus in the UNGA Plenary. This information was obtained by scraping the United Nations Bibliographic Information System.⁷ In addition, we also collected data on the voting date, issuing body session, agenda information, related documents (e.g., the committee report), the subjects of the resolutions, their title and UN code.⁸ Figures B.1 and B.2 in Appendix B.1 provides an example of the structure of the web pages in the United Nations Digital Library. Figure B.3 in Appendix B.2 instead shows the first page of a sample UNGA resolution. The final data contain 2016 resolutions for which a total of 203 countries voted on between 1976 and 2018, spanning 43 sessions of the UNGA (from the 31st to the 72nd, including the 10th Emergency Special Session).⁹ Table 3.1 reports descriptive statistics for the data.

Table 3.1: Descriptive Statistics

	N	Mean	Std. Dev.
<i>Countries</i>	203		
<i>Resolutions</i>	2016		
<i>Votes</i>	366910		
<i>Proportion of Sponsors</i>		0.178	0.146
<i>Proportion of Yes Votes</i>		0.738	0.147
<i>Resolutions (per Session)</i>		57	20

Table 3.1 shows that, on average, approximately 18% of the countries sponsor a resolution. The mean coverage in the dataset is 57 resolutions per session. The average proportion of votes in favour of a given proposal is around 0.74. Next, we look at how these statistics vary over time, to gain a better understanding of the voting patterns in the UNGA and coverage of our dataset.

Panel *a* in figure 3.1 shows two dramatic falls in the proportion of sponsorship. The first one took place around the 36th session (1981-1982), and it coincides with a major shift in US international


⁷The original website was <http://unbisnet.un.org/> (see <https://digitallibrary.un.org/record/409378>). However, at the time of writing, both the Bibliographic Records and Voting Records were moved to the United Nations Digital Library (<https://digitallibrary.un.org/>).

⁸The UN code is a unique identifier for documents in the UN system. It is made up by different components that generally indicate the issuing organ, possible subsidiary bodies, the nature of the document and the year or session of issue. For example, the resolution number 247 approved in the 74th session of the General Assembly is identified by the UN code A/RES/74/247.


⁹Unfortunately, even if older drafts were uploaded on a regular basis (which we managed to retrieve), the full coverage of the Bibliographic Records is officially from 1979 onward, limiting the temporal scope of our analysis. The final dataset does not include sessions for which we collected information on less than 10 drafts, to ensure a sufficient number of observations for every session considered.

Figure 3.1: Evolution of Sponsorship, *Yes* Votes, Resolutions and Membership over Time
Sessions 31st to 72nd (1976-2018)


(a) Average Proportion of Sponsors by Session




(b) Average Proportion of *Yes* Votes by Session



(c) Number of Resolutions Considered by Session



(d) Evolution of UNGA Membership by Session



Note: the 10th Emergency Special Session is not included in the figures above, as it spans over several years (1997-2018).

relations: the *Reagan Doctrine*. The new administration openly supported anti-Marxist insurgencies across developing countries, in an attempt to push back Soviet influence (Johnson 1988). This change in US foreign policy, with its aggressive stance towards international law (especially with respect to military interventions (Malawer 1988)), might have caused a more politically divisive environment in the UNGA. Such polarisation, in turn, could have hindered cooperation in the form of co-sponsorship of draft resolutions among member states. The second plunge occurred around the 45th session (1990-1991), concomitant with the final years of the USSR. It would not be surprising that the dissolution of the USSR caused a structural break in the dynamics of international cooperation. Around the same years, in panel *b* we observe an important drop in the mean proportion of *yes* votes, which could signal, once again, a more fragmented political landscape. It also appears that in the most recent sessions considered (starting from the 58th), both the trends in sponsorship and *yes* votes do not exhibit large fluctuations.

Panel *c* shows the coverage of the dataset over time. There, we notice that the number of resolutions included in the study increases significantly after the 52nd session, whereas the sample size is the

smallest in the previous period (between 1990 and 1998). Panel *d* reports the temporal evolution of UNGA membership. The most significant enlargement (from 159 to 179 members) occurred between the 45th and 46th session, with the collapse of the Soviet Union.

3.4 Methodology

3.4.1 Empirical Strategy

Following the institutional characteristics of the UNGA, we propose a strategy to empirically assess the strength and effectiveness of the equality principle. The strategy builds on a set of assumptions that are reasonable in this setting.

First, we expect that the intensity of preferences towards resolution issues is heterogeneous across the member states. Given the diverse nature of the proposals, the salience of a given resolution varies according to the political, geographical and cultural context of a country (Brazys and Panke 2017a). In addition, some decisions are targeted to specific members (e.g, developing countries) or geographical areas (e.g., the Middle East), so we might expect that they carry a different political weight for the actors directly involved and their allies.¹⁰ Moreover, despite the common effort to reach consensus on supranational issues (see more in the next subsection), many resolutions still trigger a vote, stressing important cleavages within the UNGA.

The second assumption entangles the three main types of strategic interactions considered in this institutional setting: vote buying, coercion and vote trading. It asserts that vote buying and coercion should crowd out vote trading. The rationale being that, if a country has the power to coerce or the resources to buy out a vote, it is likely that this form of interaction would be systematically preferred over the cost of exchanging its vote for another. Thus, these crowding-out dynamics should predominantly take place between countries with power asymmetries, which leads to our last assumption.

Third, there exist strong power asymmetries between the member states. It is evident that the unequal distribution of resources (e.g., military, economic) grants disproportionate influence to specific actors, which can shape world politics in many extralegal ways (Lockwood 2013). In addition, in the context of the UNGA, the seniority of a member and its familiarisation with the “rules of the game” can play

¹⁰One of the most notable examples are the drafts concerning the question of Palestine, which represent major political priorities for certain countries, as shown by the intense lobbying of the US in favour of Israel (e.g., Wikileaks 2009).

a major role in the ability of a country to develop political priorities and promote its own interests (Alderson 2001).

Before integrating these assumptions under a common framework, let us examine the “currency of exchange” that would be necessary to engage in each of the different strategic interactions that we consider. First, vote buying can use any type of currency (e.g., financial aid, trade, military assistance, etc.). Second, by definition, coercion does not need any currency. Third, vote trading has a unique exchange currency: votes.

Taking the three assumptions and the exchange currencies into account, we develop the following logic. In the presence of heterogeneous preferences, member states try to promote their favourite resolutions by engaging in strategic interactions (because the one country-one vote rule acts as a constraint). If the equality principle holds and every member state values its own vote the same, then vote buying and coercion should not be systemic. However, since the intensity of preferences is heterogeneous, the countries resort to vote trading, so a structural pattern of “atypical” votes should be observable. Furthermore, this pattern should expose the beneficiaries of such “deviations” from the expected voting behaviour. If there is systematic reciprocity between beneficiaries and deviators, then an assertion can be made about those votes acting as the currency of the trade.

Now, the second assumption implies that, if member states are buying votes or coercing other countries, then deviations should be systematically *not* reciprocated. What we would observe in this case is just one side of the exchange: some member states deviate from the expected behaviour to benefit other countries, but the latter ones do not return the favour by mean of a vote. Hence an opposite, anti-reciprocal pattern, should emerge from the voting data. In this situation, the equality principle is weak because member states have decided to overcome the one country-one vote constraint through other types of strategic interactions. Thus, it can be said that vote trading has been crowded out by vote buying or coercion.

Consequently, in order to assess the effectiveness of the equality principle in the UNGA, we develop a *Reciprocity Index* that measures which one of these alternative forms of strategic behaviour is the most prevalent.

3.4.2 Estimation of the *Reciprocity Index*

First, note that vote trading is the expected form of strategic voting given our institutional setting (i.e., heterogeneity in the intensity of preferences, simple majority rule, unweighted voting) if the equality principle holds. Hence, the proposed framework builds on the vote trading method developed by Matter and Guerrero (2017). Their approach starts from a fairly simple assumption: vote trading is a *quid pro quo* relationship. In our context, this translates into member state A 's representative deviating from her expected vote in order to support resolution Z favoured by member state B . In return, member state B 's representative will do the same for resolution X , which is preferred by member state A . Hence, countries A and B have strong preferences for resolutions X and Z respectively and are ready to trade their votes on a decision that they deem less relevant in order to promote their own interests. In this setting, the cost of the trade is represented by refraining from voting according to the actual preference on a proposal, whereas the benefit is given by the support we receive for our resolution. From this stylised example it is clear that the interaction between the countries in the model is characterised by two main elements: the deviation from the expected voting behaviour is *directed*, as it benefits the partner, and has to be *reciprocated* by the other agent. When accounting for reciprocal behaviour between more than two agents, network analysis is necessary.

Even if countries' representatives might deviate from their expected vote for a variety of (random) reasons, an aggregate pattern of reciprocity between agents is suggestive of systematic cooperation and thus vote trading. These strategic interactions can be represented as edges in a directed graph, where the nodes are the member states and the edges are deviations that benefit specific actors. In order to build such a network, we need three different types of information.

First consider N UNGA members voting on M different resolutions. Then, matrix \mathbb{V} (with dimensions $N \times M$) represents the pattern of actual votes, so each entry \mathbb{V}_{ij} takes the value of 1 if the observed vote for country i on resolution j is a *yes*, and 0 otherwise.

The second component is a matrix \mathbb{Q} ($N \times M$) where the entries capture the probability of country i voting *yes* on resolution j (i.e., the expected voting behaviour). We explain our approach to predict votes in section 3.5. A deviation of country i on resolution j is characterised by observing an actual vote that is inconsistent with our predictions; that is, country i voting *yes* on a resolution in which we would have expected a *no* according to the corresponding entry in \mathbb{Q} . Hence, we define a deviation of country i on resolution j every time the vote of the member state is positive, but the predicted

probability of a *yes* is less than 0.5. We call this probability threshold μ .¹¹ In other words, a country deviates from the predicted behaviour when it votes *yes* on a resolution where it was expected to vote *no* most of the times. The deviation outcome is coded as a binary variable where 1 represents a deviation and 0 a vote that was carried out as expected. These binary outcomes are stored in a matrix \mathbb{D} .

The last piece of information that we need in order to build a directed-deviations network (DDN), is to identify the beneficiaries of the deviations. In the UNGA context, it is reasonable to assume that the sponsorship of a resolution reveals an intense preference of the sponsor towards the proposal. So when a country i sponsors a resolution j , it is signalling a strong preference for that draft. We store information on countries' signals in a matrix \mathbb{S} ($N \times M$), where the entry \mathbb{S}_{ij} takes the value of 1 if country i sponsors the draft of resolution j and 0 otherwise.

Next, we construct an adjacency matrix \mathbb{A} describing a DDN between deviators and sponsors. This matrix is obtained from the dot product between the deviations matrix (\mathbb{D}) and the transpose of the signalling matrix (\mathbb{S}):¹²

$$\mathbb{A} = \mathbb{D} \cdot \mathbb{S}' \quad (3.1)$$

From this network, we can then study the degree of reciprocity of the *directed* deviations. The level of reciprocity between two countries i and k is given by $a_{ik}^{\leftrightarrow} = \min[\mathbb{A}_{ik}, \mathbb{A}_{ki}] = a_{ki}^{\leftrightarrow}$. We can think about this measure as the number of reciprocated “favours” between any two UNGA members, paid in the currency of votes. For example, if country i deviated 5 times in favour of country k and k only deviated 3 times in favour of i , the number of trades between the two is assumed to be 3. Summing over all the UNGA members considered we obtain the reciprocity estimator:

$$r = \frac{\sum_i \sum_k a_{ik}^{\leftrightarrow}}{\sum_i \sum_k \mathbb{A}_{ik}} \quad (3.2)$$

We can also extract the vote trading network (VTN) from the empirical DDN, preserving only those directed edges that are reciprocated. The resulting structure describes the pattern of mutual trades

¹¹In figure B.6 in Appendix B.7 we will explore the robustness of our results to alternative probability thresholds.

¹²In order to avoid the presence of inconsistent voting behaviour that might result from the estimation procedure (i.e., a member state deviating on a resolution that it sponsored), we ensure that all the elements on the main diagonal are set to zero.

between the countries in our institution.

Next, we can test the level of reciprocity in the empirical DDN against a null hypothesis. By building an ensemble of networks generated by random deviations, we obtain the expected level of reciprocity under the null:

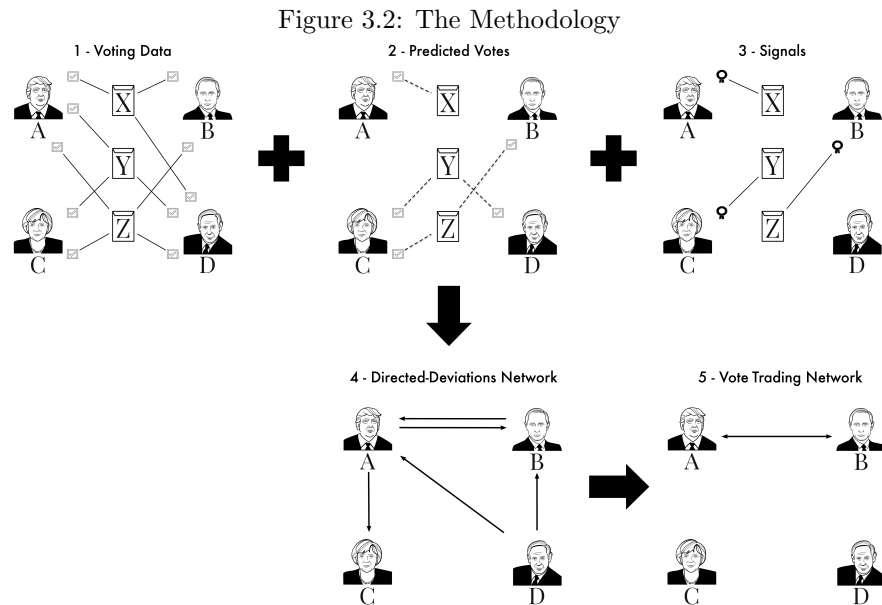
$$\bar{r}_0 = \frac{1}{N} \sum_i^N r_{0i}, \quad (3.3)$$

where N is the total number of null networks.

Finally, the *Reciprocity Index*:

$$q = \frac{r - \bar{r}_0}{1 - \bar{r}_0} \quad (3.4)$$

The measure captures the degree to which votes are used as an exchange currency for strategic interactions. It can be interpreted similarly to a correlation coefficient. A negative value of the index implies that deviations are systematically not reciprocated. That is, vote buying and coercion have crowded out vote trading and the equality principle is weak and ineffective. On the contrary, if $q > 0$ it means that vote buying and coercion are not prevalent, so member states overcome their differences in preferences by trading votes and the equality principle holds. Figure 3.2 provides a sketch of the methodology.



Before we move forward, it is important to highlight how we address a common assumption that is often made in the vote trading literature (e.g., Stratmann 1992; Cohen and Malloy 2014). That is, this form of strategic interaction occurs predominantly in *narrow* roll calls (i.e., those passed with a small margin). The rationale being that, if vote trading involves a search cost for the sponsors of a resolution, we would expect engagement only when the outcome is more uncertain and they need to seek a sufficient number of supporters in order to back up their proposal. On the other hand, when a large majority is already present, vote trading is unlikely to take place, and legislators would be free to behave according to their actual preferences. Hence, one would expect roll calls involving large margins to provide a more reliable picture of the true preferences of the legislators, whereas narrow roll calls are subject to the strategic considerations of the agents.

However, this theoretical argument of narrow margins might not hold in the context of the UNGA. Due to the non-binding nature of the resolutions, voting in the Assembly is mostly interpreted as expressive: a statement of a country's (claimed) political identity and intentions (Becker et al. 2015). Therefore, it is common practice to adopt resolutions by consensus whenever this is possible (Higgins et al. 2017), as a way to show a wide agreement on pressing international issues. As reported in the Annex IV of the *Rules of Procedure*,¹³ the Special Committee on the Rationalization of the Procedures and Organization of the General Assembly specifically considers that:

“[...] the adoption of decisions and resolutions by consensus is desirable when it contributes to the effective and lasting settlement of differences, thus strengthening the authority of the United Nations”

Hence, a resolution that is not unanimous is already a signal of political cleavages in the Assembly. In this setting, countries might exchange votes not just to pass a proposal, but to show that their own political view or course of actions have broad international support. In addition, ensuring “extra” votes is consistent with theoretical arguments about the formation of oversized majorities when policy logrolling have to be sustained over time (Volden and Carrubba 2004). Therefore, in our context, we do not limit the analysis to narrow voting outcomes, but we consider all roll calls as potentially subject to vote trading.

¹³Source: *Rules of Procedure of the General Assembly* (<https://undocs.org/en/A/520/rev.18>).

3.5 Predicting Voting Behaviour

In this section we test alternative specifications and classifiers to obtain the best possible vote predictions given the information provided to the models. We would like to emphasise that we are not trying to infer any causal mechanism driving the voting behaviour of countries in the UNGA. The primary objective here is to obtain reliable estimates of the expected votes, without compromising the underlying assumptions of the methodology. Thus, we never include information on the specific sponsors of a resolution, so that the predictions will always be oblivious of the beneficiaries of a proposal. Even if our final probabilities are going to be based on the full voting data, the method exploits the deviations from the expected votes, which result from the residuals between the observed behaviour and the factors accounted for by the predictive model. By ignoring the structure of the sponsorship across resolutions, we can exploit this “unexplained component” to find aggregate patterns that are consistent with our theoretical framework.¹⁴

A major challenge in predicting voting behaviour in our institutional setting is data availability. In order to build a network of interactions that can effectively captures reciprocity between the agents, we need to include all members of the General Assembly as nodes. That is, we have to predict the expected vote for every country on every resolution in our dataset. A sensible approach would be to employ as predictors those covariates that have been identified by the literature as important correlates of voting behaviour in the UNGA at the country level (e.g., the degree of democratisation or the political orientation of the government). However, these variables present a major shortcoming: their coverage is limited and most of them fail to include information on small or less influential countries. These countries can still play a pivotal role in the strategic interactions that we want to capture. By excluding them from the sample, we would arbitrarily cut relevant edges altering the topology of the network. This would cause a systemic bias in our estimates of the level of reciprocity.

3.5.1 Estimation of Latent Topics

To overcome this limitation we propose an alternative approach. We move our attention from the characteristics of the countries to the features of the resolutions i.e., the macro issues they deal with. The underlying rationale is clear: the topics of a resolution should be strongly correlated with a

¹⁴In figure B.9 in Appendix B.7 we replicate the analysis using a “narrow margin” approach, which is consistent with previous studies on the empirical estimation of vote trading (e.g., Stratmann 1992). The results obtained confirm the general pattern observed in the baseline.

country’s vote. By using the topics of the resolutions instead of the traditional covariates identified by the literature, we will be able to leverage all the information available in the voting and sponsorship data. To infer the latent issues of the resolutions we proceed in the following way. From the dedicated web page of a given resolution in the UN portals, we retrieve the labels representing the main subjects of the text. These labels broadly describe the themes, actors involved and reference documents related to the resolution (see figure B.1 in Appendix B.2).¹⁵ Tables 3.2 and 3.3 report descriptive statistics and the labels with the highest frequency for different time periods. The time intervals are defined by the Soviet–Afghan War (1979) and the dissolution of the USSR (1991).

Table 3.2: Descriptive Statistics for Labels

	Mean	St. Dev.	Min.	Max.	N
<i>Labels</i>	14.239	13.152	1	122	3363

There are over 3300 unique labels that could act as independent variables. In order to reduce such high dimensionality and infer the latent macro topics of the resolutions, we employ Structural Topic Models (STMs) on the collection of labels (Margaret E. Roberts et al. 2014).

Under STMs, every document d can be described as an array of proportions θ_d , where each element θ_{dk} is the fraction of words in document d allocated to topic k .¹⁶ In addition, STMs allow for correlation between topic proportions and inclusion of covariates that can arbitrarily affect topical prevalence (i.e., how much a topic is discussed) and topical content (the words used to define a given topic). We use this feature of STMs to account for changes in the agenda and priorities of the UNGA over the years. In our application, topical prevalence (how frequently a macro issue is dealt with in the UNGA) is a flexible function of time.¹⁷

Prior to conducting any analysis, we preprocess the collection of labels. First, we split the labels into tokens stemming each word via Porter’s algorithm. Then, we remove stopwords,¹⁸ numbers, punctuation and all words with less than three characters. In addition, when creating our document-feature matrix (i.e., the matrix that describes the frequency of terms in the corpus), we consider both

¹⁵As an alternative example, the labels attached to the resolution A/RES/67/121 *Israeli practices affecting the human rights of the Palestinian people in the Occupied Palestinian Territory, including East Jerusalem* are: “Human Rights in Armed Conflicts”, “Palestine Question”, “Palestinians”, “Territories Occupied by Israel”, “Israel”, “East Jerusalem State of Palestine”, “Settlement Policy”, “Separation Barriers”, “Freedom of Movement”, “Geneva Convention Relative to the Protection of Civilian Persons in Time of War 1949”.

¹⁶In our analysis, a document is represented by the combination of the labels attached to a given resolution.

¹⁷More specifically, topical prevalence is a function of the year in which the resolutions were passed, where the functional form is a spline with 10 degrees of freedom.

¹⁸These are terms that carry no information for the purpose of identifying a topic, for example, articles and prepositions. The Snowball collection of English stopwords is used for this task.

Table 3.3: Labels Frequency (Top 12)

1946-1978		1979-1990	
<i>Apartheid</i>	22	<i>South Africa</i>	158
<i>South Africa</i>	22	<i>Apartheid</i>	91
<i>Information Dissemination</i>	18	<i>Israel</i>	79
<i>Self-determination of Peoples</i>	12	<i>Conferences</i>	75
<i>Israel</i>	10	<i>Information Dissemination</i>	72
<i>Members</i>	9	<i>Palestine Question</i>	67
<i>Namibia</i>	9	<i>Middle East Situation</i>	54
<i>Namibia Question</i>	9	<i>Treaties</i>	48
<i>Military Assistance</i>	8	<i>Military Relations</i>	44
<i>Troop Withdrawal</i>	8	<i>UN. Special Committee against Apartheid</i>	43
<i>Middle East Situation</i>	7	<i>Territories Occupied by Israel</i>	42
<i>Palestine Question</i>	7	<i>Palestinians</i>	39
1991-2018		1946-2018	
<i>Territories Occupied by Israel</i>	302	<i>Israel</i>	382
<i>Israel</i>	293	<i>Territories Occupied by Israel</i>	348
<i>Palestine Question</i>	221	<i>Palestine Question</i>	289
<i>Report Preparation</i>	196	<i>Conferences</i>	254
<i>Disarmament Agreements</i>	184	<i>Middle East Situation</i>	226
<i>Accessions</i>	181	<i>Information Dissemination</i>	202
<i>Ratifications</i>	181	<i>Accessions</i>	201
<i>Signatures</i>	180	<i>Ratifications</i>	201
<i>Conferences</i>	173	<i>Palestinians</i>	200
<i>Human Rights in Armed Conflicts</i>	168	<i>Report Preparation</i>	199
<i>International Obligations</i>	167	<i>Signatures</i>	199
<i>Middle East Situation</i>	165	<i>South Africa</i>	195

unigrams and bigrams, in order to capture meaningful combinations of words that are informative in our context (e.g., the bigram “human_rights”).

One of the key issues in using STMs is the choice of the fixed number of topics, that has to be selected *a priori* (Grimmer and Brandon M. Stewart 2013). To identify the most appropriate number of topics given the corpus at hand, we adopt the following strategy. We first define a lower (K_L) and upper bound (K_U) for the number of topics K . Next, we estimate an individual STM for every value K_i in

the range $[K_L, K_U]$.¹⁹ For the lower bound, we set $K_L = 6$, which reflects the six Main Committees in the General Assembly. This number of macro issues is also consistent with previous research on voting in the UNGA (Bailey, Strezhnev, and Voeten 2017). For the upper bound K_U , we use a selection procedure based on the algorithm proposed by Mimno and Lee (2014), which automatically infers the number of topics.²⁰

We rank the set of estimated models using different criteria (Margaret E Roberts, Brandon M Stewart, and Tingley 2014): semantic coherence (based on the co-occurrence of top topic words within documents), exclusivity (the probability that top words for a given topic are unlikely under other topics), and residuals dispersion (Taddy 2012).²¹ The model that achieves the highest combined rank across the different criteria has $K = 19$ and it is chosen as the preferred one.

We qualitatively assess a selected number of topics from the preferred model in figure 3.3. The first panel presents histograms showing the terms with the highest probability for each of the topics. Below, we report the titles of their most representative document.²² The inferred clusters of terms seem to define three major issues: racial discrimination, the question of Palestine and non-proliferation of nuclear weapons.

As further validation of the qualitative sanity check performed before, we conduct a more quantitative assessment in figure 3.4. Namely, we analyse the relevance of a given issue during the period 2000-2015 and its sensitivity to UN endogenous processes. For the inferred topics to be meaningful, their prevalence in the corpus should reflect the salience of an issue within the UN. For instance, racial discrimination has increasingly gained relevance in the documents adopted in the early 2000s, peaking around a major UN initiative: the *Durban Review Conference*, which took place in 2009. The meeting assessed the implementation of the *Durban Declaration and Programme of Action* (DDPA) adopted at the *World Conference against Racism* in 2001.²³

Moving to the topic “question of Palestine”, we see that, starting from 2009 with the *Goldstone Report*,²⁴ the issue has become more and more salient within the UNGA. The expected topic proportion

¹⁹In STMs, the posterior distribution is likely to be multimodal, resulting in estimates that are sensitive to the starting values of the parameters. Thus, we initialise the models with the spectral algorithm proposed by Arora et al. (2013), a deterministic approach which is stable and outperforms alternative initialisations in terms of convergence speed (Margaret E Roberts, Brandon M Stewart, and Tingley 2014).

²⁰The resulting number of topics obtained through this selection strategy is 72, so that the final range of values is $[6, 72]$.

²¹For this last criterion, we rank the models according to their fitted overdispersion.

²²By *most representative*, we mean the document d that ranked first according to its value of $\hat{\theta}_{dk}$ (our estimate of θ_{dk}) for the given topic k .

²³The DDPA advocates action-oriented measures to combat racism and other forms of discrimination and intolerance.

²⁴The *Goldstone Report* was an independent international team established by the United Nations Human Rights

Figure 3.3: Word Frequencies and Most Representative Documents for Selected Topics

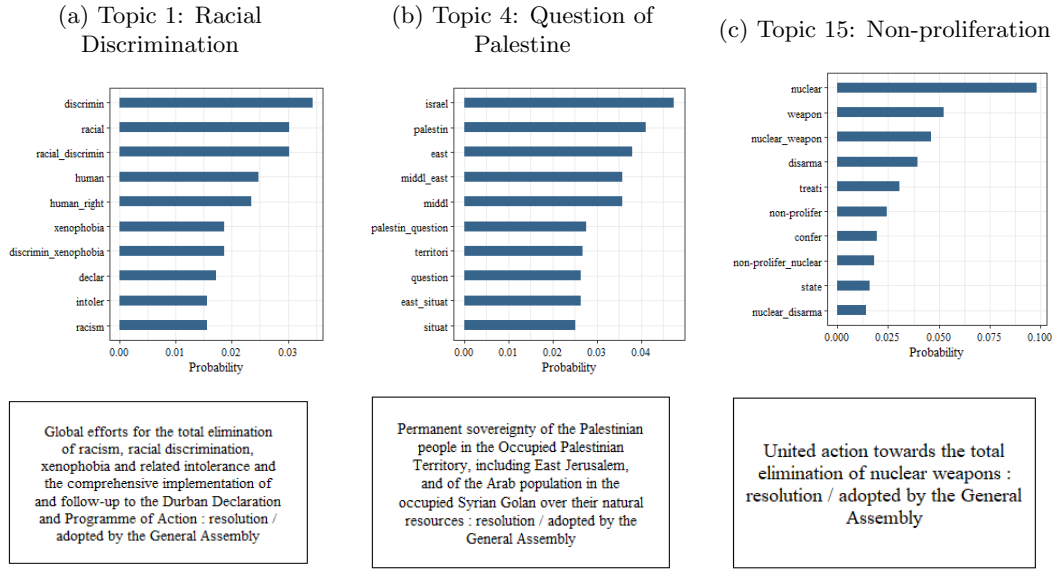
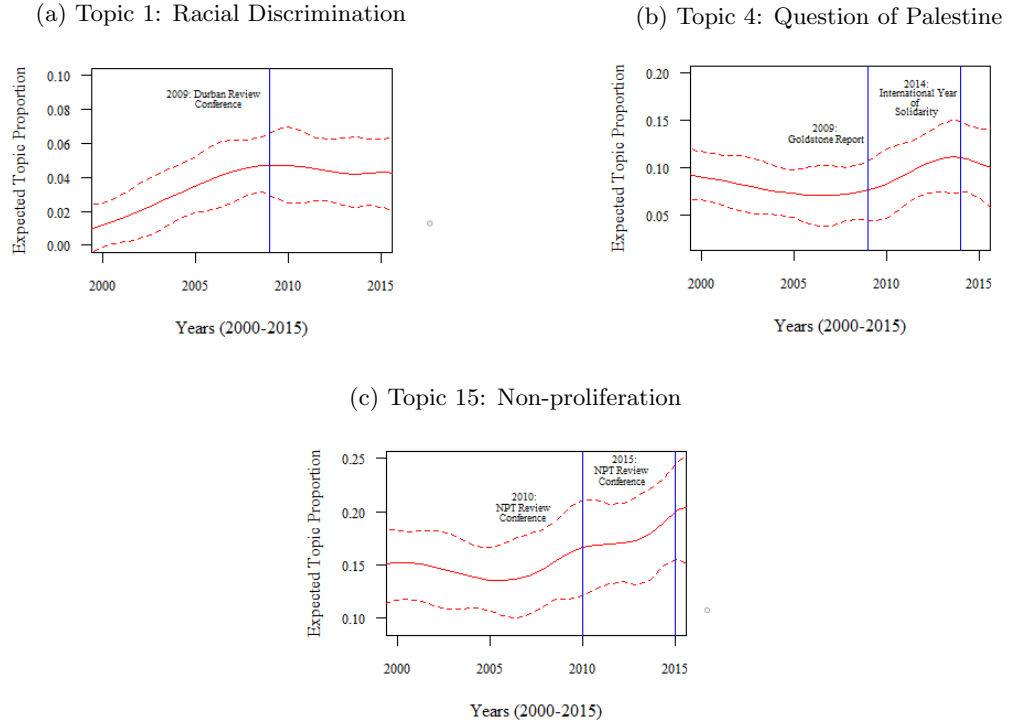


Figure 3.4: Prevalence of Selected Topics over Time (2000-2015)



Note: dashed lines denote 95% confidence intervals.

peaks around 2014, with the *International Year of Solidarity with the Palestinian People* proclaimed by the General Assembly. In the same year, the *Israel–Gaza Conflict* was taking place.

Council. Its purpose was to investigate potential violations of international human rights after the Gaza War.

We can observe an even stronger positive trend for the non-proliferation issue. The topic has become particularly relevant over the past decades. We observe two peaks around 2010 and 2015, with the 8th and 9th *Review Conference of the Parties to the Treaty on the Non-proliferation of Nuclear Weapons* (NPT).²⁵ It is also interesting to mention the international context in which these two conferences took place. In 2010 the first *Nuclear Security Summit* was held. The aim of the event was to promote nuclear security and prevent terrorism around the world. It occurred every two years until 2016. In 2015 instead, the *Joint Comprehensive Plan of Action* was signed. This agreement between Iran, the permanent members of the Security Council, Germany and the European Union imposed important limitations to the Iranian nuclear programme in order to avoid the risk of nuclear proliferation.

The estimated topic proportions are then combined into a linear regression as leading covariates, together with country and year fixed effects. We also add the fraction of sponsors for a given resolution as independent variable, to take into account possible herd behaviours (Banerjee 1992) that might result when countries do not have sincere preferences on a particular issue. Formally, the specification can be written as:

$$y_{ijt} = \alpha + \beta \text{Sponsor}_j + \hat{\theta}_j \boldsymbol{\delta} + \mathbf{x}_i \boldsymbol{\gamma} + \mathbf{z}_t \boldsymbol{\tau} + \epsilon_{ijt}, \quad (3.5)$$

Where the dependent variable y_{ijt} takes the value of 1 if country i voted *yes* on resolution j in year t and 0 otherwise. The variable Sponsor_j is the proportion of countries sponsoring resolution j . The vectors \mathbf{x}_i and \mathbf{z}_t contain dummies at the country and year level respectively. The vector $\hat{\theta}_j$ contains the MAP estimates of topic fractions for resolution j obtained through the STM with $K = 19$.²⁶

3.5.2 Topic-based Model vs. Traditional Approaches

To validate our approach, we compare the predictive power of our topic-based model with a specification that leverages the covariates identified by the literature. More specifically, the latter model includes a measure of aid relationships (Carter and Stone 2015; Brazys and Panke 2017b) captured by Net Aid Transfer (NAT) flows and a dummy variable for the country being a recipient or not;²⁷ a proxy for the degree of democratisation (Kim and Russett 1996; Voeten 2004; Carter and Stone 2015; Bailey,

²⁵This conference is held every five years since the NPT came into force (1970). It is meant to assess the Treaty and further strengthen its provisions.

²⁶As an intercept is included in the estimation, we omit one of the topic proportions to avoid perfect collinearity.

²⁷Information on NAT flows comes from the dataset compiled by David Roodman (2006), available from his personal website: <https://davidroodman.com/data>.

Strezhnev, and Voeten 2017; Brazys and Panke 2017b; Pauls and Cranmer 2017) for which we employ the combined score of the standard indicators from Freedom House and the Polity IV Project; the political orientation of the chief executive’s party (Dreher and N. M. Jensen 2013; Carter and Stone 2015; Bailey, Strezhnev, and Voeten 2017), her time in office and if legislative or executive elections were held, to consider the potential impact of leadership turnover and regime change (Dreher and N. M. Jensen 2013; Smith 2016; Brazys and Panke 2017b). We also use the natural logarithm of GDP per capita as a proxy for state capacity (Brazys and Panke 2017b), GDP per capita growth and the natural logarithm of population. To take into account trade openness and dependence (Kim and Russett 1996; Carter and Stone 2015; Bailey, Strezhnev, and Voeten 2017), we include the sum of exports and imports as a share of GDP and an index of economic globalisation, which considers also financial flows.²⁸

In addition, since much of the literature on voting behaviour in the UNGA is framed relative to US preferences (e.g., Voeten 2004; Carter and Stone 2015), we create an indicator variable for all those resolutions that were classified as crucial for the American interests by the Department of State and on which it is known that the country lobbied extensively.²⁹ Finally, we add the proportion of members sponsoring a given resolution, together with country and year fixed effects.

Table 3.4 presents the comparison between the topic-based model (1) and the one that includes the traditional covariates (2). In model 3 we simply replace all the predictors that vary over time at the country level with country-year fixed effects, keeping only the proportion of sponsors and the indicator for the US-relevant resolutions as additional covariates. The sample size is defined by the coverage of the variables employed in model 2. It consists of 132,544 observations at the country-resolution level over 34 sessions of the UNGA.³⁰

To assess the predictive power of the three different models we perform K -fold cross-validation and report a range of performance metrics.³¹ These are the F-score (the weighted harmonic mean of

²⁸All the variables mentioned above come from the Quality of Government (QoG) Standard dataset, available from the QoG Institute website: <https://qog.pol.gu.se>.

²⁹Key resolutions are identified in the Congressional report *Voting Practices in the United Nations*, first published in 1984. The reports can be accessed directly through the archived content available from the US Department of State website: <https://www.state.gov>.

³⁰More specifically, we can only consider those resolutions in our dataset that were passed between 1983-2015 and predict the vote for just a subset of countries (125). In addition, we exclude those sessions for which we have information on less than 10 resolutions.

³¹In this procedure we sequentially split the sample in K equally sized parts. Then we leave out the k^{th} subsample and fit the model to the remaining $K-1$, whereas the k^{th} part is used to compute the performance statistics. The process is repeated for $k = 1, 2, 3, \dots, K$ and the results are then averaged. We choose five fold cross-validation (i.e., $K = 5$) as it represents a good compromise between the bias and variance of the results (Hastie, Tibshirani, and Friedman 2009). Compared to the validation set approach, K -fold cross-validation offers a number of advantages over some major shortcomings of the former method. In particular, the validation set approach can be highly sensitive to the observations included into the two subsamples (the training and test sets) and the results obtained might

precision and recall), the share of false negatives (which in our framework represent the deviations from the expected voting behaviour), accuracy, and the AUC (the area under the ROC curve).³²

We observe that the topic-based model (1) outperforms the others with respect to AUC (0.760), F-score (0.860), and accuracy (0.769). Hence, our approach can improve predictions compared to a specification that relies only on traditional covariates. This allows us to extend the analysis to all available voting data, without constraining it to the subsample defined by the variables included in model 2.

Table 3.4: Model Comparison (I) with Five Fold Cross-validation

	OLS (1)	OLS (2)	OLS (3)
<i>F-score</i>	0.860 (0.002)	0.855 (0.002)	0.856 (0.002)
<i>Share False Negatives</i>	0.038 (0.002)	0.032 (0.002)	0.050 (0.002)
<i>Accuracy</i>	0.769 (0.003)	0.758 (0.003)	0.764 (0.003)
<i>AUC</i>	0.760 (0.002)	0.743 (0.003)	0.753 (0.004)
<i>N Train</i>	106036	106036	106036
<i>N Test</i>	26508	26508	26508
<i>Country FE</i>	Yes	Yes	No
<i>Year FE</i>	Yes	Yes	No
<i>Country-year FE</i>	No	No	Yes

Note: the table shows the average values of the metrics. Standard deviations are reported in parentheses.

3.5.3 Alternative Classifiers

Once we have benchmarked our topic-based model against more traditional approaches, we can fully leverage the information available in the voting data. The sample now includes 366,910 observations at the country-resolution level spanning 43 sessions of the UNGA (1976-2018).³³ With such amount of data, it is not necessary to restrict the prediction task to a linear model, so we test alternative classifiers to further improve predictive power. Table 3.5 reports performance metrics for four different classifiers assessed through five fold cross-validation: a linear probability model (LPM), a lasso regression, a

display a large variance. In addition, only a part of the observations is used to fit the model (James et al. 2013).

³²Precision is computed as: $\frac{TP}{TP+FP}$; recall as: $\frac{TP}{P}$; the share of false negatives as: $\frac{FN}{P}$ and accuracy as: $\frac{TP+TN}{P+N}$.

Where P is the positive sample, N is the negative sample, TP is the number of true positives, TN is the number of true negatives, FP is the number of false positives and FN is the number of false negatives.

³³As for the previous analysis, sessions for which we have information on less than 10 resolutions are dropped.

random forest (as in the original formulation of Breiman (2001a)),³⁴ and a neural network with a single hidden layer (Venables and Ripley 2002).³⁵

Both the random forest and the neural network outperform the LPM and the lasso regression in terms of F-score, accuracy and AUC. Note that the neural network performs slightly better than the random forest in terms of F-score (0.911 compared to 0.910) and accuracy (0.865 compared to 0.860). However, the latter still achieves the highest AUC (0.914) and it is more conservative in terms of false negative rate (0.044 against 0.064), so we choose it as our preferred classifier.³⁶

Table 3.5: Model Comparison (II) with Five Fold Cross-validation

	LPM	Lasso	Random Forest	Neural Network
<i>F-score</i>	0.849 (0.001)	0.850 (0.001)	0.910 (0.001)	0.911 (0.002)
<i>Share False Negatives</i>	0.054 (0.001)	0.040 (0.001)	0.044 (0.001)	0.064 (0.002)
<i>Accuracy</i>	0.753 (0.001)	0.751 (0.001)	0.860 (0.001)	0.865 (0.003)
<i>AUC</i>	0.762 (0.001)	0.762 (0.001)	0.914 (0.001)	0.896 (0.002)
<i>N Train</i>	293528	293528	293528	293528
<i>N Test</i>	73382	73382	73382	73382

Note: the table shows the average values of the metrics. Standard deviations are reported in parentheses.

Once we choose the random forest, we retrain the learning method with the whole sample. Then, we construct matrix \mathbb{Q} containing all the predictions on individual votes by every member state on each resolution. Note once again that the specific structure of signals underlying a resolution (i.e., the identity of the sponsors) is not included among the features of the model, so its predictions are oblivious of the potential beneficiaries of a *yes* vote. This is a key component in exploiting the deviations from the expected votes (i.e., our predictions) when building the DDN.³⁷ Now we move to the last step

³⁴For the choice of the hyperparameters, we set the node size to 0.1% of the observations in the training data. The main tuning parameter (the number of trees) is set to 100, a choice that is consistent with the range suggested by Oshiro, Perez, and Baranauskas (2012) (64-128) and it represents a good compromise between predictive power and processing time. However, we notice that, after 70 trees, the performance of the classifier exhibits only marginal improvements.

³⁵The size of the hidden layer and the weight decay are chosen with nested five fold cross-validation, as suggested by Varma and R. Simon (2006) and in line with Hastie, Tibshirani, and Friedman (2009). In the inner loop, we perform a grid search across four different values for the number of neurons (5, 10, 15, 20) and five values for the decay of the weights (0.1, 0.2, 0.3, 0.4, 0.5). The model with the highest average accuracy in the inner loop is selected and passed to the outer loop. Higher number of nodes in the hidden layer has only marginal effects on the performance of the classifier with a substantial cost in terms of processing time.

³⁶In figure B.5 in Appendix B.7 we assess the robustness of the analysis to alternative classifiers and we reproduce the main results using the neural network to obtain the matrix of predictions \mathbb{Q} .

³⁷When all the data is used to train the model (i.e., all the 366,910 observations), the random forest achieves a F-score of 0.927, a false negative rate of 0.033, an accuracy of 0.888 and a value of the AUC of 0.952.

of our methodology where we compute the *Reciprocity Index* to assess the extent of strategic voting within the UNGA.

3.6 Analysis of Strategic Voting in the UNGA

To compute the null hypothesis on which the *Reciprocity Index* from equation 3.4 is built, we follow the suggestion in Matter and Guerrero (2017) and simulate random deviations that are unrelated to strategic voting. That is, for every entry $Q_{ij} < 0.5$, we perform a Bernoulli trial to determine if member state i casts a *yes* vote for resolution j .³⁸ The intuition is that, since i is expected to vote *no* with a probability higher than 50%, a *yes* vote is considered a deviation from her predicted behaviour. Because there is no strategic element in these trials, the directed-deviations network (DDN) constructed from these deviations describes reciprocal patterns that result from random errors. When repeating this process multiple times, we obtain an ensemble of null networks used to compute the reciprocity estimator 3.3 under the null which, in turn, allows obtaining the *Reciprocity Index*.³⁹

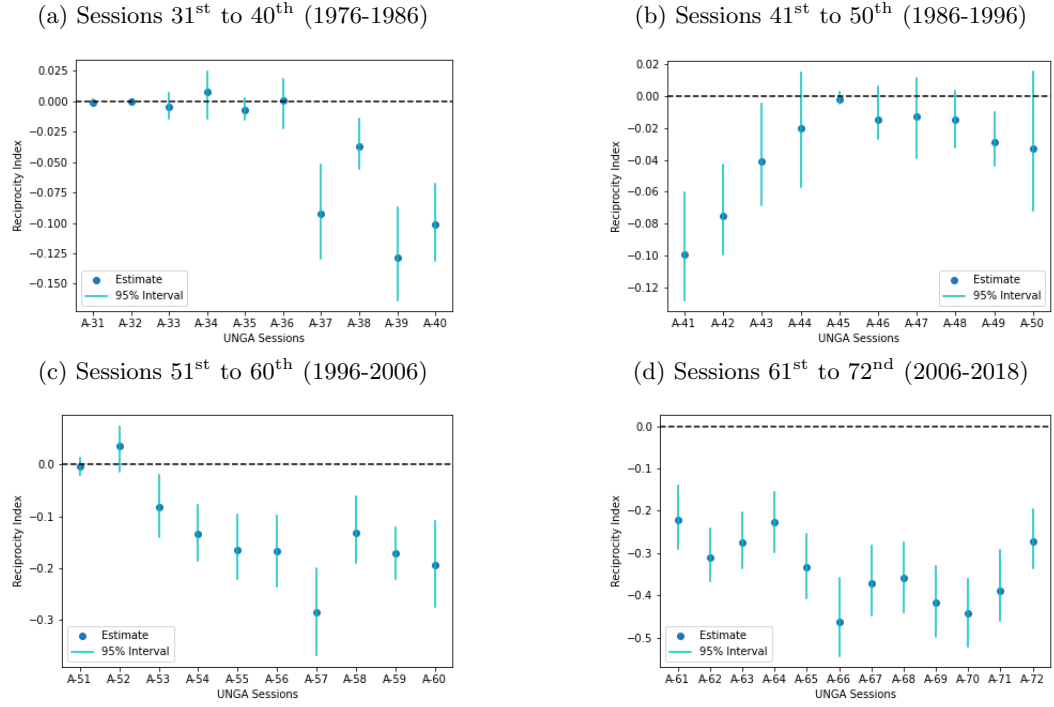
We compute indices for each session held during the sampling period. For this level of analysis (session level), we assume that deviations are reciprocated within the same annual meeting of the UNGA.⁴⁰ This, of course, might not always be the case as hidden agreements between countries could stretch across different sessions due to the recurrence and persistence of certain topics (Brazys and Panke 2017b). However, we address this concern in figures B.7 and B.8 in Appendix B.7 and show that our findings are robust across different levels of analysis.

Figure 3.5 reports the evolution of the *Reciprocity Index* over time. Between 1976 and 1982, the value of the index is not statistically different from zero. However, in the following years, we observe a generally negative trend, with the only exception being the 52nd session. Moreover, starting from the 53rd (1998), the index becomes strongly negative and significant. It is important to mention that, even in sessions for which we have a number of resolutions that is quite low, we are still able to capture indices that are significantly negative (e.g., between the 39th and 43rd session). In addition, in recent years (after 2010), the large magnitude of the index suggests strong anti-reciprocal patterns (values ranging from -0.27 in the 72nd session to -0.46 in the 66th). This is unlikely the consequence of major

³⁸In these trials the probability of success is proportional to the predicted probability of a *yes* vote on that resolution (i.e., the entry Q_{ij}).

³⁹We compute 500 null networks. For the confidence intervals (CIs) of the estimated indices, we adopt a bootstrapping procedure detailed in Appendix B.6.

⁴⁰We exclude the 10th Emergency Special Session, as it spans over different years (1997-2018). It is thus inconsistent with the assumption of “favours” (i.e., deviations) being reciprocated during the same annual meeting.

Figure 3.5: Evolution of the *Reciprocity Index* over Time

Note: point estimates and 95% confidence intervals are reported. The CIs are constructed using 500 bootstrap samples. The dashed line indicates a value of 0 for the index. The classifier used to predict voting behaviour is the random forest proposed in section 3.5. The units of analysis are individual sessions.

changes in the *observed* voting or sponsorship behaviour at the aggregate level, as both had a rather flat trend during the same period (figure 3.1).

3.7 Implications of Anti-reciprocal Deviations

Let us recall the logic behind the idea of using vote trading to assess the effectiveness and strength of the equality principle.

Given that preference intensities are highly skewed across member states in UNGA and the simple majority rule is generally adopted, we should expect the emergence of strategic behaviour. If the equality principle holds, then votes are a valuable exchange currency and we should observe systematic patterns of vote trading. If not, other forms of strategic interactions (e.g., vote buying and/or coercion) might crowd out vote trading. Since these other behaviours are based on power asymmetries, it would mean that this unbalance, and not the equality principle, is the actual driver of the “rules of the game” in the UNGA.

This can also be explained in terms of the anti-reciprocal DDNs. Our results suggest that, at the aggregate level, member states tend to deviate from the predicted voting behaviour in a non-random way. However, reciprocity in terms of directed vote deviations is systematically avoided. That is, initial deviations are not paid back in “the same currency”, i.e., with votes. Let us explore potential explanations for this phenomenon.

Power Asymmetries

The typical theoretical setting of vote trading assumes that actors are homogeneous in terms of bargaining power. Each side of the agreement has to deviate from its (weak) preferences on a specific draft in order to pursue its own interests. In the presence of power asymmetries instead, an influential sponsor could approach a potential deviator with an inferior bargaining position, so the latter may not be able to refuse the request of the former. In the extreme, this could even take the form of unilateral coercion. Let us assess this argument from different perspectives.

3.7.1 The Role of UNSC Permanent Members

We perform the analysis one more time, but excluding the few actors that could benefit from a position of extreme authority within the UNGA. In the context of the UN system, a sensible choice are the permanent members of the Security Council, an executive organ that is closely related to the UNGA. These countries are the United States, the Russian Federation (after the dissolution of the USSR), China, France and the United Kingdom. Through their veto power, these actors can block any substantive decision of the Council. This asymmetry heavily influences power relations between the member states and the decision-making process in the UN. We examine if our results are the reflection of the disproportionate power that this group of countries could exercise over the Assembly (Caron 1993; O’Neill 1996; Eldar 2008).

Before excluding these actors, we want to understand if the permanent members of the UNSC have a central role in the structure of the DDNs that we estimated, so that their removal would cause a major change in the topology of the networks. We run a set of Kolmogorov-Smirnov tests (one for each session) in which we compare both the indegree and outdegree distributions of the DDNs with and without these influential actors.⁴¹ Only in one case out of 42 sessions ($\approx 2\%$) the test is significant

⁴¹The null hypothesis of the test is that the two sets of data are drawn from the same distribution.

at the 5% level for the outdegree distributions. When looking at the indegree distributions, only three statistics ($\approx 7\%$) are significant at the 5%.⁴² Thus, it appears that these countries are not pivotal for the general topology of our networks.

We then explore their local level of reciprocity and we compare it with the reciprocity estimator for the whole General Assembly. To make the two statistics comparable, we normalise the level of reciprocity for a permanent member of the UNSC i by the sum of the incoming deviations:

$$r_i = \frac{\sum_k a_{ik}^{\leftrightarrow}}{\sum_k \mathbb{A}_{ki}} \quad (3.6)$$

If the permanent members drove the negative values of the *Reciprocity Index*, we should expect their level of reciprocity being far lower than the one observed at the aggregate level. Figure 3.6 compares the two statistics for each of the permanent members separately. We observe that, although with an irregular trend, the reciprocity exhibited by the US, UK and France is generally higher compared to the values for the whole UNGA. This does not appear to be the case for China and the Russian Federation (in the most recent sessions), for which we observe values that are consistently lower than the aggregate.

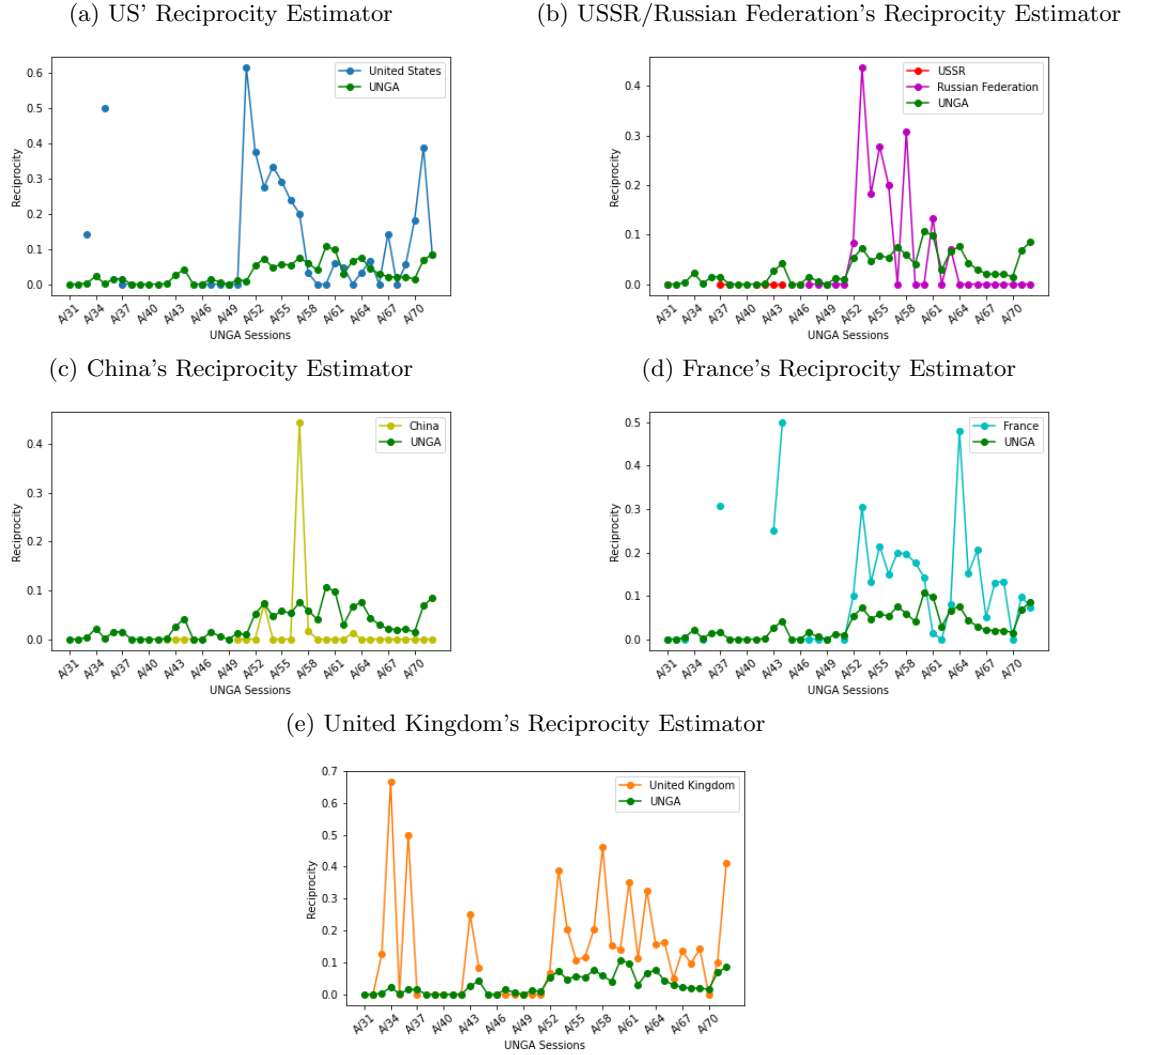
In addition, these nations seem to differ also with respect to their rate of sponsorship. From figure 3.7 we can see that the US in particular and, to a lesser extent, the Russian Federation (and the USSR before its dissolution), tend to sponsor a number of resolutions below the average. On the other hand, if we look at the most recent sessions (i.e., from the 54th), it appears that China, France and the United Kingdom drafted a number of proposals above the average. However, they are never among the top ten sponsors in any of the sessions considered.

We now proceed to the removal of the UNSC permanent members from matrices \mathbb{D} , \mathbb{S} and \mathbb{Q} , to understand if the behaviour of the index is mainly driven by this small group of countries. Notwithstanding their leverage, from figure 3.8 we observe that our baseline results are fundamentally unaffected by the omission of these actors: the *Reciprocity Index* still exhibits significant negative values in the large majority of the sessions analysed.

Thus, it appears that the anti-reciprocal patterns that we observe are not solely determined by few influential actors. This does not exclude that the permanent members of the UNSC might signal their

⁴²Table B.1 in Appendix B.3 reports the full set of tests.

Figure 3.6: Reciprocity Estimator for the Permanent Members of the UNSC



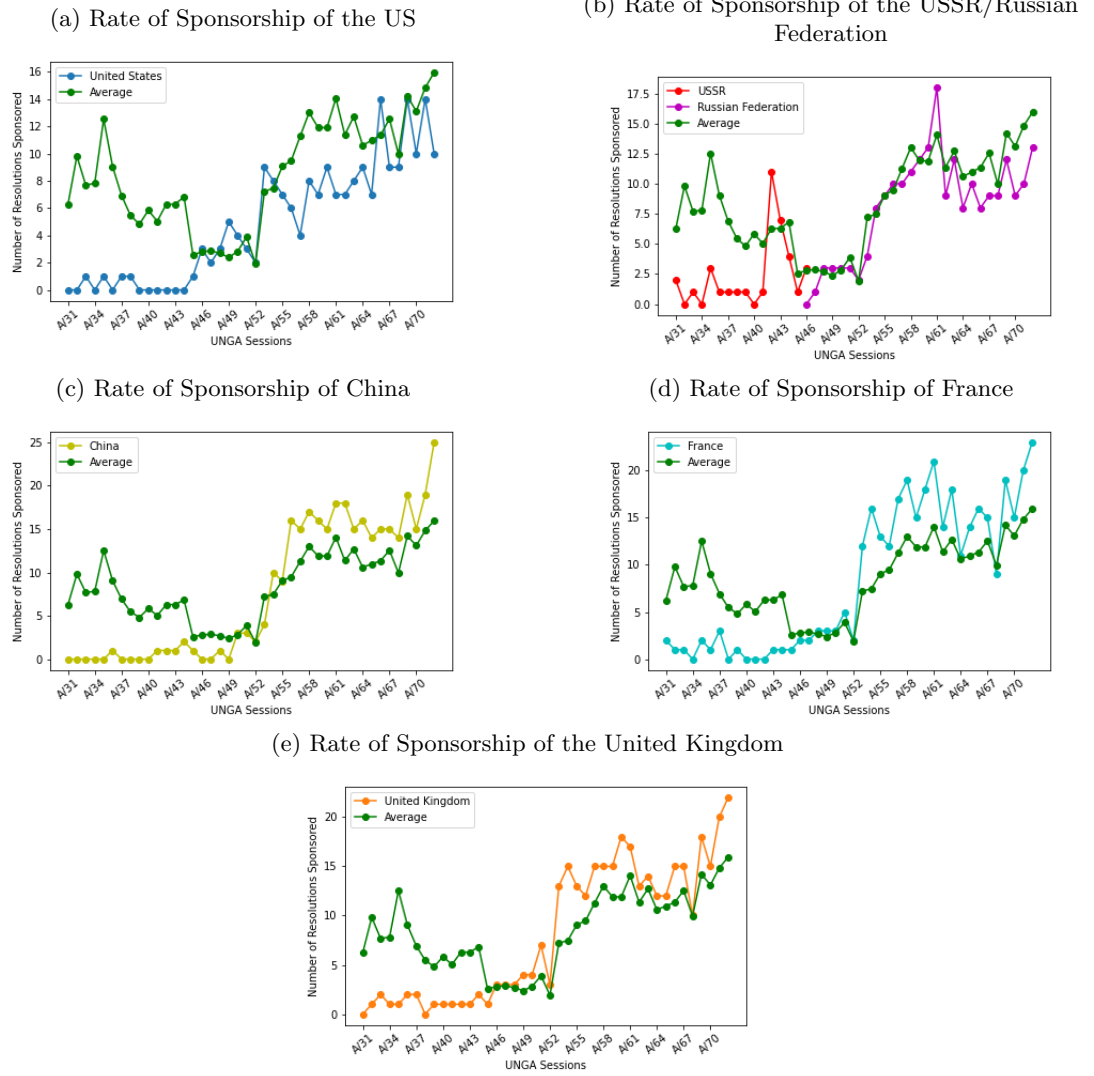
Note: in every graph the reciprocity estimator for the UNSC permanent member is reported together with the one for the whole UNGA. Missing values represent sessions in which the number of incoming deviations for the permanent member is equal to 0.

priorities through different channels (e.g., their speeches addressing the Assembly during the General Debate) that are not leveraged by our methodology.

3.7.2 Generalised Inequality in Resources

Let us focus on the whole set of countries. We may be inclined to think that, among the member states, some can systematically rely on their financial, military or political resources to influence the decision-making process. In this scenario, these members could pursue their own agenda at the expense of the less powerful, driving the decisions of the latter ones. Thus, one would expect nations

Figure 3.7: Rate of Sponsorship of the Permanent Members of the UNSC

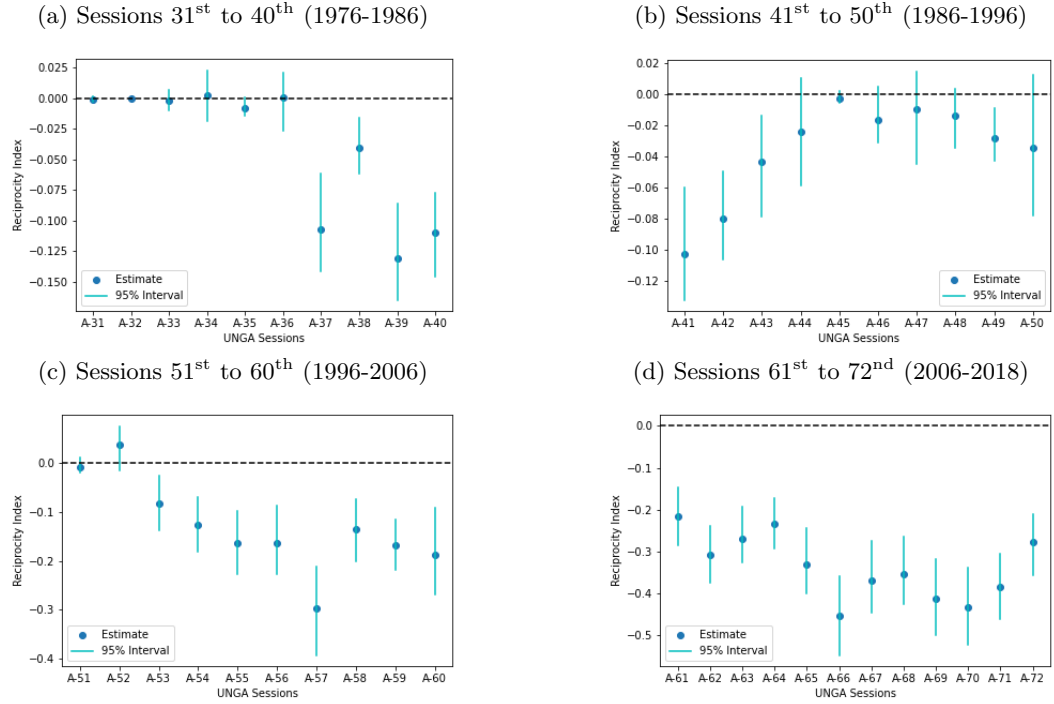


Note: in every graph the rate of sponsorship of the UNSC permanent member is reported together with the average for all countries in the UNGA.

in a weaker position to exhibit more unpredictable voting behaviour (i.e., deviating more), as subject to the external influence of the others. According to this logic, in the DDN, the number of incoming and outgoing deviations of a country should systematically relate to specific attributes. In the analysis that follows we are going to test this argument.⁴³

⁴³For descriptive purposes, tables B.2 and B.3 in Appendix B.4 report the three countries with the highest outdegree (outgoing deviations) and indegree (incoming deviations) for each session analysed.

Figure 3.8: Evolution of the *Reciprocity Index* over Time
Permanent Members of the UNSC Omitted



Note: point estimates and 95% confidence intervals are reported. The CIs are constructed using 500 bootstrap samples. The dashed line indicates a value of 0 for the index. The classifier used to predict voting behaviour is the random forest proposed in section 3.5. The units of analysis are individual sessions. The United States, the Russian Federation (and USSR), China, France and the United Kingdom are excluded from the analysis.

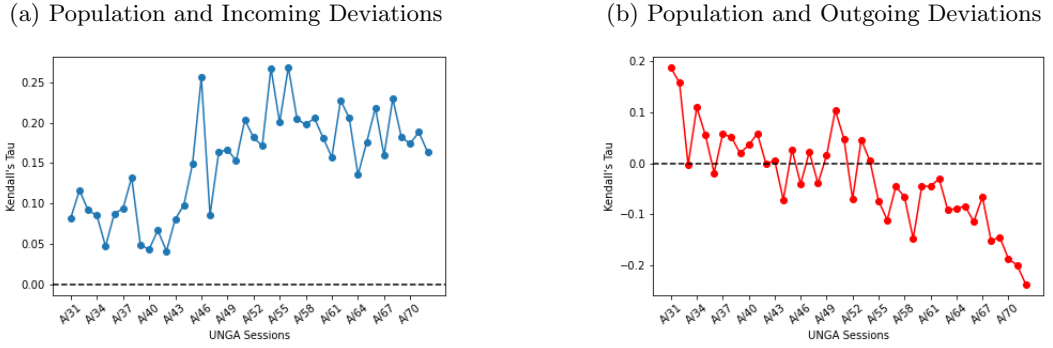
Population Size

A factor that can potentially affect the bargaining between nations is their size. By representing a higher share of the world's population, larger countries could be enabled to promote their own interests and sway the opinion of smaller states (as suggested by Brazys and Panke (2017b)). We formalise this hypothesis in the following way:

H1: the size of a country is positively related to the number of incoming deviations and negatively associated to the number of outgoing deviations.

We first address this hypothesis by studying the monotonic relationship between a country's size and its deviations over the different sessions considered. For both graphs in figure 3.9 the vertical axis shows the Kendall's τ (to capture possible non-linear associations). However, in the left panel this is computed between a country's population and the number of incoming deviations, whereas in the right panel we substitute the latter with the number of outgoing edges.

Figure 3.9: Evolution of the Correlation between Population and Deviations



Note: observations are constrained by the information available on countries' population. See panel *a* of figure B.4 in Appendix B.5 for the data coverage over the different sessions. Data come from the Quality of Government dataset: <https://www.gu.se/en/quality-government/qog-data/data-downloads>.

Focusing on the most recent sessions (i.e., after the 54th), we observe a positive relationship between size and incoming deviations, whereas the outgoing ones display a negative association. The magnitude of the coefficient, however, displays a large variability. Looking at the period after the 54th session in the right panel, the statistic ranges from -0.030 in the 62nd session to -0.236 in the 72nd, and 8 out of these 18 coefficients are statistically significant at the 5% level.

Economic Size

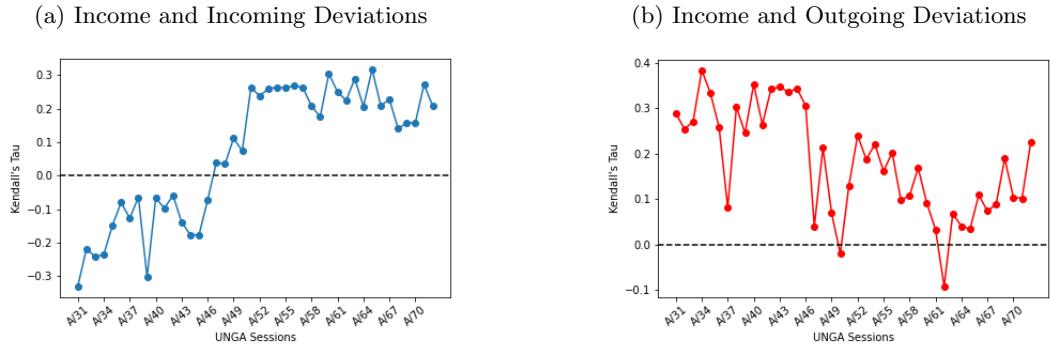
Economic resources are an alternative factor driving disparities in bargaining relationships. These are not just a proxy for diplomatic capacity (Brazys and Panke 2017b), but they also capture a nation's ability to resist specific forms of external pressures such as vote buying. In fact, the absence of systematic vote trading is consistent with this narrative, because bought deviations are not paid in votes.

The empirical literature studying the UNGA suggests that vote buying can take the form of development aid flows or even financial benefits for governing elites (Eldar 2008). In both cases, one would expect less affluent countries to be more prone to these economic incentives and their vote swinging more frequently in favour of the most wealthy ones. We formalise this argument in the following hypothesis:

H2: the income of a country is positively related to the number of incoming deviations and negatively associated to the number of outgoing deviations.

We analyse again the rank correlation between the variables over time. Figure 3.10 presents the two trends, providing mixed evidence for $H2$. In panel *a*, starting from the 47th session, we observe a positive correlation. In the years after 1992, the value of the coefficient ranges from 0.035 in the 48th session to 0.317 in the 65th, with 23 out the 26 coefficients statistically significant at the 5% level. However, we also observe, on average, a positive association in panel *b*.

Figure 3.10: Evolution of the Correlation between Income and Deviations



Note: income is operationalised through a country's GDP per capita at current prices. Observations are constrained by the information available on countries' income. See panel *b* of figure B.4 in Appendix B.5 for the data coverage over the different sessions. Data come from the Quality of Government dataset: <https://www.gu.se/en/quality-government/qog-data/data-downloads>.

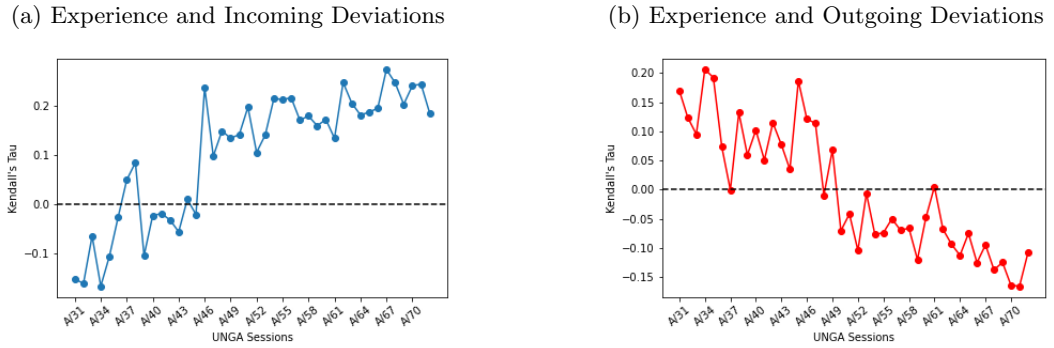
Seniority and Socialisation

Nations may differ in their degree of state socialisation (Alderson 2001), and so in their familiarisation to international norms. Countries that joined the UNGA earlier had more time to develop their own agenda in the global debate. Less experienced members, on the other hand, might have not yet formed strong preferences and stable voting decisions on several supranational issues (Brazys and Panke 2017b). In such scenario, old members could take advantage of the newcomers by influencing their votes, resulting in a more unpredictable behaviour of the latter ones. Thus, a third hypothesis can be stated as follows:

H3: the experience of a country in the UNGA is positively related to the number of incoming deviations and negatively associated to the number of outgoing deviations.

Figure 3.11 presents the correlation between countries' seniority in the UNGA and their incoming and outgoing deviations. The two panels support $H3$ when focusing on the second half of the sample period. Seniority is positively associated with incoming deviations; with values spanning from 0.098 in the 47th session to 0.275 in the 67th. It is interesting to observe how the coefficient becomes steadily positive

Figure 3.11: Evolution of the Correlation between Experience and Deviations



Note: experience is operationalised through the number of years of UN membership. Data come from the list of current members, in the UN website: <https://www.un.org/en/member-states>.

after the 45th session (1990-1991), which corresponds to the years preceding the collapse of the Soviet Union. This might suggest a structural break in the norms and priorities originated in the international system and their internalisation within the UNGA. When looking at the right panel, we observe that, starting from the 50th session, the correlation between experience and outgoing deviations exhibits a negative trend (with the only exception being the 61st session). The magnitude of the coefficients, however, is smaller (in absolute terms); ranging from -0.006 in the 53rd session to -0.165 in the 71st, and only 8 out of these 22 (negative) coefficients are statistically significant at the 5% level.

Deals in Other UN Organs

Systemic anti-reciprocal voting patterns in the UNGA may happen because the other side of a deal materialises in a different place. Considering the plurality of interdependent organs that constitute the UN system, it is possible that reciprocated deviations between two agents (i.e., countries) occur *across* different bodies. In this case, the *Reciprocity Index* estimated would fail to capture such exchanges.⁴⁴

To assess the above proposition we focus on the UN Security Council (UNSC).⁴⁵ This is the most related UN organ to the UNGA, as its non-permanent members are directly elected by the General Assembly. We hypothesise that a country holding a temporary seat at the UNSC can leverage its position to obtain “favours” (i.e., incoming deviations) from other states in the UNGA, in exchange

⁴⁴Even if the literature suggests that this cross-institutional trades could occur (Malone 2000), one needs to consider that the cognitive and coordination costs that they would imply are non-negligible.

⁴⁵It is theoretically possible to include UNSC voting data in our methodology. However, given the different structure (limited membership, veto players) and nature of the decisions (legally binding for UN members) it is not clear how to determine the “exchange rate” between votes in the two institutions. Every weighting system would result in an arbitrary choice. Thus, we prefer to keep the focus on the General Assembly, given that our main aim is to assess the “rules of the game” in this specific UN organ.

for preserving their interests within the Council.⁴⁶ More specifically, we want to test the following hypothesis:

H4a: being a non-permanent member of the Security Council is positively related to the number of incoming deviations.

We operationalise UNSC membership through a binary variable that takes the value of 1 if the country holds a temporary seat during the session of the UNGA considered, and 0 otherwise. We exclude the five permanent members from the current analysis. Since we are now dealing with a dichotomous variable (i.e., UNSC membership) we apply a log-transformation to the number of incoming deviations to reduce the skewness of the indegree distribution and then compute the point-biserial correlation between the two variables.⁴⁷ Panel *a* in figure 3.12 shows that the correlation coefficient between UNSC membership and the natural log of the incoming deviations is generally positive (with the only exception being the 37th session), providing some suggestive evidence in favour of *H4a*. However, we should point out that the magnitude of these positive coefficients is quite weak, ranging from 0.032 in the 39th session to 0.216 in the 51st. In addition, only 9 out of the 41 (positive) coefficients are statistically significant at the 5% level.

An alternative, but related, hypothesis is that being a member of the UNSC raises your international profile and increases your chances of being re-elected again in the Council. Former UNSC members could then leverage this prestige to gain “favours” from the other countries, with their reputation raising as the number of times the country held a seat in the UNSC increases. We can formalise this hypothesis as follows:

H4b: the number of times that a country has been a non-permanent member of the Security Council is positively related to the number of incoming deviations.

We operationalise former membership in the UNSC with the cumulative number of sessions during which a country was holding a seat in the Security Council. Panel *b* in figure 3.12 shows the rank correlation (Kendall’s τ) with the number of incoming deviations. We observe again a positive relationship between the two variables (only in the 37th session the coefficient is negative), which is consistent with *H4b*. However, even in this case, the magnitude is not particularly large, ranging from 0.007 in the 68th session to 0.201 in the 51st and 9 out of these 41 (positive) coefficients are significant at the 5%


⁴⁶The ten non-permanent members of the Security Council are elected on a regional basis for a two-year term, and in every session of the UNGA five new members are elected.

⁴⁷When we apply the logarithmic transformation to the number of incoming deviations, we add 0.1 to the argument to preserve observations on countries that do not receive any deviation in a given session.


level.

Figure 3.12: Evolution of the Correlation between
SC Membership and Incoming Deviations

(a) UNSC Membership and Incoming Deviations



(b) Cumulative UNSC Membership and Incoming Deviations



Note: UNSC membership is operationalised through a binary variable that takes the value of 1 if the country holds a temporary seat during the session of the UNGA considered, and 0 otherwise. Former membership in the UNSC is operationalised through the cumulative number of sessions during which the country was holding a seat in the Security Council. The five permanent members of the Security Council are excluded from the analysis. In panel *a* we take the natural logarithm of the number of incoming deviations, adding 0.1 to the argument. In panel *b* we use the raw numbers. Data come from the list of countries elected members, in the UNSC website: <https://www.un.org/securitycouncil/content/countries-elected-members>.

Testing All Hypotheses

To properly test the relationship between all the factors previously explained (i.e., the size of a country, its economic resources, its experience in the UNGA, being a current member of the UNSC and the number of sessions that a country held a seat in the Security Council) and the inferred deviations, we run six different linear models. All regressions in table 3.6 include session and country fixed effects to take into account confounders that might be common to all countries in a given session or specific to a nation throughout the sampling period. We cluster the errors at the state level.

In models 1 and 2, the dependent variable is the natural logarithm of the number of outgoing deviations, whereas in models 3 to 8 is the log of the incoming deviations. In model 1, we observe that, when the different factors are simultaneously taken into account, population and UNGA experience are not statistically significant. The coefficient on income instead is statistically significant and show the expected sign. In model 2 we allow for experience having an impact through a non-linear learning process. We notice now that both experience and its square value are statistically significant and a Wald test on both variables reveals that they are also jointly significant at 5%. The inverted U shape described by the signs of the coefficients suggest that UNGA socialisation might require some time

Table 3.6: Deviations' Covariates

	OLS (1) <i>Outgoing</i>	OLS (2) <i>Outgoing</i>	OLS (3) <i>Incoming</i>	OLS (4) <i>Incoming</i>	OLS (5) <i>Incoming</i>	OLS (6) <i>Incoming</i>	OLS (7) <i>Incoming</i>	OLS (8) <i>Incoming</i>
<i>Experience</i>	0.0015 (0.0075)	0.0304** (0.0129)	0.0308*** (0.0081)	-0.0106 (0.0130)		-0.0073 (0.0130)		-0.0072 (0.0130)
<i>Experience</i> ²		-0.0004** (0.0001)		0.0005*** (0.0001)		0.0005*** (0.0001)		0.0005*** (0.0001)
$\ln(\text{Population})$	0.3010 (0.3617)	0.2425 (0.3561)	-1.3538*** (0.4434)	-1.2700*** (0.4440)		-1.2039*** (0.4394)		-1.1944*** (0.4387)
$\ln(\text{Income})$	-0.3606** (0.1570)	-0.3720** (0.1555)	0.4824*** (0.1689)	0.4987*** (0.1640)		0.4508*** (0.1629)		0.4525*** (0.1619)
<i>SC Membership</i>					0.1435** (0.0719)	0.093 (0.076)		
<i>SC Experience</i>							0.0366*** (0.0112)	0.0195* (0.0104)
<i>N Countries</i>	193	193	193	193	197	188	197	188
<i>N Observations</i>	6848	6848	6848	6848	7242	6654	7242	6654
<i>Country FE</i>	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>Session FE</i>	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: the dependent variable in models 1 and 2 is the natural logarithm of outgoing deviations. In models 3 to 8 is the log of the incoming deviations. In both cases we add 0.1 to the argument of the logarithm. Standard errors clustered at the country level are reported in parentheses. Models 5 to 8 do not include the permanent members of the UNSC. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

before a member can effectively leverage its experience in this international arena to resist external pressures.

Model 3 shows that the coefficient on the log of population is statistically significant at the 1% level, however its sign contradicts *H1*: an increase of 1% in a country's population is predicted to *reduce* the number of incoming deviations by 1.3%. The coefficient on income has the expected (positive) sign, and it is statistically significant at the 1% level. These results, together with the findings from models 1 and 2, provide strong support for *H2*. The coefficient on experience also shows a positive association with incoming deviations and it is statistically significant at 1%. However, when we allow for a non-linear effect, its sign switches and the variable is no longer statistically significant. The square value instead has a positive sign and it is highly significant. When we perform a Wald test on the two terms we discover that they are jointly significant at the 1% level. In this case, the U shape suggests that UNGA members might need some time before being able to exploit their seniority to pursue their own interests in the Assembly. These results, coupled with the insights from the first two models seem to suggest that the role of UNGA experience might be more nuanced than the simple monotonic relationship stated in *H3*.

Model 5 presents a linear specification where the only regressor used (in addition to the fixed effects) is a dummy for being temporary member of the UNSC. We see that the variable is statistically significant at the 5% level and its effect is large. A temporary seat in the Security Council is predicted to increase the number of incoming deviations by around 15.4%, on average and *ceteris paribus*; a result consistent with *H4a*. However, once we include the other covariates in model 6 (i.e., experience, its square

value, log-income and log-population), the size of the effect decreases and it is no longer statistically significant. The other coefficients remain significant, but their magnitude is slightly reduced. Thus, once we take into account UNGA seniority and a country's income and population, holding a temporary seat in the UNSC is no longer relevant to explain the number of incoming deviations.

In model 7 we use as only covariate (together with the fixed effects) the cumulative number of sessions during which a country held a seat in the UNSC. The estimated effect is positive and statistically significant at the 1%, in accordance with *H4b*. On average and *ceteris paribus* a unit increase in the variable is predicted to raise incoming deviations by (approximately) 3.7%. Interestingly, when we add the other covariates in model 8 the effect of *SC Experience* still remains positive and significant at the 10% level (even if the magnitude decreases). The significance of the other variables stay the same, even if their effect is smaller compared to model 4.

Cross-institutional vote trading might be one of the concurring explanations of the inferred behaviour, but the role played by the UNSC does not limit to current membership. Our findings seem to suggest that the reputation gained through past UNSC membership might be a more relevant factor. In addition, the results point towards a narrative that is consistent with both UNGA socialisation and disparities in economic resources.

To conclude, the systemic anti-reciprocal patterns observed provide strong evidence *against* vote trading as predominant form of strategic voting in the UNGA. This means that member states tend to not consider votes a valuable exchange currency. Thus, the equality principle, under which we should expect vote trading, does not seem to hold and the actual "rules of the game" appear to be quite different. Our results imply that disparities in UNGA socialisation, economic resources and past UNSC membership are the main correlates of the inferred deviations. These factors seem to underlie the unbalanced bargaining power that our negative *Reciprocity Index* suggests and so the nature of strategic voting in the UNGA.

3.8 Conclusion

In this paper we have provided an innovative framework to measure the extent and structure of strategic voting in the United Nations General Assembly (UNGA). This represents a necessary exercise in order to assess the suitability of the institutional rules adopted by the Assembly. The literature on

public choice suggests that strategic voting might arise in any deliberative organ due to the mismatch between the intensity of agents’ preferences and the features of the decision-making process. Without entering the normative debate, we develop a methodology to quantify and characterise the presence (or absence) of strategic voting in the UNGA. We consider alternative forms of strategic interactions that are plausible in our institutional setting (i.e., vote trading, vote buying and coercion) and the “currency of exchange” that they would imply. For instance, vote buying is a trade of one’s vote for an economic benefit, whereas vote trading is a plain exchange of votes. We assume that these hidden behaviours leave patterns in the roll call data, that can be captured through the analysis of the deviations from the expected votes. By leveraging a newly assembled dataset, Structural Topic Models and other machine learning techniques, we predict the voting behaviour of UNGA members. We then combine our predictions with information on the beneficiaries of eventual deviations from the expected votes. This allows us to build a network that represents the structure of directed deviations in the institution under analysis. The graph can be further exploited to compute an aggregate statistic: the *Reciprocity Index*. This measure quantifies the degree to which deviations from the expected votes are reciprocated by the agents and can be used to assess the extent of different forms of strategic interactions (e.g., vote trading, vote buying/coercion).

Through our methodology, we find that deviations from the expected voting behaviour are systematically *not* reciprocated. The results are robust to the use of a different classifier for predicting votes, an alternative probability cutoff for defining deviations and the choice of longer time frames to build the directed-deviations networks (DDNs). The anti-reciprocal pattern that the *Reciprocity Index* captures seems to suggest a structural imbalance in the bargaining power of UNGA members, a reasonable conclusion given the high heterogeneity of this international arena. We test a series of hypotheses, finding only weak evidence for vote trading *across* UN organs. The results instead seem to point towards a narrative of disparities in terms of UNGA socialisation and economic resources. This last consideration in particular is consistent with a large body of literature that documents the extent of vote buying practices in the UNGA. The vast majority of these studies frames vote buying in terms of development aid flows. Overseas Development Assistance (ODA) can be included in the construction of the DDNs, as well as other forms of financial flows (e.g., FDI). In the present paper, we do not extend the methodology to investigate further this relationship, but it appears to be a promising avenue for future research.

To conclude, our findings challenge the leading institutional features of the UNGA: the equality principle and the resulting unweighted voting system. These criteria do not seem to take into account

the structural heterogeneity of the political actors involved in the decision-making process, potentially undermining the effectiveness of the deliberative organ.

However, the current study is subject to a number of limitations. First, in the analysis we do not discriminate between *no* votes, abstentions and absences, that could be the result of a strategic choice too (Dreher and N. M. Jensen 2013). We follow Voeten (2000) by assuming that, due to the non-binding nature of UNGA resolutions, the relevant feature of a vote is the willingness of a country to go on the record and support a given proposal. In addition, incorporating the different outcomes into the proposed methodology is not straightforward and it would imply further assumptions. A limitation of the method is that it only considers intense preferences for passing a given proposal and not for blocking it. However, this limit is partly due to the nature of the UNGA: most of the resolutions in the General Assembly are passed and voting information on those that are not is unrecorded. Besides, coding negative signals is not as clear-cut as for the positive ones, for which we can rely on a resolution's sponsors and it would require additional assumptions. Moreover, in our framework we assume that countries signal their intense preferences through draft sponsorship, but they might express their key priorities in different ways. For instance, when they address the Assembly during the General Debate held at the beginning of each session. A thorough investigation of these speeches is beyond the scope of our study, but it could provide valuable insights into countries' policy preferences and it would be an interesting extension of the present research.

The General Assembly is the only organ within the UN system in which each country has *theoretically* equal representation. However, our analysis suggests a systemic disparity in the bargaining power between these political actors, where differences seem to be based on inequalities in economic resources and experience in the international sphere. The UNGA would then just be another political arena where Thucydides' words come to life and “[...] *the strong do what they can and the weak suffer what they must*”.

In a world where an increasing number of global challenges can be no longer tackled at the national level and require collective action, it is crucial to ensure that every country has truly equal opportunities within this international *agorá*, irrespective of its resources or history.

Chapter 4

Information Processing and Policy Overreactions: Evidence from Mexico

Joint work with Omar A. Guerrero

4.1 Introduction

“[...] in an information-rich world, the wealth of information means a dearth of something else: a scarcity of whatever it is that information consumes. What information consumes is rather obvious: it consumes the attention of its recipients. Hence a wealth of information creates a poverty of attention and a need to allocate that attention efficiently among the overabundance of information sources that might consume it.”

Simon, Herbert A. (1971). “Designing Organizations for an Information-Rich World”.

In these few lines written almost 50 years ago, Herbert A. Simon highlight a key aspect of human decision-making that could not be more relevant in today’s world, where we are constantly overloaded by endless streams of information striving for our attention. Moreover, these cognitive constraints that we face when taking decisions based upon a surplus of inputs are not just a peculiar feature of single

individuals. They can also materialise within an organisation, but the way in which the structure of this one interacts with the limited attention of its members remains an open empirical question.

In this paper we address the issue of how incoming flows of information are processed by one of the most pivotal decision-making systems in many societies: the government. Starting from the model of choice for public policy of B. D. Jones and Baumgartner (2005), we empirically test two alternative processing mechanisms and their implications for policymakers' decisions, considering the role that limited attention might play. We narrow the focus of the research on the policy domain of budget allocation. In particular, we look at short term adjustments to the expenditure on individual budget programmes. This policy outcome has been relatively disregarded by the literature mainly due to the lack of available data on such a granular level of analysis. We also consider a distinct source for the informational signals that enter the decision-making system: the media. More specifically, we examine opinion pieces from newspapers, given the political role that they play in the institutional context under study (Mexico).

In order to create a mapping between the informational inputs and the policy outputs, we rely on Latent Dirichlet Allocation -LDA- (Blei, Ng, and Jordan 2003) to devise a *News Index*. This index exploits our newspapers corpus of more than 35000 articles and relevant textual data that we collected on the single budget programmes. It captures the main topics discussed in the news together with their salience and looks at how close the budget programmes are to these public issues. We use the index to assess two different descriptive models of information processing: parallel and serial. We find that changes in the *News Index* are associated with overreactions of the government (measured as the gap between approved and paid expenditure on a programme) which are consistent with the serial processing hypothesis.

We then take a step forward, and try to analyse the causal relationship between information shocks and government's disproportionate response. We leverage the different media coverage of two earthquakes that occurred in Mexico in 2017 to frame a natural experiment with a Difference-in-Differences (DID) design. Overall, our results pass a series of robustness checks and are consistent with a narrative of disproportionate information processing in government's decisions, which is in line with the theoretical framework of B. D. Jones and Baumgartner (2005).

We contribute to the literature in several ways. To the best of our knowledge, we provide one of the first empirical studies on alternative mechanisms of information processing in government's decisions. We employ text analysis techniques and rely on a newly assembled dataset to devise a novel approach

that links informational signals to policy outputs. This allows the study of the relationship between media coverage and budget allocation for individual expenditure programmes, which is a level of disaggregation that has not been analysed so far. Finally, by conceiving a natural experiment, we explore further the implications of attention-based policy responses within a causal framework.

The rest of the paper is articulated as follows. The next section covers the relevant literature on the leading model of policy process and the role of media as inputs for policy outcomes. It lays the basis for testing alternative information processing mechanisms within the governmental decision-making system. In section 4.3, we present the expenditure data and the newspapers corpus, providing details about the institutional context under analysis, both in terms of the budget cycle and the characteristics of opinion journalism in Mexico. Next, in section 4.4, we first formally define two alternative descriptive models of information processing and then move to explaining the construction of the *News Index* and presenting the results of our estimations. In section 4.5, we introduce the setting of our natural experiment, the findings obtained and we test their robustness by assessing competing hypotheses. Finally, section 4.6 briefly concludes.

4.2 Literature Review

In the last two decades, the dynamic model of choice proposed by B. D. Jones and Baumgartner (2005) has emerged as the dominant decision-making model of policy change among public policy and administration scholars. It unifies the standard framework of incrementalism (Wildavsky 1964; Lindblom 1959), the leading view of policy dynamics until the 1970s, with the one of punctuated equilibrium (Baumgartner and B. D. Jones 1993). Its theoretical underpinnings have been successfully applied to study the dynamics of both agenda-setting (e.g., Alexandrova, Carammia, and Timmermans 2012) and public budgeting (e.g., Breunig, Koski, and Mortensen 2010). In particular, the framework offers a robust explanation for the well-established empirical generalisation that yearly changes in government budget tend to follow a power-law distribution. Such result has been observed consistently across different political systems and policy domains (B. D. Jones, Baumgartner, et al. 2009). More recently, J. L. Jensen, Mortensen, and Serritzlew (2016) have developed a sound mathematical formalisation of the dynamic component of the model and proposed a few revisions that improve its empirical fit. In particular, they emphasise how the frictions of the decision-making system and its efficiency in reacting to informational signals might vary across policy issues.

Still, how information is processed by the political system remains a key feature of the framework. In their theory of policy dynamics, B. D. Jones and Baumgartner (2012) assert that policy changes cannot be uniquely produced by shifts in the preferences of legislators caused by electoral cycles. They also result from variations in the information environment in which policymakers act. They must collect, assemble, interpret, and prioritise signals from the environment (Workman, B. D. Jones, and Jochim 2009). However, they face an oversupply of information that conflicts with the bounded rationality of individuals (H. A. Simon 1996), which in turn interacts with the operating rules of their organisational structure. The limited attention of both policymakers and political institutions force them to prioritise and rank issues by their relative importance.

Generally, organisations can process the incoming flow of information in two different ways: parallel and serial processing (Workman, B. D. Jones, and Jochim 2009). During parallel processing, decision systems are effectively able to tackle multiple issues at the same time. This mechanism is thus characterised by incremental adjustment and proportional response in policy choices. Serial processing instead usually occurs when the level of attention is very high and focused on a subset of problems. This mechanism stresses attention limits at both individual and institutional level, causing overresponse to issues that appear particularly salient and underreaction to those that seem less relevant. Such reactions can be further exacerbated by the decision costs and frictions inherent in the institution, and only large changes in the information environment will elicit a response from the system. Hence, serial processing is marked by disproportional updating, which is a central component in the dynamic choice model of B. D. Jones and Baumgartner (2005). Underreactions to low informational signals cumulate over time and when they overcome cognitive and institutional constraints attention is dramatically shifted, and the system overreacts. Such error accumulation and correction dynamics should generate those leptokurtic distributions of policy outcomes that have been consistently observed in empirical work.

In our study, we approach the issue of government information processing and public budgeting from a narrower perspective. We will focus on a single budget cycle, but we will conduct our analysis at the level of individual budget programmes. The importance of unpacking the budget to uncover the dynamics of policymaking at a more disaggregated level has already been highlighted by the comparative analysis of Breunig, Koski, and Mortensen (2010). It is also in line with the theoretical developments proposed by J. L. Jensen, Mortensen, and Serritzlew (2016). By examining a single budget cycle, we will not consider the dynamics of error accumulation and correction, but only the mechanism that governs policymakers' reactions to a set of informational signals of different intensity.

Besides, since we only look at a single fiscal year, our approach allows us to assume that the institutional costs of decision-making are fixed, at least within the same policy area, so that we can focus on how information is processed by the legislators.

The temporal framework is further justified by the outcome of interest: discrepancies between the approved budget and actual expenditure on individual programmes. This policy outcome has not received much attention in the literature, mainly due to the lack of available data on such a fine-grained level of analysis. Nevertheless, these corrections are of great interest. Their analysis can allow deeper insights into the dynamics of short term information processing and the way in which policymakers might adjust initial decisions given a changing information environment. In studying how the government processes information flows, we will focus on a specific type of signal: opinion pieces from newspapers.

The link between media and policy outputs has been poorly analysed in the literature, especially in political communications studies (Wolfe, B. D. Jones, and Baumgartner 2013). Policy process scholars instead, underline the pivotal role that media might play in the allocation of attention within the information processing framework. Media coverage can be viewed as a weighting mechanism for informational signals, highlighting specific attributes that define a policy problem and its possible solutions. Thus, by filtering relevant information, media can shift the attention of legislators, but the effects on the policy process and its outputs are not straightforward.

Within this strand of research, most of the literature has been focusing on the role of mass media as agenda-setter. For instance, in the cross-country study of Vliegthart et al. (2016), the authors analyse the strategic interplay between political elites and media, emphasising the role of the latter one as a crucial source of information for politicians. However, they find that the effect on parliamentary agenda is mediated by the idiosyncratic institutional features of the political system considered. B. D. Jones, Thomas III, and Wolfe (2014) move to specific policy outputs and argue that media coverage can trigger positive feedback that is strongly related to the emergence and development of policy bubbles. These are defined by the authors as the systematic overinvestment in specific policy instruments beyond their instrumental value, that become self-sustained over time. Wolfe (2012) instead, shows that media attention might also engender negative feedback that constraints policy change. This can occur not just through favouring the status quo, but also by allowing new issue attributes and related interest groups to enter the political arena, increasing conflicts and slowing down the policy process. In her analysis of lawmaking in the US Congress, the author finds that high media attention increases the

length of time it takes for a bill to become law.

However, all these studies call for a better understanding of how governments process the informational signals that they receive from the media, especially in the short term, when the effect of mass media on the political system is expected to be stronger (Vliegenthart et al. 2016).

Hence, in the first part of the paper, we try to close this gap. We provide a formal assessment of two different processing mechanisms that might describe the link between government’s actions and media coverage in the short run: parallel and serial processing. In the second part of the paper instead, we take a step forward, and explore the causal relationship between increased media attention triggered by specific events and government’s disproportionate policy reactions.

4.3 Institutional Setting and Data Description

4.3.1 Institutional Framework

Mexico’s political system consists of a federation of 32 Federal Entities.¹ The head of state and government is the President, who is elected every six years, with no opportunity to run for election (re-election of the President is banned by the Constitution). In contrast with many other federations, Mexico’s states have little tax-collection capacity. Instead, it is the federal government the one concentrating most of the public revenue, which is redistributed every fiscal year across the states according to the *Fiscal Coordination Law*.

The fiscal year follows the calendar year, so it runs from the 1st of January to the 31st of December. The agency in charge of proposing and spending the federal budget is the Ministry of Finance and Public Credit or SHCP for its name in Spanish (Secretaría de Hacienda y Crédito Público). The SHCP has to present the proposed federal budget to the Congress no later than the 8th of September for discussion, revisions, and approval. Two of the submitted documents are the *Draft of the Federal Revenue Law* and the *Draft Budget of Expenditures of the Federation*. The former provides the legal framework to enable the government to generate income, and the latter describes the allocation and scope of public spending across the Executive, Legislative, and Judicial powers (and across autonomous publicly funded bodies) and has to be approved by the Chamber of Deputies before the 15th of November.² Our study focuses

¹The 32 Federal Entities are the 31 states plus the Federal District of Mexico City.

²In addition, throughout the fiscal year, the SHCP produces reports on the performance and progress of the economy. These include public revenues and expenditures, and are submitted to the Congress every quarter.

on this last document, as Mexico’s transparency laws have enabled the creation of fine-grained datasets on public spending.

The Mexican federal budget can be classified according to different criteria; the relevant one for this study is the so-called classification by expenditure object. At the most aggregate level, public spending is organised into 48 broad categories called *branches* (literal translation from its term in Spanish: *ramos*).³ Some examples of branches are “Public Education”, “Agriculture and Rural Development” and “Environment and Natural Resources”. At the most disaggregated level, there are 751 *budget programmes* (*programas presupuestarios*, in Spanish).⁴ Let us abbreviate budget programmes as BPs from here onwards. A BP represents a homogeneous and integrated set of processes, activities and services with the same purpose. Their objectives are so specific that those agencies receiving the funding from a BP (to implement public policies) are required to report performance indicators according to the guidelines provided by Mexico’s Performance Evaluation System. Thus, BPs enable the government to achieve specific development goals by allocating resources to diverse policy interventions. Two examples of BPs are the “National Programme for Financing Micro-entrepreneurs and Rural Women” and the “Programme for the Prevention and Control of Overweight, Obesity and Diabetes”.

4.3.2 The Importance of Opinion Journalism

Scholars of Mexican media argue that the process of democratisation and awakening of the civil society started in the 1980s in the country, together with inner changes in the organisational culture, contributed to a progressive modernisation of news outlets, especially newspapers (Hughes 2003, 2006). Both the development of a competitive media market, with its pressures to gain readership, and changes in journalistic ethical norms, bringing considerations about the role that press should play in the society, were conducive to this transformation (Lawson 2002). From an authoritarian media institution completely subordinated to the regime, the development of new, civic-oriented, models of journalism helped the creation of a public sphere that granted citizens access to a more diverse information environment.

Since the 1988 general election, the press started to play a major role in the definition of the democratic demand, providing an open arena for those policy issues most heatedly debated in the country. By drawing attention to specific societal problems, it tried to actively influence the decision-making

³The number of total branches refers to the fiscal year of 2017, the one considered in the analysis.

⁴The number of total budget programmes refers to the fiscal year of 2017, the one considered in the analysis.

system (Santillán 2013). In particular, opinion pieces and editorials became real platforms of political negotiations (Santillán 2016). A leading medium for the representation of public issues and the promotion of possible solutions, spurring the mobilisation of the political actors involved. As a consequence, nowadays in Mexico, there is a significant number of columnists that exert great influence on political parties and authorities (Santillán 2013).

Thus, in our analysis, we focus on opinion pieces, which are then supposed to embody the public sphere where the most salient political and social issues in the country are discussed. These are going to be the sources of the informational signals to which the policymakers in our decision-making system are exposed.

4.3.3 Data on Public Expenditure

The Mexican Treasury reports two types of expenditure: the one approved by the Congress (before the fiscal year begins) and the paid one (at the end of the fiscal year). The difference between the approved and paid budget reflects shifts in policy priorities by the government through a total or partial cancellation of payment obligations. Thus, we focus on the absolute difference between approved and paid expenditure.⁵

In our study we consider 614 BPs, spanning 35 branches during the fiscal year of 2017.⁶ Table 4.1 reports the average and standard deviation of the absolute difference between approved and paid expenditure at the level of each branch. Just by looking at a subset of these macro-categories, we can already detect systematic disparities across major policy issues.

These are coupled with quite large spending variability for the programmes within specific categories, as for “Environment and Natural Resources”, “Government” and “Public Education”. Tables C.1 and C.2 in Appendix C.1 show the values for the full list of branches (both in levels of spending and as proportions of the total budget approved/paid), together with the distribution of the programmes across the expenditure categories.

In order to link expenditure data to the news, it is necessary to have further information about the BPs, not just their monetary amounts. These data are provided by the SHCP through its fiscal

⁵We look at *absolute* differences as our proxy for informational inputs will not convey the polarity of the signals. However, we will also provide results for non-negative changes in expenditure.

⁶The set of BPs considered is mainly determined by the availability of relevant textual data on the programmes (see later).

Table 4.1: Absolute Difference between Expenditure Approved and Paid for Selected Branches

Branch	Absolute Difference
<i>Economy</i>	0.5610 (0.6227)
<i>Work and Social Security</i>	0.6774 (0.9527)
<i>Energy Regulatory Commission</i>	1.1383 (0.8468)
<i>National Hydrocarbons Commission</i>	1.1856 (0.8945)
<i>National Council for Science and Technology</i>	2.2088 (1.9383)
<i>Environment and Natural Resources</i>	3.7828 (8.2770)
<i>Government</i>	4.5178 (10.9795)
<i>National Defense</i>	6.6448 (9.1195)
<i>Marine</i>	9.6050 (11.0651)
<i>Public Education</i>	12.9781 (41.1075)

Note: the table shows the average absolute difference between expenditure approved and paid. Standard deviations are reported in parentheses. Values are reported in millions of Mexican pesos.

transparency portal.⁷ We collect textual data on the government department in charge of managing the budget allocated to the BP, the socio-economic objectives of the BP, the actions carried out to achieve these goals and the indicators constructed to evaluate the department’s performance in achieving such goals, to mention a few features. Tables C.3 and C.4 in Appendix C.2 provide the list with all the BP features from which we obtain textual data.

4.3.4 Data on Opinion Columns

We construct a dataset of opinion columns containing 35822 articles from the major 9 newspapers in Mexico.⁸ These pieces are obtained from the *Opinion* section of each newspaper (which includes editorials and op-eds as well). Table 4.2 shows the distribution of the opinion columns across the

⁷Source: *Transparencia Presupuestaria* (<https://www.transparenciapresupuestaria.gob.mx/>).

⁸All these newspapers have national circulation.

newspapers, together with their average circulation (per issue).⁹

In our data, *El Universal* has the highest share of articles, with 7958 items. Founded in 1916, it is one of the oldest newspapers in Mexico and one of the most popular, as shown by its average circulation per issue (130307), which is the second highest among the media outlets that we consider in the analysis.

Table 4.2: Distribution of Articles across Newspapers and their Average Circulation

	Number of Articles	Average Circulation
<i>El Economista</i>	1482	40600
<i>El Excelsior</i>	5548	121039
<i>El Financiero</i>	2727	91230
<i>El Herald de México</i>	4288	40232
<i>El Universal</i>	7958	130307
<i>La Jornada</i>	5103	69752
<i>La Razón de México</i>	3085	92758
<i>Milenio</i>	3134	84646
<i>Reforma</i>	2497	132262

Note: information on average circulation (per issue) comes from the website of the National Register of Printed Media (<https://pnmi.segob.gob.mx/reporte>, consulted in October 2020).

The articles were published between the 1st of January 2017 and the 31st of December 2017, with an average daily coverage of 98 opinion pieces, ranging from a minimum of 30 to a maximum of 155 per day. Table 4.3 presents more detailed descriptive statistics on daily coverage across the different newspapers.

4.4 Government’s Information Processing and Budget Allocation

4.4.1 Parallel and Serial Processing

To formulate our hypotheses, we rely on two straightforward formalisations of government’s information processing based on the work of B. D. Jones and Baumgartner (2005). However, we would like to emphasise the descriptive nature of these formalisations. They do not result from solving analytically

⁹We obtain the newspapers’ circulation data from the National Register of Printed Media (<https://pnmi.segob.gob.mx/reporte>, consulted in October 2020).

Table 4.3: Coverage across the Newspapers

	Mean	Std. Dev.	Min.	Max.
<i>El Economista</i>	4	5	0	24
<i>El Excelsior</i>	15	2	7	21
<i>El Financiero</i>	7	5	0	17
<i>El Heraldo de México</i>	12	10	0	26
<i>El Universal</i>	22	10	0	37
<i>La Jornada</i>	14	3	0	21
<i>La Razón de México</i>	8	4	0	19
<i>Milenio</i>	9	2	0	10
<i>Reforma</i>	7	1	0	10
<i>All Newspapers</i>	98	23	30	155

an equilibrium model, where agents maximise an objective function. Nevertheless, they provide a useful descriptive framework to interpret the patterns observed in the data.

In the first scenario, government's decisions are guided by parallel processing. Hence, each signal is processed through an independent channel, so governments are able to cope with multiple sources of information without experiencing bottlenecks. In theory, this allows policymakers to adapt their policy responses in information-efficient ways. Here, the information processing costs C are separable from the signal S . Then, the response in terms of policy P is proportional to such signal. B. D. Jones and Baumgartner (2005) formalise this model through:

$$\Delta P = \beta S - C \quad (4.1)$$

Where ΔP denotes the change in policy P (i.e., the response) and β represents the proportional adjustment of the response to the incoming stream of information. The information processing costs C can be interpreted as institutional constraints imposed by organisational practices, cultural legacies, lack of capacity, transaction costs, or other factors that may lead to technical inefficiencies (i.e., inefficiencies arising from the policymaking process). For a single fiscal cycle, where the budget has already been approved, it is reasonable to assume that such constraints, together with the influence from the political system (e.g., congressional negotiations), are fixed. Therefore, we assume that C captures mainly institutional frictions that are persistent and vary across broad policy areas. This is consistent with the theoretical framework of J. L. Jensen, Mortensen, and Serritzlew (2016) and

with the systematic issue-area differences in budgeting emphasised by Breunig, Koski, and Mortensen (2010).

The second model, serial processing, posits that governments are victims of the cognitive constraints of their bureaucracies. Under bounded rationality, the policymakers' actions are attention-driven, and this attention can only operate under a reduced number of information processing channels. Therefore, an over-saturation of information generates bottlenecks and considerable discrepancies in policy responses, for example, overreactions. In contrast to the parallel approach, the serial model assumes non-separable costs and non-linear policy responses. Formally B. D. Jones and Baumgartner (2005) specify this model as:

$$\Delta P = \beta S^\gamma C \quad (4.2)$$

Where the exponent γ introduces non-linearities due to the over-saturation of information processing channels and, in particular, overreactions when $\gamma > 1$. The case of $\gamma > 1$ is consistent with a salient feature of behavioural rationality (B. D. Jones 2017): the tendency to stick to a previous set of decision rules, while an overreaction takes place only when the intensity of the signal is strong enough.

Taking the natural logarithm on both sides, we obtain:

$$\ln(\Delta P) = \ln(\beta) + \gamma \ln(S) + \ln(C) \quad (4.3)$$

When we narrow down the focus on individual policy choices (e.g., adjustments to budget programmes), the attention-driven process results in a selective bias and a subsequent disproportionate policy reaction to those informational signals that are more extreme. Notice that the functional form implied by model 4.2 leads (in principle) to those leptokurtic distributions of policy outcomes that have been consistently observed in the empirical literature (B. D. Jones, Baumgartner, et al. 2009).

4.4.2 Building a *News Index*: Methodology

To empirically test the alternative models of government's information processing proposed in the previous section, we construct a *News Index* that captures the signals sent by opinion journalists to

the government about specific societal issues. The idea is to take into account not only the information encoded in the opinion columns, but also how relevant this information is to each budget programme. Thus, the index does more than just quantifying news; it effectively provides a link to the policy instruments used by the government (because they are funded by highly specific BPs). Our methodology is divided into two steps. In the first one, we consider as informational signals the topics discussed in the opinion pieces together with their salience, inferring them from the newspapers corpus. In the second step, we link these signals to features from the BPs, obtained through their textual data. The outcome of this linkage is an index quantifying how close the budget programmes are to the different issues identified.

To infer the latent topics in the opinion pieces, we employ Latent Dirichlet Allocation -LDA- (Blei, Ng, and Jordan 2003). LDA is a three-level hierarchical Bayesian model for topic modelling. The main idea is that documents (opinion pieces in our application) can be described as mixtures over latent topics, and in turn every topic is characterised by a distribution of words.

Once we apply LDA to our preprocessed corpus,¹⁰ we can describe every article d with a vector of proportions $\theta_d = \{\theta_{d1}, \theta_{d2}, \dots, \theta_{dK}\}$, where each element θ_{dk} represents the fraction of words in article d belonging to topic k . Next, for each article d , we identify the main topic discussed ($\hat{\theta}_d^*$), i.e., the topic inferred with the highest proportion ($\hat{\theta}_d^* = \max\{\hat{\theta}_{dk}\}_{k=1}^K$). We compute, for each topic k , the fraction of articles in the corpus that talk mainly about that topic (p_k); that is, those articles for which $\hat{\theta}_d^* = \hat{\theta}_{dk}$. These proportions $\{p_k\}_{k=1}^K$ are going to be the first component of the *News Index* and describe the salience of the topics in the corpus.

In the second step, we apply the topic model trained on the newspapers corpus on the textual data associated with the BPs' features.¹¹ This means that we can now describe each BP j as a distribution over the K topics inferred from the articles. Intuitively, these proportions $\{\hat{\theta}_{jk}\}_{k=1}^K$ capture how close the BP j is to the different topics that emerge from the opinion pieces, creating the link between policy outputs and informational signals.

Formally, the *News Index* for BP j is computed as the sum of the signals from the different topics:

¹⁰The corpus is preprocessed in the following way: we remove stopwords (i.e., common terms with extremely high frequency), make all tokens lowercase, discard those with less than 2 characters, create bigrams and trigrams (discarding those that occur less than 5 times), consider only nouns, proper nouns, verbs, adverbs, adjectives, and other relevant part-of-speech as foreign words. Finally, we lemmatise the resulting tokens. Once we represent the corpus in a Bag of Words (BoW) format, we also discard those tokens that occur in more than 50% of the opinion pieces.

¹¹We apply the same preprocessing steps described in footnote 10, but on the textual data on the BPs.

$$NewsIndex_j = p_1\hat{\theta}_{j1} + p_2\hat{\theta}_{j2} + \dots + p_K\hat{\theta}_{jK} \quad (4.4)$$

Where the signal for topic k ($p_k\hat{\theta}_{jk}$) is the product between the fraction of articles that mainly cover topic k (p_k) and its proportional presence in BP j ($\hat{\theta}_{jk}$).

The LDA model that we train on the newspapers corpus has $K = 9$. This number of topics is optimised according to the combined rank across three different performance metrics: the statistic proposed by Mimno, Wallach, et al. (2011), normalised pointwise mutual information -NPMI- (Bouma 2009) and the combined coherence measure C_V of Röder, Both, and Hinneburg (2015).¹² Figure 4.1 reports word clouds for two of the inferred topics. On the left, we can observe that the words mainly refer to the issues related to crime and corruption, as we can deduce from terms as *violencia* (violence), *impunidad* (impunity) and *investigación* (investigation). On the right instead, the focus is on foreign affairs and US-Mexico relations, as suggested by words like *amenazar* (to threaten), *negociación* (negotiation), *Donald Trump* and *inmigrante* (immigrant).

Figure 4.1: Word Clouds for Selected Topics

(a) Crime and Corruption

(b) US and Foreign Affairs



Note: in the word clouds the size of the terms is proportional to their weight in the topic distribution.

4.4.3 *News Index*: Analysis and Results

We now turn to a regression framework and employ our *News Index* to assess the two formalisations of government's information processing presented in subsection 4.4.1.

¹²More specifically, we select the number of topics in the model (which is a hyperparameter that has to be chosen *a priori*), by assessing topic coherence across 5 possible values of K : {3, 5, 7, 9, 11}. The best performing model is selected according to the rank on the three different performance measures.

Parallel processing and proportional updating are captured by the following specification, which reflects equation 4.1:

$$|Exp_j^{paid} - Exp_j^{approved}| = \alpha + \beta NewsIndex_j + \mathbf{x}_j \boldsymbol{\delta} + \epsilon_j \quad (4.5)$$

Where the dependent variable is the absolute difference between the expenditure paid (Exp_j^{paid}) and the expenditure approved ($Exp_j^{approved}$) for BP j (i.e., the change in policy: ΔP). The variable $NewsIndex_j$ described in the previous subsection quantifies the informational input S for BP j . Finally, the vector \mathbf{x}_j includes dummy variables for the branch categories, which are meant to capture systematic differences in formal institutional costs related to specific policy areas (i.e., the term C in equation 4.1).

To assess serial processing and disproportionate updating instead, we estimate the regression below, which is consistent with equation 4.3:

$$\ln(|Exp_j^{paid} - Exp_j^{approved}|) = \alpha + \gamma \ln(NewsIndex_j) + \mathbf{x}_j \boldsymbol{\delta} + \epsilon_j \quad (4.6)$$

Testing each of these two hypotheses boils down to parameters β and γ . An ambiguous result would be non-zero statistically significant values for both models. In contrast, decisive findings would suggest significance for one of the parameters, but not for the other.

In the rest of this subsection, we show that the latter is the case. In particular, we find strong evidence in favour of the serial information processing hypothesis.

Main Results

From left to right, columns 1 and 2 in table 4.4 report the estimates for specifications 4.5 and 4.6 respectively. First, the results suggest a significant coefficient ($p < 0.05$) for the serial model ($\ln(NewsIndex)$) and a non-significant value for the parallel one ($NewsIndex$). Thus, these results favour the hypothesis of a limited information processing capacity, which implies overreactions and disproportionate policy updating. Second, the magnitude of the effect of media on policy adjustments is large: on average and *ceteris paribus*, an increase of 1% in the index is associated with an increase of approximately 11% in the absolute difference between approved and paid expenditure.

In columns 3 and 4, we use the expenditure as proportion of the total budget (so we compute absolute changes in the approved and paid proportions). The results are robust to this alternative operationalisation of the dependent variable: the coefficient of *NewsIndex* (column 3) is still not significant, whereas the one of $\ln(\text{NewsIndex})$ (column 4) is ($p < 0.01$). However, the magnitude of the effect decreases: on average, raising the index by 1% is associated to an increase of the dependent variable by around 6%.

Table 4.4: Parallel and Serial Information Processing

	OLS (1) $\Delta Level$	OLS (2) $\ln(\Delta Level)$	OLS (3) $\Delta Prop$	OLS (4) $\ln(\Delta Prop)$	OLS (5) $\Delta Level$	OLS (6) $\ln(\Delta Level)$	OLS (7) $\Delta Prop$	OLS (8) $\ln(\Delta Prop)$
<i>Intercept</i>	-7.901 (18.159)	33.799*** (9.693)	-0.098 (0.284)	-0.044 (4.294)	-8.216 (17.470)	28.074*** (9.062)	-0.107 (0.264)	-0.776 (3.744)
<i>NewsIndex</i>	58.182 (124.838)		0.753 (1.952)		59.253 (120.066)		0.783 (1.815)	
$\ln(\text{NewsIndex})$		10.992** (5.015)		6.173*** (2.221)		9.139* (4.723)		4.696** (1.895)
<i>Exp^{approved}</i>					0.035*** (0.009)		0.060*** (0.010)	
$\ln(\text{Exp}^{\text{approved}})$						0.145*** (0.043)		0.210*** (0.041)
<i>N</i>	614	614	614	614	614	614	614	614
<i>Branch FE</i>	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>Adj R²</i>	0.022	0.117	0.014	0.126	0.038	0.152	0.072	0.246

Note: the dependent variable in models 1, 3, 5 and 7 is the absolute difference between the expenditure paid and the expenditure approved. In models 2, 4, 6 and 8 is the natural logarithm of the aforementioned variable. In models 1, 2, 5 and 6 we use levels of expenditure to compute the policy response. In models 3, 4, 7 and 8 we use expenditure proportions. Expenditure in levels is reported in millions of Mexican pesos. Expenditure proportions are reported in per thousand of the budget. Heteroscedasticity-consistent standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

In columns 5 to 8, we repeat the estimations but we include as additional control the approved expenditure (or its natural logarithm) for the given programme. This variable is meant to capture incentives to adjustment in the budget that might come from the previous cycle (i.e., any high level of discrepancy between approved and paid expenditure should be incorporated by policymakers in the formulation of the next budget). Our main findings are robust to the addition of the control. When we look at the models representing serial information processing (models 6 and 8) the coefficient of interest is still positive and significant, both when considering the level of spending ($p < 0.10$) and expenditure as proportion of the total budget ($p < 0.05$). However, its magnitude and statistical significance are somewhat reduced. When considering proportional updating instead (models 5 and 7), the coefficient of interest is never statistically different from 0.

Robustness Checks

In table 4.5 we assess the robustness of our previous findings to variations in the underlying newspapers corpus. We reproduce the estimations of table 4.4, but excluding the opinion columns from *El*

Universal and *Reforma*, which were the two newspapers with the highest average circulation in the sample.¹³ The purpose of this exercise is to test if our main results are mainly driven by these highly popular newspapers. Again, we find no evidence of parallel processing (the coefficient of *NewsIndex* is not significant across the different models), whereas variable $\ln(\text{NewsIndex})$ is always statistically significant, with a magnitude of the effect even larger than before, especially in columns 4 and 8. When considering the serial processing models, the inclusion of the additional control (approved expenditure) increases the magnitude of the effect of interest and its statistical significance (columns 6 and 8).

Table 4.5: Parallel and Serial Information Processing
Subsample of Opinion Pieces

	OLS (1) $\Delta Level$	OLS (2) $\ln(\Delta Level)$	OLS (3) $\Delta Prop$	OLS (4) $\ln(\Delta Prop)$	OLS (5) $\Delta Level$	OLS (6) $\ln(\Delta Level)$	OLS (7) $\Delta Prop$	OLS (8) $\ln(\Delta Prop)$
<i>Intercept</i>	-22.044 (17.224)	30.403*** (8.288)	-0.402 (0.279)	1.224 (4.019)	-20.930 (16.783)	28.613*** (7.539)	-0.370 (0.260)	3.787 (3.578)
<i>NewsIndex</i>	116.064 (88.413)		2.122 (1.431)		109.530 (86.147)		1.937 (1.336)	
$\ln(\text{NewsIndex})$		10.894** (5.055)		8.055*** (2.446)		11.159** (4.693)		8.286*** (2.104)
<i>Exp^{approved}</i>					0.035*** (0.009)		0.060*** (0.010)	
$\ln(\text{Exp}^{\text{approved}})$						0.151*** (0.043)		0.217*** (0.040)
<i>N</i>	614	614	614	614	614	614	614	614
<i>Branch FE</i>	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>Adj R²</i>	0.022	0.118	0.016	0.141	0.039	0.156	0.073	0.270

Note: the dependent variable in models 1, 3, 5 and 7 is the absolute difference between the expenditure paid and the expenditure approved. In models 2, 4, 6 and 8 is the natural logarithm of the aforementioned variable. In models 1, 2, 5 and 6 we use levels of expenditure to compute the policy response. In models 3, 4, 7 and 8 we use expenditure proportions. Expenditure in levels is reported in millions of Mexican pesos. Expenditure proportions are reported in per thousand of the budget. Heteroscedasticity-consistent standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The narrative of budgetary changes as policy responses to societal signals could be contested for the case of public spending shrinkage. For instance, a government may decide to decrease its expenditure in a particular BP because the goal has been met. Alternatively, a BP may have an expiration date. In both situations, budgetary reductions have nothing to do with societal signals, but with administrative procedures.¹⁴ Therefore, it is important to test whether our results also hold for a dataset that only considers non-negative budgetary changes.

Hence, we replicate the estimations of table 4.4, but considering only non-negative changes in expenditure (either in the levels of spending or as proportions of the budget). The results are shown in table 4.6. We find no evidence in support of parallel information processing: the coefficient of the variable *NewsIndex* is not statistically significant. In contrast, the coefficient of the variable $\ln(\text{NewsIndex})$

¹³This new sample is composed by 25367 opinion pieces. We apply the same steps described in footnote 10 to preprocess the articles. We employ the same selection procedure outlined in footnote 12 to choose the optimal number of topics K , which is again 9.

¹⁴We could also argue that the withdrawal of funds from a BP might generally incur more political constraints within the decision-making system.

is positive and significant across the different models (columns 2, 4, 6 and 8), and its effect is larger in magnitude compared to the corresponding estimates in table 4.4. Focusing on the second column, the estimate implies that, on average and *ceteris paribus*, a 1% increase in the index raises the (positive) difference between paid and approved expenditure by around 18.5%. As in table 4.4, the inclusion of the expenditure approved for the BP reduces only marginally the effect of interest, that stays large and significant in the different specifications (columns 6 and 8). These findings are in line with the cross-country empirical generalisation of B. D. Jones, Baumgartner, et al. (2009), who observe that exceptional expansions in the budget tend to occur more often than dramatic cutbacks.

Table 4.6: Parallel and Serial Information Processing
Non-negative Changes

	OLS (1) $\Delta Level$	OLS (2) $\ln(\Delta Level)$	OLS (3) $\Delta Prop$	OLS (4) $\ln(\Delta Prop)$	OLS (5) $\Delta Level$	OLS (6) $\ln(\Delta Level)$	OLS (7) $\Delta Prop$	OLS (8) $\ln(\Delta Prop)$
<i>Intercept</i>	-28.566 (38.199)	48.850*** (15.140)	-0.843 (0.786)	16.674** (8.266)	-29.321 (37.532)	47.398*** (14.953)	-0.723 (0.763)	15.331* (8.132)
<i>NewsIndex</i>	199.667 (256.574)		5.981 (5.451)		203.763 (252.020)		5.046 (5.298)	
$\ln(NewsIndex)$		18.523** (7.883)		14.331*** (4.215)		17.960** (7.786)		13.124*** (4.178)
<i>Exp^{approved}</i>					0.036*** (0.013)		0.190*** (0.046)	
$\ln(Exp^{approved})$						0.025 (0.046)		0.104*** (0.039)
<i>N</i>	314	314	232	232	314	314	232	232
<i>Branch FE</i>	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>Adj R²</i>	-0.022	0.243	-0.059	0.231	-0.014	0.242	-0.033	0.252

Note: only non-negative changes in expenditure are considered. The dependent variable in models 1, 3, 5 and 7 is the absolute difference between the expenditure paid and the expenditure approved. In models 2, 4, 6 and 8 is the natural logarithm of the aforementioned variable. In models 1, 2, 5 and 6 we use levels of expenditure to compute the policy response. In models 3, 4, 7 and 8 we use expenditure proportions. Expenditure in levels is reported in millions of Mexican pesos. Expenditure proportions are reported in per thousand of the budget. Heteroscedasticity-consistent standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Overall, we find robust evidence supporting the serial processing hypothesis, which is consistent with government's budgetary overreactions. In addition, the sensitivity of policymakers' short term decisions to changes in the strength of informational signals from the media seems strong.¹⁵

However, despite the important insights provided by the previous analysis, the setting employed cannot be leveraged to make any substantial causal claim. In particular, reverse causality might be a source of bias in the estimated relationship. For instance, we could argue that what we observe is not the government which is responding with its interventions to the societal problems raised by columnists. It is the other way round. In fact, the government already wants to prioritise a policy area and influences

¹⁵In table C.5 of Appendix C.3 we test the robustness of the results on serial processing shown in table 4.4 using an alternative topic model to build the *News Index*. We choose the LDA with $K = 7$, which was the second best performing model in the selection procedure described in footnote 12. The coefficient of interest is slightly less significant across the different specifications, but the main findings are basically unaffected by this test. In addition, in table C.6 of Appendix C.4 we report descriptive statistics for the *News Index* across the different sets of regressions that we run.

the journalists to promote it (possibly, through bribes).

Under this scenario, corruption or other forms of clientelism would be an important factor leading the choice of the topics covered by the media outlets that we consider.

To tackle this issue, we reconduct the analysis using only opinion columns from *Reforma*. This civic-oriented newspaper is renowned for the sound professional practices of its journalists and having a zero-tolerance policy towards payoffs, gifts, perks or anything that could jeopardise its editorial independence or integrity (Márquez Ramírez 2014a,b).¹⁶ Results are reported in table 4.7 for the full sample of budget programmes. In table 4.8 instead, we consider only non-negative changes in the BPs.

Table 4.7: Parallel and Serial Information Processing
(Only *Reforma*)

	OLS (1) $\Delta Level$	OLS (2) $\ln(\Delta Level)$	OLS (3) $\Delta Prop$	OLS (4) $\ln(\Delta Prop)$	OLS (5) $\Delta Level$	OLS (6) $\ln(\Delta Level)$	OLS (7) $\Delta Prop$	OLS (8) $\ln(\Delta Prop)$
<i>Intercept</i>	-11.359 (15.210)	18.910*** (3.023)	-0.155 (0.237)	-9.318*** (1.609)	-11.967 (14.943)	16.882*** (2.752)	-0.172 (0.227)	-6.911*** (1.566)
<i>NewsIndex</i>	27.156 (34.636)		0.380 (0.539)		28.178 (34.025)		0.408 (0.517)	
$\ln(NewsIndex)$		7.581** (3.583)		3.169* (1.893)		7.817** (3.317)		3.437** (1.691)
<i>Exp^{approved}</i>					0.035*** (0.009)		0.060*** (0.010)	
$\ln(Exp^{approved})$						0.151*** (0.043)		0.216*** (0.041)
<i>N</i>	614	614	614	614	614	614	614	614
<i>Branch FE</i>	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>Adj R²</i>	0.022	0.117	0.015	0.119	0.038	0.155	0.072	0.247

Note: the dependent variable in models 1, 3, 5 and 7 is the absolute difference between the expenditure paid and the expenditure approved. In models 2, 4, 6 and 8 is the natural logarithm of the aforementioned variable. In models 1, 2, 5 and 6 we use levels of expenditure to compute the policy response. In models 3, 4, 7 and 8 we use expenditure proportions. Expenditure in levels is reported in millions of Mexican pesos. Expenditure proportions are reported in per thousand of the budget. Heteroscedasticity-consistent standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Note that, despite the substantial reduction in the size of the underlying newspapers corpus, the overall pattern is consistent with our baseline findings. The coefficient of $\ln(NewsIndex)$ is positive and significant, even if its magnitude is decreased. The effect of variable *NewsIndex* instead is not statistically different from 0 in any of the specifications (columns 1, 3, 5 and 7 in both tables).

The results provide additional evidence for the serial processing hypothesis, addressing corruption as a possible (unobserved) factor affecting the relationship under analysis. However, to further address the issue of causality, we design in the next section a more robust empirical setting.

The framework that we propose tries to deal with the potential confounders that may drive the government's overreaction on how it allocates resources across the relevant BPs as a consequence of

¹⁶The sample we consider now is composed by 2497 opinion pieces. After applying the preprocessing steps described in footnote 10, we select the optimal model following the procedure discussed in footnote 12. The resulting LDA has $K = 3$.

Table 4.8: Parallel and Serial Information Processing
Non-negative Changes (Only *Reforma*)

	OLS (1) $ \Delta Level $	OLS (2) $\ln(\Delta Level)$	OLS (3) $ \Delta Prop $	OLS (4) $\ln(\Delta Prop)$	OLS (5) $ \Delta Level $	OLS (6) $\ln(\Delta Level)$	OLS (7) $ \Delta Prop $	OLS (8) $\ln(\Delta Prop)$
<i>Intercept</i>	-27.637 (27.539)	28.281*** (5.172)	-0.705 (0.547)	-4.581 (2.905)	-27.645 (27.140)	27.776*** (5.121)	-0.719 (0.525)	-3.097 (2.880)
<i>NewsIndex</i>	68.887 (65.765)		1.636 (1.233)		68.555 (64.812)		1.635 (1.183)	
$\ln(\text{NewsIndex})$		16.645*** (5.672)		7.939** (3.467)		16.730*** (5.604)		8.309** (3.364)
<i>Exp^{approved}</i>					0.036*** (0.014)		0.192*** (0.046)	
$\ln(\text{Exp}^{\text{approved}})$						0.038 (0.045)		0.123*** (0.037)
<i>N</i>	314	314	232	232	314	314	232	232
<i>Branch FE</i>	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
<i>Adj R²</i>	-0.021	0.253	-0.056	0.221	-0.012	0.252	-0.030	0.251

Note: only non-negative changes in expenditure are considered. The dependent variable in models 1, 3, 5 and 7 is the absolute difference between the expenditure paid and the expenditure approved. In models 2, 4, 6 and 8 is the natural logarithm of the aforementioned variable. In models 1, 2, 5 and 6 we use levels of expenditure to compute the policy response. In models 3, 4, 7 and 8 we use expenditure proportions. Expenditure in levels is reported in millions of Mexican pesos. Expenditure proportions are reported in per thousand of the budget. Heteroscedasticity-consistent standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

the different level of signals.

4.5 News and Budget Allocation: A Natural Experiment

In this section, we exploit the fact that increased media coverage triggered by specific events can dramatically shift the attention of the politicians (Wolfe, B. D. Jones, and Baumgartner 2013), leading to potential overreactions by the government. For this, we set up a natural experiment framework and rely on two separate exogenous events to try to analyse the relationship of interest from a causal perspective.

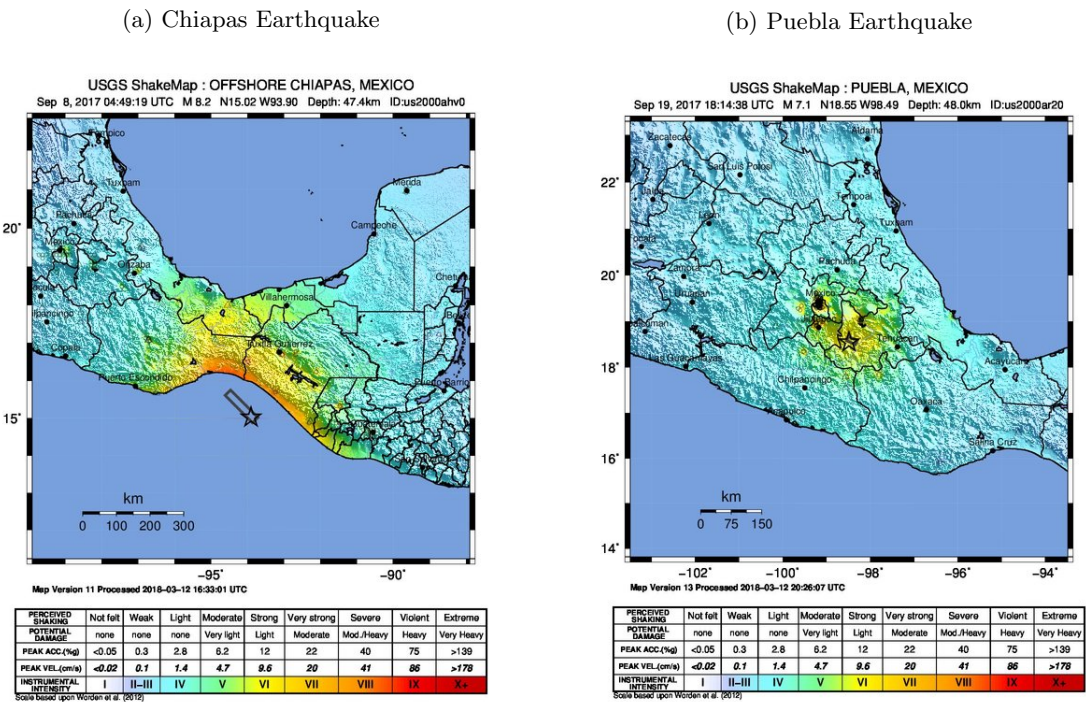
4.5.1 The Setting

We focus on two earthquakes that hit Mexico in 2017 and the different media coverage that they received through opinion pieces. Since these events represent negative outcomes, it is difficult to argue that reverse causality could take place in this setting (e.g., that the government would pay opinion journalists to give ample coverage).

Each of the two earthquakes happened in different states in the month of September. The first one occurred in the state of Chiapas on the 7th of September and had its epicenter at the Gulf of

Tehuantepec. It was the strongest recorded earthquake since the beginning of the 20th century,¹⁷ with a moment magnitude of 8.2. The Chiapas earthquake also affected the neighbouring state of Oaxaca. The second earthquake took place in the state of Puebla on the 19th of September 2017 and had its epicenter located, according to the United States Geological Survey, at 1km East from San Felipe Ayutla (Puebla).¹⁸ This earthquake had a moment magnitude of 7.1, while it caused important economic losses and several deaths in the states of Puebla and Morelos and in the Metropolitan Area of Mexico City. Figure 4.2 shows the ShakeMaps for the two events.

Figure 4.2: Earthquakes ShakeMaps



Source: United States Geological Survey
(<https://www.usgs.gov/natural-hazards/earthquake-hazards/earthquakes>).

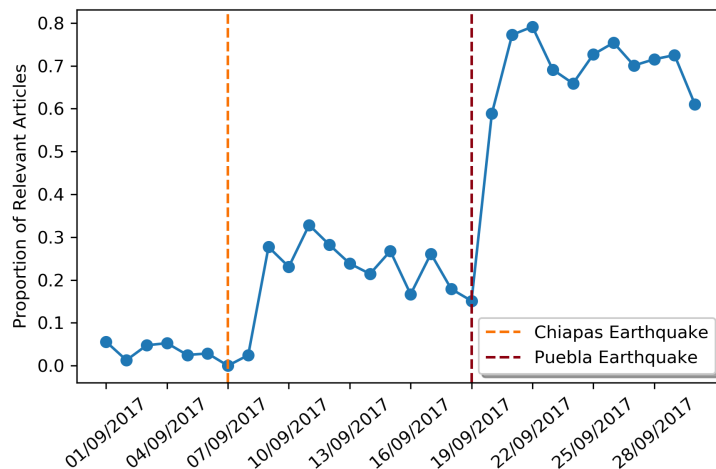
In spite of its lower magnitude, the Puebla earthquake received substantially more media attention at national and international level. A significant factor at play was the high number of victims (more than 220) and the devastating damages that the seism caused in the capital of the country.¹⁹ To top it up, it occurred during the anniversary of the 1985 Mexico City earthquake, which was one of the deadliest in the entire history of the country.

¹⁷Source: National Centers for Environmental Information (<https://www.ngdc.noaa.gov/hazel/view/hazards/earthquake/search>).
¹⁸Source: United States Geological Survey (<https://www.usgs.gov/natural-hazards/earthquake-hazards/earthquakes>).
¹⁹Source: Instituto Belisario Domínguez (<http://bibliodigitalibd.senado.gob.mx/handle/123456789/3721>).

This last event (the anniversary of the 1985 earthquake) is crucial because, every year, there is ample media coverage on this national tragedy. The 1985 seism is perceived as both a catastrophe and as a moment of national union. Therefore, this remembrance event, by itself, generates unbalanced media coverage without affecting policy responses. Thus, we could say that the data capture not only one but two exogenous events that shift media attention.

The differentiated level of media coverage is clearly seen in our newspapers corpus, as highlighted by figure 4.3. We report the daily proportion of earthquake-related opinion pieces during the month of September 2017.²⁰ We can see the proportion peaking in the aftermath of the two events. However, the peak is definitely more pronounced in the wake of the Puebla earthquake, with the fraction of earthquake-related articles exceeding 79% of the daily number of opinion pieces on the 22nd of September.

Figure 4.3: Proportion of Earthquake-related Opinion Pieces (September 2017)



Note: relevant articles are defined as those opinion pieces containing at least one of the following words: *terremoto*, *sismo*, *seísmo*, *cataclismo*, *hecatombe*, *sacudida* and *temblor*.

To combine the newspapers data on these events with budgetary information, we exploit an additional feature of the open spending database: quarterly and state-level disaggregation. For the year under study (2017), the expenditure data offer quarterly numbers on the paid budget at the level of each BP, which will allow us to examine changes in expenditure over time. Furthermore, it is possible to identify those BPs that are assigned to specific states. In this experiment, we are interested in states where the earthquakes had a similar effect in terms of economic costs and human lives. For the Chiapas

²⁰By *earthquake-related*, we mean an opinion piece containing at least one of the following words: *terremoto* (earthquake), *sismo* (seism), *seísmo* (seism), *cataclismo* (cataclysm), *hecatombe* (calamity), *sacudida* (shake), *temblor* (tremor).

earthquake, we focus on the state of Oaxaca. For the Puebla earthquake, we concentrate on the state of Morelos.

To frame this natural experiment setting, we use a Difference-in-Differences (DID) design. Since both seisms occurred at the end of the third quarter, we employ data on the last two quarters to conduct the experiment, while the previous ones can be used to assess the parallel trends assumption (Angrist and Pischke 2008). Notice that the fact that both earthquakes occurred in a reduced time frame is favourable to our design because we can discard the potential change/improvement of the public administration or official procedures to deal with natural disasters.

Oaxaca suffered 78 deaths during the Chiapas earthquake.²¹ Morelos was affected by 74 casualties during the Puebla earthquake.²² In the DID design, we consider Oaxaca to be the control and Puebla the treatment group, since the latter is the one experiencing media over-saturation. If the high media coverage of these two natural disasters drove the attention of policymakers and their subsequent decisions, then we would expect differentiated policy responses in terms of the BPs that are relevant to both Oaxaca and Morelos and to the topics related to the earthquakes.

To identify those BPs that are closer to the earthquake issue, we train different LDA models on the newspapers corpus for the relevant quarters (i.e., the 3rd and 4th), with the aim of isolating the earthquake topic.²³ We apply the preferred model to the textual data of the BPs, so that we can encode each budget programme as a vector of topic proportions (as described in subsection 4.4.2).²⁴ Finally, to identify the BPs related to the earthquake issue, we select those BPs that are common to both states where the earthquake-topic proportion is greater than the average.²⁵ The treatment group is defined as those earthquake-related BPs directed to Morelos, while the control group includes the earthquake-related BPs directed to Oaxaca.²⁶ Table 4.9 summarises the components of our DID

²¹Source: *La Jornada* (<https://www.jornada.com.mx/2017/09/13/politica/007n1pol>, retrieved on the 20th of October 2020). Oaxaca's population consists of 3,967,889 inhabitants (according to the 2015 intercensal survey).

²²Source: *Instituto Belisario Domínguez* (<http://bibliodigitalibd.senado.gob.mx/handle/123456789/3721>). Morelos' population consists of 1,903,811 inhabitants (according to the 2015 intercensal survey).

²³The subsample that we consider now contains 19352 opinion pieces and it spans the period from the 1st of July 2017 to the 31st of December 2017.

²⁴We conduct a narrower search to find the optimal number of topics K across 8 different values: {3, 4, 5, 6, 7, 8, 9, 10}. We rank the models according to the same statistics mentioned in subsection 4.4.2. We then select as our preferred model the LDA with $K = 9$, which is the second best. We favour this model instead of the best performing one ($K = 6$) as a visual inspection of the relevant topic (i.e., the one that should relate to the earthquake) reveals a cluster of words much closer to the issue of interest in the first case. However, to assess the consistency of our findings, we replicate the main estimations using the LDA with $K = 6$ in table C.7 of Appendix C.5. The size of the coefficient of the interaction term is only marginally different and the effect is still significant at 10% ($p = 0.071$). We are not able to test the findings that we get when we use the tighter definition of earthquake-related programmes, as the resulting sample size is too small (32 observations) to obtain any meaningful estimate from the DID regression.

²⁵That is, budget programme j is selected if $\hat{\theta}_{j, \text{earthquake}} > \bar{\hat{\theta}}_{\text{earthquake}}$.

²⁶Some examples of earthquake-related BPs are: *scientific research and technological development in public education, human resources' education and training for the health sector, the employment support programme, urban development policy and land management and the housing support programme*.

design.

Table 4.9: Difference-in-Differences Design

Component	Description
<i>Treatment Group</i>	Earthquake-related budget programmes directed to Morelos.
<i>Control Group</i>	Earthquake-related budget programmes directed to Oaxaca.
<i>Treatment</i>	Being hit by the earthquake that received higher media coverage (i.e., the Puebla earthquake).
<i>Period Before the Treatment</i>	The 3 rd quarter.
<i>Period After the Treatment</i>	The 4 th quarter.

4.5.2 Analysis and Results

We estimate the following DID regression:

$$Exp_{jst} = \alpha + \beta_1 Treat_s + \beta_2 Quarter_t + \beta_{DID}(Treat_s \times Quarter_t) + \beta_3 Exp_{jst}^{past} + \beta_4 Exp_{js}^{approved} + \epsilon_{jst} \quad (4.7)$$

Where Exp_{jst} denotes paid expenditure on BP j directed to state s during quarter t . The dummy variable $Treat_s$ takes the value of 1 for BPs directed to Morelos and 0 for those directed to Oaxaca. The dummy variable $Quarter_t$ instead takes the value of 1 if the paid expenditure on BP j corresponds to the 4th quarter and 0 otherwise. The interaction term $Treat_s \times Quarter_t$ captures being a BP directed to Morelos during the 4th quarter. Note that, given this setting, we can formulate a hypothesis about the direction of the interaction's effect. More specifically, the higher media coverage of the Puebla earthquake should attract government's attention and generate a disproportionate reallocation of earthquake-related resources towards Morelos compared to Oaxaca. Hence, we expect β_{DID} to be positive. In addition, we control for the spending level approved for BP j ($Exp_{js}^{approved}$) and for how much the government has already spent on it during the previous quarters (Exp_{jst}^{past}), as they represent relevant constraints on the amount of expenditure that can be paid at each quarter.

Model 1 in table 4.10 reports the estimates of the DID regression. The coefficient of the interaction term is positive and significant at 5%, which is consistent with the hypothesis of serial information processing and policy overreactions. On average and *ceteris paribus*, earthquake-related BPs directed to Morelos received in the 4th quarter almost 37 millions of Mexican pesos more. In the second column of the

table we run the same specification, but using budget data from the first two quarters. A significant coefficient of the variable $Treat \times Quarter$ would cast doubt on the parallel trends assumption on which the DID identification strategy relies. However, the coefficient of the interaction term is not statistically different from 0.

Table 4.10: Natural Experiment

	Model (1) <i>Level</i>	Parallel Trends (1) <i>Level</i>	Model (2) <i>Level</i>	Parallel Trends (2) <i>Level</i>
<i>Intercept</i>	9.717 (11.475)	6.372 (4.523)	24.779 (24.108)	14.685 (9.225)
<i>Treat</i>	-10.200 (11.148)	-3.375 (5.475)	-30.570 (25.971)	-2.740 (12.358)
<i>Quarter</i>	-29.860* (15.306)	32.335** (16.034)	-51.942 (33.048)	59.960 (46.864)
<i>Treat</i> \times <i>Quarter</i>	36.974** (18.033)	-20.258 (18.407)	75.336** (37.505)	-47.727 (52.196)
<i>Exp^{past}</i>	0.273** (0.120)	-0.144*** (0.027)	0.414*** (0.029)	-0.161*** (0.011)
<i>Exp^{approved}</i>	0.097 (0.064)	0.264*** (0.002)	0.021 (0.017)	0.261*** (0.002)
<i>N</i>	284	284	84	84
<i>Adj R²</i>	0.958	0.971	0.992	0.978

Note: all models are estimated via OLS. The dependent variable is paid expenditure in millions of Mexican pesos. In model 1 budget programmes are selected if their earthquake-topic proportion is greater than the average. In model 2 budget programmes are selected if their earthquake-topic proportion is greater than the value that we would expect under a uniform distribution across the topics. Heteroscedasticity-consistent standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

In model 2, we use a stricter criterion to identify earthquake-related BPs in terms of earthquake-topic proportions.²⁷ The coefficient of the interaction term is still positive and significant at 5% and its magnitude is even larger than before, meaning that earthquake-related BPs directed to Morelos received during the 4th quarter, on average, 75 millions of Mexican pesos more. It seems that, when we look at programmes that are closer to the issue that is driving attention through media coverage (i.e., the seism), the policy response is even greater. In the last column of the table we provide evidence in support of the parallel trends assumption for this alternative sample of BPs. Even in this case, the coefficient of the interaction term is not statistically significant.

²⁷More specifically, we consider only those BPs in which the earthquake-topic proportions are greater than the value that we would expect under a uniform distribution across the K topics (i.e., BP j is selected if $\hat{\theta}_{j, earthquake} > \frac{1}{K} = \frac{1}{9} \approx 0.11$).

4.5.3 Testing Alternative Hypotheses

In this subsection, we would like to consider alternative hypotheses that may explain the differentiated policy response in terms of public funding allocated to Morelos and Oaxaca after their respective earthquakes. Naturally, there may be other political, societal, and institutional factors that contribute to the estimated effects. Nevertheless, as we demonstrate here, the alternative hypotheses that we address would hardly explain the disproportionate responses that we find and, in some situations, they would contradict them.

Systematic Differences in Media Coverage

The first alternative hypothesis that we consider is that a systematic higher coverage of Morelos is the cause of the observed differences in government expenditure. Under this scenario, the government would channel additional resources because it constantly receives more informational signals about this state. If this hypothesis held, we would expect the effect of the dummy variable *Treat* to be positive and significant, which is not.

In addition, we formally test for differences in media coverage between Morelos and Oaxaca using the newspapers corpus. For this, we consider the opinion pieces written during the first two quarters of the year and conduct a Boolean search identifying those articles that mention either the state, its governor, the capital city or its mayor.²⁸ We then compute the differences in the daily number of opinion pieces referring to the two states and perform a paired *t*-test. The mean of the differences in daily mentions (0.1492) is not statistically different from 0 at any conventional significance level ($p = 0.3457$). Hence, not only this hypothesis is inconsistent with the econometric tests, but also there seems to be no systematic difference in media coverage between the two states.

Differentiated Economic Damages

Next, let us consider the possibility that the observed government's responses are a consequence of substantially different levels of economic damages between Morelos and Oaxaca during their respective earthquakes. To assess such differences, we employ data from Mexico's National Disasters Fund.²⁹

This is a federal fund that is mobilised in extreme circumstances like the earthquakes of our study,

²⁸Note that we only consider those articles that are not used in the training of the topic model employed in the natural experiment.

²⁹Source: *Transparencia Presupuestaria* (<https://www.transparenciapresupuestaria.gob.mx/es/PTP/fuerzamexico>).

to provide emergency support for immediate disaster relief and reconstruction. The amount of funds transferred to the states is determined by the estimated economic damages. The data indicate that Morelos received \$1,318,535,480 Mexican pesos, while Oaxaca was supported with \$1,436,112,556.³⁰ Thus, the data suggest that Oaxaca's economic damages were larger than Morelos', so this hypothesis turns out to be contradictory. Needless to say, the data on federal emergency funds are not part of the BPs that we use in our study.

Corrupt Media

We could think of a situation in which the differentiated coverage between Morelos and Oaxaca is the result of media outlets that respond to government's interests. First, as we have argued at the beginning of subsection 4.5.1, since they represent negative shocks it is not clear how earthquake outcomes could motivate the government to buy media presence. Furthermore, even if this happened, it seems like a rather weak explanation of the government's disproportionate response towards Morelos' BPs.

The most plausible scenario under which this hypothesis would work would be that the government pays the media to write about an extremely effective response in the aftermath of the Puebla earthquake in order to capitalise politically in Morelos. Thus, this would imply that there are ulterior political motives, and that buying favourable opinion pieces should be accompanied by an increase in earthquake-related BPs. Consequently, it is necessary to analyse the political context in order to verify the plausibility of this argument, something that we do next.

Political Motives

The last alternative interpretation of our findings is that the federal government channelled resources to Morelos in the aftermath of the earthquake due to political reasons. As we show in this subsection, this seems rather unlikely.

First, let us consider the case in which the party sitting in the federal office has incentives to pay political favours to friendly parties ruling the states. The party in office from 2012 to 2018 at the federal level was the *Partido Revolucionario Institucional* (PRI). It turns out that, during 2017, the

³⁰These values represent federal exercised expenditure and include the scheme of Immediate Partial Supports and the Damage Assessment Expense.

party governing Morelos was an opposition one: the *Partido de la Revolución Democrática* (PRD). Thus, it seems unlikely that the federal government would use earthquake-related BPs to pay political favours to the government of Morelos. In addition, the ruling party in Oaxaca during the same period was the PRI. This also contradicts the hypothesis, which would imply more funds for the latter state since the ruling parties are the same.

Second, an alternative political explanation for the uneven distribution of earthquake-related funds is that the federal government was using that money in political campaigns (especially in the common and illegal practice of vote buying) for the upcoming gubernatorial election (July 2018). This hypothesis would assume some certainty by the PRI about its real possibilities of winning the election in Morelos; otherwise these resources could be better used in swing states. This also seems an unlikely scenario because the 2018 election in Morelos was won by a landslide (52.59%) by the popular candidate of the opposing coalition (*Juntos Haremos Historia*), Cuauhtémoc Blanco (a public figure who was a prolific football player), whereas the candidate of the PRI got only around 6%. Hence, this hypothesis also seems rather weak.

A similar argument could be made for the general election, also held in July 2018. However, the mechanism through which the Mexican federal government has historically bought votes has been through BPs related to development and welfare. For instance, by leveraging established social assistance programmes that the authority can arbitrarily condition on the votes of the beneficiaries.

We formally test this hypothesis using the topic model trained for the natural experiment. This time, however, we select BPs related to those clusters of words that capture the topics of elections and social policy.³¹ We choose the same criterion used to identify earthquake-related BPs. That is, we select those BPs for which the relevant topic proportion is greater than the average.³² We then run the same regression specification as in 4.7 but using these BPs.

The results are reported in table 4.11. In the first column, we consider the BPs related to elections, whereas in the second one those closer to social policy. In both cases the coefficient of interest (the one of the interaction term) is not statistically different from 0, providing no evidence in support of this alternative hypothesis.

³¹Word clouds for these topics, together with the one related to the earthquake, are provided in figure C.1 of Appendix C.6.

³²Hence, BP j is selected in the first case if $\hat{\theta}_{j,elections} > \bar{\theta}_{elections}$ and in the second one if $\hat{\theta}_{j,socialpolicy} > \bar{\theta}_{socialpolicy}$. In addition, it must also satisfy the condition $\hat{\theta}_{j,earthquake} \leq \bar{\theta}_{earthquake}$, so that we do not select BPs that are considered part of the natural experiment.

Table 4.11: Natural Experiment
Alternative Hypothesis

	Model (1) <i>Level</i>	Model (2) <i>Level</i>
<i>Intercept</i>	11.128 (7.380)	17.807 (14.163)
<i>Treat</i>	-10.542 (8.029)	-15.313 (16.034)
<i>Quarter</i>	-13.311 (8.208)	-25.909 (21.251)
<i>Treat</i> \times <i>Quarter</i>	11.028 (10.135)	21.567 (20.666)
<i>Exp^{past}</i>	0.063* (0.036)	0.241** (0.113)
<i>Exp^{approved}</i>	0.218*** (0.029)	0.127* (0.071)
<i>N</i>	232	236
<i>Adj R²</i>	0.961	0.804

Note: all models are estimated via OLS. The dependent variable is paid expenditure in millions of Mexican pesos. In model 1 budget programmes are selected if their elections-topic proportion is greater than the average. In model 2 budget programmes are selected if their social policy-topic proportion is greater than the average. Heteroscedasticity-consistent standard errors are reported in parentheses.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

4.6 Conclusion

In this paper we explored how a government processes information and how this is reflected into its decisions. Our policy domain of interest is budget allocation. More specifically, we examine short term adjustments to expenditure on individual budget programmes (BPs). The type of informational signals that we consider are opinion pieces from newspapers, given their political relevance as a proxy for societal issues in Mexico, which is the institutional setting under analysis.

We exploit a newspapers corpus of more than 35000 articles and textual data on the features of the BPs to create a *News Index*, that links informational signals to policy outputs. Our index captures how close the single programmes are to the main issues discussed in the opinion columns. We use the index to assess two different descriptive models of government's information processing: parallel and serial. We find that changes in the *News Index* are associated with overreactions by the government, measured through the difference between approved and paid expenditure on the BPs. The results

pass a series of robustness checks, including variations in the underlying newspapers corpus and a different operationalisation of the dependent variable. Our findings are consistent with a model of serial information processing, which is characterised by selective attention to a subset of policy areas.

We further investigate the relationship between news coverage and government's overreactions within a causal framework. We devise a natural experiment with a Difference-in-Differences design, exploiting the different level of media attention to two earthquakes that occurred in Mexico in 2017. With our design, we try to address the potential confounders that might bias the effect of newspapers coverage on government's policy response. The results show how media can focus the attention of policymakers and influence the way in which a government allocates economic resources. Our natural experiment suggests that federal spending was relatively higher towards the state hit by the earthquake that received more coverage. To substantiate our claim, we address a number of alternative hypotheses that turn out to be inconsistent with the data. Overall, our findings are in line with the theory of disproportionate information processing that underlies the decision-making model proposed by B. D. Jones and Baumgartner (2005).

The first part of the analysis however is limited to a single budget cycle. Hence, it does not take into account how a poor focus on specific issues might cumulate over time, causing overreactions by the system when attention is dramatically shifted. In addition, we do not consider the dynamics of interplay between policymakers and media outlets. Policy responses through budget adjustments can be incorporated later on in the media, where they are addressed and debated, spurring subsequent budgetary reactions from the political system. Further research on the topic should extend our methodology to incorporate these complex feedback loops across budget cycles in a consistent way.

Moreover, the index that we devise is solely based on a particular type of informational signal (which is nevertheless of high interest in our context), but it does not encompass the whole spectrum of different inputs to which policymakers are exposed. Besides, our *News Index* only captures the main issues discussed in newspapers' opinion pieces and their salience, but it is agnostic about their polarity. This is an important caveat and future work should investigate how distinct sentiments can attract government's attention in different ways.

It is also important to underline that the design of our natural experiment is highly dependent on the institutional context under analysis, both in terms of the political system and media characteristics, which might threaten the external validity of our conclusions.

Despite the limitations, we hope with this study to provide an original approach based on text analysis to link media and policy outputs, together with new evidence on how a government tends to focus its attention and overreact to informational signals.

An issue increasingly important in a world where everyone is constantly exposed to a growing flow of information, but where we all face the same cognitive constraints, even the government, possibly the most impactful decision-making system in many societies.

Chapter 5

Conclusion

In this thesis I provide three empirical studies where I combine methods developed in other fields, such as Computer Science and Computational Linguistics, with more traditional Econometrics techniques. The underlying rationale of this exercise is to fully leverage the potential of large and unstructured data, especially texts.

In the three papers the general theme addressed is the behaviour of policymakers, either when exposed to exogenous shocks (chapter 2), informational signals (chapter 4) or when they strategically interact in a broader institutional setting (chapter 3).

In chapter 2, I analyse the effect of terrorist attacks on the immigration rhetoric of British Members of Parliament (MPs) on Twitter. I employ text analysis and machine learning techniques to collect and identify those Tweets that are related to the topic of immigration. The panel structure of the dataset that I build is leveraged to frame an event study design. I find that, during the general election, the 2017 Manchester bombing caused a significant reduction in the number of immigration-related Tweets posted by the MPs. In addition, I provide suggestive evidence that this “muting effect” might result from the risk-averse behaviour of the politicians during the election campaign. Moreover, this counterintuitive finding displays interesting variations according to the socio-economic characteristics of the MPs.

However, it is important to highlight that, at the time of the Manchester attack, there was no politician belonging to any far-right or anti-establishment parties holding a seat in the Parliament. We might well expect that these politicians would be more willing to expose themselves on such a risky topic

given their firm position on the issue of immigration.

In addition, the paper does not explore a possible link between the observed behaviour of the MPs and the electoral response, which seems an interesting research question that could be addressed in the future.

Moreover, even if the literature suggests that politicians do not use Twitter to promote common party policies (Adi, Erickson, and D. G. Lilleker 2014), exploring coordinated action among cluster of accounts after the attack might be a relevant extension of the work. Besides, a deep analysis of different topics that are more directly related to terrorist attacks (e.g., multiculturalism and Islamophobia) might provide further insights on how and if such events can affect the social media agenda of policymakers.

An important caveat of my study is that, mainly due to the research design chosen, it does not consider the polarity of the Tweets, but only their topic. Hence, a valuable extension of the work would address this issue by providing a sound framework to capture the causal effect of exogenous events on policymakers' sentiment towards immigration.

In chapter 3, we investigate the nature of strategic voting in the United Nations General Assembly (UNGA). By considering different forms of strategic interactions (i.e., vote trading, vote buying and coercion) and their respective “currency of exchange”, we propose a methodology to capture if deviations from the expected votes are systematically reciprocated or not. For this purpose, we first rely on Structural Topic Models (Margaret E. Roberts et al. 2014) and other machine learning techniques to predict the voting behaviour of member states in the UNGA. We identify the beneficiaries of eventual deviations from the expected votes through draft sponsorship. We then combine the information that we have to build a directed-deviations network (DDN), which describes the structure of possible exchanges between members of the Assembly. Through the graph, we can compute an aggregate statistic: the *Reciprocity Index*. This measure quantifies the degree to which UNGA members exchange votes and can be used to assess the prevalence of different forms of strategic voting. By applying our methodology, we find that deviations from the expected votes are systematically *not* reciprocated. Further analyses suggest that our results are consistent with a narrative of state socialisation (Alderson 2001) and vote buying (e.g., Dreher, Nunnenkamp, and Thiele 2008; Carter and Stone 2015).

One of the limitations of the methodology employed is that it does not consider that countries might have strong preferences not just for passing a resolution, but also for blocking it. However, such

limitation is mainly due to data availability: voting information on UN resolutions that are not passed is not recorded.

Another shortcoming of the study is that we do not differentiate between *no* votes, abstentions and absences. In fact, these alternatives might well be the result of different strategic considerations (Dreher and N. M. Jensen 2013). Hence, a possible addition to our work could integrate these different voting outcomes within our current framework through a weighting scheme.

One could also enrich the set of signals through which countries express their preferences for certain proposals, for instance by considering the speeches that they deliver during the General Debate held at the beginning of each UNGA session.

In addition, a relevant extension of our work could also include development aid and financial flows in the construction of the DDNs. One might then analyse if voting deviations are systematically matched by the beneficiaries with such flows, providing further evidence for the vote buying hypothesis.

Finally, in chapter 4, we address the issue of how a government processes and responds to incoming flows of information. More specifically, we focus on how the federal Mexican government adjusts its expenditure on individual budget programmes (BPs). The informational signals that we consider are newspapers' opinion columns, given their political relevance as proxies for societal issues in our context. We employ Latent Dirichlet Allocation (Blei, Ng, and Jordan 2003) on a newspapers corpus of more than 35000 articles and textual data describing the main features of the BPs to build a *News Index*. The index provides a mapping between the informational signals and the policy outputs, capturing how close the BPs are to the main issues covered in the media. We use the index to assess two different descriptive models of information processing: parallel and serial. Our results provide evidence for the serial processing hypothesis, which is characterised by selective attention and disproportionate policy response. These conclusions are further supported by our natural experiment, in which we try to provide a causal framework to study the relationship between newspapers coverage and government's overreactions. For this, we exploit the different media attention received by two earthquakes that occurred in Mexico in 2017 to frame a Difference-in-Differences research design.

However, the setting of our natural experiment heavily relies on the institutional context considered, which might be problematic for the external validity of our conclusions. Hence, other case studies would provide a valuable addition to test the robustness of the general insights provided by the analysis.

I should also emphasise that, since we focus on a single fiscal year, the methodology that we employ

to link informational signals with government's response does not take into account the complex loops of negotiation that occur between policymakers and the media. Further research on the topic should incorporate these feedback dynamics into our methodology, to extend the study to subsequent budget cycles.

A general remark is that we restrict the set of informational signals considered to opinion columns. Even if this choice is motivated by their relevance in the Mexican context, a possible addition to our work could explore how the government processes and reacts to alternative sources of information.

One last limitation, is that our *News Index* does not incorporate the tone of the opinion pieces, which might well elicit a differentiated policy response. Hence, future research could investigate how informational inputs with different polarities capture government's attention in different ways, resulting in specific policy interventions.

While providing new insights into the individual research questions that they address, the main message that these different studies (and my thesis as a whole) would like to convey is the value of keeping an open dialogue with different methodologies and disciplines.

The importance of embracing a multidisciplinary approach in Economics already resonated in the words of John Maynard Keynes:

"[...] the master-economist must possess a rare combination of gifts. He must reach a high standard in several different directions and must combine talents not often found together. He must be mathematician, historian, statesman, philosopher-in some degree. He must understand symbols and speak in words. He must contemplate the particular in terms of the general, and touch abstract and concrete in the same flight of thought. He must study the present in the light of the past for the purposes of the future. No part of man's nature or his institutions must lie entirely outside his regard. He must be purposeful and disinterested in a simultaneous mood; as aloof and incorruptible as an artist, yet sometimes as near the earth as a politician."

Keynes, John M. (1924). "Alfred Marshall, 1842-1924". In: *The Economic Journal* 34.135, pp. 311-372.

Although this represents an almost utopian task, the few lines above show us the path that we have to pursue. A path that requires time and effort. The patience to learn and understand the language that different disciplines may speak. The ability, and willingness, to appreciate the work of others.

Bibliography

- Adi, Ana, Kristofer Erickson, and Darren G. Lilleker (2014). “Elite tweets: Analyzing the Twitter communication patterns of Labour party peers in the House of Lords”. In: *Policy & Internet* 6.1, pp. 1–27 (cit. on pp. [9](#), [10](#), [114](#)).
- Aksoy, Deniz (2012). “Institutional arrangements and logrolling: Evidence from the European Union”. In: *American Journal of Political Science* 56.3, pp. 538–552 (cit. on pp. [43](#), [47](#)).
- Alderson, Kai (2001). “Making sense of state socialization”. In: *Review of International Studies*, pp. 415–433 (cit. on pp. [3](#), [44](#), [52](#), [74](#), [114](#)).
- Alexandrova, Petya, Marcello Carammia, and Arco Timmermans (2012). “Policy punctuations and issue diversity on the European Council agenda”. In: *Policy Studies Journal* 40.1, pp. 69–88 (cit. on p. [84](#)).
- Allen, W. and S. Blinder (2013). *Migration in the news: Portrayals of immigrants, migrants, asylum seekers and refugees in national British newspapers, 2010-2012*. Tech. rep. Compas, pp. 1–31 (cit. on pp. [2](#), [5](#), [11](#), [12](#)).
- Angrist, Joshua D. and Jörn-Steffen Pischke (2008). *Mostly harmless econometrics: An empiricist’s companion*. Princeton University Press (cit. on p. [104](#)).
- Arora, Sanjeev et al. (2013). “A practical algorithm for topic modeling with provable guarantees”. In: *Proceedings of the 30th International Conference on International Conference on Machine Learning*. Vol. 28. ICML’13. JMLR.org, pp. 280–288 (cit. on p. [60](#)).
- Athey, Susan and Guido W. Imbens (2015). “Machine learning methods for estimating heterogeneous causal effects”. In: *Stat* 1050.5, pp. 1–26 (cit. on p. [11](#)).

- Bailey, Michael A, Anton Strezhnev, and Erik Voeten (2017). “Estimating dynamic state preferences from United Nations voting data”. In: *Journal of Conflict Resolution* 61.2, pp. 430–456 (cit. on pp. [45](#), [60](#), [62](#), [63](#)).
- Bajari, Patrick et al. (2015). “Machine learning methods for demand estimation”. In: *American Economic Review* 105.5, pp. 481–485 (cit. on p. [11](#)).
- Banerjee, Abhijit V. (1992). “A simple model of herd behavior”. In: *The Quarterly Journal of Economics* 107.3, pp. 797–817 (cit. on p. [62](#)).
- Barbieri, Francesco, Francesco Ronzano, and Horacio Saggion (2015). “How topic biases your results? A case study of sentiment analysis and irony detection in Italian”. In: *Proceedings of the International Conference Recent Advances in Natural Language Processing*, pp. 41–47 (cit. on p. [14](#)).
- Baumgartner, Frank R. and Bryan D. Jones (1993). *Agendas and instability in American politics*. Chicago, IL: University of Chicago Press (cit. on p. [84](#)).
- Becker, Raphael N. et al. (2015). “The preoccupation of the United Nations with Israel: Evidence and theory”. In: *The Review of International Organizations* 10.4, pp. 413–437 (cit. on p. [56](#)).
- Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen (2014). “Inference on treatment effects after selection among high-dimensional controls”. In: *The Review of Economic Studies* 81.2, pp. 608–650 (cit. on p. [11](#)).
- Blei, David M., Andrew Y. Ng, and Michael I. Jordan (2003). “Latent dirichlet allocation”. In: *Journal of Machine Learning Research* 3, pp. 993–1022 (cit. on pp. [3](#), [83](#), [94](#), [115](#)).
- Bouma, Gerlof (2009). “Normalized (pointwise) mutual information in collocation extraction”. In: *Proceedings of GSCL*, pp. 31–40 (cit. on p. [95](#)).
- Bozzoli, Carlos and Cathérine Müller (2011). “Perceptions and attitudes following a terrorist shock: Evidence from the UK”. In: *European Journal of Political Economy* 27, pp. 89–106 (cit. on p. [8](#)).
- Brader, Ted, Nicholas A. Valentino, and Elizabeth Suhay (2008). “What triggers public opposition to immigration? Anxiety, group cues, and immigration threat”. In: *American Journal of Political Science* 52.4, pp. 959–978 (cit. on pp. [2](#), [5](#), [9](#)).
- Brazys, Samuel and Diana Panke (2017a). “Analysing voting inconsistency in the United Nations General Assembly”. In: *Diplomacy & Statecraft* 28.3, pp. 538–560 (cit. on p. [51](#)).

- Brazys, Samuel and Diana Panke (2017b). “Why do states change positions in the United Nations General Assembly?” In: *International Political Science Review* 38.1, pp. 70–84 (cit. on pp. 45, 62, 63, 66, 72–74).
- Breiman, Leo (2001a). “Random forests”. In: *Machine learning* 45.1, pp. 5–32 (cit. on p. 65).
- (2001b). “Statistical modeling: The two cultures (with comments and a rejoinder by the author)”. In: *Statistical Science* 16.3, pp. 199–231 (cit. on p. 11).
- Breunig, Christian, Chris Koski, and Peter B. Mortensen (2010). “Stability and punctuations in public spending: A comparative study of budget functions”. In: *Journal of Public Administration Research and Theory* 20.3, pp. 703–722 (cit. on pp. 84, 85, 93).
- Bright, Jonathan et al. (2020). “Does campaigning on social media make a difference? Evidence from candidate use of Twitter during the 2015 and 2017 UK elections”. In: *Communication Research* 47.7, pp. 988–1009 (cit. on p. 9).
- Buchanan, James M. and Gordon Tullock (1962). *The calculus of consent*. Vol. 3. Ann Arbor, MI: University of Michigan Press (cit. on p. 46).
- Cameron, A. Colin and Pravin K. Trivedi (2013). *Regression analysis of count data*. Vol. 53. Cambridge, UK: Cambridge University Press (cit. on pp. 14, 34).
- Caron, David D. (1993). “The legitimacy of the collective authority of the Security Council”. In: *The American Journal of International Law* 87.4, pp. 552–588 (cit. on p. 68).
- Carter, David B. and Randall W. Stone (2015). “Democracy and multilateralism: The case of vote buying in the UN General Assembly”. In: *International Organization* 69.1, pp. 1–33 (cit. on pp. 2, 3, 43, 46, 47, 62, 63, 114).
- Casella, Alessandra, Aniol Llorente-Saguer, and Thomas R. Palfrey (2012). “Competitive equilibrium in markets for votes”. In: *Journal of Political Economy* 120.4, pp. 593–658 (cit. on p. 46).
- Casella, Alessandra, Thomas Palfrey, and Sébastien Turban (2014). “Vote trading with and without party leaders”. In: *Journal of Public Economics* 112.C, pp. 115–128 (cit. on p. 46).
- Casella, Alessandra and Sébastien Turban (2014). “Democracy undone. Systematic minority advantage in competitive vote markets”. In: *Games and Economic Behavior* 88, pp. 47–70 (cit. on p. 46).

- Ceobanu, Alin M. and Xavier Escandell (2010). “Comparative analyses of public attitudes toward immigrants and immigration using multinational survey data: A review of theories and research”. In: *Annual Review of Sociology* 36, pp. 309–328 (cit. on p. 8).
- Cohen, Lauren and Christopher J. Malloy (2014). “Friends in high places”. In: *American Economic Journal: Economic Policy* 6.3, pp. 63–91 (cit. on pp. 46, 56).
- Cohu, Medhi, Christelle Maisonneuve, and Benoit Testé (2016). “The “Charlie-Hebdo” effect: Repercussions of the January 2015 terrorist attacks in France on prejudice toward immigrants and North-Africans, social dominance orientation, and attachment to the principle of laïcité”. In: *International Review of Social Psychology* 29.1, pp. 50–58 (cit. on p. 8).
- Coupe, Tom (2017). “The impact of terrorism on expectations, trust and happiness—the case of the November 13 attacks in Paris, France”. In: *Applied Economics Letters* 24.15, pp. 1084–1087 (cit. on p. 8).
- Deriu, Jan Milan et al. (2017). “Potential and limitations of cross-domain sentiment classification”. In: *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media*, pp. 17–24 (cit. on p. 14).
- Downs, Anthony (1972). “Up and down with ecology: The issue-attention cycle”. In: *The Public Interest* 28, pp. 38–50 (cit. on p. 15).
- Dreher, Axel and Nathan M. Jensen (2013). “Country or leader? Political change and UN General Assembly voting”. In: *European Journal of Political Economy* 29.C, pp. 183–196 (cit. on pp. 45, 63, 81, 115).
- Dreher, Axel, Peter Nunnenkamp, and Rainer Thiele (2008). “Does US aid buy UN General Assembly votes? A disaggregated analysis”. In: *Public Choice* 136.1, pp. 139–164 (cit. on pp. 3, 47, 114).
- Dreher, Axel, Jan-Egbert Sturm, and James Raymond Vreeland (2009a). “Development aid and international politics: Does membership on the UN Security Council influence World Bank decisions?” In: *Journal of Development Economics* 88.1, pp. 1–18 (cit. on p. 47).
- (2009b). “Global horse trading: IMF loans for votes in the United Nations Security Council”. In: *European Economic Review* 53.7, pp. 742–757 (cit. on p. 47).
- Einav, Liran and Jonathan Levin (2014). “Economics in the age of big data”. In: *Science* 346.6210 (cit. on p. 1).

- Eldar, Ofer (2008). “Vote-trading in international institutions”. In: *European Journal of International Law* 19.1, pp. 3–41 (cit. on pp. [2](#), [43](#), [46](#), [47](#), [68](#), [73](#)).
- Evans, Heather K., Victoria Cordova, and Savannah Sipole (2014). “Twitter style: An analysis of how house candidates used Twitter in their 2012 campaigns”. In: *PS: Political Science & Politics* 47.2, pp. 454–462 (cit. on pp. [10](#), [17](#)).
- Gadarian, Shana Kushner and Bethany Albertson (2014). “Anxiety, immigration, and the search for information”. In: *Political Psychology* 35.2, pp. 133–164 (cit. on pp. [6](#), [9](#), [38](#)).
- Galeotti, Andrea and Sanjeev Goyal (2010). “The law of the few”. In: *American Economic Review* 100.4, pp. 1468–1492 (cit. on p. [9](#)).
- Garlaschelli, Diego and Maria I. Loffredo (2004). “Patterns of link reciprocity in directed networks”. In: *Physical Review Letters* 93.26, p. 268701 (cit. on p. [48](#)).
- Gottfried, Jeffrey and Elisa Shearer (2016). *News use across social media platforms 2016*. Tech. rep. Pew Research Center (cit. on p. [9](#)).
- Graham, Todd et al. (2013). “Between broadcasting political messages and interacting with voters: The use of Twitter during the 2010 UK general election campaign”. In: *Information, Communication & Society* 16.5, pp. 692–716 (cit. on p. [10](#)).
- Grimmer, Justin and Brandon M. Stewart (2013). “Text as data: The promise and pitfalls of automatic content analysis methods for political texts”. In: *Political Analysis* 21.3, pp. 267–297 (cit. on p. [59](#)).
- Hainmueller, Jens and Daniel J. Hopkins (2014). “Public attitudes toward immigration”. In: *Annual Review of Political Science* 17.1, pp. 225–249 (cit. on pp. [8](#), [29](#)).
- Hanes, Emma and Stephen Machin (2014). “Hate crime in the wake of terror attacks: Evidence from 7/7 and 9/11”. In: *Journal of Contemporary Criminal Justice* 30.3, pp. 247–267 (cit. on p. [8](#)).
- Hanretty, Chris (2017). “Areal interpolation and the UK’s referendum on EU membership”. In: *Journal of Elections, Public Opinion and Parties* 27.4, pp. 466–483 (cit. on pp. [17](#), [31](#)).
- Hastie, Trevor, Robert Tibshirani, and Jerome Friedman (2009). *The elements of statistical learning: Data mining, inference, and prediction*. Springer Science & Business Media (cit. on pp. [63](#), [65](#)).
- Higgins, Rosalyn et al. (2017). *Oppenheim’s international law: United Nations*. Oxford, UK: Oxford University Press (cit. on p. [56](#)).

- Hilbe, Joseph M. (2011). *Negative binomial regression*. Cambridge, UK: Cambridge University Press (cit. on p. 14).
- Hu, Mei-Chen, Martina Pavlicova, and Edward V. Nunes (2011). “Zero-inflated and hurdle models of count data with extra zeros: Examples from an HIV-risk reduction intervention trial”. In: *The American Journal of Drug and Alcohol Abuse* 37.5, pp. 367–375 (cit. on p. 14).
- Huddy, Leonie et al. (2005). “Threat, anxiety, and support of antiterrorism policies”. In: *American Journal of Political Science* 49.3, pp. 593–608 (cit. on p. 8).
- Hughes, Sallie (2003). “From the inside out: How institutional entrepreneurs transformed Mexican journalism”. In: *Harvard International Journal of Press/Politics* 8.3, pp. 87–117 (cit. on p. 88).
- (2006). *Newsrooms in conflict: Journalism and the democratization of Mexico*. University of Pittsburgh Press (cit. on p. 88).
- Jackson, Nigel and Darren Lilleker (2011). “Microblogging, constituency service and impression management: UK MPs and the use of Twitter”. In: *The Journal of Legislative Studies* 17.1, pp. 86–105 (cit. on pp. 10, 17).
- James, Gareth et al., eds. (2013). *An introduction to statistical learning: With applications in R*. Springer Texts in Statistics 103. New York, NY: Springer (cit. on p. 64).
- Jensen, Jens Ledet, Peter B. Mortensen, and Søren Serritzlew (2016). “The dynamic model of choice for public policy reconsidered: A formal analysis with an application to US budget data”. In: *Journal of Public Administration Research and Theory* 26.2, pp. 226–238 (cit. on pp. 84, 85, 92).
- Johnson, Robert H. (1988). “Misguided morality: Ethics and the Reagan Doctrine”. In: *Political Science Quarterly* 103.3, pp. 509–529 (cit. on p. 50).
- Jones, Bryan D. (2017). “Behavioral rationality as a foundation for public policy studies”. In: *Cognitive Systems Research* 43, pp. 63–75 (cit. on p. 93).
- Jones, Bryan D. and Frank R. Baumgartner (2005). “A model of choice for public policy”. In: *Journal of Public Administration Research and Theory* 15.3, pp. 325–351 (cit. on pp. 3, 4, 83–85, 91–93, 111).
- (2012). “From there to here: Punctuated equilibrium to the general punctuation thesis to a theory of government information processing”. In: *Policy Studies Journal* 40.1, pp. 1–20 (cit. on p. 85).

- Jones, Bryan D., Frank R. Baumgartner, et al. (2009). “A general empirical law of public budgets: A comparative analysis”. In: *American Journal of Political Science* 53.4, pp. 855–873 (cit. on pp. [84](#), [93](#), [99](#)).
- Jones, Bryan D., Herschel F. Thomas III, and Michelle Wolfe (2014). “Policy bubbles”. In: *Policy Studies Journal* 42.1, pp. 146–171 (cit. on p. [86](#)).
- Jones, Edward E. and Thane S. Pittman (1982). “Toward a general theory of strategic self-presentation”. In: *Psychological Perspectives on the Self* 1.1, pp. 231–262 (cit. on p. [10](#)).
- Jungherr, Andreas (2016). “Twitter use in election campaigns: A systematic literature review”. In: *Journal of Information Technology & Politics* 13.1, pp. 72–91 (cit. on p. [9](#)).
- Kendall, Chad, Tommaso Nannicini, and Francesco Trebbi (2015). “How do voters respond to information? Evidence from a randomized campaign”. In: *American Economic Review* 105.1, pp. 322–353 (cit. on p. [9](#)).
- Kim, Soo Yeon and Bruce Russett (1996). “The new politics of voting alignments in the United Nations General Assembly”. In: *International Organization* 50.4, pp. 629–652 (cit. on pp. [46](#), [62](#), [63](#)).
- Kober, Thomas and David Weir (2015). “Optimising agile social media analysis”. In: *Proceedings of the 6th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pp. 31–40 (cit. on p. [12](#)).
- Kreiss, Daniel (2016). “Seizing the moment: The presidential campaigns’ use of Twitter during the 2012 electoral cycle”. In: *New Media & Society* 18.8, pp. 1473–1490 (cit. on p. [10](#)).
- Kuziemko, Ilyana and Eric Werker (2006). “How much is a seat on the Security Council worth? Foreign aid and bribery at the United Nations”. In: *Journal of Political Economy* 114.5, pp. 905–930 (cit. on p. [47](#)).
- Lawson, Chappell J. H. (2002). *Building the fourth estate: Democratization and the rise of a free press in Mexico*. University of California Press (cit. on p. [88](#)).
- Legewie, Joscha (2013). “Terrorist events and attitudes toward immigrants: A natural experiment”. In: *American Journal of Sociology* 118.5, pp. 1199–1245 (cit. on pp. [2](#), [5](#), [8](#), [16](#), [27](#)).
- Lindblom, Charles E. (1959). “The science of ”muddling through””. In: *Public Administration Review*, pp. 79–88 (cit. on p. [84](#)).

- Lockwood, Natalie J. (2013). “International vote buying”. In: *Harvard International Law Journal* 54, pp. 97–156 (cit. on pp. [43](#), [45](#), [47](#), [51](#)).
- Lucas, Michael and Doug Downey (Aug. 2013). “Scaling semi-supervised naive Bayes with feature marginals”. In: *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 343–351 (cit. on p. [12](#)).
- Malawer, Stuart S. (1988). “Reagan’s law and foreign policy, 1981-1987: The Reagan corollary of international law”. In: *Harvard International Law Journal* 1, pp. 85–110 (cit. on p. [50](#)).
- Malone, David M. (2000). “Eyes on the prize: The quest for nonpermanent seats on the UN Security Council”. In: *Global Governance* 6.1, pp. 3–23 (cit. on pp. [2](#), [43](#), [46](#), [75](#)).
- Márquez Ramírez, Mireya (2014a). “Post-authoritarian politics in a neoliberal era: Revising media and journalism transition in Mexico”. In: *Media systems and communication policies in Latin America*. Springer, pp. 272–292 (cit. on p. [100](#)).
- (2014b). “Professionalism and journalism ethics in post-authoritarian Mexico: Perceptions of news for cash, gifts and perks”. In: *The ethics of journalism: Individual, institutional and cultural influences*. I.B. Tauris, pp. 55–64 (cit. on p. [100](#)).
- Matter, Ulrich and Omar Guerrero (2017). *Uncovering vote trading through networks and computation*. Tech. rep. Oxford, UK: Saïd Business School, University of Oxford (cit. on pp. [44](#), [53](#), [66](#)).
- Mimno, David and Moontae Lee (2014). “Low-dimensional embeddings for interpretable anchor-based topic inference”. In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. Doha, Qatar: Association for Computational Linguistics, pp. 1319–1328 (cit. on p. [60](#)).
- Mimno, David, Hanna Wallach, et al. (2011). “Optimizing semantic coherence in topic models”. In: *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pp. 262–272 (cit. on p. [95](#)).
- Montalvo, José G. (2010). “Voting after the bombings: A natural experiment on the effect of terrorist attacks on democratic elections”. In: *The Review of Economics and Statistics* 93.4, pp. 1146–1154 (cit. on p. [8](#)).
- O’Neill, Barry (1996). “Power and satisfaction in the United Nations Security Council”. In: *Journal of Conflict Resolution* 40.2, pp. 219–237 (cit. on pp. [45](#), [68](#)).

- Oshiro, Thais Mayumi, Pedro Santoro Perez, and José Augusto Baranauskas (2012). “How many trees in a random forest?” In: *International Workshop on Machine Learning and Data Mining in Pattern Recognition*. Springer, pp. 154–168 (cit. on p. 65).
- Parmelee, John H (2014). “The agenda-building function of political tweets”. In: *New Media & Society* 16.3, pp. 434–450 (cit. on p. 10).
- Pauls, Scott D. and Skyler J. Cranmer (2017). “Affinity communities in United Nations voting: Implications for democracy, cooperation, and conflict”. In: *Physica A: Statistical Mechanics and its Applications* 484, pp. 428–439 (cit. on p. 63).
- Peysakhovich, Alexander and Jeffrey Naecker (2017). “Using methods from machine learning to evaluate behavioral models of choice under risk and ambiguity”. In: *Journal of Economic Behavior & Organization* 133, pp. 373–384 (cit. on p. 11).
- Philipson, Tomas J. and James M. Snyder (1996). “Equilibrium and efficiency in an organized vote market”. In: *Public Choice* 89.3, pp. 245–265 (cit. on p. 46).
- Roberts, Margaret E, Brandon M Stewart, and Dustin Tingley (2014). “stm: R Package for structural topic models”. In: *Journal of Statistical Software*, pp. 1–40 (cit. on p. 60).
- Roberts, Margaret E. et al. (2014). “Structural topic models for open-ended survey responses”. In: *American Journal of Political Science* 58.4, pp. 1064–1082 (cit. on pp. 3, 58, 114).
- Röder, Michael, Andreas Both, and Alexander Hinneburg (2015). “Exploring the space of topic coherence measures”. In: *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*, pp. 399–408 (cit. on p. 95).
- Roodman, David (2006). *An index of donor performance*. Tech. rep. Washington, DC: Center for Global Development (cit. on p. 62).
- Saif, Hassan et al. (2016). “Contextual semantics for sentiment analysis of Twitter”. In: *Information Processing & Management* 52.1, pp. 5–19 (cit. on p. 14).
- Santillán, José Ramón (2013). “Campañas de papel. La construcción de la democracia en México”. In: *Global Media Journal México* 5.9 (cit. on p. 89).
- (2016). “Prensa y democratización en México. Tres miradas: Excélsior, La Jornada y Reforma”. In: *Observatorio (OBS*)* 10.4, pp. 77–96 (cit. on p. 89).

- Schüller, Simone (2016). “The effects of 9/11 on attitudes toward immigration and the moderating role of education”. In: *Kyklos* 69.4, pp. 604–632 (cit. on p. 8).
- Signorino, Curtis and Jeffrey Ritter (1999). “Tau-b or not tau-b: Measuring the similarity of foreign policy positions”. In: *International Studies Quarterly* 43.1, pp. 115–144 (cit. on p. 45).
- Simon, Herbert A. (1996). *The sciences of the artificial*. 3rd. Cambridge, MA: The MIT Press (cit. on p. 85).
- Smith, Alastair (2016). “Leader turnover, institutions, and voting at the UN General Assembly”. In: *The Journal of Conflict Resolution* 60.1, pp. 143–163 (cit. on pp. 45, 46, 63).
- Stratmann, Thomas (1992). “The effects of logrolling on congressional voting”. In: *The American Economic Review* 82.5, pp. 1162–1176 (cit. on pp. 44, 46, 56, 57).
- Taddy, Matt (2012). “On estimation and selection for topic models”. In: *Artificial Intelligence and Statistics*, pp. 1184–1193 (cit. on p. 60).
- Teng, Zhiyang, Duy-Tin Vo, and Yue Zhang (2016). “Context-sensitive lexicon features for neural sentiment analysis”. In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 1629–1638 (cit. on p. 14).
- Tullock, Gordon (1959). “Problems of majority voting”. In: *Journal of Political Economy* 67.6, pp. 571–579 (cit. on p. 46).
- (1970). “A simple algebraic logrolling model”. In: *The American Economic Review* 60.3, pp. 419–426 (cit. on p. 46).
- Van Rijsbergen, C. J. (1979). *Information retrieval*. Publisher: Newton MA, USA. Newton MA, USA: Butterworth-Heinemann (cit. on p. 13).
- Varian, Hal R. (2014). “Big data: New tricks for econometrics”. In: *Journal of Economic Perspectives* 28.2, pp. 3–28 (cit. on p. 13).
- Varma, Sudhir and Richard Simon (2006). “Bias in error estimation when using cross-validation for model selection”. In: *BMC Bioinformatics* 7.1, p. 91 (cit. on p. 65).
- Vasilopoulos, Pavlos (2018). “Terrorist events, emotional reactions, and political participation: The 2015 Paris attacks”. In: *West European Politics* 41.1, pp. 102–127 (cit. on p. 8).
- Venables, W. N. and B. D. Ripley (2002). *Modern applied statistics with S*. 4th ed. Statistics and Computing. New York, NY: Springer-Verlag (cit. on p. 65).

- Vliegthart, Rens et al. (2016). “Do the media set the parliamentary agenda? A comparative study in seven countries”. In: *European Journal of Political Research* 55.2, pp. 283–301 (cit. on pp. 86, 87).
- Voeten, Erik (2000). “Clashes in the Assembly”. In: *International Organization* 54.2, pp. 185–215 (cit. on pp. 45, 46, 81).
- (2001). “Outside options and the logic of Security Council action”. In: *American Political Science Review* 95.4, pp. 845–858 (cit. on p. 45).
- (2004). “Resisting the lonely superpower: Responses of states in the United Nations to U.S. dominance”. In: *The Journal of Politics* 66.3, pp. 729–754 (cit. on pp. 46, 62, 63).
- Volden, Craig and Clifford J. Carrubba (2004). “The formation of oversized coalitions in parliamentary democracies”. In: *American Journal of Political Science* 48.3, pp. 521–537 (cit. on pp. 43, 56).
- Wager, Stefan and Susan Athey (2018). “Estimation and inference of heterogeneous treatment effects using random forests”. In: *Journal of the American Statistical Association* 113.523, pp. 1228–1242 (cit. on p. 11).
- Whitaker, Richard and Philip Lynch (2011). “Explaining support for the UK Independence Party at the 2009 European Parliament elections”. In: *Journal of Elections, Public Opinion & Parties* 21.3, pp. 359–379 (cit. on p. 8).
- Wibberley, Simon, David Weir, and Jeremy Reffin (Aug. 2014). “Method51 for mining insight from social media datasets”. In: *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: System Demonstrations*, pp. 115–119 (cit. on p. 15).
- Wikileaks (2009). *Report on African delegations at the 64th General Assembly*. Wikileaks Public Library of US Diplomacy 09USUNNEWYORK1131.a. New York, NY: United Nations, p. 48 (cit. on p. 51).
- Wildavsky, Aaron B. (1964). *Politics of the budgetary process*. Boston, MA: Little, Brown (cit. on p. 84).
- Williams, Matthew L. and Pete Burnap (2016). “Cyberhate on social media in the aftermath of Woolwich: A case study in computational criminology and big data”. In: *British Journal of Criminology* 56.2, pp. 211–238 (cit. on pp. 14, 16).

- Wolfe, Michelle (2012). “Putting on the brakes or pressing on the gas? Media attention and the speed of policymaking”. In: *Policy Studies Journal* 40.1, pp. 109–126 (cit. on p. 86).
- Wolfe, Michelle, Bryan D. Jones, and Frank R. Baumgartner (2013). “A failure to communicate: Agenda setting in media and policy studies”. In: *Political Communication* 30.2, pp. 175–192 (cit. on pp. 86, 101).
- Wooldridge, Jeffrey M. (2016). *Introductory econometrics: A modern approach*. Nelson Education (cit. on p. 38).
- Workman, Samuel, Bryan D. Jones, and Ashley E. Jochim (2009). “Information processing and policy dynamics”. In: *Policy Studies Journal* 37.1, pp. 75–92 (cit. on p. 85).

Appendix A

Appendix to Chapter 2

A.1 List of Words for Boolean Search

List of words used to conduct the Boolean search:

- Migrant
- Asylum Seeker
- Refugee
- Migration
- Influx
- Wave
- Not Native
- Deportation
- Border
- Foreigner
- Exodus
- Free Movement
- Confine
- Expatriate
- Displacement
- Non-native
- Flee
- Frontier

A.2 Examples of Tweets



Figure A.1: Examples of Irrelevant Tweets Picked Up with the Boolean Search



Figure A.2: Examples of False Positives for the *Relevant* CategoryFigure A.3: Examples of False Positives for the *Irrelevant* Category

Figure A.4: Examples of Tweets with Different Polarities

(a) Example of Tweet with Positive Stance


**Chuka Umunna** 
@ChukaUmunna

Follow

We should be helping integrate immigrants, not demonise them. You can read @IntegrationAPPG's alternative here:
[d3n8a8pro7vhmx.cloudfront.net/themes/570513f ...](https://d3n8a8pro7vhmx.cloudfront.net/themes/570513f...)

6:40 AM - 6 Sep 2017

16 Retweets 44 Likes



(b) Example of Tweet with Negative Stance

**Charlie Elphicke** 
@CharlieElphicke

Follow

Making the case on @VictoriaLIVE to take back control of or our borders and end uncontrolled EU immigration.



3:52 AM - 6 Sep 2017

2 Retweets 2 Likes



(d) Example of Tweet with Ambiguous Stance

(c) Example of Tweet with Neutral Stance

**Tom Newton Dunn** 
@tnewtondunn

Segui

Net migration down to 248,000 in 2016 - a drop of 84,000 from 2015. Largely to less EU8 (Eastern Europeans) coming and more going home.

01:37 - 25 mag 2017

21 Retweet 22 Mi piace



**Sarah Champion** 
@SarahChampionMP

Segui

Immigration target is 'an ambition', say Conservatives. My ambition's 2b a Power Ranger, that's probably more likely



Tories: Migration target is 'an ambition'
Tories deny confusion over immigration after the PM signals she wants it down to 'tens of thousands'.
bbc.co.uk

02:25 - 2 glu 2017

71 Retweet 121 Mi piace



A.3 Additional Descriptive Statistics

Here I present additional descriptive statistics on the MPs included in the analyses and those never considered. Table A.1 shows the distribution by political affiliation. The large majority (74.5%) of the MPs excluded in the study belongs to the Conservative Party. Table A.2 presents the comparison of demographic characteristics between the two groups. It appears that the politicians in my sample are younger than the excluded ones, and women are more represented. Finally, table A.3 shows the reasons for the exclusion of some MPs from the analyses. Most of them were not considered as they did not have a Twitter account at the time of the collection (76.5%). Four politicians had a protected account, whereas one MP was using the Commons Leader account. Three accounts were never considered because the limit on the collection from their timeline (i.e., 3,200 Tweets) was reached before the day that represents the upper temporal bound in my analyses (the 29th of May 2017).

Table A.1: MP Distribution by Political Party

Party	MPs Included		MPs Excluded	
	Freq.	Percent	Freq.	Percent
<i>Conservative</i>	257	46.56	73	74.49
<i>Democratic Unionist Party</i>	7	1.27	1	1.02
<i>Green</i>	1	0.18	-	-
<i>Independent</i>	3	0.54	2	2.04
<i>Labour</i>	184	33.34	19	19.39
<i>Labour Co-operative</i>	26	4.71	1	1.02
<i>Liberal Democrats</i>	9	1.63	-	-
<i>Plaid Cymru</i>	3	0.54	-	-
<i>Scottish National Party</i>	53	9.60	1	1.02
<i>Sinn Féin</i>	4	0.73	-	-
<i>Social Democratic and Labour Party</i>	3	0.54	-	-
<i>Ulster Unionist Party</i>	2	0.36	-	-
<i>Speaker</i>	-	-	1	1.02
<i>Total</i>	552	100.00	98	100.00

Table A.2: Comparison Demographic Characteristics

	MPs Included		MPs Excluded	
	Mean	Std. Dev.	Mean	Std. Dev.
<i>Male</i>	0.67	0.47	0.85	0.36
<i>Age</i>	51.28	10.17	58.08	12.12

Table A.3: Reasons for Exclusion

Reason	Freq.	Percent
<i>Too Active</i>	3	3.06
<i>No Personal Account</i>	1	1.02
<i>Protected</i>	4	4.08
<i>No Account</i>	75	76.53
<i>Not Active</i>	15	15.31
<i>Total</i>	98	100.00

Appendix B

Appendix to Chapter 3

B.1 Web Pages in the UN Digital Library

Figure B.1: Examples of Web Pages in the United Nations Digital Library (1)

The screenshot displays a web page for a United Nations resolution. At the top, the title 'Agricultural technology for sustainable development : resolution / adopted by the General Assembly' is shown, along with the adopting body 'UN General Assembly (72nd sess. : 2017-2018)' and the year '2018'. Below this, there are three main action buttons: 'Download', 'Formats', and 'Add to List'. A 'Files' section follows, containing a table with columns for Action, Filename, Size, Access, Description, and License. The table lists five PDF files in Arabic, Other, English, Spanish, and French. Below the files, a 'Details' section provides comprehensive information about the resolution, including its symbol (A/RES/72/215), title, access links in multiple languages, action note date (2017-12-20), draft status (A/C.2/72/L.33/Rev.1), committee report (A/72/420), meeting record (A/72/PV.74), authors (UN General Assembly), agenda information (A/72/251), imprint (New York, 19 Jan. 2018), description (7 p.), notes (Issued in GAOR, 72nd sess., Suppl. no. 49), and collections (Resource Type > Documents and Publications > Resolutions and Decisions; UN Bodies > General Assembly > General Assembly Plenary). At the bottom, a 'Browse Subjects' section lists related topics like '2030 Agenda for Sustainable Development', 'AGRICULTURAL DEVELOPMENT', 'AGRICULTURAL ENGINEERING', 'SUSTAINABLE AGRICULTURE', 'AGRICULTURAL INNOVATIONS', and 'AGRICULTURAL POLICY'. A 'PDF Reader' button is located at the very bottom.

Action	Filename	Size	Access	Description	License
Download	A_RES_72_215-ar.pdf	405.5 kB	Public	العربية	-
Download	A_RES_72_215-ot.pdf	297.5 kB	Public	Other	-
Download	A_RES_72_215-en.pdf	392.9 kB	Public	English	-
Download	A_RES_72_215-es.pdf	390.1 kB	Public	Español	-
Download	A_RES_72_215-fr.pdf	365.2 kB	Public	Français	-


Details	
Symbol	A/RES/72/215
Title	Agricultural technology for sustainable development : resolution / adopted by the General Assembly
Access	English: A_RES_72_215-EN - PDF ; Español: A_RES_72_215-ES - PDF ; Français: A_RES_72_215-FR - PDF ; Other: A_RES_72_215-OT - PDF ; Русский: A_RES_72_215-RU - PDF ; العربية: A_RES_72_215-AR - PDF ; 中文: A_RES_72_215-ZH - PDF ;
Action note	2017-12-20
Draft	A/C.2/72/L.33/Rev.1
Committee report	A/72/420
Meeting record	A/72/PV.74
Authors	UN General Assembly (72nd sess. : 2017-2018)
Agenda information	A/72/251 19 Sustainable development. SUSTAINABLE DEVELOPMENT
Imprint	[New York] : UN, 19 Jan. 2018
Description	7 p.
Notes	Issued in GAOR, 72nd sess., Suppl. no. 49
Collections	Resource Type > Documents and Publications > Resolutions and Decisions UN Bodies > General Assembly > General Assembly Plenary

Browse Subjects	
2030 Agenda for Sustainable Development	AGRICULTURAL DEVELOPMENT
AGRICULTURAL ENGINEERING	SUSTAINABLE AGRICULTURE
AGRICULTURAL INNOVATIONS	AGRICULTURAL POLICY

PDF Reader


Note: the web page belongs to the United Nations Digital Library. However, we scraped the information from the previous web portal (UNBISnet), which contained the same type of information for UNGA resolutions.


Figure B.2: Examples of Web Pages in the United Nations Digital Library (2)



Agricultural technology for sustainable development : resolution / adopted by the General Assembly

2017

 Formats

 Add to List

Details

Title Agricultural technology for sustainable development : resolution / adopted by the General Assembly

Agenda A/72/251 19 Sustainable development: SUSTAINABLE DEVELOPMENT

Resolution A/RES/72/215

Meeting record A/72/PV.74

Draft resolution A/C.2/72/L.33/Rev.1

Committee report A/72/420

Note RECORDED - No machine generated vote

Vote summary Voting Summary
Yes: 152 | No: 1 | Abstentions: 29 | Non-Voting: 11 | Total voting membership: 193


Vote date 2017-12-20

Vote
A AFGHANISTAN
Y ALBANIA
A ALGERIA
Y ANDORRA
Y ANGOLA
Y ANTIGUA AND BARBUDA
Y ARGENTINA
Y ARMENIA
Y AUSTRALIA
Y AUSTRIA
Y AZERBAIJAN
Y BAHAMAS
A BAHRAIN
A BANGLADESH
Y BARBADOS
Y BELARUS
Y BELGIUM
Y BELIZE
Y BENIN
Y BHUTAN
A BOLIVIA (PLURINATIONAL STATE OF)
Y BOSNIA AND HERZEGOVINA
Y BOTSWANA
Y BRAZIL
A BRUNEI DARUSSALAM
Y BULGARIA
Y BURKINA FASO
Y BURUNDI
Y CABO VERDE
Y CAMBODIA
Y CAMEROON
Y CANADA
Y CENTRAL AFRICAN REPUBLIC
CHAD
Y CHILE

Note: the web page belongs to the United Nations Digital Library. However, we scraped the information from the previous web portal (UNBISnet), which contained the same type of information for UNGA resolutions.

B.2 Example of UNGA Resolution

Figure B.3: Sample of a UNGA's Resolution Text

	<p>United Nations</p> <p>General Assembly</p>	<p>A/RES/72/215</p> <p>Distr.: General 19 January 2018</p>
---	--	--

Seventy-second session
Agenda item 19

**Resolution adopted by the General Assembly
on 20 December 2017**

[on the report of the Second Committee (A/72/420)]

72/215. Agricultural technology for sustainable development

The General Assembly,


Recalling its resolution [70/198](#) of 22 December 2015,


Reaffirming its resolution [70/1](#) of 25 September 2015, entitled “Transforming our world: the 2030 Agenda for Sustainable Development”, in which it adopted a comprehensive, far-reaching and people-centred set of universal and transformative Sustainable Development Goals and targets, its commitment to working tirelessly for the full implementation of the Agenda by 2030, its recognition that eradicating poverty in all its forms and dimensions, including extreme poverty, is the greatest global challenge and an indispensable requirement for sustainable development, its commitment to achieving sustainable development in its three dimensions – economic, social and environmental – in a balanced and integrated manner, and to building upon the achievements of the Millennium Development Goals and seeking to address their unfinished business,


Reaffirming also its resolution [69/313](#) of 27 July 2015 on the Addis Ababa Action Agenda of the Third International Conference on Financing for Development, which is an integral part of the 2030 Agenda for Sustainable Development, supports and complements it, helps to contextualize its means of implementation targets with concrete policies and actions, and reaffirms the strong political commitment to address the challenge of financing and creating an enabling environment at all levels for sustainable development in the spirit of global partnership and solidarity,

Welcoming the Zero Hunger Challenge initiative launched by the Secretary-General at the United Nations Conference on Sustainable Development as a vision for a future free from hunger, and recalling the Rome Declaration on Nutrition, adopted at the Second International Conference on Nutrition,¹ the Decade of Action on Nutrition (2016–2025)² and the recommendations contained in the report on

17-23284 (E) 240118



Please recycle 



¹ World Health Organization, document EB 136/8, annex I.

² See resolution [70/259](#).

B.3 Kolmogorov-Smirnov Tests

Table B.1: Kolmogorov-Smirnov Tests
for Degree Distributions

Session	K-S Test (Indegree)	<i>p</i> -value	K-S Test (Outdegree)	<i>p</i> -value
<i>A/31</i>	0.188	0.010	0.025	1.000
<i>A/32</i>	0.288	0.000	0.026	1.000
<i>A/33</i>	0.118	0.230	0.023	1.000
<i>A/34</i>	0.085	0.605	0.015	1.000
<i>A/35</i>	0.133	0.123	0.018	1.000
<i>A/36</i>	0.140	0.085	0.029	1.000
<i>A/37</i>	0.148	0.058	0.249	0.000
<i>A/38</i>	0.166	0.024	0.024	1.000
<i>A/39</i>	0.065	0.866	0.014	1.000
<i>A/40</i>	0.095	0.442	0.015	1.000
<i>A/41</i>	0.063	0.886	0.053	0.965
<i>A/42</i>	0.061	0.905	0.015	1.000
<i>A/43</i>	0.040	0.998	0.022	1.000
<i>A/44</i>	0.097	0.410	0.036	1.000
<i>A/45</i>	0.040	0.999	0.010	1.000
<i>A/46</i>	0.050	0.970	0.015	1.000
<i>A/47</i>	0.042	0.995	0.019	1.000
<i>A/48</i>	0.069	0.744	0.015	1.000
<i>A/49</i>	0.031	1.000	0.010	1.000
<i>A/50</i>	0.054	0.934	0.014	1.000
<i>A/51</i>	0.090	0.414	0.037	0.999
<i>A/52</i>	0.072	0.690	0.044	0.989
<i>A/53</i>	0.065	0.805	0.037	0.999
<i>A/54</i>	0.067	0.754	0.056	0.907
<i>A/55</i>	0.046	0.980	0.058	0.888
<i>A/56</i>	0.077	0.593	0.061	0.840
<i>A/57</i>	0.042	0.991	0.131	0.069
<i>A/58</i>	0.069	0.723	0.047	0.975
<i>A/59</i>	0.038	0.998	0.049	0.967
<i>A/60</i>	0.039	0.997	0.079	0.554
<i>A/61</i>	0.053	0.929	0.051	0.951
<i>A/62</i>	0.025	1.000	0.028	1.000
<i>A/63</i>	0.043	0.989	0.051	0.950
<i>A/64</i>	0.044	0.988	0.047	0.975
<i>A/65</i>	0.076	0.597	0.043	0.990
<i>A/66</i>	0.038	0.998	0.098	0.287
<i>A/67</i>	0.062	0.820	0.040	0.995
<i>A/68</i>	0.033	1.000	0.054	0.923
<i>A/69</i>	0.064	0.801	0.042	0.991
<i>A/70</i>	0.060	0.853	0.055	0.914
<i>A/71</i>	0.037	0.998	0.090	0.386
<i>A/72</i>	0.057	0.890	0.081	0.517

Note: the Kolmogorov-Smirnov tests compare the degree distributions of the DDNs with and without UNSC permanent members. Values in red are statistics significant at the 5% level.

B.4 Countries with Highest Degree by Session

Table B.2: Countries with Highest Outdegree and Indegree by Session (31st to 50th)

Session	Country	Outdegree	Country	Indegree
A/31	GERMANY, FEDERAL REPUBLIC OF	231	NIGERIA	27
A/31	ISRAEL	202	GHANA	22
A/31	UNITED KINGDOM	116	ALGERIA	22
A/32	GERMANY, FEDERAL REPUBLIC OF	464	INDIA	49
A/32	UNITED STATES	357	YUGOSLAVIA	46
A/32	UNITED KINGDOM	279	NIGERIA	40
A/33	CANADA	249	ALGERIA	91
A/33	JAPAN	240	GUINEA	88
A/33	NETHERLANDS	225	AFGHANISTAN	87
A/34	ITALY	174	INDIA	48
A/34	AUSTRALIA	147	PAKISTAN	45
A/34	GERMANY, FEDERAL REPUBLIC OF	142	MALAYSIA	44
A/35	BELGIUM	125	PAKISTAN	38
A/35	ISRAEL	120	GUINEA	34
A/35	LUXEMBOURG	102	TUNISIA	32
A/36	GERMANY, FEDERAL REPUBLIC OF	141	TUNISIA	45
A/36	CANADA	129	GUINEA	39
A/36	ISRAEL	100	MADAGASCAR	37
A/37	GERMANY, FEDERAL REPUBLIC OF	250	YUGOSLAVIA	101
A/37	ISRAEL	220	ALGERIA	68
A/37	DOMINICA	149	MALI	58
A/38	FRANCE	224	GUYANA	50
A/38	ITALY	162	YUGOSLAVIA	48
A/38	LUXEMBOURG	135	PAKISTAN	48
A/39	DOMINICA	267	TUNISIA	61
A/39	GERMANY, FEDERAL REPUBLIC OF	224	ALGERIA	57
A/39	CANADA	205	MAURITANIA	51
A/40	ISRAEL	164	ALGERIA	31
A/40	GERMANY, FEDERAL REPUBLIC OF	160	YUGOSLAVIA	30
A/40	EQUATORIAL GUINEA	103	CUBA	30
A/41	DOMINICA	255	SUDAN	41
A/41	ISRAEL	169	ALGERIA	39
A/41	SAINT KITTS AND NEVIS	138	INDIA	39
A/42	GERMANY, FEDERAL REPUBLIC OF	336	VIET NAM	66
A/42	PARAGUAY	259	LIBYA	60
A/42	EQUATORIAL GUINEA	164	SUDAN	59
A/43	SAINT KITTS AND NEVIS	383	LIBYA	65
A/43	CANADA	234	SUDAN	63
A/43	GERMANY, FEDERAL REPUBLIC OF	196	CUBA	58
A/44	DOMINICA	477	PAKISTAN	61
A/44	SAINT KITTS AND NEVIS	268	CUBA	59
A/44	GERMANY, FEDERAL REPUBLIC OF	183	MADAGASCAR	50
A/45	DOMINICA	177	CUBA	44
A/45	AUSTRALIA	74	UKRAINIAN SSR	35
A/45	IRELAND	65	AFGHANISTAN	34
A/46	DOMINICA	235	CUBA	64
A/46	SAINT KITTS AND NEVIS	225	AFGHANISTAN	59
A/46	MARSHALL ISLANDS	149	VIET NAM	59
A/47	SAINT KITTS AND NEVIS	244	CUBA	125
A/47	DOMINICA	213	MOROCCO	70
A/47	ISRAEL	145	MAURITANIA	70
A/48	DOMINICA	228	CUBA	64
A/48	DEMOCRATIC PEOPLE'S REPUBLIC OF KOREA	102	YEMEN	60
A/48	MICRONESIA (FEDERATED STATES OF)	75	SENEGAL	52
A/49	DOMINICA	237	CUBA	58
A/49	ISRAEL	204	MALAYSIA	38
A/49	SAINT KITTS AND NEVIS	131	SENEGAL	37
A/50	SOUTH AFRICA	98	CUBA	83
A/50	BOSNIA AND HERZEGOVINA	87	JORDAN	69
A/50	VANUATU	82	DJIBOUTI	69


Table B.3: Countries with Highest Outdegree and Indegree by Session (51st to 72nd)

Session	Country	Outdegree	Country	Indegree
A/51	DOMINICA	220	MALAYSIA	71
A/51	BOSNIA AND HERZEGOVINA	145	TUNISIA	71
A/51	VANUATU	128	INDONESIA	70
A/52	LIBERIA	167	MALTA	70
A/52	VANUATU	138	GREECE	60
A/52	ISRAEL	135	EGYPT	53
A/53	DEMOCRATIC REPUBLIC OF THE CONGO	362	EGYPT	148
A/53	SAINT KITTS AND NEVIS	299	FIJI	129
A/53	VANUATU	299	BANGLADESH	126
A/54	EQUATORIAL GUINEA	296	CUBA	107
A/54	MARSHALL ISLANDS	271	EGYPT	95
A/54	DOMINICA	251	NETHERLANDS	84
A/55	VANUATU	346	ALGERIA	126
A/55	NAURU	327	BANGLADESH	104
A/55	EQUATORIAL GUINEA	223	EGYPT	94
A/56	SEYCHELLES	290	BANGLADESH	87
A/56	NAURU	157	GERMANY	86
A/56	EQUATORIAL GUINEA	142	SPAIN	86
A/57	SAO TOME AND PRINCIPE	725	BANGLADESH	150
A/57	SEYCHELLES	343	SOUTH AFRICA	133
A/57	EQUATORIAL GUINEA	312	JORDAN	129
A/58	SEYCHELLES	452	BANGLADESH	116
A/58	NAURU	373	MOROCCO	113
A/58	ISRAEL	292	MALAYSIA	100
A/59	SEYCHELLES	591	QATAR	189
A/59	NAURU	559	MALAYSIA	170
A/59	EQUATORIAL GUINEA	439	ALGERIA	167
A/60	PALAU	298	BANGLADESH	132
A/60	DEMOCRATIC REPUBLIC OF THE CONGO	293	JORDAN	108
A/60	ISRAEL	238	MALTA	106
A/61	CHAD	495	CUBA	149
A/61	PALAU	393	NAMIBIA	147
A/61	ISRAEL	379	SOUTH AFRICA	137
A/62	EQUATORIAL GUINEA	1075	INDONESIA	107
A/62	DEMOCRATIC REPUBLIC OF THE CONGO	418	CUBA	104
A/62	PALAU	369	VENEZUELA (BOLIVARIAN REPUBLIC OF)	102
A/63	PALAU	671	SWEDEN	160
A/63	EQUATORIAL GUINEA	497	CUBA	157
A/63	ISRAEL	364	BRAZIL	131
A/64	DEMOCRATIC REPUBLIC OF THE CONGO	390	VENEZUELA (BOLIVARIAN REPUBLIC OF)	111
A/64	EQUATORIAL GUINEA	295	COMOROS	100
A/64	PALAU	277	MALAYSIA	98
A/65	SAO TOME AND PRINCIPE	323	MALTA	103
A/65	ISRAEL	314	AUSTRIA	91
A/65	TUVALU	268	IRELAND	88
A/66	SAO TOME AND PRINCIPE	443	EGYPT	125
A/66	LIBERIA	333	CUBA	113
A/66	PALAU	276	BANGLADESH	103
A/67	SOUTH SUDAN	897	ECUADOR	118
A/67	ISRAEL	519	SOUTH AFRICA	115
A/67	VANUATU	261	CUBA	111
A/68	KIRIBATI	779	CUBA	98
A/68	SAO TOME AND PRINCIPE	587	EGYPT	98
A/68	NAURU	533	INDONESIA	97
A/69	KIRIBATI	805	ECUADOR	121
A/69	SAO TOME AND PRINCIPE	651	YEMEN	119
A/69	EQUATORIAL GUINEA	433	EGYPT	118
A/70	KIRIBATI	807	ECUADOR	79
A/70	DOMINICA	299	SOUTH AFRICA	78
A/70	ISRAEL	273	CUBA	71
A/71	KIRIBATI	919	ECUADOR	175
A/71	ISRAEL	653	SWEDEN	173
A/71	NAURU	604	EGYPT	156
A/72	KIRIBATI	757	GEORGIA	162
A/72	PALAU	702	ROMANIA	128
A/72	NAURU	459	POLAND	112


B.5 Data Coverage over Time

Figure B.4: Coverage of Population and Income Data over Time
Sessions 31st to 72nd (1976-2018)

(a) Coverage of Population Data



(b) Coverage of Income Data



Note: in the two panels the orange bars show the number of current members of the UNGA, whereas the blue ones show the number of members for which we have information in each dataset.

B.6 Bootstrapping Procedure

We can assess the statistical significance of the index by computing confidence intervals (CIs). We adopt a bootstrapping procedure that involves resampling with replacement the set of member countries, that are the rows of matrices \mathbb{V} , \mathbb{S} and \mathbb{Q} . At each iteration the new matrices are used to compute a different *Reciprocity Index* q_i^* :

$$\left(\frac{r_1 - \bar{r}_{01}}{1 - \bar{r}_{01}}, \frac{r_2 - \bar{r}_{02}}{1 - \bar{r}_{02}}, \dots, \frac{r_n - \bar{r}_{0n}}{1 - \bar{r}_{0n}} \right) = (q_1^*, q_2^*, \dots, q_n^*) \quad (\text{B.1})$$

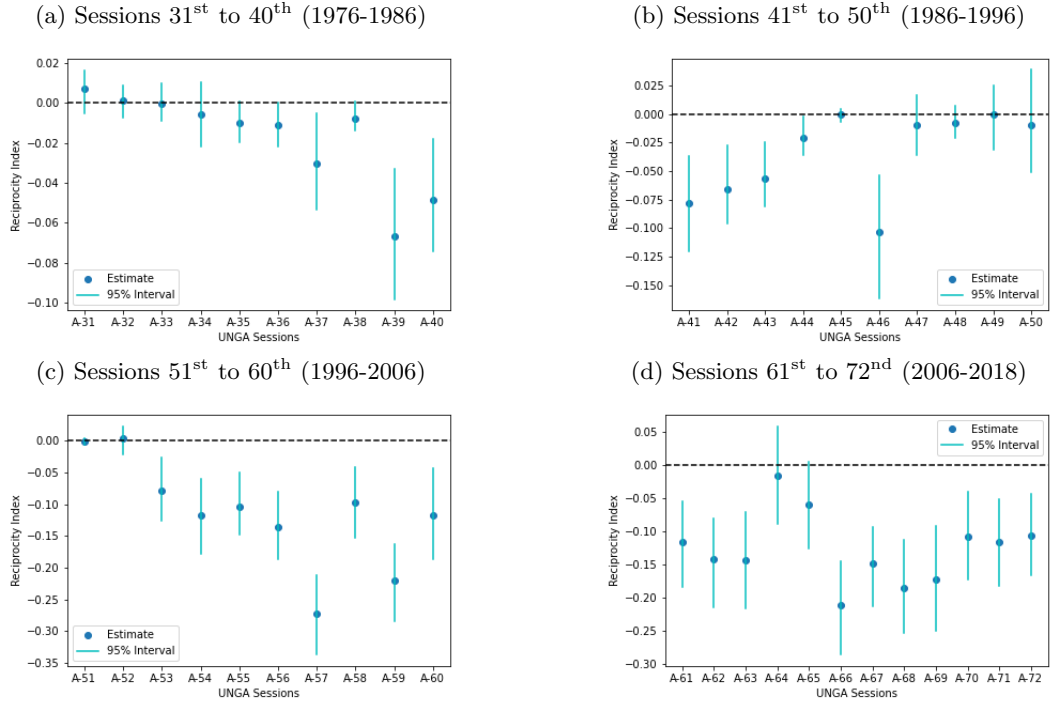
Where n stands for the number of bootstrap samples. To build the confidence intervals, we first obtain the distribution of Δ_i^* , which is defined as the difference between each bootstrap index and the one computed using the observed data ($\Delta_i^* = q_i^* - q$). The empirical distribution of Δ_i^* is then sorted in ascending order to construct the 95% bootstrap CIs: $[q - \Delta_{0.025}^*, q - \Delta_{0.975}^*]$.¹ The procedure is meant to capture the (theoretical) uncertainty arising from alternative representations in the institutional setting considered.

¹The value $\Delta_{0.025}^*$ represents the element at the 97.5th percentile and $\Delta_{0.975}^*$ the element at the 2.5th percentile.

B.7 Robustness Checks: *Reciprocity Index*

A first component of the methodology that we might want to test is the consistency of our results with the use of a different classifier. The choice of the predictive model is a key component in the specification of matrix \mathbb{Q} , which subsequently informs the topology of the DDN. Alternative predictions will determine a distinct pattern of deviations that might lead to values of the index that substantially diverge from our baseline results. To address this issue, we re-implement the methodology using the neural network proposed in section 3.5, which was the best predictive model in terms of F-score and accuracy. The analysis is still conducted at the session level and the new set of indices and CIs are reported in figure B.5.

Figure B.5: Evolution of the *Reciprocity Index* over Time
Alternative Classifier



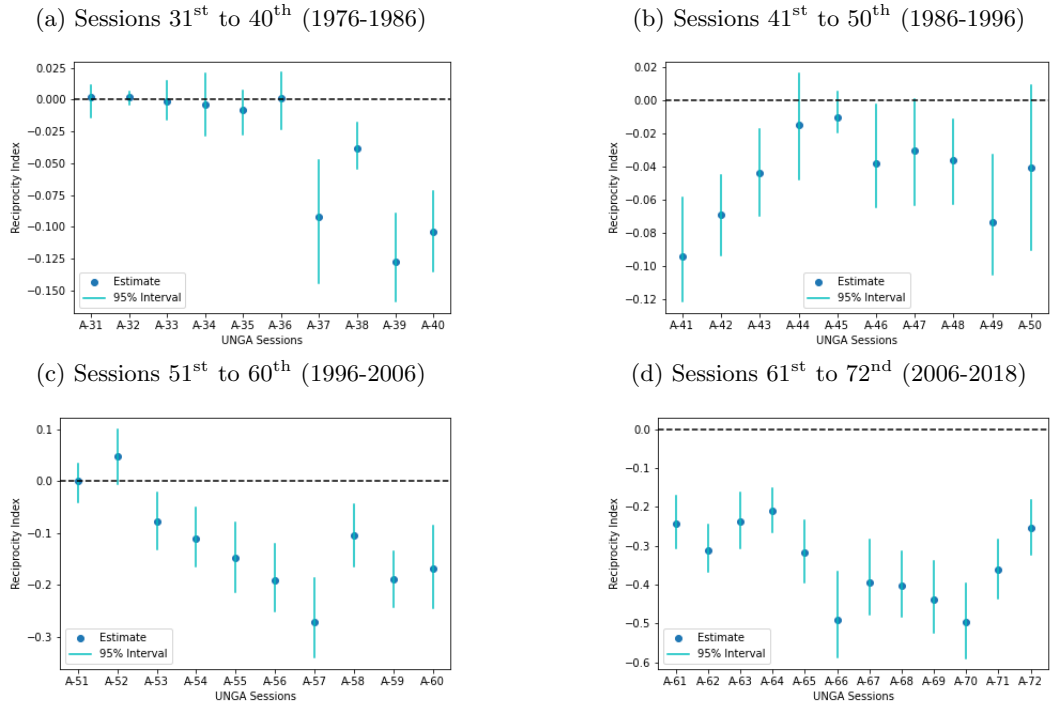
Note: point estimates and 95% confidence intervals are reported. The CIs are constructed using 500 bootstrap samples. The dashed line indicates a value of 0 for the index. The classifier used to predict voting behaviour is the neural network proposed in section 3.5. The units of analysis are individual sessions.

The pattern that we observe is not qualitatively different from the one we get using our preferred classifier (figure 3.5). The index is not statistically different from 0 in the first period of the analysis (until 1984), then it starts exhibiting a generally negative trend with a break between the 45th and 52nd session. We notice however that the most recent sessions, even if still characterised by a systematic negative *Reciprocity Index* (the only exceptions being the 64th and 65th sessions), display values that

are on average lower in their magnitude compared to the baseline results. If we focus on the years after 2010 (65th session), the index ranges from -0.10 in the 72nd session to -0.21 in the 66th.

A second robustness check that we perform involves the choice of the probabilities' threshold μ mentioned in section 3.4. This cutoff defines our deviations, the entries of matrix \mathbb{D} and therefore the structure of the DDN. Due to the lack of theoretical guidance on justifying the choice of a specific threshold in this institutional context, we decide to assess the stability of the baseline results by using the value of μ that maximises the accuracy of our preferred classifier (i.e., the random forest presented in section 3.5 and trained with the full dataset). More specifically, we set $\mu = 0.559$ which results in an accuracy of 0.893. The indices and CIs obtained using this alternative cutoff are reported in figure B.6.

Figure B.6: Evolution of the *Reciprocity Index* over Time
Alternative Cutoff (μ)



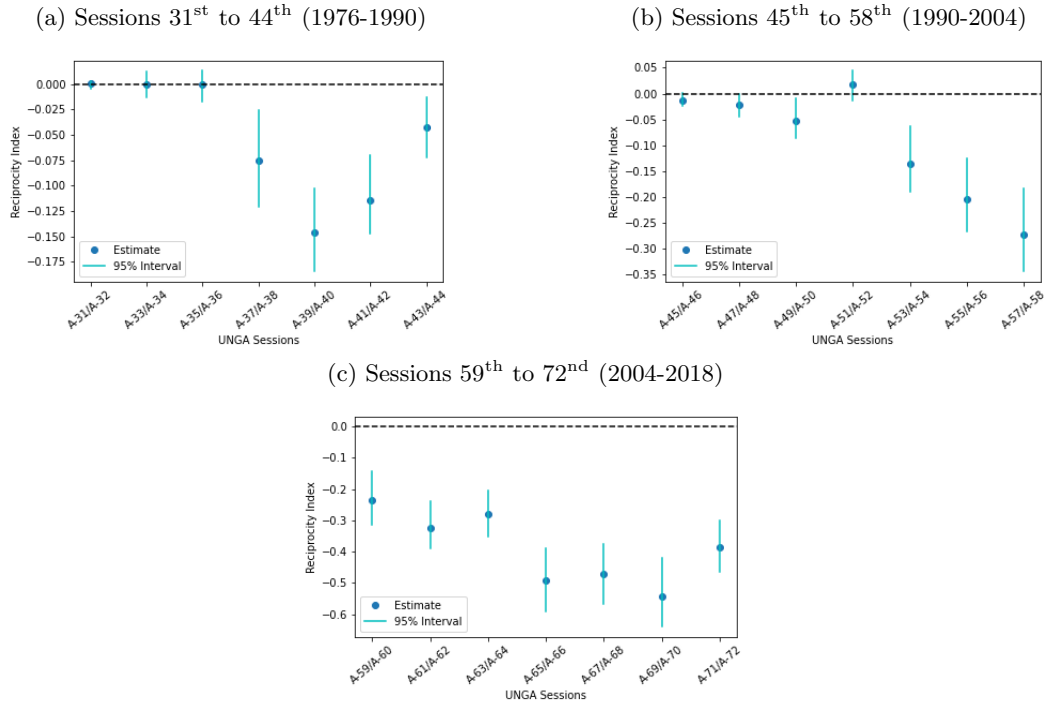
Note: point estimates and 95% confidence intervals are reported. The CIs are constructed using 500 bootstrap samples. The dashed line indicates a value of 0 for the index. The classifier used to predict voting behaviour is the random forest proposed in section 3.5. The cutoff for the probabilities is set to $\mu = 0.559$. The units of analysis are individual sessions.

The trend we observe now is remarkably similar to our baseline results: the *Reciprocity Index* is not statistically significant until the 37th session, followed by systematic negative values, especially from 1998 (53rd session). As in figure 3.5, we notice that in the most recent years of the analysis the Assembly seems to be characterised by indices that are even more negative, reaching a minimum of

-0.50 in the 70th session.

Next, we implement the methodology using different time frames. By specifying a given unit of analysis (i.e., a temporal window) we might lose exchanges of votes that occur between two different periods. We thus decide to recompute the set of indices pooling consecutive meetings of the Assembly. First, we combine the sessions in blocks of two (figure B.7), then in groups of three (figure B.8). Thus, we are now allowing for deviations being reciprocated in subsequent sessions.

Figure B.7: Evolution of the *Reciprocity Index* over Time
Alternative Time Frame (2 Sessions)

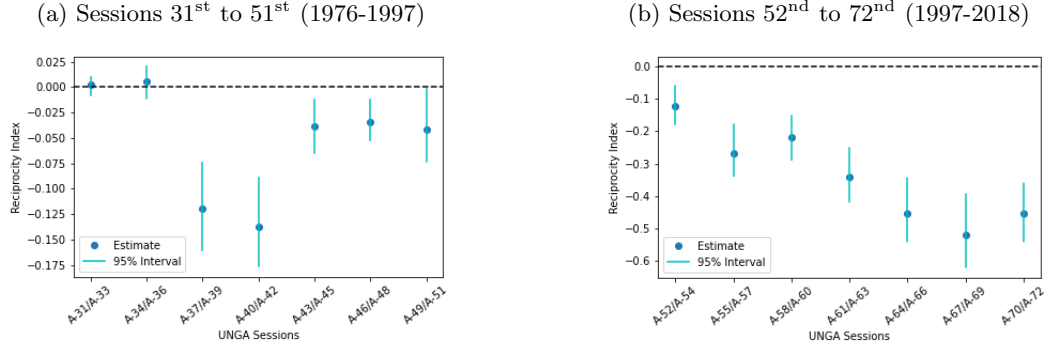


Note: point estimates and 95% confidence intervals are reported. The CIs are constructed using 500 bootstrap samples. The dashed line indicates a value of 0 for the index. The classifier used to predict voting behaviour is the random forest proposed in section 3.5. The units of analysis are groups of two subsequent sessions.

Changing the unit of analysis does not substantially alter the main insights that we get from our baseline results. In both cases we observe a similar behaviour of the index: in the first years considered it is not statistically different from 0, but then it tends to take negative values, which increase in magnitude in the most recent sessions. It is interesting to notice that, when we move from single sessions to longer time frames, the *Reciprocity Index* reaches minimum values that are larger (in absolute terms), peaking at -0.54 when we pool two consecutive meetings of the Assembly and -0.52 when we analyse groups of three.

To further test the robustness of our baseline results, we reimplement the methodology using a “narrow


Figure B.8: Evolution of the *Reciprocity Index* over Time
Alternative Time Frame (3 Sessions)



Note: point estimates and 95% confidence intervals are reported. The CIs are constructed using 500 bootstrap samples. The dashed line indicates a value of 0 for the index. The classifier used to predict voting behaviour is the random forest proposed in section 3.5. The units of analysis are groups of three subsequent sessions.

margin” approach that is consistent with the assumption outlined in section 3.4. That is, vote trading is more likely to occur in narrow roll calls. We proceed in the following way. We first select the resolutions for which we observe a percentage of *yes* votes (over the total members) that is less than 60%. We then look at the sessions during which the documents are voted on and retrieve all other resolutions passed in these meetings (i.e., those passed by a large margin). We also exclude the sessions for which we have information on less than 10 narrow roll calls. Next, we use the sample of resolutions passed by large margin as training dataset for our preferred classifier (i.e., the random forest presented in section 3.5). We employ the model to obtain the matrix of predictions \mathbb{Q} for the narrow roll calls, which are then used in the computation of the *Reciprocity Index* at the session level. The underlying rationale is to train the model with the votes that are assumed to be expression of the “true preferences” of the countries. Subsequently, we use the classifier to make predictions for the roll calls in which we would expect vote trading (and so our guesses to be less accurate), in order to identify deviations that more clearly suggest instances of trades. Figure B.9 presents the results of this exercise. We can see that the major insights provided by our analysis in section 3.6 are confirmed. The indices computed for the set of sessions considered are generally negative and 10 out of 15 are statistically significant at the 5% level. However, their magnitude is substantially lower (in absolute terms) compared to the baseline results.

Figure B.9: Evolution of the *Reciprocity Index* over Time
Narrow Margin Approach



Note: point estimates and 95% confidence intervals are reported. The CIs are constructed using 500 bootstrap samples. The dashed line indicates a value of 0 for the index. The classifier used to predict voting behaviour is the random forest proposed in section 3.5. The units of analysis are individual sessions. The index is computed using only narrow roll calls, which are defined as those resolutions passed with a percentage of *yes* votes (over the total members) that is less than 60%. Sessions with less than 10 narrow roll calls are excluded from the analysis.

Appendix C

Appendix to Chapter 4

C.1 Branches: Descriptive Statistics

Table C.1: Descriptive Statistics for All Branches (I)

Branch	Number of Programmes	Absolute Difference (Level)	Absolute Difference (Proportion)
<i>Office of the Presidency of the Republic</i>	6	3.3714 (5.2259)	0.0496 (0.0782)
<i>Government</i>	32	4.5178 (10.9795)	0.0578 (0.1320)
<i>External Relations</i>	8	9.0218 (13.1380)	0.1354 (0.2044)
<i>Finance and Public Credit</i>	34	7.4832 (11.3176)	0.1119 (0.1728)
<i>National Defense</i>	21	6.6448 (9.1195)	0.1174 (0.1724)
<i>Agriculture and Rural Development</i>	22	5.5444 (10.8947)	0.1100 (0.2007)
<i>Communications and Transports</i>	35	9.3655 (23.7782)	0.1459 (0.3640)
<i>Economy</i>	23	0.5610 (0.6227)	0.0113 (0.0138)
<i>Public Education</i>	39	12.9781 (41.1075)	0.2067 (0.6398)
<i>Health</i>	28	8.1178 (25.5067)	0.1060 (0.2773)
<i>Marine</i>	8	9.6050 (11.0651)	0.1166 (0.1237)
<i>Work and Social Security</i>	11	0.6774 (0.9527)	0.0106 (0.0158)
<i>Agrarian, Territorial and Urban Development</i>	18	5.3639 (10.4154)	0.0872 (0.1700)
<i>Environment and Natural Resources</i>	41	3.7828 (8.2770)	0.0604 (0.1357)
<i>Attorney General of the Republic</i>	13	0.7836 (1.8117)	0.0236 (0.0522)
<i>Energy</i>	11	3.7154 (6.9005)	0.0561 (0.1068)

Note: the table shows the average absolute difference between expenditure approved and paid, both in level of spending and as proportion of the total budget approved/paid. Standard deviations are reported in parentheses. Values are reported in millions of Mexican pesos (levels) and as per thousand of the budget (proportions).

Table C.2: Descriptive Statistics for All Branches (II)

Branch	Number of Programmes	Absolute Difference (Level)	Absolute Difference (Proportion)
<i>Social Security Contributions</i>	24	6.4299 (11.9451)	0.3306 (0.6682)
<i>Welfare</i>	25	4.1710 (11.9327)	0.1060 (0.2960)
<i>Tourism</i>	13	4.3526 (11.0467)	0.0690 (0.1699)
<i>Salary and Economic Provisions</i>	56	60.0014 (200.3587)	0.9372 (3.1094)
<i>Provisions and Contributions for Education Systems</i>	5	10.9884 (12.6734)	0.2485 (0.2889)
<i>Public Function</i>	7	0.5745 (0.9557)	0.0086 (0.0137)
<i>Agrarian Courts</i>	4	0.1313 (0.1130)	0.0015 (0.0009)
<i>Federal Court of Administrative Justice</i>	1	0.5845 (0.0000)	0.0332 (0.0000)
<i>Federal Contributions for States and Municipalities</i>	12	10.7798 (23.6049)	0.3738 (0.5189)
<i>National Commission for Human Rights</i>	22	0.0854 (0.1097)	0.0018 (0.0029)
<i>Legal Counsel of the Federal Executive</i>	3	0.0292 (0.0360)	0.0002 (0.0001)
<i>National Council for Science and Technology</i>	11	2.2088 (1.9383)	0.0369 (0.0299)
<i>Energy Regulatory Commission</i>	4	1.1383 (0.8468)	0.0169 (0.0125)
<i>National Hydrocarbons Commission</i>	5	1.1856 (0.8945)	0.0179 (0.0135)
<i>Non-sectorized Entities</i>	13	1.0670 (1.5682)	0.0199 (0.0252)
<i>Culture</i>	15	0.9688 (1.4844)	0.0122 (0.0192)
<i>Mexican Social Security Institute</i>	18	14.2025 (23.4404)	0.4398 (0.6748)
<i>Institute of Social Security and Services for State Workers</i>	22	16.3820 (28.2467)	0.3295 (0.7575)
<i>Federal Electricity Commission</i>	4	127.7971 (171.7930)	1.5213 (2.2814)

Note: the table shows the average absolute difference between expenditure approved and paid, both in level of spending and as proportion of the total budget approved/paid. Standard deviations are reported in parentheses. Values are reported in millions of Mexican pesos (levels) and as per thousand of the budget (proportions).

C.2 Textual Data on Public Expenditures

Table C.3: Textual Data on the Budget Programmes (I)

Feature	Description
<i>Responsible Unit</i>	Description of the responsible unit assigned to the branch that coordinates the recording of information on the performance of the budget programme.
<i>Functional Group</i>	Description of the first level (or digit) of the Functional Classification of Expenditure, the budget's classification that groups expenses according to the purposes or socio-economic objectives pursued by the different public entities. It allows identifying the activities carried out by the State to fulfill its purposes of social development, economic development and governance.
<i>Function</i>	Description of the second level (or digit) of the Functional Classification of Expenditure. It allows identifying the actions carried out by the responsible units to comply with the legal regulations, in accordance with each one of the functional groups.
<i>Subfunction</i>	Description of the third level (or digit) of the Functional Classification of Expenditure. It identifies more precisely the activities carried out by agencies and entities within a function.
<i>Institutional Activity</i>	Description of the substantive or support actions carried out by the administrative structure in order to comply with the objectives and goals contained in the budget programmes, in accordance with the attributions stated in their respective organic law or the legal system that is applicable to them.
<i>Branch</i>	Description of the branch according to the programmatic structure of the Budget of Expenditures of the Federation in force for each budget cycle.
<i>Modality</i>	Description of the classification that allows identifying the budget programmes according to the type of services/products they provide or their specific nature.
<i>Budget Programme</i>	Name of the federal budget programme according to the current programmatic structure for each budget cycle.
<i>National Development Plan (NDP)</i>	Description of the relevant national (or transversal) goal established in the National Development Plan (NDP) 2013-2018.

Table C.4: Textual Data on the Budget Programmes (II)

Feature	Description
<i>NDP Objective</i>	Description of the objective that seeks to be achieved through the relevant national (or transversal) goal.
<i>NDP Programme</i>	Description of the programme derived from the 2013-2018 NDP, which specifies the objectives, priorities and policies that will govern the performance of the administrative sector in question.
<i>NDP Programme Objective</i>	Description of the objective to be achieved through the different sector programmes.
<i>Strategic Objective</i>	Description of the objectives (of the agency or entity) which are intended to be achieved.
<i>Objective</i>	Description of the objectives of the Matrix of Indicators for Results that the budget programme intends to achieve. These are identified at the following levels: <i>goal</i> (higher objectives to which the budget programme seeks to contribute), <i>purpose</i> (objectives to be achieved with the budget programme), <i>component</i> (goods or services that the budget programme intends to generate) and <i>activity</i> (actions carried out through the budget programme).
<i>Bases</i>	Description of the conditions external to the responsible unit that must be considered to achieve the objectives of the budget programme.
<i>Indicator</i>	Description of the indicator for each level of the Matrix of Indicators for Results. This allows to measure the achievement of the objectives of the programmes, in addition to being a reference for the monitoring of the progress and evaluation of the results achieved by the budget programme.
<i>Indicator Definition</i>	Description of what is to be measured of the objective associated to the indicator. It helps to understand its utility, purpose or use.

C.3 Additional Robustness Checks: News Index

Table C.5: Serial Information Processing with Alternative Topic Model

	OLS (1) $\ln(\Delta Level)$	OLS (2) $\ln(\Delta Prop)$	OLS (3) $\ln(\Delta Level)$	OLS (4) $\ln(\Delta Prop)$
<i>Intercept</i>	31.405*** (10.144)	-0.359 (4.279)	27.385*** (9.769)	0.212 (3.835)
$\ln(NewsIndex)$	10.815* (5.818)	6.664*** (2.452)	9.761* (5.636)	5.761*** (2.158)
$\ln(Exp^{approved})$			0.148*** (0.043)	0.213*** (0.041)
<i>N</i>	614	614	614	614
<i>Branch FE</i>	Yes	Yes	Yes	Yes
<i>Adj R²</i>	0.114	0.123	0.151	0.247

Note: the underlying topic model has $K = 7$. The dependent variable is the natural logarithm of the absolute difference between the expenditure paid and the expenditure approved. In models 1 and 3 we use levels of expenditure to compute the policy response. In models 2 and 4 we use expenditure proportions. Expenditure in levels is reported in millions of Mexican pesos. Expenditure proportions are reported in per thousand of the budget. Heteroscedasticity-consistent standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

C.4 News Index: Descriptive Statistics

Table C.6: Descriptive Statistics for the *News Index*

<i>News Index</i>	Mean	Std. Dev.	Min.	Max.
<i>Full Sample</i>	0.137	0.015	0.093	0.178
<i>Restricted Newspapers Corpus</i>	0.190	0.021	0.119	0.246
<i>Non-negative Changes (Levels)</i>	0.138	0.015	0.093	0.178
<i>Non-negative Changes (Prop.)</i>	0.138	0.015	0.093	0.175
<i>Only Reforma</i>	0.478	0.063	0.273	0.638
<i>Only Reforma Non-negative Changes (Levels)</i>	0.477	0.065	0.273	0.634
<i>Only Reforma Non-negative Changes (Prop.)</i>	0.473	0.068	0.273	0.634
<i>Alternative LDA</i>	0.164	0.015	0.116	0.207

Note: the table reports the mean, standard deviation, minimum and maximum for the *News Index* across the different sets of regressions that we run. The first row refers to the *News Index* employed in the regressions of table 4.4. The second row to the one used in table 4.5. The third row refers to columns 1, 2, 5 and 6 in table 4.6. The fourth row to columns 3, 4, 7 and 8 in table 4.6. The fifth row refers to table 4.7. The sixth row refers to columns 1, 2, 5 and 6 in table 4.8. The seventh row refers to columns 3, 4, 7 and 8 in table 4.8. The eighth row refers to table C.5.

C.5 Additional Robustness Checks: Natural Experiment

Table C.7: Natural Experiment with Alternative Topic Model

	Model (1) <i>Level</i>	Parallel Trends (1) <i>Level</i>
<i>Intercept</i>	18.914 (13.134)	11.143 (8.221)
<i>Treat</i>	-13.393 (13.563)	-7.979 (10.359)
<i>Quarter</i>	-34.902** (16.966)	4.363 (10.225)
<i>Treat</i> \times <i>Quarter</i>	35.161* (19.385)	-2.486 (13.544)
<i>Exp</i> ^{past}	0.225* (0.120)	-0.114*** (0.032)
<i>Exp</i> ^{approved}	0.128* (0.065)	0.259*** (0.004)
<i>N</i>	328	328
<i>Adj R</i> ²	0.953	0.981

Note: all models are estimated via OLS. The underlying topic model has $K = 6$. The dependent variable is paid expenditure in millions of Mexican pesos. Budget programmes are selected if their earthquake-topic proportion is greater than the average. Heteroscedasticity-consistent standard errors are reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

(a) Earthquake



(b) Elections



(c) Social Policy



Note: in the word clouds the size of the terms is proportional to their weight in the topic distribution.