Speech as Data for the Politics and Partisanship of Legislators

Matthew Blayney

St Antony's College

University of Oxford

Thesis submitted for the Degree of Doctor of Philosophy

Trinity 2025

45,969 words.

Abstract

How can we empirically prove complex theories about legislative behaviour? Parliamentary speeches provide large quantities of data about legislators' behaviour, daily. Until recent developments in natural language processing (NLP), however, researchers could not make full use of the richness of this data. Introducing new methods, I automatically code millions of statements made in parliaments into data about legislators' partisan loyalties and political positions. I use these methods to yield new evidence about old questions in the study of legislatures, which so far have had long-standing theoretical debate and mixed empirical evidence. These include the effects of electoral system (change), of constituency-level opinion and electoral threat, and of legislators' ethnicity and gender. Firstly, I use NLP methods to identify the similarity of legislators' speeches to their party's relevant manifesto chapters based on the topic they are discussing. This similarity is used to measure strength of partisanship, and determine the effect of electoral system change in New Zealand on party discipline. The results show greater partisanship after introducing closed party lists. Secondly, I use a large language model (LLM) as an 'expert coder' to label over 2 million speeches in the UK parliament for their left-right political leaning on two axes. I compare these positions to estimates of constituency opinion and voting intention polls. I show that MPs' positions are not significantly correlated with those of their constituents, but where UKIP, the Brexit Party, or Reform UK gained support in their constituency, MPs took more right-wing positions, especially on non-economic issues. Finally, I add to this dataset LLM-coded topics identifying relevance to ethnic minorities' and women's interests, and measure positions on these topics separately. I show that ethnic minority and female MPs' speeches are generally more left-wing (especially on ethnic and gender topics), except for ethnic minority MPs in the Conservative Party.

Contents

Acknowledgements

I can't offer thanks high enough for the mentorship and guidance Stephen Fisher has provided throughout this DPhil course. Beyond his excellent advice throughout drafting and re-drafting these papers, he has provided incredible opportunities to work on projects, and insights into the academic world that I would be lost without.

I'm also grateful to Geoffrey Evans and the Nuffield Politics Research Centre, who I have worked for on research outside of the DPhil throughout this course. This has enriched my development as a researcher, and inspired a series of projects which I hope we can continue.

I likely would not have applied for postgraduate study if it wasn't for the support of Giovanni Capoccia, who supervised my undergraduate thesis, which provided the basis for the first paper of this, my doctoral thesis.

This research, my tuition, and the past four years of my life in Oxford were funded by the ESRC. I'm grateful to them for the vital funding they provide, without which most people, including myself, could never afford to undertake postgraduate study. By making the pathway to academia accessible to all, all of academia benefits.

My family has been vital support throughout the past three years of study. Ben's insights into data science and machine learning have been a sounding board for my ideas throughout. My parents' enthusiasm for the work has seen a house transformed into a control room monitoring select committees and House of Commons debates around the clock: an addiction I don't think they'll thank me for, but I'm not sure they'll be abandoning any time soon, either.

Finally, Sorcha has been there for me every single step of the way. Her kindness, patience, and positivity have managed to bring endless joy, even whilst studying a subject as pessimistic as the behaviour of British MPs.

Chapter 1: Introduction

Legislatures exist to scrutinise, and sometimes reject, legislation. They can do this only insofar as they can freely criticise and vote against proposals of the government. And they can do this in the public interest only insofar as the basis of their criticism is representative of the views or interests of the people they represent. Crucial to understanding how (well) representative democracy is functioning, therefore, is understanding how often, and why, legislators diverge from their party leadership. Given these questions are so important, they have engaged much theoretical debate and qualitative analyses. Only with recent developments in the computational analysis of natural language, however, have robust empirical analyses become possible. Rebellious voting data was the main source for empirical analyses until now, but rebellious votes are too rare and provide too little information, especially to test a theory which is asking *why* legislators behave a certain way. I solve this problem by applying new methods to analyse speech data: a rich data source of millions of sentences, containing daily information about MPs' political positions. These methods include 'classic' n-gram models to identify speeches that echo partisan manifesto chapters. They also include an application of large language models (LLM) which is new to social science: improving their accuracy in coding tasks by leveraging 'chain-of-thought' methods. Through these new methods, the contribution of this thesis is a set of new findings on legislative behaviour. Firstly, that introducing an electoral system with closed lists increases the cohesiveness and loyalty of legislative parties. Secondly, that MPs in the UK are uniquely responsive to UKIP, Brexit Party, and Reform UK's (polled) vote share in their constituency. Thirdly, that although ethnic minority or female MPs are more left-wing than white or male MPs respectively, this does not hold true for ethnic minority MPs of the Conservative Party.

A core premise of all chapters is that an MP's speeches can be used to measure their effective positions or loyalties. A speech in parliament plays four important functions: signalling, legitimisation, expression, and construction.

1. It signals group identities and belongings: using political or partisan language and framing debates from the perspectives of particular factions or social groups.
2. It justifies and explains their other behaviours: it can legitimise disloyal behaviour to their party whips or constituents by explaining their behaviour as constituency service, or more true to the core party ideology than the current party leadership, or based on their own personal convictions and political beliefs. It can also legitimise following the whip to their constituents or fellow MPs.
3. It expresses preferences: it references aspects or implications of a bill which specifically appeal to (or offend) the legislator, their constituents, or their faction, and in doing so express what policy priorities and beliefs a legislator holds. It also expresses approval for aspects of the party platform.
4. It constructs a personal platform: through the above three processes, speeches define a legislator's own priorities and preferences for the reference of others. As a politician advances their career, decisions and speeches in their past affect their prospects in the future: whether a new party leadership looking for good ideological fits for their reshuffle, or the potential for journalists to dig up and question their past conduct. Legislators, even if dishonest, want their record in parliament to be reflective of a platform that is useful to their political aims, consistent, and defensible.

Whilst legislators can be dishonest, these functions require legislators to choose their words such that they justify their behaviour in parliament, in a way which advances their legislative career in a direction which is desirable to them. That means expressing loyalties and political positions which appeal to constituents who elect them (Carey &

Shugart, 1995), partisans who select them, leaders who can promote them (Rush & Giddings, 2011), or other legislators who can work with them (Bowler, Farrell, & Katz, 1999). Whether pursuing 'electoral benefits' or 'intralegislative benefits' (Cox & McCubbins, 1994), or their own convictions: the political positions they espouse in parliament reflect their interests, and the policy goals they (currently) pursue to further those interests. Legislators may also situate themselves strategically for the sake of power over policy, to strengthen their role as an individual veto player[1] (Tsebelis, 2002): in which case such positions would have to be expressed to be used for bargaining, and would have to be frequently followed through with their voting behaviour when bargaining fails[2]. Similarly, where bargaining fails or the whip is strong, MPs may use speeches as a way to vocalising dissent without risking exit by rebelling or harming their own party (Hirschman, 1970). There are therefore many strategic (and sincere) reasons why MPs espouse political positions that accurately reflect the policies they will support. By analysing speeches, you therefore can capture legislators' group identities, sources of legitimacy, and policy preferences.

What is studied, therefore, is a legislator's behaviour as a representative, in line with Goffman's (1956) approach of treating social interaction itself as the focus of sociological study. This is unlikely to be a perfect reflection of their personal beliefs, but that is also not the target of these analyses. The intention is to analyse how they perform as representatives: regardless of their personal beliefs, the degree to which they align with their constituents, conform to their party line, adopt positions strategically to win elections, or indeed follow their own beliefs. It is therefore not necessary to measure their

[1] Or, as part of a faction, strengthen their faction's role as a collective veto player.
[2] Making this behaviour potentially more identifiable through speech data than voting data.

inner thoughts, or what positions they personally support: only to measure the beliefs they represent when they are acting as representatives.

I therefore talk throughout the thesis of 'political positions' as the positions a legislator takes through their behaviour, and likewise of 'partisan loyalty' as the loyalty a legislator expresses through their behaviour. A legislator's political positions may therefore be in contradiction with their personal beliefs, and their behaviour may be loyal to their party without reflecting a deeply held personal loyalty.

Central to all chapters of the thesis is the concept of competing principals (Strøm, 2003; Carey, 2007): the division of an MP's loyalties between their party, their constituents, and their own interests. The core research question of this thesis is therefore how MPs navigate their principals: each of which has both a normative and a strategic reason they might be expected to guide an MP's behaviour.

Each chapter of the thesis analyses MPs' responsiveness to a different principal. In Chapter 2 the dependent variable is strength of partisanship, measuring the degree to which an MP is beholden to their party's interests: a party-centred type of behaviour. Unexplained variance in this chapter is therefore likely to be a mix of constituency and personal interests, both of which may be pressuring them to be more rebellious or independent in their voting and speech patterns. In Chapter 3, the dependent variable is the political position of their speeches, and the independent variable represents the political positions of their constituents. This tests the degree to which an MP is beholden to their constituents' interests: a delegate-type behaviour. Unexamined variance in this chapter is therefore likely to be a mix of partisanship and personal interests. Chapter 4 analyses gender and ethnicity. Depending on interpretation, this may either be regarded as part of a legislator's individual interests, or as an additional principal: group interests of co-ethnic or female voters across the country, whose own MP may not share their

ethnicity or gender. This may therefore be interpreted as a trustee-type behaviour, or a kind of delegate-type behaviour.

The main contribution I'm offering is a new toolset for measuring the expressed political positions and loyalties of legislators. Legislative behaviour is already an expansive literature, and exists as a sub-literature within many others. For quantitative analysis, however, the only data sources have been voting records and speeches. Rebellious voting is too rare, and the amount of information each instance provides is too small, for confident conclusions to be drawn even about simple hypotheses, let alone deep and nuanced theoretical frameworks. Speeches, however, are rich data, and only relatively recently have more advanced computational methods been adopted in social science. I outline the shortcomings of existing methods, and demonstrate how mine will improve on them.

Chapter 2 uses more traditional natural language processing methods, and compares the outcomes of these methods to an analysis of voting records. Findings are comparable, demonstrating the robustness of the method and the conclusions, as well as enabling a more nuanced interpretation of the results, given that the two methods study different behaviours and have slightly different results. By using manifestos as training data, this first chapter is able to analyse partisan loyalty specifically, as the method is specialised to identify the deployment of buzzwords, slogans, and flagship policies that are central to the party's platform.

In chapters 3 and 4, I introduce the more novel method of using LLMs for the role of a human coder, which is necessary in order to capture the more subjective and general concepts of 'left-' and 'right-wing' positioning. Models reliant on training data would be too sensitive to over-specification (capturing highly specific instances of partisanship based on terminology), lacking adequate appreciation of context (capturing generally left-

or right-wing terminology, but not sensitive enough to context-specific vocabulary like the names of policies), and not having strong enough appreciation of complex grammar and semantics (like identifying when a sentence is in the negative or not, and therefore whether right-wing terminology is being deployed to criticise right-wing positions or to support them). When chain-of-thought methods are used, LLMs are surprisingly capable of recognising context-specific vocabulary, identifying meaning, and even recognising subtext and subtle implications from a text.

Thesis Outline

The thesis is comprised of three substantive chapters (chapters 2, 3, and 4), each written to be published as individual papers.

Whilst all chapters use speeches as data, and analyse legislative behaviour, there are variations in the methods and the specific aspects of behaviour being studied in each. They therefore differ in how they respond to existing literature, and the types of methodological problems they solve using these new methods, as well as variations in the methods themselves. Each substantive chapter is therefore best understood in the context of its own question, and its own methodological response to that question. The first uses clustering methods to analyse partisan loyalty. The second uses LLMs with chain-of-thought methods to analyse political positions in two dimensions over time in response to constituency-level changes. The third uses LLMs with chain-of-thought methods to analyse political positions on specific policy interests cross-sectionally in comparison with legislators' ethnicity and gender.

Chapter 2

Chapter 2 focuses on electoral system change, using the case of New Zealand. It's the only example of a stable post-war democracy using a first-past-the-post electoral system, and switching to closed-list proportionality. It's therefore an informative case for the

question of electoral reform in first-past-the-post countries. Closed-list system have been found in theoretical terms and in cross-national empirical studies to be linked to greater party discipline (Carey & Shugart, 1995). Comparisons within the same country, however, provide a more robust way to study these differences, by avoiding the many unobserved and unobservable differences between countries' political cultures and systems. Highly cohesive partisan voting trends in New Zealand have made it a difficult case study for existing methods, but new methods of analysing speeches provide a new way to leverage this case.

This chapter uses older machine learning methods, as measuring the strength of partisanship is more compatible with the underlying assumptions of n-gram models. Partisan influence[3] is expected to be evidenced through the use of partisan slogans and 'buzzwords', name-dropping of flagship policies, and adoption of a 'partisan lexicon' of common ways party members talk about certain concepts and ideas, alongside the possibility MPs deliver speeches scripted and provided by the central party that they want to hear delivered. All of these would have a direct effect on word frequencies, in a way that would be easily picked up by n-gram models.

The findings of Chapter 2 are that New Zealand's electoral system change from 'first-past-the-post' (FPTP) to closed-list mixed-member proportionality (CL-MMP) caused a large increase in partisanship. This effect was observable for both 'Electorate MPs' and 'list MPs' in the new system – with even MPs who represent a constituency being more loyal to the party line in the CL-MMP system. I confirm the findings of the natural language

[3] The subject of this chapter is the power of central party leaderships to constrain the behaviour of MPs. Partisanship, partisan loyalty, and partisan influence, are all therefore used in this chapter to refer to the degree to which an MP's behaviour aligns with the interests of the incumbent party leadership. It does not refer to a longer-term concept of partisanship which may include commitment to the party's core/founding doctrines regardless of the interests of the current leadership. It also does not refer simply to which party an MP is a member of or loyal to.

processing analyses by also examining conscience votes, as a site where voting is nominally free, but partisan pressures can still exist. Using a breakpoint analysis, a statistically significant 'break' is found at the date of the electoral system change, where MPs begin becoming more partisan over time.

Chapter 3

Chapter 3 focuses on MPs' responsiveness to the local political context in their constituencies. Do MPs adopt positions that track their constituencies' preferences, as the theoretical principles of representative democracy, or strategic theories about delegation and accountability might suggest? If so, they ought to present arguments in parliament that align with the perspectives of their constituents, or which respond to strategic demands of electoral threats. By analysing positions expressed in MPs' speeches, and comparing to constituency-level opinion and voting intention data, I find new answers to these questions.

This chapter introduces novel methods using LLMs. I use the open source LLM Llama to play the role of an 'expert coder' in order to code the political leanings of every speech in parliament from the 2015 UK election until the 2024 UK election. This is necessary because deriving political leaning is far more complex than can be captured with older methods. It requires interpretation of subtext, and deriving unspoken logical conclusions and implications from a speaker's statements. This is something LLMs are demonstrably capable of when applied appropriately. I enhance the LLM's ability to draw logical conclusions from the texts by applying a chain-of-thought methodology. Many applications of LLMs as coders in social science specify strictly limited output formats, in the form that would be requested of a human coder, such as Ziems et al. (2024), and papers evaluated by Barrie, Palaiologou, and Törnberg (2025) and by Barrie, Palmer, and Spirling (2025). Given how LLM outputs are generated, however, their performance improves when they output longer texts, and provide answers in natural language format

(Nye, et al., 2021; Wei, et al., 2022; Kojima, et al., 2022). I leverage the strengths of LLMs by asking for a fully articulated explanation of the political leanings of each speech, and then request a strictly formatted summary of that explanation's conclusions. The estimated positions of MPs over time are then compared to measurements of constituency political leanings (which are derived by applying MRP methods to British Election Study public opinion data), and to recent party popularity data according to a mix of commercial MRPs and election data (whichever is most recent).

I find that MPs of both parties move to the right when UKIP, the Brexit Party, or Reform UK increase in popularity in their constituency, but that there is no such response to increasing popularity of the Liberal Democrats. I also find that public opinion within an MP's constituency has minimal, but identifiable and measurable, effects on MPs' expressed political positions in parliament. I also find partisan differences, with different levels of responsiveness and different patterns of response in each party.

Chapter 4

Chapter 4 addresses questions about descriptive and substantive representation in parliament, and their implications for legislative behaviour. These suggest that marginalised groups' interests are unlikely to be represented in electoral democracy, which in principle seeks to serve majority interests, and in practice is run by a personnel skewed towards majority and privileged groups (Hughes, 2013). Marginalised groups' interests are also most likely to be represented in the legislature (substantive representation) when members of those groups are, themselves, in the legislature (descriptive representation) (Pitkin, 1967).

Given the vital role played by female and ethnic minority MPs in protecting and advancing civil rights and protections, especially in the face of underrepresentation, an additional principal makes demands of their behaviour. An ethnic minority MP, for example, may feel obligations to represent the interests of their local constituency, their

party, their own interests, *and* the interests of co-ethnic voters across the country. This can be seen through the role of the Congressional Black Caucus in the USA (Rivers, 2012), and of various All-Party Parliamentary Groups in the UK which organise marginalised groups' interests (UK Parliament, 2025). To what degree, therefore, does this obligation (or interest) of ethnic minority and female MPs affect their behaviour? Furthermore, recent theories suggest that parties may play a role in 'gatekeeping' entry of ethnic minorities and women into elite politics, selecting them for views unrepresentative of the groups they belong to, or simply to use their loyalty as a 'reputational shield' (Dancygier 2017). To what extent do these theories about marginalised group candidacies have effects on the behaviour of MPs once elected? And to what extent does it subvert the expectations that descriptive representation leads to substantive representation?

To answer these questions, I expand the dataset created for Chapter 3, creating two new axes of political positioning: one representing positions on issues relevant to ethnic minority interests, the other positions on issues relevant to women's interests. These axes are created by again using an LLM as a coder, this time taking debate titles and introductory speeches to outline the broad themes of a debate, and then to code it as 'yes' or 'no' for whether it is relevant to each group's interests. The same is done for speeches in case those groups' interests are raised by an MP on a debate that is otherwise not directly related. The average of MPs' left-right positions when speaking on these topics represents their left-right positions on ethnic and gender issues.

Using this dataset, I test for the basic expectations of descriptive representation leading to substantive representation, and whether results differ at the party- or individual-level, to identify which (if any) of the mechanisms subverting descriptive representation can be observed. I find that the expectations of descriptive representation hold true for female MPs in both parties, but only for ethnic minority MPs in the Labour Party. Whereas ethnic minority and female MPs are generally more left-wing, especially on issues related to

their group interests, ethnic minority Conservative MPs are more right wing generally, and indistinguishable on ethnic issues.

Chapter 2: Electoral Systems and Legislative Cohesion: Comparing First-Past-the-Post and Closed List Proportional Representation

Abstract

Can electoral systems affect the behaviour of legislators between elections? Scholars have compared the effects of first-past-the-post (FPTP) to the effects of closed-list proportional representation (CLPR), in which parties rank candidates to determine their electoral chances in large multi-member districts. By giving the party leadership control over chances of re-election, CLPR incentivises legislators to toe the party line and increases legislative party cohesion. This literature, however, suffers from endogeneity problems, and its substantive findings have been criticised. To address these issues, I leverage the 1993 electoral reform in New Zealand, when FPTP was replaced with a closed-list mixed-member proportional system (CL-MMP), in which MPs are elected in part in single-member districts and in part with closed-list PR. I conduct two analyses. First, I perform an interrupted time-series analysis of parliamentary votes between 1983 and 2013, comparing legislative cohesion before and after the introduction of CL-MMP. Second, I use machine learning for natural language processing to compare MPs' speeches in Parliament to the language of their party manifesto before and after the introduction of CL-MMP. Both analyses show that under CL-MMP, MPs are overall more obedient to the party line. This effect is larger for MPs elected from closed lists vis-à-vis MPs elected in single-member districts. This evidence shows that electoral systems can have an independent effect on legislative behaviour.

Introduction & Literature Review

Electoral systems shape not only the electoral arena, but the legislative arena too. They not only determine allocation of seats between parties, but between individuals within parties, and so shape the incentives of individual legislators with respect to their parties and their electorates. Analysis on these inter-electoral effects of electoral systems is scarce, but there is a prevailing view that closed lists bias partisan behaviour by legislators. This chapter tests this hypothesis using evidence from New Zealand's parliament before and after its transition from Single Member District Plurality (SMDP) to a mixed system of single member districts and Closed List Proportional (CLP) elements.

Though the literature on electoral systems is extensive, the bulk of the analysis focuses on seat allocation and government formation. Far less examines how the legislatures constructed by the systems then operate between elections. The analysis of these effects of electoral systems has yielded strong theoretical theses about incentive structuring, and micro-level theses about delegation and responsibility. Where empirical analyses have been conducted, however, they have compared CLP to other proportional systems, demonstrating that closed lists create more partisan cohesion relative to open ones (Carroll and Nalepa 2020), or that pure CLP creates more partisan cohesion relative to a mixed system (Chiru 2021). With the public discourse around electoral system change being centred in countries with SMDP, where the public is dissatisfied with disproportionate seat allocation and the two party

system, it is important to also test these theories in a case which compares SMDP to a closed list system.

The theoretical intuition for these effects traces back to rational choice models of incentive structuring. These typically explain legislative behaviour based on maximising their prospects of keeping their seats. Cain, Ferejohn and Fiorina (1987) propose such an explanation. They argue that 'constituency service', including backbench rebellion on its basis, is only of direct value to constituents themselves, whilst only instrumentally valuable to the legislator to the extent that serving constituents helps them keep their seat. Therefore, by having legislators who lack constituencies and are selected only through the party list, CLP has legislators whose incentives are structured entirely by their party. This suggests that CLP should create more cohesive parties with more obedient legislators.

Another theoretical basis for these effects comes from delegation theories. These regard legislators as having 'dual agency', as being 'agents' delegated power by two 'principals': constituents and parties. Incentive to rebel is created when the principals have conflicting interests. Principal-agent theories are common in the literature on institutional design, but often neglect the role of parties themselves as principals. A traditional example is Strøm's (2003) work. Strøm's models, however, make legislators direct agents only of voters. When parties are considered, according to Carey's (2007) theory, legislators have competing principals. The amount of sway principals hold over their agents can be altered by the institutional design, based on

what tools principals are given to hold their agents to account. Tools which SMDP gives to constituents to hold legislators to account, such as constituency-level elections and local party organisations' role in candidate selection, are missing for legislators elected by party lists. Meanwhile the party list gives additional tools to the party to hold legislators to account, through the carrot and stick of raising and lowering legislators' list ranking respectively. Furthermore, MPs who are elected through party lists may be expected to conceptualise their representative role differently. How a representative is elected – why they are in parliament – serves as their source of legitimacy. If they are in their representative roles solely due to the voters' support for the party, and without any role played by the representative's own personality or characteristics, then it is more difficult to legitimise diverging from the obligations of their party, its manifesto, its leadership, and its platform. MPs elected this way may therefore not conceptualise their role as one with 'dual agency' at all: their sole principal is the party through which they were elected. These theories therefore agree that CLP should create cohesive parties with obedient legislators.

Microlevel analyses have found evidence for these effects. Carey and Shugart's (1995) analysis examines 13 electoral formulae and the 'incentive to cultivate a personal vote'. This incentive, in their theory, includes all non-partisan incentives, including winning the support of a local constituency and being favoured on an open list. Their findings show that CLP systems (including mixed systems) deviate from all others by devaluing personal reputation as district magnitude increases. In other words, standing out from the crowd as a candidate harms electoral prospects under CLP, and

not in other systems. Similar findings were reported by Carroll and Nalepa (2020) in their analysis of the incentives of party leaderships. In a comparison of open- and closed-list proportional systems, they find that open lists encourage 'preference cohesion', whilst closed lists encourage 'discipline'. Whereas national party success partly depends on individual candidates' appeal under open lists (and therefore there are costs to disciplining popular legislators), under closed lists votes for the national party are what deliver seats for the party, meaning a unified platform and discipline which guarantees delivery on promises are more important. Disciplinary strategies have therefore been found to be stronger and more strategically appealing under CLP.

Country-level case studies of mixed systems yield some further evidence. Fishel (1972) compared constituency candidates to list candidates in the West German SPD. They found that, of constituency-winning candidates, 10 were 'left' and 1 was 'centre', whilst of list-winning candidates, 11 were 'left' and 19 were 'centre'. The central party of the SPD was traditionally more moderate (Hunt, 1964) relative to its electoral support base of labour and union interests. Fishel's (1972) findings, however, are mostly descriptive, and not an explicit test of the hypothesis. Sieberer (2010) subsequently found that MPs with SMDP electorates were more likely to deviate from their party line in their voting behaviour than List MPs. A similar descriptive case is provided by Ward (1998) in examining the claim that List MPs were 'second-class MPs' to Electorate MPs[4] – a popular claim contemporarily. They therefore do

[4] In New Zealand, legislators elected by party list are called 'List MPs', and those elected by constituency are called 'Electorate MPs'.

not compare ideologies but qualifications, roles, and aptitude. They found little difference between the two types, except that List MPs were more disproportionally middle class. In their analysis, they make this a point about proficiency: List MPs had more highly skilled jobs prior to entering parliament. With regard to accountability, however, this could equally be indicative of a decreased selection power for the electorate to elect working class legislators.

The most rigorous and direct tests for the effects of electoral systems on legislative party cohesion have varied in their operationalisations of the two concepts, with each facing their own shortcomings. Hix (2004), for example, focuses on the European Parliament, as a single legislative assembly containing multiple electoral institutions for each country's MEPs. By comparing MEPs' votes to their co-partisans, and grouping MEPs by electoral system, they find that national parties more effectively enforce loyalty with CLP, even within the same assembly. Being within the same supernational assembly, however, does not alleviate the issue of varying national contexts and party organisations, and adds an issue of low public engagement (YouGov, 2019)[5]. Ceron (2015) operationalises both preference cohesion and candidate selection in order to test the relationship posited by Carroll and Nalepa (2020). They measure preference cohesion using 'wordfish' on Italian party conferences: a tool which places speeches on a left-right axis based on word frequencies. Their findings confirm the relationship: however strong the preference

[5] Equivalent surveys don't seem to have been conducted for other EU states.

cohesion of the party, CLP parties are more disciplined in parliament. Italy, however, has had electoral system changes which have usually been paired with replacement of the major parties involved, creating a confounding factor of party organisations changing at the same time.

Theory, Case Presentation & Research Design

Theory

The prevailing view of the literature is that closed lists hand significant coercive power to parties to select and deselect legislators. This gives parties greater power to control legislators in parliament. This takes place through multiple mechanisms. As well as the direct influence (candidate selection and de-selection being easier for the party) there is the underlying change to parliamentary culture. Via the new selection processes, List MPs realise a legitimacy deficit in their capacity to act as anything but agents of the party. They were elected through votes for a party, unrelated to themselves as individuals. Furthermore, the direct power has externalities: more demandingly partisan expectations of MPs' behaviour in general. The latter more gradual changes are labelled the 'indirect influences' on legislator behaviour. Through an electoral system which institutionalises parties as what voters elect and what legislators represent (rather than being merely coalitions of directly elected representatives) partisan demands on representatives' behaviour are given greater legitimacy. The 'direct influences', on the other hand, work through selection, party discipline and the threat thereof. Whereas the indirect influences work through a gradually changing political culture which slowly changes MPs' approach to legislative dilemmas, the direct influences work through selection and strategic response to a new institutional framework. Both mean introduction of closed lists ought to increase legislative party cohesion.

Under SMDP, legislators have incentive to rebel on behalf of constituency interests. There is much variation in parties' candidate selection systems, but the role of 'local parties' is a feature common to candidate selection systems in SMDP. Due to both the direct role of selecting or nominating candidates, and the indirect role of organising and staffing a candidate's local campaign, legislators have a strong incentive to consider their constituency's interests when acting in parliament. Furthermore, a particularly strong constituency grounding can protect a candidate from deselection even with a persistent record of rebellion. This dynamic is verified through candidates' electoral promises to rebel against their parties on prominent constituency-level issues, such as UK Conservative Zac Goldsmith's pledge to rebel against any bill proposing to expand Heathrow, an issue important to his local party. Failing to meet this promise ultimately led him to resign his seat, and lose it to the Liberal Democrats in a by-election (BBC News, 2016). Similar incentives also lead to large-scale backbench rebellions over local issues, such as the proposed HS2 railway line in the UK, and in the USA pork-barrel politics to appease local interests when voting on issues of national partisan importance. The protection role of local party organisations is seen archetypally through Dennis Skinner, who despite accumulating 322 rebellions in parliament from 1997 to 2019 (The Public Whip, 2019), was immovable from his Labour candidacy.

Under a pure CLP system, incentive to serve local-level interests is theoretically gone. Even within a mixed system, however, whilst constituencies can still protect an Electorate MP to an extent, the central party can provide better protection (backed by

more reliable promises) using the safety net of the party list. Furthermore, the makeup of parliament is shaped by voters' party votes (rather than electorate-level votes for their local MP), which are fought for at the national level. Seeking re-election through a constituency, therefore, is not necessary for an incumbent MP if they can secure instead a place on the party list. A conflict of interests which may lead to rebellion, therefore, can be resolved with the promise of a high list position (or threat of a low one). Parties, therefore, are far more capable of incentivising votes on party lines and of protecting electoral prospects of loyal MPs than constituencies are under CL-MMP, even for Electorate MPs. Party cohesion, therefore, ought to increase under systems involving closed lists.

As well as these strategic incentives created by the institutional setup which directly influence legislative behaviour, there is expected to be an indirect effect via parliamentary culture. The way legislators justify rebellions against the party line is usually with respect to conflicting allegiance to their constituency or their personal conscience. As democratic representatives, in navigating these allegiances they are conscious of the implications of the formal systems of representation. Debates, dating back as far as Burke (1774), pivot on the means by which citizens elect representatives as what justifies legislative behaviour: the implicit contract and style of delegation taking place between voters and legislators. A closed party list electoral system is quite a radical deviation on that formal system of representation: representatives from the party list are elected solely in virtue of being a member of their party selected by their party's leadership. Their mandate, therefore, is unrelated

to their individual conscience as a candidate, and there is no subnational constituency to speak of. This is expected to have an effect on how legislators – in particular List MPs – perceive their own roles as representatives. Under Closed-List MMP, as they navigate the dilemmas of allegiance between constituency, individual conscience and party, the legitimacy of party allegiance ought to be stronger relative to the other two.

There are therefore two distinct types of new influences on legislative behaviour which arise from Closed-List MMP (direct and indirect) which affect two distinct types of MP (List MPs and Electorate MPs) uniquely relative to SMDP MPs.

	MMP Electorate MPs	MMP List MPs
Direct Influences	Coercion through new selection power held by the central party: threats of lower list position (weaker safety net) or promises of higher list position (stronger safety net). Explicit provision of speeches, buzzwords & talking points by central party for potential preferential treatment in list position.	Coercion through sole selection power held by central party: threats of lower list position (deselection) or promises of higher list position (re-election prospects). Explicit provision of speeches, buzzwords & talking points by central party for potential preferential treatment in list position.
	Moderate effect	Strong effect
Indirect Influences	Pre-emption of new selection power held by central party, alongside awareness of local constituency party role persisting. Diffusion of political culture: possible legitimacy deficit of non-partisan behaviour for non-List MPs also.	Pre-emption of sole selection power held by central party. Legitimacy deficit of non-partisan behaviour: elected through party ballot to represent party.
	Small effect	Strong effect

Table 2.1: Summary breakdown of indirect and direct influences over MMP Electorate and MMP List MPs' partisan behaviour relative to SMDP MPs, as outlined above.

If and only if these mechanisms that alter the incentives of legislators do exist, then it ought to be possible to confirm the following thesis and hypotheses:

T: Closed-List MMP causes increased legislative party cohesion because MPs are more partisan in their behaviour.

H1a: Partisanship due to indirect influences increases gradually under Closed-List MMP.

H1b: Partisanship due to direct influences increases immediately after Closed-List MMP is introduced.

H2a: Under Closed-List MMP, List MPs are more partisan than Electorate MPs.

H2b: Under Closed-List MMP, Electorate MPs are more partisan than MPs under SMDP.

The main thesis, T, is demonstrable through the hypotheses H1 and H2. By H1 it would be demonstrated that there are identifiable and separable mechanisms by which Closed-List MMP increases legislative party cohesion, and by H2 that types of MP are uniquely affected based on their relationships to those mechanisms, a causal relationship between the electoral system and partisan cohesion can be established.

H1 presents the hypothesis about the differing mechanisms. The influences on legislative behaviour come in two distinct types: indirect and direct. If each of H1a and H1b can be validated independently then the causal thesis T is made stronger, as

there would not only be an increase in partisanship, but a specific pattern of increase across two different measurements which is consistent with the hypothesised consequences of an empowered electoral system whip through CLP.

Furthermore, the combination of H2a and H2b would similarly strengthen the causal thesis T by demonstrating another specific pattern of increase on different types of MP based on their relationship to the electoral system. H2a and H2b would relate differently to the mechanisms of H1a and H1b (summarised in Table 2.1). H2a and H2b would be demonstrated with relation to the mechanisms of H1a and of H1b, and would therefore demonstrate that these two specific patterns – one relating to the mechanisms associated with the electoral system whip, and the others to the types of MP – not only exist, but interact with each other as hypothesised. H2b would furthermore determine whether the effects hypothesised in H1 are only because of the presence of List MPs, or whether they affect all MPs under CL-MMP.

To create the measurements involved in analysing T, a clear definition of the type of party cohesion being analysed is required. What matters for the thesis presented here is the political activity of legislators within the party. The focus is therefore on how legislators vote and speak within parliament. Given the coercive nature of the electoral system whip proposed; the political activities of legislators and the degree of

loyalty to the party line they demonstrate are what connect to its strengths[6]. This can be in the public expression of views that are ideologically proximate to the party ideal point, or in loyalty to the party's position in votes. Sincere agreement with the party leadership's platform or policy positions is therefore importantly not a condition for the type of loyalty being studied in this chapter. The motivations for loyalist behaviour are not the focus of this chapter: an MP may be expressing loyalty or voting loyally because they agree with the party leadership, because they see their personal electoral success as tied to the success of the party, because of career ambitions, or any number of other reasons. Any of these motviations, however, may be influenced by the electoral rules: as the electoral rules give the party leadership additional powers to select candidates who sincerely agree with their platform, alter the degree to which the party leadership can provide an electoral safety net, and as a result of these, alter the degree to which factionalism and challenges to the leadership can be a likely method for furthering intra-party career ambitions.

The research design must avoid as many potential confounding variables as is possible. For this reason, an analysis of a single country which hypothesises a change (increased cohesion) at a certain point in time (introduction of new electoral system) is preferable to comparing between countries. Between countries, party organisations themselves are differently centralised; political culture is more or less candidate-

[6] Exit from the party, however, would not be suitable way of measuring the strength of the electoral system whip with respect to cohesion. Exit can indicate either a strong or a weak whip. It can also result mechanically from electoral system change because of the increase in fractionalisation in the party system rather than changes in whipping power.

centred; the constitution gives the legislature a different role; the executive's power over the legislature is different. An exhaustive list is outside the scope of this thesis. By limiting the analysis to a single country where the electoral system has changed, those (relatively few) confounding variables that do exist can be identified and controlled for.

Whilst cases of electoral system change which would be suitable for analysis are few, New Zealand stands out as ideal for this analysis. Those countries that adopted proportional systems during, or shortly after, the expansion of suffrage are not suitable. The expansion of suffrage triggers a strong, long-term adjustment process of its own, making it difficult to disentangle its effects from those of closed lists. The only cases of 'long-standing democracies changing the basic principles of their electoral rules' are: France in 1986, Italy in 1993, Japan in 1994 and New Zealand in 1996 (Renwick, 2018). Each of these has their merits and shortcomings, which are discussed in turn, but New Zealand stands out as the most appropriate for analysis. France, however, is discarded immediately as a possible case for analysis, as it only used its proportional system for one election before reverting to its original two-round system, meaning there is not enough information about the French legislature under PR.

The research design must disentangle two variables which both influence party discipline in parliament through the selection process: the electoral system and the party's own methods of candidate selection and discipline. Ceron's (2015) analysis of

these two ‘whipping resources’ risked confounding these two variables by using Italy as their selected case. Italy’s changes in electoral system have, in each case, been combined with a total overhaul of the parties in the system. 1994’s switch from pure proportionality to MMP saw Christian Democracy and the Democratic Party of the Left replaced with Pole of Freedoms and Good Government and Alliance of Progressives. By 2006, after the next change, these were replaced by The Union and the House of Freedoms. The two largest parties were again replaced when the electoral system was changed for 2018. The electoral system whip, therefore, has never changed whilst the candidate selection whip has been constant. This means that no parties have remained stable enough to conduct an interrupted time series analysis of the same party before and after any of the electoral system changes. Comparing different parties before and after the reform would lead to too many confounding variables. Parties all have entirely different selection processes, parliamentary whipping procedures, ideologies, cultures and agendas. The differences cannot be exhaustively listed, much less systematically controlled for in a way which accounts for all of these factors.

Another concern with the design of the research is what closed lists ought to be compared to. Other works in the literature (Carroll and Nalepa, 2020; Ceron, 2015; Carey and Shugart, 1995) have focused on comparing open and closed-list systems. This alone does not isolate the effects of closed lists specifically, however. Just as there is a hypothesised incentivisation to follow the party line under closed lists, there is also a hypothesised incentivisation to differentiate under open lists. To isolate the

effects of closed lists, a comparison must also be made to SMDP. Comparing a mixed system to SMDP within the same country is ideal for this, because features of candidate selection other than the presence of a closed list are equivalent. This is what this chapter looks to contribute, and why Japan is not chosen as the subject of the analysis. Before introducing a closed-list mixed system, Japan had a multimember district system with open lists. A candidate is therefore incentivised to differentiate from the party line in order to compete with copartisans (Carey & Shugart, 1995). New Zealand, therefore, remains the most appropriate case for analysis in order to contribute a new point of comparison and isolate the effects of closed lists, as it maintained the same major parties, and their selection processes for single-member district candidacies remained the same (Stephens & Leslie, 2011).

Case Presentation

The analysis I present addresses the concerns given above through the use of New Zealand as its case study. New Zealand switched from SMDP to CL-MMP in 1996. It had universal suffrage under SMDP for many decades before, and has operated under CL-MMP for several decades since, providing plenty of data either side. The two main parties since the Second World War also remained constant through the 1996 transition. Changes to the parties' own selection criteria were minimal at the time of the reform, changing only to add new processes for the selection and ranking of members for the new 'party lists.' Its challenges as a case are few and manageable.

New Zealand, especially before the electoral reform of 1996, was seen as the archetypal case of a majoritarian Westminster democracy. Lijphart (1989) noted pre-

MMP New Zealand as even more majoritarian than the UK, given its unicameral parliament. It has no federalism, no written constitution, no upper house and no powerful judiciary, making the only control on government the general election. This is even more true when parties are considered. Whereas, at the time, New Zealand and the UK were nearly identical institutionally, the way the systems operated in the legislature was substantially different. The UK parliament regularly sees government defeats due to backbench rebellion. In the period from the Second World War to New Zealand's electoral reform, UK governments had suffered 101 defeats in parliament. In that same time, a New Zealand government bill was never defeated (New Zealand Parliament, 2010). Barker and McLeay (2000), in their analysis of the New Zealand parliament before the change to CL-MMP, noted how 'independence of MPs, supposed to be a characteristic of simple-plurality single-member-constituency electoral systems, was not true of New Zealand'[7]. One institutional difference they recognise as a cause is the size of their parliament. At 95 members before the introduction of CL-MMP, there's little room in the parliamentary party for intra-party factionalism. This all shows how cohesive legislative parties always have been in New Zealand, which may make New Zealand a 'least likely' case to observe changes in cohesion due to the electoral system change (especially when relying on parliamentary votes as a measurement).

[7] A challenge addressed in the Research Design below.

The electoral system changed after over a decade of active debate on the issue. It was promised at elections by both the Labour and National parties in 1987 and 1990 respectively. After a set of non-binding referendums, a binding referendum in 1993 saw 53.9% of voters agreeing to the adoption of CL-MMP (Roberts 2001). New Zealand's SMDP system and its adopted CL-MMP system are summarised in the online appendices. Important features of the system beyond the introduction of List MPs are the still-small size of parliament, the 'fill-down' method for filling vacancies in list seats, and the potential to run for constituencies whilst being on the party list. Under SMDP, the Labour and National parties usually had 30-50 MPs each. The 35-member increase in the size of parliament means that even with increased fractionalisation, the number of Labour and National MPs has stayed each in the 30-50 range. This meant relative continuity of the faces making up the two main parliamentary parties between the 44th and 45th Parliaments[8], and that internal factionalism was just as unlikely. The fill-down method for list seat vacancies and the ability to run for both a constituency and the party list contribute to party power, the former through the threat of deselection (discussed further below) and the latter through diminishing the importance of constituency service.

Coincident with the electoral reform was a related change to how legislative votes were held, with the introduction of 'party votes', in which whips would vote on behalf of MPs under the assumption no legislators would rebel. Whilst any MP can,

[8] The last SMDP-elected and first CL-MMP-elected parliaments respectively.

by request to the speaker, force a 'personal vote', where legislators divide into the voting lobbies in person, this does change the dynamics of *how* legislators rebel, and the costs involved. This makes analysis of ordinary legislative votes before and after the electoral reform, such as Williams' (2012), likely fraught, as MPs typically cannot rebel unexpectedly, ad-hoc, or without explicitly informing their whip[9]. This is a challenge of the case which is addressed below in the research design.

The electoral reform also coincided with new rules for allocating speaking time to members. Speaking time was, since the reform, allocated by the Business Committee, which is typically attended by the party whips and the Leader and Shadow Leader of the House (Osnabrügge, 2021). The Committee allocates speeches according to rules that give parties speeches in parliament roughly proportional to their size. Within a party's allocation of speeches, party leaderships can play a role in selecting which MPs specifically get to speak. Osnabrügge yields findings that suggest that this gave party leaderships significant power over who speaks in parliament. This change, however, is not entirely independent of the electoral system change. It is a change that is made possible by the formalisation of parties into the electoral and legislative rules so that list seats can be allocated. It is, however, a feature that is likely to make any observed changes more extreme. I partially address the effects of this change in the research design below.

[9] Williams accounts for bills which are passed without holding a vote at all by limiting to votes where a 'recorded vote' was requested: but 'recorded votes' include party votes.

Given how cohesive parties already were before CL-MMP was introduced, it's not surprising that little was observed to have changed in the legislative cohesion of parties. MPs were obedient to the party line beforehand, and continued to be afterwards. Barker and McLeay (2000) note only that "a further incentive was added in that coalition deals require followers on whom leaders can rely." Furthermore, worries that were raised at the time about the qualifications, skills and duties of the new List MPs were unfounded according to Ward's (1998) analysis.

The legitimacy of List MPs has been a major source of debate since the introduction of CL-MMP, however. Ward (1998) also notes the tradition that arose almost immediately in the first parliament of 'shadow MPs.' List MPs were systematically assigned an opposing party's constituency to 'shadow' by setting up an office and serving as a point of contact within parliament for constituents who don't support their sitting Electorate MP. This, however, served to further bring into question the democratic legitimacy of List MPs. Barker and Levine (1999) goes so far as to describe the country's and parties' attitudes toward List MPs as regarding them as "55 MPs in search of a constituency." Instead of finding another source of demographic legitimacy for List MPs, their deficit of legitimacy has been highlighted, giving them less power and less incentive to distinguish themselves or act against the party line.

The democratic legitimacy of List MPs to act independently of their party was further challenged through the legal issue of List MPs resigning from their parties. Alamein

Kopu, an Alliance List MP, resigned her party whip in 1997 (Ward 1998). The Alliance party had had all its MPs sign a written pledge that if they resign the party, they resign from parliament. For List MPs, this meant handing their seat directly to the next-highest Alliance candidate from the previous election's list. The party attempted and failed in a legal case to have the speaker regard her notice of resignation from the party whip as an implied notice of resignation from parliament due to the agreement. Legislation enacted in 2001, however, not only reaffirmed party pledges like Alliance's, but enshrined the premise in law. The Electoral (Integrity) Amendment Act (2001) made it a legal requirement that List MPs resign from parliament if they are no longer party members. It failed to be renewed in 2005, and has been a subject of ongoing debate, being re-enacted in 2018 and currently still being in force (Electoral (Integrity) Amendment Act 2018). This all undermines the legitimacy of List MPs in acting independently of the party line in New Zealand, and suggests that confirming or denying the hypothesis that List MPs are less independent would be interesting for understanding the impact the electoral system change has had on New Zealand's parliament.

New Zealand is therefore a uniquely appropriate case for a single-country case study analysis of the effects of closed lists on a country's legislative politics. It provides plenty of data either side of the electoral system change, and maintains significant continuity in its institutions and party organisations before and after the reform. A glance at the country's political history since the reform also yields some promising insights, with the partisanship and accountability of List MPs having been

commented on by academic and contemporary news sources alike. It is, however, not without its challenges: the high degree of partisanship even before the reform, and a change to how it carries out legislative votes at the same time make finding an operationalisation for legislative party cohesion difficult.

Research Design

For a comparative study of legislative behaviour with and without closed lists, New Zealand is an ideal case in many ways, outlined above, yet also a 'least likely' case in others. It's a rare example of electoral system change, providing good data both before and after the transition. It presents one central challenge, however: that roll call votes are not a viable operationalisation of legislative party cohesion. Firstly, New Zealand is noted for its exceptionally cohesive parliamentary parties, even before the introduction of CL-MMP, with rebellions rarely ever seen[10]. This would make data on rebellions too sparse to be meaningful. Secondly, the method by which legislators vote in parliament changed when the electoral system did. Since the reform, roll call votes and personal divisions have been rare, with whips voting on their MPs' behalf. This not only makes data on individual MPs' rates of rebellion inaccessible; it makes it inappropriate by introducing a confounding variable to backbench rebellion. These challenges can be addressed, however, through two methods. Firstly, through the study of 'conscience votes' and speeches in parliament. The whip is formally removed from all MPs in conscience votes, yet partisan trends

[10] This is to the extent that it was noted as one of the countries most accurately described as an 'elective dictatorship' (Gold 1989).

persist, and party interests in seeing a certain outcome are apparent[11], as parties which appear divided may suffer electoral penalties (Greene & Haber, 2015). Secondly, through the study of speeches as a source of data on partisan loyalty, which may reveal variations in partisan loyalty even if voting records were entirely partisan.

Furthermore, there is an additional challenge for the use of speeches: the method for allocating speaking time to MPs changed after the electoral reform. The introduction of the Business Committee gave party leaderships more power to select speaking MPs on each debate (Osnabrügge, 2021). This is addressed through three measures.

Firstly, the methods using speeches are validated by comparison to the methods using conscience votes. An MP's ability to vote independently during a conscience vote is not conditional on their ability to speak during the debate. The results of the conscience votes analysis therefore provide another means of measuring and understanding the effect of electoral system change.

Secondly, the methods using speeches measure partisan loyalty at the MP-level rather than the speech-level. This avoids the newly introduced selection effect of party leaders selecting loyalist MPs to speak more often than less-loyal MPs.

[11] As they tend to be on civil rights or post-materialist social issues where there are often clear partisan agendas: National more socially conservative, Labour more progressive.

Thirdly, as well as comparing post-reform (CL-MMP) MPs to pre-reform (SMDP) MPs, I compare post-reform List MPs to post-reform Electorate MPs. Both of these types of MP are beholden to the same rules about how speaking time is allocated. List MPs' re-election chances, however, are more contingent on their party list ranking introduced by the new electoral rules. Electorate MPs, meanwhile, have their constituency as well as their list placement. These MPs exist in otherwise the same context, but have different relationships to the electoral system.

Effects between types of MPs after the reform therefore cannot be attributed to the Business Committee[12]. Given I avoid issues of selection effects at the MP level, it is also unlikely that effects between MPs before and after the reform can be attributed to the Business Committee. The Business Committee's control over speaking time affects how often different MPs speak, but not what their speeches contain. The content of speeches would only be affected if MPs have an incentive to signal party loyalty in order to be rewarded with speaking time. Speaking time, however, is only valuable insofar as it helps an MP develop their personal reputation in a direction they desire. Winning additional speaking time in order to make more loyalist speeches would only be desirable for an MP who wishes to make loyalist speeches, and an MP

[12] For cautious interpretation, a more conservative version of the analysis presented in this chapter may consider only the effects on List MPs relative to pre- and post-reform SMDP MPs. If List MPs are consistently more loyalist, this is indicative of the electoral rules playing an incentive-structuring role within the electoral reform. These effects are observed here in a country which used to have a SMDP electoral system, and even if isolated to List MPs, have an effect in aggregate on the level of backbench independence in the legislature.

who wishes to make rebellious or constituency-serving speeches would still use their time to make non-loyalist speeches. There is also the possibility that the party selects speakers conditional on the content of their proposes speeches, but similarly, this will only mean that less-loyal MPs will speak less often. I therefore expect any such effects on speech content to be minor, and the primary reason for MPs to have an incentive to signal greater loyalty to the party leadership to be the party leadership's control over their re-election chances under the new electoral rules, more so than their control over speaking time.

The operationalisation of legislative party cohesion employed as a test for H1a is how MPs vote during conscience votes. There is a long tradition in New Zealand of conscience voting. The whip is removed for MPs to 'vote with their consciences' on specific policy areas determined by parliamentary precedent, meaning that there is minimal variation over time in the topics voted on by this method (Lindsey, 2008)[13]. Since 1990, 93 bills have been put to conscience votes. Furthermore, commentators have observed increased partisanship at conscience votes, especially among List MPs (O'Sullivan, 2018). Parliamentary debates on conscience votes also betray the growing use of pressure on MPs to vote as party blocs despite the whip being formally removed, especially on List MPs:

[13] The only 'new' topic introduced as a subject of conscience votes in the period studied was 'life and death', with a single bill on euthanasia in the 2000s. The main change in the distribution of topics in the period was the decline of alcohol and marriage/family/children, and the rise of gambling and business/employment topics.

Tracey Martin (New Zealand First): "So New Zealand First will be casting all their votes in favour of this bill."

Tracey Martin (New Zealand First): "[…]they didn't vote our consciences here; they voted our party manifestos."

Simeon Brown (National): "[…] A deal made with the Green Party whereby they would support this bill in a bloc, […] The party which considers itself the conscience of the nation refuses its members to be able to exercise a conscience vote on a bill of this magnitude,"
(Hansard 2019)

Party cohesion at conscience votes, then, serves as a measurable operationalisation of legislative party cohesion. Whereas party cohesion in regular votes has always been high in New Zealand, the expectation of the hypotheses being tested would be that increased party power from closed lists would enable partisan voting blocs even on conscience issues. Due to rules and precedents restraining explicit whipping activities during conscience votes, the direct influences are expected to play less of a role compared to the indirect, cultural ones. This makes it an appropriate way of disentangling H1a (a gradual increase in legislative party cohesion associated with indirect influences) from H1b (a sudden increase in legislative party cohesion associated with direct influences).

The operationalisation of legislative party cohesion employed to test H1b is the partisanship of legislators' statements in parliament. Given the formal power of the party whip on voting, an act of direct rebellion risks constituting an 'exit' in Hirschman's terms (1970) by openly risking deselection. Further, given that the role of 'voice' in politics is as an outlet which reduces pressure on exit, either through negotiation or simply by mild dissent as coping; speeches in parliament which deviate from the party line would indicate a lack of cohesion. A stronger electoral system whip, however, which gives parties more powers of selection and coercion over legislators, would lead to greater partisanship of speeches, as even voice could be constrained. This can be through a variety of methods, direct or indirect[14]: ranging from tacitly rewarding openly loyal MPs by consideration for higher party list rankings and so encouraging others to demonstrate open loyalty; to invoking selection and deselection powers to coerce members directly. Whatever the leverage is, an important strategy employed by whips is feeding prewritten speeches and reinforcing repetition of partisan buzzwords and phrases. Feeding speeches and phrases to MPs to be spoken in parliament is standard practice (BBC2 England 2015) as a way parties can control debate and ensure partisan cohesion. These speeches are particularly partisan, and drop references to key partisan promises, policies and issues.

To establish whether this is something which may take place in the New Zealand parliament, a rudimentary analysis was done to establish expectations by running a

[14] As per Table 2.

phrase matching function built by Reeve (2020) to spot instances where Labour MPs have directly quoted their manifesto. In the 44th (SMDP) Parliament, it detected 20 instances of matched phrases; 3 of which were accurate and significant enough to be reported in full, and only 1 of which was a substantial quotation:

Suzanne Sinclair (Labour): Amalgamation has meant a local government unit of sustainable size that can be a true partner with central government in providing better services and opportunities for its citizens.

Labour Manifesto 1993: […] New Zealand now has city and local government units of substantial size that can be true partners with central government in providing better services and opportunities for citizens.

In the 45th (CL-MMP) parliament, 65 matches were detected in the cursory search, which filtered down to 14; most of which appear substantial, including:

Trevor Mallard (Labour): […] If central government acts with indifference towards that, or even with bias against Government workers – as we heard from David Carter earlier – as a result of the Government pursuing very short- term fiscal objectives, then the impact on the productivity of the country will be with us for a very long time.

Labour Manifesto 1996: If the central government structure continues to decline as a result of indifference towards, or even bias against government workers, or as a result of government pursuing short term fiscal objectives, then the impacts on the productivity and competitiveness of New Zealand will be very serious.

It is also worth noting that these are only the cases in which the quotation is long enough and the OCR of the manifesto was accurate enough to be picked up by the function. Many more may also exist, and use of shorter buzzwords and phrases aren't detected. These excerpts, however, give confidence that party manifestos do influence the wording of speeches. Whether through coercion or implied promise of reward or protection, it is an indicator partisan cohesion, making it a valid operationalisation.

The untrained natural language processing tool 'wordfish' (Slapin & Proksch, 2008) has been used in similar analyses of party cohesion using transcripts of speeches at party conferences (Greene & Haber, 2014) rather than parliamentary speeches. Wordfish is a statistical tool which analyses word frequencies in transcripts to estimate the ideological position of speakers on a unidimensional axis, but is limited through both its unidimensionality and its lack of specificity to the sample. This method can be greatly improved upon by using a training sample which could inform its placement of speeches, and could recognise different topics as their own dimensions. New Zealand Hansard provides full transcripts of all debates in parliament as far back as 1854. The New Zealand Labour Party has extensive manifestos published for each election, with reasonably consistent chapters which divide it into topics. The Labour party is chosen as the focus of this analysis because the National party produces far shorter manifestos, and in 1996 governed in a coalition, meaning a coalition agreement was drafted in replacement of their manifesto, which complicates the analysis by providing conflicting documents that may be cited, and introducing coalition dynamics to party strategy. Labour, on the

other hand, was in opposition under the same leader for both the 44th and 45th Parliaments[15], with most of the front bench being consistent, too.

Analysing just an opposition party also helps avoid the complex dynamics of majority size (Kam, 2009; Patty, 2008) and multipartism (Huber & Powell, 1994; Bowler, Farrell, & Katz, 1999) which are expected to primarily affect governing parties, who may have to alter their policy platforms for the sake of coalition (or minority government) deals, or be under increased pressure to deliver policy promises. In the specific case of New Zealand's 44th and 45th parliaments, change in these regards is at least minimal. Though the New Zealand parliament increased in size after the electoral reform, the size of the two parties remained relatively constant. Before the 1996 election, Labour and National had 41 MPs each. After the election, Labour had 37 and National had 44.

The scope is limited to these parliaments due to the demandingness of the method of analysis (involving machine learning analysis of over 100,000 speeches even when limited to these two parliaments), and to maintain these reasonably constant conditions described above. By training a machine learning model using the 1993 and 1996 manifestos[16], and comparing the language features of MPs' speeches to that model, the strength of the whip can be measured. MPs whose speeches are closer to

[15] Those elected by the 1993 and 1996 elections respectively; i.e. immediately preceding and succeeding the switch from SMDP to CL-MMP.
[16] Those used for the last SMDP election and the first CL-MMP election

the manifesto – borrowing words, phrases and ways of talking about the given topic – are more likely to be whipped. This operationalisation disentangles H1b and H1a by limiting its scope to the period 1993-1999, meaning it tests for the immediate, direct changes of H1b, and not the gradual changes of H1a.

Methodology and Analysis

Conscience Votes: Methodology

H1a is tested by measuring an operationalisation of legislative party cohesion that focuses on the indirect influences on legislative behaviour. It uses voting at conscience votes as a site of legislative behaviour which theoretically precludes direct coercive methods by central parties[17]. An interrupted timeseries analysis applied to a measure of cohesion at those votes then determines whether the date of the introduction of Closed-List MMP constitutes a 'breakpoint' where legislative party cohesion begins to change positively.

An interrupted timeseries analysis is used to determine the existence of a 'breakpoint' in 1996, when the electoral system changes. Following Campbell and Ross (1968), the data is to be treated as a 'quasi-experiment' with a 'pretest' of data from 1983 to 1996, an 'experimental treatment' in 1996, and a 'posttest' of data from 1996 to 2013. More recent analysis of policy outcomes uses a more formal method of testing for 'breakpoints' (Musazzi, et al., 2018; Ko, 2015), which occur where there is a significant change in the values or the slope of the data, such that segmenting the data at the hypothesised breakpoint creates a new model with a higher R^2 than the unsegmented model. The expected difference between the cohesion values before and after the hypothesised breakpoint in 1996 would be a change in slope, given that this

[17] In practice these methods may still be attempted, but if so, they would suffer from a deficit in legitimacy given the premise of a conscience vote. In practice, attempts to censure or punish an MP for 'rebellious' behaviour in a conscience vote may prove difficult. The party may face backlash from its own personnel or the public if the reasoning is explicit, and otherwise its options may be limited by finding another excuse to punish the 'rebellious' MP.

is designed as a test of the indirect and gradual hypothesis H1a. H1a would therefore be confirmed if a segmented regression model, segmented at the hypothesised date, exhibits a positive change in slope and is statistically significant relative to a simple linear model.

Legislative cohesion at conscience votes is then measured by application of a size-and-closeness-weighted version of the Rice (1925) index of partisan cohesion to conscience votes from 1983 to 2013. A complete list of conscience votes was compiled on advice of the New Zealand Parliamentary Information Service using a combination of their own records for conscience votes since 1990, and the list compiled by David Lindsey (2011, 239-268) for those before. The Rice index is a measure of partisan cohesion at votes obtained by the absolute (non-negative) value of the subtraction of the fraction of the party voting against from the fraction of the party voting in favour. The result is a number between 0 and 1. 1 represents the party being totally united in favour or against, whilst 0 represents a 50/50 split.

The Rice index has been criticised for ignoring abstentions, and not weighting votes based on how 'close' they are, prompting the creation of a closeness-weighted version of the index (Carey, 2009). Abstentions, whilst important to assessing the strength of party whips in general roll call votes, are not relevant to conscience vote analysis because there is no formal party line. Abstentions are common in conscience votes, and unlike a rebellious abstention in a normal whipped vote, are unlikely to represent active rebellion. Weighting by closeness, however, is important. At votes where the outcome is narrow, defying the party line is more significant. Without this, given the

formal nonpartisanship of the votes, many of the votes that show most party cohesion are those where there's little disagreement on the vote even between parties, whereas the results that matter are those where there is disagreement which forms on partisan lines. Furthermore, to analyse the cohesion of multiple parties, sets of parties, or whole parliaments, Rice index scores for each party must be weighted by party size and summed. A 'size-and-closeness-weighted' version of the Rice index, $SCRice_{kj}$, is therefore used for the Conscience Vote Cohesion statistic. The calculations for this index are detailed in the online appendices.

The Conscience Vote Cohesion values are calculated for all of the conscience votes data for the breakpoint analysis. To confirm the hypothesis H1a: the segmented model should exhibit a positive change in slope, and be more significant than the alternative linear model. This is determined by p-score significance and having a higher R^2. The software package 'segmented'[18], provides the tools for this analysis (Muggeo, 2020). The function pscore.test can be fed a regression model of the data and carry out a more robust significance test. It approximates the optimum breakpoint and compares the hypothesis that there is a breakpoint to the null hypothesis that there is either no breakpoint or more than one. If this test rejects the null hypothesis and uses a breakpoint in the region of the hypothesised date, this would also confirm the hypothesis H1a. The package can also provide a segmented regression model with the given breakpoint to compare to the linear regression for the other means of

[18] Built for the programming language R.

comparison. This provides two means of verifying the results which ought to agree with one another.

Conscience Votes: Analysis

The full timeseries data for Conscience Vote Cohesion is provided for the conscience votes data in the period 1983-2013 in Figure 2.1. The data is very noisy, which is expected as it describes partisan cohesion on formally nonpartisan votes. For short term variation, variables to do with the bills' content are far more relevant. This doesn't impact the relevance of breakpoint analysis, however, as long term trends can still be identified when noise is mitigated.

Before handling the data and performing regression analyses, the noise of the data is managed by using a 6-point rolling average to smooth the data (represented by the bold line in Figures 2.1 and 2.2).

The hypothesis that there is a breakpoint around the time of the first CL-MMP election is first tested by comparing a linear regression model to a segmented one. Figure 2.1 and Figure 2.2 visualise these models with dashed lines, and the regressions are formally presented in Table 2.2 and Table 2.3. The simple linear regression has a p-score of over 0.1[19], making it statistically insignificant, and a very low R^2 of 0.02, meaning that there's a high degree of unexplained variance from the fitted linear model. The segmented model was created using the software package 'segmented'.

[19] p=0.239

Based on an input breakpoint of the 1996 date of the first CL-MMP election in New Zealand, the segmented function selected as the optimum breakpoint a date in early 1997, just before the first conscience vote to take place after the electoral reform. The model is statistically significant in terms of p-score at $p<0.01$[20], and has a much higher R^2 of 0.4. This means that by segmenting the data between the last SMDP and first CL-MMP parliaments, time becomes a stronger explanatory variable for partisan cohesion at conscience votes. Given that it wouldn't otherwise be expected that time alone plays a major role in explaining variation in legislative party cohesion, this suggests a role of adjustment to the new electoral system as an explanation of increasing party cohesion at conscience votes after 1996.

[20] $p=3.87e^{-4}$

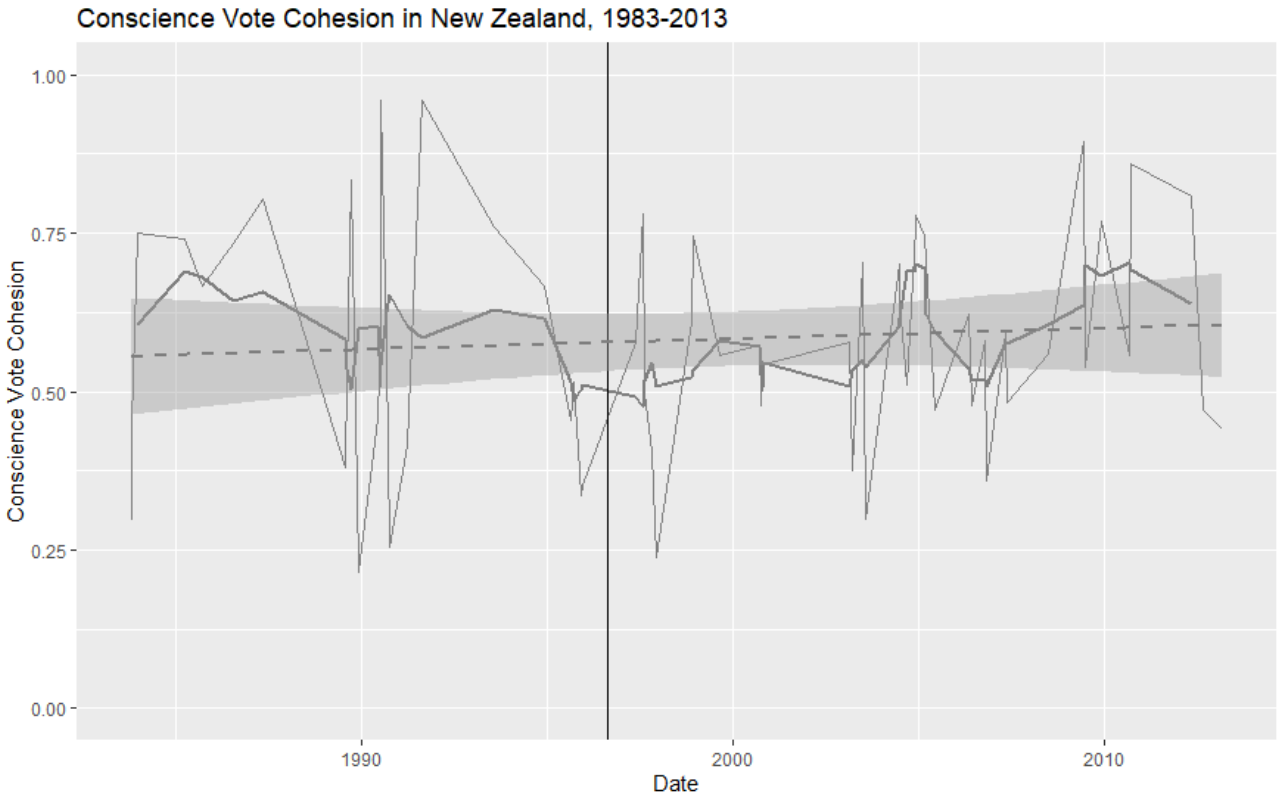

Figure 2.1: Full timeseries data of Conscience Vote Cohesion (solid), including 6-month rolling average (bold) linear regression (dashed) and hypothesised breakpoint (black, vertical)

Results: Linear Model

	Dependent variable:
	SCRicekj
Date	$3.680e^{-6}$
	$(3.094e^{-6})$
Constant	$5.458e^{-1***}$
	$(3.434e^{-2})$
Observations	62
R^2	0.02303
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 2.2: Linear regression analysis of 6-month rolling average Conscience Vote Cohesion (dependent) and Date (independent)

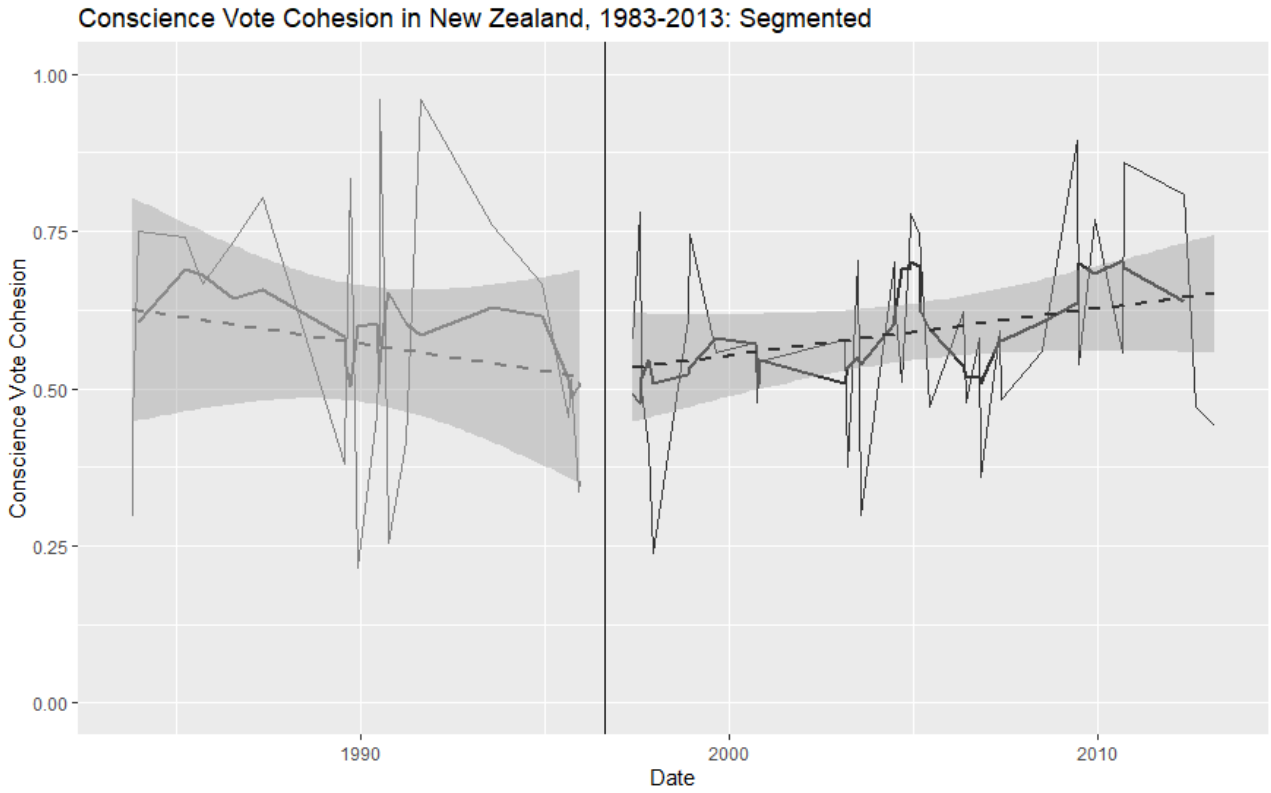

Figure 2.2: Segmented timeseries data of Conscience Vote Cohesion: pretest (left), posttest (right), respective 6-month rolling average (bold), linear regression (dashed), and hypothesised breakpoint (black, vertical)

Results: Segmented Model

	Dependent variable:
	SCRicekj
Date	$-3.104e^{-5***}$
	$(8.239e^{-6})$
Constant	$8.192e^{-1***}$
	$(6.404e^{-2})$
Observations	62
R^2	0.4225
Input Breakpoint	1996-10-12 (9781)
Predicted Breakpoint	1997-04-29 (9980)
First MMP Election	1996-10-12 (9781)
Standard Error	454.6
Within Std Error?	Yes
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 2.3: Segmented linear regression analysis of Conscience Vote Cohesion (dependent) and Date (independent)

This does not yet, however, exclude the possibility that there are other breakpoints that may be statistically significant. The presence of other statistically significant breakpoints would not threaten this analysis: multiple factors can be explaining variation in partisan cohesion at conscience votes without being mutually exclusive. Testing for other breakpoints[21], however, proves the breakpoint at the time of the electoral reform to be most compelling. Table 2.12 to Table 2.15 (Appendix) present the segmented regression analyses for 4 different potential breakpoints identified by eye in 1990, 1991, 2003, and 2006. The only models that were significant at the $p<0.01$ level were ones where the strongest breakpoint detected based on the given input date was within 1 standard deviation of the electoral reform date: these were the 1991 and 2003 inputs, which could not find a suitable breakpoint in the input year, landing instead on breakpoints in 1997, both around the date of the first conscience vote of the first CL-MMP parliament. The other two were less significant and had lower R^2 values.

Also with the 'segmented' package, a different type of hypothesis test can be carried out comparing the hypothesis that there is a breakpoint in 1996 to a two-sided null hypothesis that there are no breakpoints or that there is more than one breakpoint. This function – 'pscore.test' – is unsupervised: it detects a breakpoint automatically

[21] By running them through the comparative breakpoint analysis described above.

(albeit to a lower degree of specificity than the 'segmented' function). It used a breakpoint in late 1997 as the basis of the hypothesis test, and was significant at the $p<0.01$ level[22]. This is a highly significant result, which not only rejects the hypothesis that there is no breakpoint in the data, but also rejects the hypothesis that there are any additional significant breakpoints other than the one around the time of the electoral reform.

This confirms the hypothesis H1a: after the implementation of CL-MMP in 1996, the operationalisation for indirect influences on partisan cohesion exhibits an upward trend which isn't present before 1996. This finding also cannot be explained easily by variations in topics voted on, as such a variation would have to not be a simple decline in certain (partisan-cohesive) topics and the rise in other (partisan-divisive) topics, but such a trend until the 1996 electoral reform, and a reverse of that trend after the 1996 electoral reform[23].

[22] $p = 6.66e^{-6}$

[23] Such a theory would also have to confront the possibility that the topic selection is the causal consequence of the strengthened party leaderships of the new electoral system, and in fact a part of the theory I present here.

Parliamentary Speeches: Methodology

The parliamentary speeches analysis looks to test H1b: the hypothesis that legislative cohesion as measured by an operationalisation that includes direct influences on legislative behaviour increased after the electoral reform of 1996. The operationalisation to be used to measure this is the comparison of legislators' speeches in parliament to their party manifesto.

Legislative cohesion in parliamentary speeches is measured by using a machine learning model to compare the wording of each MP's corpus of speeches in 7 topics to the wording of the prior election's Labour Party Manifesto chapters relevant to that topic.

The categories were defined using a machine learning tool called LDA Mallet from gensim (Řehůřek & Sojka, 2010), a natural language processing package for python which performs advanced topic modelling. By identifying clustering of certain sets of rarer (subject-specific) words, it predicts topic categories, which can then be used to predict which speeches belong to which topic. This allows the corpus of over 100,000 speeches to be categorised automatically with categories tailored specifically to the political context and the party's platform. A recent working paper by Osnabrügge, Ash & Morelli (2020) presents a supervised model of categorisation which they apply to the New Zealand parliament, also using Hansard transcripts, and using relevant manifesto chapters. My categorisation uses an unsupervised model, however, because training the categorisation using the manifestos would make the comparison to

manifestos self-referential[24]. The categorisations created by the unsupervised model are also more interpretable and of a higher quality than was expected by Osnabrügge, Ash & Morelli. LDA Mallet output two sets of nine categories independently for speeches in the 44th and 45th parliament, and the sets agreed on the categories. In both of the outputs, two of the nine categories grouped speeches which didn't reveal the topic of debate and contained mostly technical parliamentary vocabulary, based around words such as: 'raise', 'agree', 'paper', 'introduce', 'provisions', and so forth. These were used to exclude those more procedural statements in parliament which weren't part of topical debate. The remaining seven, bar one, aligned clearly: creating the near-identical categorisations presented in Table 2.4, with categories in each for the topics of civic/local government, education, elections, fiscal policy, healthcare and monetary policy. Some variance in the important topics which are important contemporarily for each parliament is to be expected, so the replacement of a global affairs topic in the 44th Parliament with a civil rights topic in the 45th is not surprising. That an unsupervised model output near-identical topics in two separate corpuses is itself a good sign. Except for speeches which were especially short (which were excluded in the following analysis), the topic-matching is also accurate for individual speeches. Examples of speeches it categorised are displayed in the online appendices.

[24] Measuring similarity between speeches categorised based on manifesto chapters to manifesto chapters would only measure categorisation quality.

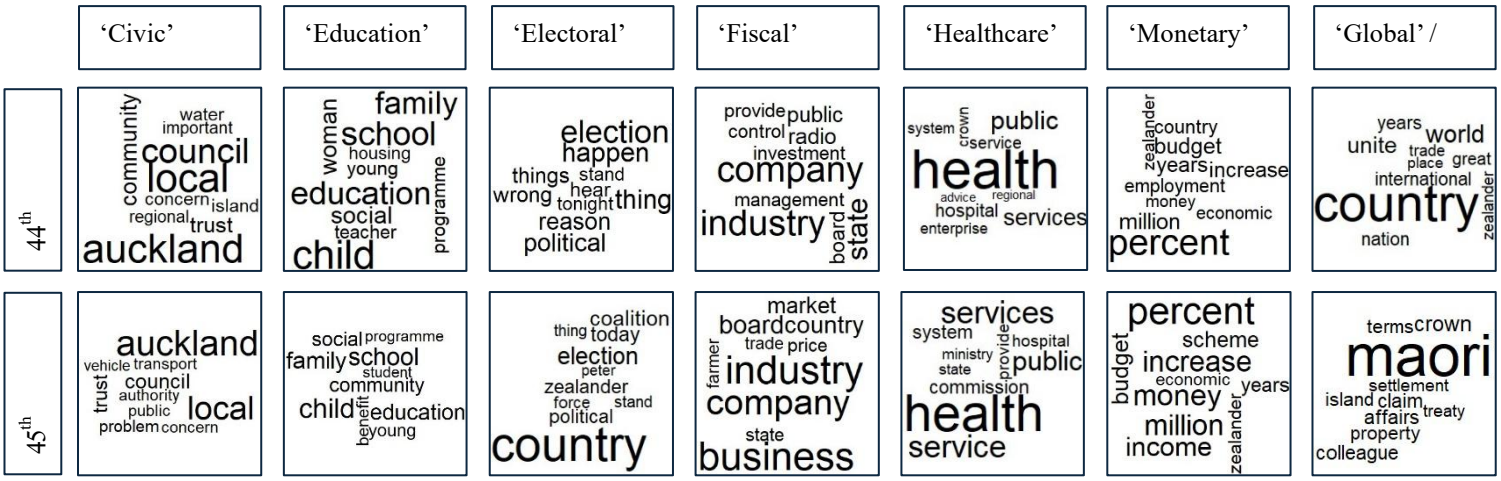

Table 2.4: 'Word clouds' for the 7 categories identified in the speeches of the two parliaments studied.

The comparison between speeches and manifestos is then done using an unsupervised machine learning technique for authorship attribution originally developed by AICBT (2015), altered for the purpose of this analysis (explained in the appendices). It clusters based on language features of the speech, such as vocabulary. Once clustered, those which belong in the same cluster as the pre-defined examples of text by each author are predicted to have been written by the respective author. The 'author' in this context is the party as the collective author of the manifesto, with more partisan 'authorship' as evidence of greater partisanship.

The output of this analysis is, for each MP on each category, a prediction as to whether they belong in the more partisan group or the less partisan group on that category represented as a 1 (whipped) or 0 (not). The mean across those categories then represents what I have termed their 'Speech Partisanship Score'. For example, being in the more partisan cluster on three of seven topics would yield a Speech Partisanship Score of 0.43.

To test the significance of the findings of this analysis, the populations of the 44th (SMDP) and 45th (CL-MMP) parliaments are separated, and a Wilcoxon signed-rank test is performed with the partisanship score as the dependent variable. The Wilcoxon signed-rank test is a variation on the t-test which dispenses with the assumption that the data is normally distributed. It interprets the mean and distribution of a dependent variable (Speech Partisanship Score) in two samples, and tests against a null hypothesis. In this case, the null hypothesis is that the two samples are taken from the same population (i.e. that the sampling might as well be randomly chosen). If the test yields a significant p-score, then the sampling is relevant to the dependent variable (Speech Partisanship Score): whether you were speaking in an SMDP parliament or an CL-MMP parliament is relevant to how likely you are to be more partisan in your speeches, confirming H1b.

Parliamentary Speeches: Analysis

The distribution of Speech Partisanship Scores given by the analysis is represented in the form of a box-and-whisker plot overlayed with a dotplot in Figure 2.3. At a glance, the distribution of Speech Partisanship Scores has stark differences before and after the electoral reform. Under SMDP, the vast majority of MPs are clustered at a Speech Partisanship Score of 0, meaning they were not predicted to be influenced by the manifesto in any of their speeches. Under CL-MMP, though a Speech Partisanship Score of 0 is still (by a narrow margin) the most common score; there are similar sized clusters at 0.14 and 0.29, as well as some distributed at higher levels between 0.3 and 0.6.

The results of the two-tailed, unpaired Wilcoxon signed-rank test performed on these two samples is given in Table 2.5. They confirm the hypothesis that the sampling is relevant to the dependent variable (Speech Partisanship Score), providing strong evidence (at the $p<0.01$ level) against the null hypothesis that the MPs are selected randomly. The integrity of these results is also strengthened by a Wilcoxon test statistic (W) below the critical value for the sample size and significance level, suggesting that this level of significance to the difference between the two samples is highly unlikely to be random.

This suggests that, as per H1b, legislative party cohesion according to an operationalisation that includes direct and coercive influences on legislative behaviour, significantly increases when Closed-List MMP is introduced. This suggests the role of influences such as invoking the selection power of the party list by the central party in order to maintain obedience of MPs, and of issuing speeches to be read by MPs which uphold the party line, making the party whip stronger under Closed-List MMP.

Hypothesis H1b is therefore confirmed by this analysis: the implementation of CL-MMP in 1996 had an effect on the partisanship of speeches given by MPs in parliament, leading to a greater degree of influence of the manifesto (as a representation of the party line) over the wording of MPs' speeches.

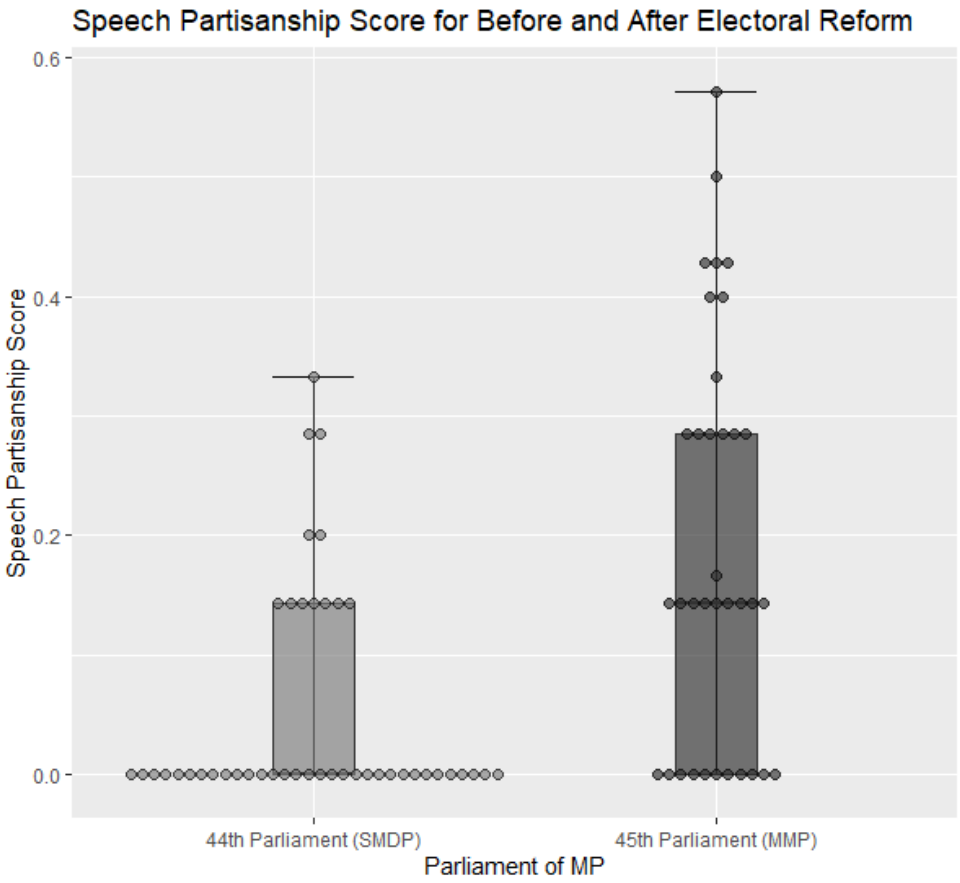

Figure 2.3: Box-and-whisker plot overlayed with dotplot for quartiles and data points in the distribution of SPSs for Labour MPs of the SMDP 44th (left) Parliament and the CL-MMP 45th (right) Parliament.

Wilcoxon Test Results

	44th & 45th Parliaments:
	p-value
Speech Partisanship Score	6.819×10^{-5}***
N	44, 35
W	402
Critical Value	535
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 2.5: Output of the Wilcoxon signed-rank test comparing SPS of MPs in the 44th and 45th Parliaments.

Application to Types of MPs: Methodology

For each of these analyses, as well as comparing the parliaments before and after the change of the electoral system, List and Electorate MPs under the CL-MMP system are compared. This serves as a test for H2a and H2b by determine the extent to which the changes since 1996 have affected the behaviour of List and Electorate MPs differently. It is expected that the application of the Speech Partisanship Score measure created for the Parliamentary Speeches analysis to confirm H1b would be the most relevant to the difference between List and Electorate MPs (H2a). It is the through more direct mechanisms which influence MPs' behaviour via the candidate selection system where the greatest difference between the two new types of MP would be expected. Changes to political culture over time, on the other hand, are likely to influence all MPs in a more even way.

To test the H2a and H2b predictions, MPs are sampled as 'SMDP MPs' (all MPs before the reform in 1996), 'MMP Electorate MPs' (MPs with a single member constituency after the reform) and 'MMP List MPs' (MPs from party lists without a single member constituency after the reform). The three populations are then compared to one-another using a Wilcoxon signed-rank test, firstly using the Conscience Vote Cohesion[25] as the dependent variable, then using the Speech Partisanship Score of MPs' speeches in parliament as the dependent variable. If the tests yield significant p-scores, then the samplings are relevant to the dependent

[25] Adapted so that, in the post-reform period, there are two 'parliaments': one of only Electorate MPs, and one of only List MPs. $SCRice_{kj}$ scores are calculated for each.

variables: the type of MP you are is relevant to how cohesive you are to the party line, such that all MPs under CL-MMP, and especially List MPs, are more partisan than SMDP MPs, confirming H2a and H2b (and also evidencing H1).

Application to Types of MPs: Analysis

Comparing how the two effects impact list and Electorate MPs independently yields results which are represented in the form of box-and-whisker plots overlayed with dotplots in Figure 2.4 and Figure 2.5. As anticipated in H2a and H2b, relative to SMDP MPs, there is a marginal increase in partisanship for MMP Electorate MPs, and a more substantial one for MMP List MPs. It also shows that the differentiation between List and Electorate MPs is far more substantial in Speech Partisanship Score than in Conscience Vote Cohesion, reflecting the 'directness' of the mechanisms behind each operationalisation as expected.

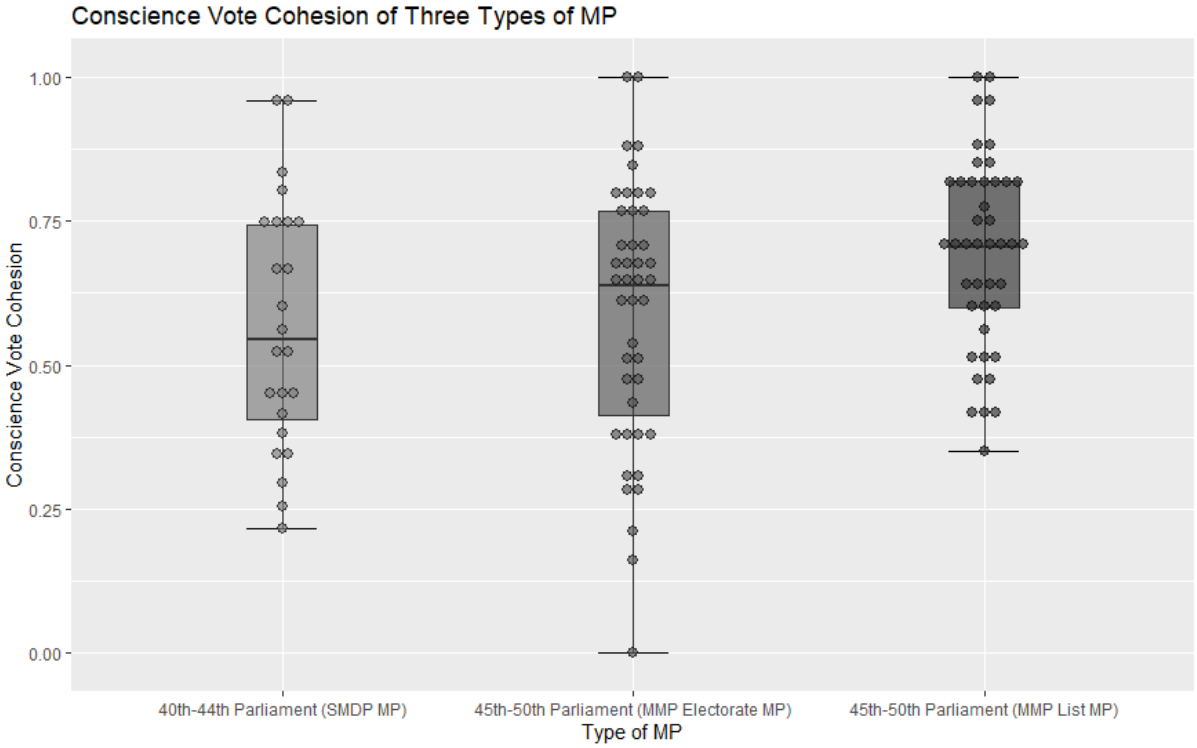

Figure 2.4: Box-and-whisker plot overlayed with dotplot representing quartiles and data points for Conscience Vote Cohesion of SMDP (left), MMP Electorate (centre) and MMP List (right) MPs at conscience votes. Each dot represents a conscience vote.

Wilcoxon Test Results	
	SMDP & MMP List MPs:
	p-value
SCRicej	0.01640065**
N	24, 43
W	700
Critical Value	390
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 2.6: For SMDP and MMP List MPs

Wilcoxon Test Results	
	SMDP & MMP Electorate MPs:
	p-value
SCRicej	0.6518629
N	24, 43
W	551
Critical Value	546
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 2.7: For SMDP and MMP Electorate MPs

Wilcoxon Test Results	
	MMP Electorate & MMP List MPs:
	p-value
SCRicej	0.0006237332***
N	43
W	166
Critical Value	282
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 2.8: Wilcoxon test results comparing Conscience Vote Cohesion scores for MMP Electorate and MMP List MPs

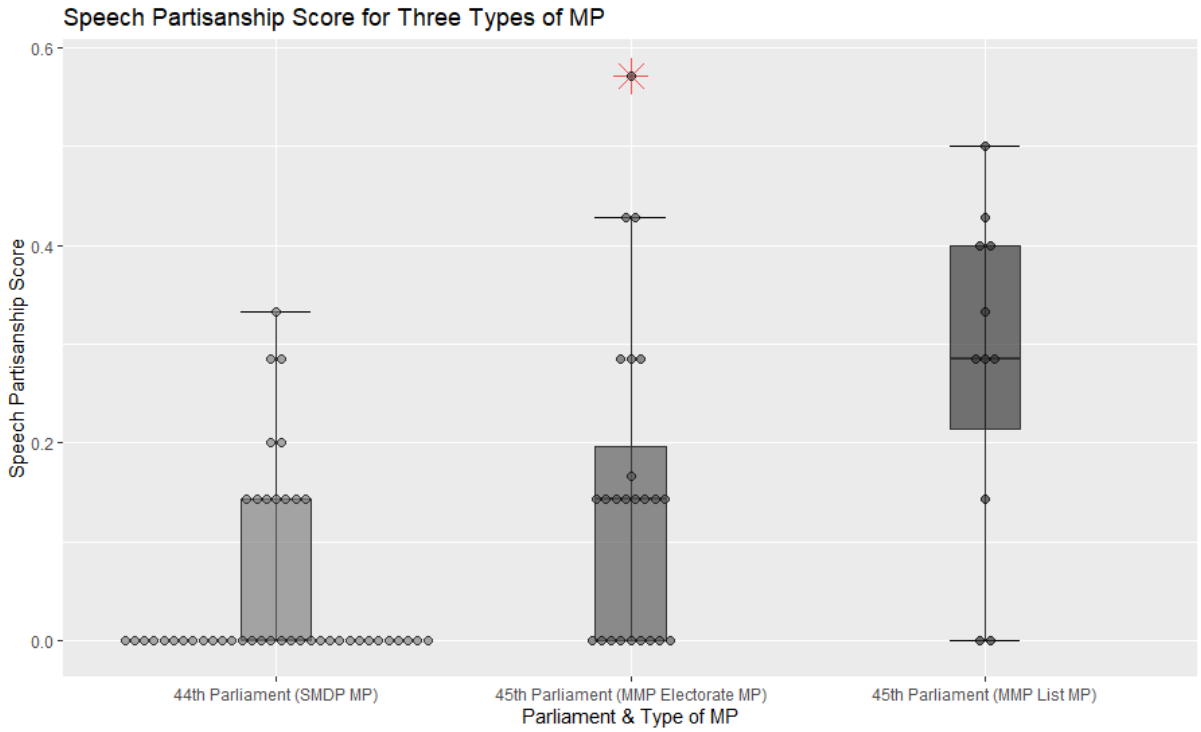

Figure 2.5: Box-and-whisker plot overlayed with dotplot representing quartiles and data for SPS of SMDP (left), MMP Electorate (centre) and MMP List (right) MPs' corpuses of speeches. Each dot represents a Labour MP's corpus.

Wilcoxon Test Results

	SMDP & MMP List MPs:
	p-value
Speech Partisanship Score	3.870×10^{-5}***
N	44, 11
W	71
Critical Value	133
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 2.9: For SMDP and MMP List MPs

Wilcoxon Test Results

	SMDP & MMP Electorate MPs:
	p-value
Speech Partisanship Score	0.004050898***
N	44, 24
W	331
Critical Value	348
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 2.10: For SMDP and MMP Electorate MPs

Wilcoxon Test Results

	MMP Electorate & MMP List MPs:
	p-value
Speech Partisanship Score	0.0408512**
N	24, 11
W	75.5
Critical Value	86
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 2.11: Wilcoxon test results comparing Speech Partisanship Score for MMP Electorate and MMP List MPs

Table 2.8 and Table 2.11 present the Wilcoxon signed-rank test results for Conscience Vote Cohesion and Speech Partisanship Score respectively, comparing MMP List MPs to MMP Electorate MPs. Both are significant (at the $p<0.01$ and $p<0.05$ level respectively) rejections of the null hypothesis that the MPs are sampled randomly with respect to the independent variable: the sampling as List or Electorate MPs is relevant to how cohesive they are to the party line. They are also assigned Wilcoxon test statistics which are below the critical value for the sample sizes and significance level, suggesting the results are reliable. This confirms H2a with both measurements.

Comparing MMP List MPs and MMP Electorate MPs each independently to SMDP MPs yields more varied results, shown in Table 2.7, Table 2.6, Table 2.10 and Table 2.9. For the Speech Partisanship Score measure of the direct mechanisms of increased legislative party cohesion, the two tests reject the null hypothesis with a significance at the $p<0.01$ level and have Wilcoxon test statistics below their critical values. This suggests that, for the Speech Partisanship Score measure, all three samples are distinct: all three respond differently to their electoral system conditions. SMDP MPs are the least cohesive; MMP Electorate MPs are marginally more cohesive, and MMP List MPs are highly cohesive. The results for equivalent tests on Conscience Vote Cohesion as a measurement of the indirect, political culture mechanisms of increased legislative party cohesion are less clear. The difference in Conscience Vote Cohesion between SMDP and MMP Electorate MPs is not significant enough to reject the null hypothesis, and though the difference in Conscience Vote Cohesion between SMDP and MMP List MPs is significant in its p-score, the Wilcoxon test statistic exceeds the

critical value, meaning that the chance of selecting these samples from the population at random are too high to be confident in the result. Given how similar the distributions are (visualised in Figure 2.4) and that Conscience Vote Cohesion measures the indirect effects that would gradually impact all MPs, this was to be expected. This therefore confirms H2b with the measurements of H1b, but not with the measurements of H1a.

This means that the results in terms of H1a's measurements confirm H2a but not H2b; whilst the results in terms of H1b's measurements confirm both H2a and H2b. An interpretation of these results is that H2b (that even Electorate MPs behave more cohesively after CL-MMP is introduced) is true only through the direct influences over legislator behaviour: coercion using the new selection powers and provision of speeches to MPs does have a stronger effect, but the effect of political culture and preemption of coercion is small for Electorate MPs. This was what was anticipated in Table 2.1, and also makes sense in the more specific terms of the operationalisations used: Electorate MPs are still elected directly, so they are more likely to feel legitimate in voting based on their own personal conscience rather than following the party line in a conscience vote, but the positive selection power given to the whip in being able to provide Electorate MPs with a safety net via the list is still a strong influence over their behaviour. Only a small effect of indirect influences on Electorate MPs was expected, and a small one is visible to the eye in Figure 2.4, but is unsurprisingly not statistically significant.

The results therefore confirm the hypotheses of H2a and H2b: whilst the effects of CL-MMP on legislative party cohesion do impact both list and Electorate MPs, with a greater effect on list MPs, and in particular through the more direct mechanisms that are measured by Speech Partisanship Score.

Conclusion

The findings of the three analyses presented above are significant and confirm the hypotheses outlined in the research design. Legislative party cohesion increased after 1996 according to both operationalisations, confirming H1b and H1a. There was a marked difference between SMDP, MMP List and MMP Electorate MPs to the extent anticipated for both operationalisations, which not only confirms H2a and H2b, but gives even greater confidence in H1: that the electoral system change is the variable which caused the increasing legislative party cohesion.

This confirms that closed party lists do have a positive effect on legislative party cohesion: whether through indirect influences such as legitimacy deficits and political culture, or more direct influences like coercion and (implicit) promises of party list rankings; the electoral system gives more power to central parties to rely on loyalty of MPs in parliament.

These results should contribute to an understanding of how electoral systems affect not only the composition, but the operation of parliaments. Closed party lists can be regarded as providing strong whipping resources to party leaderships through the power to reorder candidates on the basis of loyalty. This can amount to the provision or removal of a safety net for Electorate MPs, or to effective deselection of List MPs, on the basis of their legislative behaviour. This effect can be intuitively inferred, but has not been rigorously tested in comparison to SMDP systems. The research design

presented above provides an effective way of comparing the whipping power in SMDP and CLP systems using the case of New Zealand as one of the few cases of electoral system change where there is significant data on the country's democratic history both before and after the transition.

Whilst the results focus on a single country case study, the analyses are valid and important and confirm one-another's findings, all supporting the hypotheses that closed party lists positively affect legislative party cohesion. They combine long-established quantitative methods in the conscience vote analysis with novel, advanced methods in the parliamentary speeches analysis. The Rice index of voting cohesion in an interrupted time series analysis is a novel application of a tried and tested operationalisation of cohesion, whilst the machine learning techniques used to compare wording of parliamentary speeches to party manifestos is an entirely original method of analysis. These two very different operationalisations of legislative cohesion produced mutually supportive results, reaffirming one-another's validity.

Whether the partisan loyalty measured in this chapter is the result of increased preference cohesion due to selection effects, or increased party control to restrict the expression of diverse preferences within the legislature, is not identified here. Both selection effects at elections and coersive whipping activites that alter behaviour between parliaments are mechanisms by which MPs may behave more loyally due to the change in electoral rules. That the electoral rules are the cause, however, is what I

identify, given the breakpoint coincident with the electoral system change, and the much stronger effect sizes for List MPs as for Electorate MPs.

Deeper qualitative research will be required to determine exactly how closed lists affect MPs' behaviour: for example whether party whips invoke the list directly as a way to coerce MPs, or if the list acts as a passive deterrent that frames MPs' choices indirectly. The analyses presented above nonetheless appreciate that there are a number of ways closed lists could affect MPs' behaviour which are not mutually exclusive, and respects that different operationalisations may be more relevant to different mechanisms.

This therefore contributes to the findings of Carey and Shugart (1995) as well as recent findings by Chiru (2021) which used evidence from Romania's transition from full CLPR to a mixed proportional system to demonstrate that the mixed system enhanced constituency service. These findings confirm the same dynamic using evidence from New Zealand as a case of transition in the opposite direction: from full SMDP to a mixed system, which is shown to enhance party cohesion.

Appendices

Appendix 1: Regression Tables for Alternative Breakpoints:

Results: Earlier Breakpoint Segmented Model

	Dependent variable:
	SCRicekj
Date	$-2.525e^{-5}$***
	($1.059e^{-5}$)
Constant	$7.780e^{-1}$***
	($7.778e^{-2}$)
Observations	62
R^2	0.4068
Input Breakpoint	1990-06-28 (7483)
Predicted Breakpoint	1998-10-12 (10511.6)
First MMP Election	1996-10-12 (9781)
Standard Error	602.8
Within Std Error?	No
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 2.12: Segmented linear regression analysis of Conscience Vote Cohesion (dependent) and Date (independent) with an earlier potential breakpoint in 1990.

Results: Early Breakpoint Segmented Model

Dependent variable:
SCRicekj
$-3.104e^{-5}$***
($8.239e^{-6}$)
$8.192e^{-1}$***
($6.404e^{-2}$)
62
0.4225
1991-08-22 (7903)
1997-04-29 (9980)
1996-10-12 (9781)
454.6
Yes
*p<0.1; **p<0.05; ***p<0.01

Table 2.13: With an early potential breakpoint in 1991.

Results: Late Breakpoint Segmented Model

Dependent variable:
SCRicekj
$-3.066e^{-5}$***
($8.872e^{-6}$)
$8.165e^{-1}$***
($6.794e^{-2}$)
62
0.4225
2003-07-30 (12263)
1997-05-28 (10009.2)
1996-10-12 (9781)
480.0
Yes
*p<0.1; **p<0.05; ***p<0.01

Table 2.14: With a late potential breakpoint in 2003.

Results: Later Breakpoint Segmented Model

Dependent variable:
SCRicekj
$-3.333e^{-6}$
($3.575e^{-6}$)
$6.073e^{-1}$***
($3.686e^{-2}$)
62
0.1897
2006-11-08 (13460)
2006-04-07 (13245.1)
1996-10-12 (9781)
712.3538
No
*p<0.1; **p<0.05; ***p<0.01

Table 2.15: With a later potential breakpoint in 2006.

Appendix 2: Summary of Changes in New Zealand's Electoral Reform

Feature	SMDP	MMP[9]
MPs	95 MPs	120 MPs
Breakdown	95 single-member districts 4 Maori seats superimposed 0 national closed list seats	60 single-member districts 5 Maori seats superimposed 55 national closed list seats MPs can contest electorates and be on the national list.
Ballot	All candidates by simple plurality	Electors cast two votes on one ballot: for constituency candidate and national party.
Redrafting	South Island seats fixed at 25 Maori seats fixed at 4 Total seats increase with population Neutral boundary commission.	Neutral boundary commission. St Lague formula for list seats.
Vacancies	By-elections	By-elections for electorate seats. Ranked party list for list seats.
Voter Registration	Compulsory registration. Maori choose to register in local constituency or Maori seat.	Compulsory registration. Maori choose to register in local constituency or Maori seat.
Legislative Voting	Always personal vote	Party vote by default (whips vote on behalf of party's MPs) Personal vote if requested by any MP (rebellions and conscience votes)

Appendix 3: Conscience Vote Cohesion ($SCRice_{kj}$) Calculation:

The closeness score is obtained by the absolute (non-negative) value of the subtraction of the fraction of all votes in favour from 0.5, doubled and subtracted from 1. The result is represented as a number between 0 and 1, where 1 represents a perfect tie, and 0 represents a unanimous vote either for or against. From the recorded conscience votes, then, the Rice index for each party (i) for that bill (j), along with the closeness (Carey 2009) for the bill are calculated. Both $Rice_{ij}$ and $Close_j$ are on a scale of 0 to 1.

$$Rice_{ij} = \left|\%Aye_{ij} - \%Nay_{ij}\right|$$
$$Close_j = 1 - 2 \times \left|0.5 - \%Aye_j\right|$$

In order to aggregate the $Rice_{ij}$ results into a single result for the bill as a whole (or a set of parties on that bill), rather than a single party, the $Rice_{ij}$ results for each party in a set (k) must be weighted by each party's size and summed. The size-weighted Rice index ($SRice_{kj}$) for the bill is therefore the sum of Rice indexes of each party in k, each multiplied by the party's size, all divided by the summed size of all parties in k.

This summarises cohesion for multiple parties, where a score of 1 means all parties were perfectly cohesive, and a score of 0 means parties were on average split 50/50.

$$SRice_{kj} = \frac{\sum_{i=1}^{k}(Rice_{ij} \times Size_{ij})}{\sum_{i=1}^{k}(Size_{ij})}$$

The size-and-closeness-weighted Rice index ($SCRice_{kj}$) is calculated by subtracting 0.5 from $SRice_{kj}$ and multiplying that by $Close_j$, before adding 0.5 again. This means that, for bills which weren't close, the lower $Close_j$ value means $SCRice_{kj}$ is brought closer to 0.5, the midpoint of potential values. Where bills were close, the higher $Close_j$ value means $SCRice_{kj}$ is brought to either extreme, depending on whether the bill saw relative cohesiveness or in-cohesiveness from parties. A very low score means parties were evenly split on average, and the vote itself was close (highly nonpartisan). A very high score means parties were unified and the vote itself was close (highly partisan).

$$SCRice_{kj} = \left(SRice_{kj} - 0.5\right) \times Close_j + 0.5$$

$SCRice_{kj}$ is then used as the Conscience Vote Cohesion statistic.

Appendix 4: Further Categorisation Examples:

> ***CIVIC – Rob Storey (National):*** *What reports, if any, has he received on community benefits from the introduction of speed cameras and tougher drinking/driving rules?*
> ***CIVIC - John Luxton (National):*** *There were 50 fewer deaths on the roads following the introduction of compulsory breath-testing; […] There is certainly good evidence that strongly suggests major community benefits from the introduction of speed cameras and tougher drinking/driving rules.*
> ***HEALTHCARE - John Luxton (National):*** *Hospital and long-term health costs associated with motor accident victims have been reduced, allowing the accident compensation levy, which accompanies the registration of cars, to be reduced by 30 percent or $40 for a motorcar […]*
> ***MONETARY - Michael Cullen (Labour):*** *Has he seen a statement by the Alliance spokesperson on social security that restoring benefits to levels applying before the 1991 benefit cuts would cost in the region of $200 million a year; if so, has he received any advice as to the accuracy of that figure?*

> ***MONETARY - Wyatt Creech (National):*** *[…] I have received the estimated net cost for the 1994-95 year of restoring benefits to the real level of March 1991, and that would be $515 million on 1 April 1994. […]*

The topic model successfully tracks how the exchange moves focus from Storey's question about community benefits from road policy – a civic issue – to Luxton's firstly on-topic reply in terms of community benefits and road usage, to his healthcare-related addendum about hospitals and health costs, to the next question from Cullen about benefits cuts posed in monetary terms[26], which is responded to by Creech also in monetary terms. Where the topic model mis-categorises speeches is usually where the speeches are especially short (which are removed from the dataset), where the speeches are in vague terms (meaning direct reference to a manifesto promise is unlikely anyway), and where the speech addresses multiple topics (which is unfortunate, but largely unavoidable, and would also be a hazard for a trained model).

Appendix 5: Alterations to AICBT Authorship Attribution

I adapted this code in several ways to make it appropriate for determining the influence of a (written) document over another (spoken) document, rather than direct attribution of authorship. I removed the lexical and punctuation features, as they were not appropriate for a cross-comparison between written manifestos and spoken speeches in parliament. Sentence length and quantitative breadth of vocabulary (lexical features) are fundamentally changed in a spoken, rather than written, form, and whilst useful for directly attributing authorship, they are not useful for detecting

[26] The label 'monetary' refers to economic dialogue in technical, monetary terms; not necessarily on monetary policy.

influence of one document over another. Punctuation features are even less relevant, as they are not direct features of speech, meaning the use of commas and semicolons are features of the transcriber's writing style more than the speaker's. Furthermore, the OCR extraction of text from the scanned manifesto documents, whilst accurate for words (and spell-checked), is less reliable for punctuation, meaning the reading of commas, full stops, colons, semicolons and quotation marks is far from perfect, causing problems for both the lexical and punctuation features.

Finally, I replaced the k-means clustering model used to create the categories with one which includes a minimum size constraint. In the original code, it was known there were only two authors, meaning two distinct clusters were expected, and it was known they were the sole authors, so the clustering would be strong. In this analysis, there are 45 different authors in the 44th Parliament and 36 in the 45th, and the prediction is of influence of one document's authorship over others through the borrowing of words and phrases. The one document (the manifesto) is also, furthermore, of a different form and in a very different context to all the rest. It is therefore likely to be isolated from the parliamentary speeches when clustering, because of parliamentary language and features of verbal speech creating similarity across all speeches. The case must be avoided, therefore, where the two clusters created by the k-means function are a single-document cluster containing the manifesto, and another cluster containing all the parliamentary speeches. A version of the k-means function where there is a minimum size constraint (Babaki 2020; Bradley, Bennett and Demiriz 2000) is used to coerce the clusters so that there are at

least 2 documents in each cluster (at the minimum, the manifesto and an MP predicted to be the most whipped by the analysis).

Chapter 3: Strategic Representation: the Responsiveness of Legislators to Constituency Interests in the United Kingdom

Abstract

What positions legislators take when they speak in parliament is vital information for understanding their degree of accountability, their strategies, and what can cause their positions to change. I develop a method for measuring the positions they adopt in parliament using their speeches as data, and use those measurements to analyse their responsiveness to their constituents' positions, and the popularity of third party threats in their constituency. My method involves using a large language model (LLM) as a coder to classify over 2.5 million statements in the UK Parliament as left- or right-wing on a 7-point scale, covering the period between the General Elections of 2015 and 2024. I improve on previous attempts at LLM coding in social science by introducing chain-of-thought methods, which improve the LLM's performance on complex tasks and enable it to interpret subtext and unspoken conclusions or implications of a statement. I also measure average political positions in constituencies by conducting an MRP using British Election Study data on left-right and libertarian-authoritarian values. I test for the effects of local opinions and party popularities on MPs' annual speech positions in two policy dimensions using fixed effects models with lagged dependent variables. I find that MPs are highly partisan, and not significantly correlated with their constituents' positions in general. Conservative MPs' non-economic positions are responsive to constituent positions if their seat is highly marginal. All MPs, however, are consistently responsive to the popularity of UKIP and the Brexit Party, adopting more right-wing positions as these parties start to achieve higher poll results and vote shares.

Introduction

Legislators advocate for diverse political positions when speaking in parliament. Deep factional divides and distinctive personalities with personal political platforms exist within parties. At times this diversity in legislators' views and interests can result in rebellion, but often they are latent: revealed through how they justify their positions or what questions they raise, rather than how they vote. There are, however, conflicting understandings of what motivates this diversity in legislators' positions. There is also little consensus on the prevalence of these different motives, or what conditions may affect them. Legislators' positions may therefore primarily be the product of the personal beliefs of the legislator. They may, however, be part of a strategic attempt to construct a platform which appeals to voters in their constituency. They may be genuine attempts to represent constituents' views. Understanding the prevalence of these motives, and the degree to which they are the source of individualistic behaviour in parliament, is vital for understanding the workings of representative democracy. I utilise LLM-coded legislative speech data and MRP models of constituency opinions to construct a panel model to measure the responsiveness of UK legislators to their constituents' views and voting preferences. I identify strategic behaviour using an interaction effect between constituents' views and expected seat marginality, as well as comparing responsiveness to views to responsiveness to voting preferences.

Speeches in parliament provide rich information about a legislator's policy preferences, priorities, and group loyalties. Whether by pointing out a bill will help balance a budget as a key priority, arguing it will help deliver core manifesto

promises, or declaring it neglects the interests of marginalised communities, legislators express diverse political preferences. In doing so, they construct a personal platform of political positions, and influence ongoing debates towards preferred outcomes. Selecting these positions and outcomes entails prioritising competing interests: between the legislator's own views, those of their constituents, and those of their party and leaders. What I analyse in this chapter is, in line with this, not necessarily the legislator's personally held beliefs, but the beliefs they represent when they are in parliament. This is a set of beliefs selected based on competing strategic incentives, personal interests, and loyalties, which together construct a version of themselves as a legislator. This version of themselves is what matters for predicting legislative behaviour, and is the only version which is useful for identifying their commitment to different principals[27]. Many theories exist to account for how legislators might navigate these competing interests, but so far methodological limitations have prevented satisfactory quantitative testing of these theories.

I develop a new method using a large language model (LLM) as a coder to quantify the political positions of speeches in parliament, allowing speeches to be used as a new source of data on legislative behaviour. This approach to analysing political positions based on natural language data has already been applied extensively to the study of parties, by treating the text of manifestos as indicators of party positions, and

[27] I don't expect that MPs often change their own personal beliefs to match those of their party or constituents. Though it can, of course, happen (especially through processes such as socialisation into partisan culture, or perhaps hearing a moving account from a constituent), there are often purely strategic reasons for MPs to claim to support, and to vote for, policies they do not personally believe in. Measuring how they behave, regardless of what they 'truly' believe, therefore, is key to analysing legislative behaviour.

typically relying on expert coders to quantify political leanings. My thesis contributes by providing a way to apply natural language methods to much larger corpora of political text, so that the positions of individual legislators can be identified. Analyses of legislators' speeches, tweets, and public statements are becoming more popular, but existing methods have critical shortcomings which I resolve, enabling not just more valid, but more complex, robust, and interpretable analyses.

Legislative behaviour is an expansive literature, and exists as a sub-literature within many others. For quantitative analysis, however, the data relied upon has been low resolution: based on rebellious voting, committee memberships, quantity of speeches or written contributions, and only rarely analyses of the content of speeches. Rebellious votes are rare, and the amount of information each instance provides is too small, for confident conclusions to be drawn even about simple hypotheses, let alone deep and nuanced theoretical frameworks. Speeches are a richer source of data, but only relatively recently have more advanced computational methods been adopted in social science, and these methods are always improving. I contribute a new implementation of LLM coding methods which enables a greater degree of accuracy. I use this to contribute new insights into legislative behaviour by quantifying the effects of partisanship, constituent opinions, seat marginality, and third-party threats on legislators' positions in parliament.

Literature Review and Theory

Theoretical Background

Accountability is often modelled as a restraint on legislative behaviour. Legislators are imagined as autonomous, but constrained by re-election threats into adopting vote-seeking behaviours. Despite their own aims, therefore, they must serve constituency interests insofar as they may lose the next election if they fail to satisfy their constituents. Strøm (2003) argues that 'fear of electoral punishment is a strong incentive for incumbents to remain in tune with their voters' demands.' They furthermore argue that re-election threats can have indirect effects even on legislators who intend to (or must) resign their seat at the next election, due to pressure from their parties which also depend on re-election prospects in their seat. This relationship is fundamental to their model of representative democracy, which is treated as a chain of delegation from voters to the executive, in which legislators' accountability to voters is the first link. How strong this link is, though, and under what conditions it may be strengthened or weakened, remains a question without satisfactory quantitative answers.

Strøm recognises an argument from Palmer (1995) that there may be an alternative framework, in which the realities of voter behaviour mean that parties, rather than candidates, operate more like the direct principals of voters. Voters' choices are plausibly based more on party reputations than individual candidates', and parties (or party leaders) are therefore the agents which have to respond to public opinion in order to secure re-election chances. In this interpretation, a legislator is indirectly accountable to the public via the party, who serves as their main principal. This power

is exercised through powers over candidate selection, and other whipping resources. Often, however, constituency interests can conflict with national-level agendas. In this situation, a legislator has 'competing principals'. Understanding the accountability and responsiveness of individual legislators to their constituents, therefore, is key not only to understanding the theoretical 'first link' in representative democracy, but to understanding what happens when legislators have competing principals.

Carey (2007) has sought to answer this question using voting records to compare how different institutional arrangements affect the power of parties over legislators. They find that federal systems and open lists are both sources of disunity in parliamentary parties, suggesting that they weaken the role of parties. Their method, however, cannot identify whether this weakening of accountability to the party is the result of or coincident with the strengthening of accountability to constituents. It may be that these systems are simply less restrictive of legislative behaviour, and therefore enable legislators to act more individualistically on their own interests, rather than in the service of constituents. This could be because the party simply has less power over selection. Closed-list electoral systems give the party full control over the relative re-election chances of each of their candidates through the power to rank order them. Federal systems often also involve decentralised candidate selection procedures (Hazan & Rahat, 2010), which may create a barrier to central party intervention in (de-)selection, which may not necessarily be sensitive to a candidate's record of constituency service[28]. It could, however, be the result of increased local power over

[28] As an institution, local party offices may simply be resistant to central party interventions of all kinds, as they are territorial over their remit and authority. An attempt by the central party to de-select

selection and re-election. For open list electoral systems, the ability of voters to prioritise certain candidates over others gives legislators an incentive to perform constituency service in order to defeat other candidates, including co-partisans. Relying on the party's reputation alone is not sufficient, and therefore developing an individual record of constituency service may be valuable. This is what is implied by the theory of the 'personal vote', in which candidates succeed in open list systems by differentiating from other candidates (Carey & Shugart, 1995). It is not, however, necessarily the case that candidates are incentivised to engage in constituency service in how they differentiate. Developing a personal platform and distinctive 'brand' as a legislator has pay-offs in the form of name recognition, even if that platform does not match the views of constituents. Likewise, in federal systems, differences in political interests between the local area and the rest of the country may be more salient in the minds of voters, and lead to greater accountability based on constituency service. Alternatively, local party offices with power over candidate selection may be more active in the use of their power, and be able to use their own threats of de-selection (or control over campaigning resources) to influence legislators' behaviour. It is not, however, necessarily the case that candidates are incentivised to converge on constituents' positions. Candidates may use local-level re-election threats as a tool to justify dissent based on false premises – e.g. citing non-representative interest groups in their constituency – as a way to deter the central party from punishing or censuring their dissent. Furthermore, local party organisations may not be representative of

a candidate as punishment may be resisted, declined, or scrutinised. This friction alone may reduce the amount that central parties attempt to exercise this power.

constituency interests, and therefore their increased power would not cause legislators to converge on constituency positions. More precise measurements of the positions adopted by legislators, therefore, is a necessary but missing piece in disentangling individualism from constituency service.

There are mixed findings on whether there are real electoral consequences for the behaviour of individual legislators. In US politics, legislators who are 'out of step' have been found to be punished at elections based on their roll call voting record. Canes-Wrone, Brady, and Cogan (2002) found that more extremely liberal Democrats and more extremely conservative Republicans received fewer votes than more moderate counterparts – confirming a 'median voter' theory of strategic incentives. Carson, et al. (2010) found similarly that legislators who are more loyally partisan are punished at elections, whereas those who rebel more frequently are rewarded – confirming more to a 'personal vote' theory of strategic incentives. Canes-Wrone, Minozzi, and Revely (2011) furthermore observed patterns of individual legislators being punished and rewarded based on voting records on individual issues, contingent on the issue's salience.

These findings, however, are often framed as surprising given a stable finding in political science that voters' level of political knowledge and attention is very low. Delli Carini and Keeter (1991) showed using US data from 1946 and 1989 that roughly half of respondents can identify the majority party in the Senate, the meaning of an economic recession, and the effect of tariffs. Approximately a third can name their state's senators, or their district's congressperson. Similar surveys in the UK have shown that less than half of respondents between 1991 and 2011 were able to correctly

name their MP (Parsons, 2024)[29]. This may mean either that the rewards operate through indirect effects (such as garnering support from the local party and activists who provide campaign resources), or that the small portion of the electorate who would be sufficiently engaged to notice MPs' activities are also highly volatile in their voting behaviour. Whilst sorting effects have been found to be far stronger among more politically knowledgeable voters (Andersen, et al., 2005), and survey experiments show that voters are responsive to candidate characteristics if the information is provided to them, even when the partisanship of the candidates is also provided, at the local (Munis, 2021) and national (Flavin & Law, 2022) levels. Though it's likely they rely on partisan generalisations, there is evidence that voters can correctly predict their representative's positions on important contemporary bills being voted on, that they hold their own views on the issues, and that their approval ratings of their representative and their electoral behaviour are responsive to disagreement between themselves and their representative's behaviour (Ansolabehere & Jones, 2010).

Regardless of whether incentives are real, however, legislators also have imperfect information, and may act as though they are real anyway: they may believe in individual-level consequences or in the importance of their representative role. Cain, Ferejohn, and Fiorina (1987) suggest that regardless of whether there are 'real'

[29] Based on phone surveys. Figures vary, but multiple choice and online formats suffer from problems of guesswork and cheating. For example, the British Election Study's question on MP names is online and provides five fake names as options (Cowley, 2014). This means a respondent may just know the gender and ethnicity of their MP, and this can often be enough to narrow it down to a single option. Furthermore, a respondent may be able to recall the names of candidates in their constituency from campaign materials they received, but not know who won.

incentives, legislators may still behave strategically or idealistically. They find, by surveying them, that legislators and staffers believe that constituency service and personal reputation are of great importance in the US (and in the UK, subject to seat marginality). 34 years later, Hanretty, Mellon, and English (2021) find that legislators' opinions have changed. In their survey of legislators they find that they believe voters are minimally responsive to their own positions as a candidate – in their case, even on the contemporarily salient issue of Brexit. The question posed, however, differed. Cain, Ferejohn, and Fiorina looked broadly at the effects of 'constituency work' and effort. Hanretty, Mellon, and English looked specifically at the effect of adopting one issue position or another – and so does not include effort-signalling or symbolic acts of constituency service. It is therefore Hanretty, Mellon, and English's analysis which is most applicable to the question of whether legislators are accountable for their substantive record as a representative of their constituency in parliament[30]. It may, therefore, be the case that MPs engage in effort-signalling behaviours strategically, but do not change their issue positions in response to constituency interests. This was the pattern of behaviour identified by Fouirnaies and Hall (2022) in the US. Modelling strategic incentives using term limits, they found that legislators seeking re-election are more likely to engage in effort-signalling activities, but not likely to be any different in their political positions. Meyerrose and Watson (2024), however, did find that French senators' political positions were responsive to indicators of radical

[30] Though they only use their pre-Referendum position and therefore don't capture post-Referendum adjustment, in theory their pre-Referendum position ought to have been informed by the interests and views of their constituents. Dissonance between their pre-Referendum position and their constituency's vote is therefore a failure of responsiveness. Likewise, even if an MP changes their position post-Referendum, to the extent that constituents respond to it, they may still resent their MP for the side they took in the Referendum, and may not learn of their change in position.

right threat in their constituencies. Butler and Nickerson (2011) also found in the US that legislators who were sent opinion data about their constituents' views on an upcoming bill were more likely to vote in line with those views. There may, however, be a degree to which MPs respond to constituency opinions regardless of potential strategic payoffs, simply based on normative beliefs about their role as a representative.

A further possibility is that constituent interests serve not as a strategic restriction, but as a strategic tool for legislators. Based on Palmer's (1995) version of the delegation model, in which parties are the main principals of legislators rather than voters, competing principals can provide a way for legislators to bargain for more individualistic positions. Based on Tsebelis' (2002) modelling of legislators as strategic actors in creating legislation, issue positioning can provide strategic advantages, by making specific actors more pivotal and increasing their bargaining power. Where constituent interests and legislator interests align, therefore, the legislator may use either normative arguments about constituency representation, or potential re-election threats, to justify a position against the party line. This may either require the party to accept amendments or otherwise change the position of the party line in order to win the support of that legislator, or may accept the need to allow that legislator to rebel in order to preserve the party's shared interests in that legislator's re-election. The latter has payoffs for the legislator even in lieu of affecting legislation, as records of rebellion help build a personal platform (Carey & Shugart, 1995). Even if a personal platform has minimal effects on their appeal to voters, it may have effects on lobbyists and interest groups, or on factional politics within the parliamentary party. If this is the role of constituency interests in legislative

behaviour, this would cause a legislator to appear to be pursuing constituency representation, unless the personal interests of the legislator are controlled for. Representative behaviour which is not conditional on electoral strategic interests, therefore, is not necessarily founded on idealism or normative roles, but may also be a strategic behaviour in response to a different principal.

Strategic incentives may differ based on issue topics. Gormley (1986) and Culpepper (2010) theorise that government accountability varies by issue, due to the dynamics of public attention and intellectual resources. The public is equipped to hold governments to account when an issue is high-salience, and low-complexity, as these are the issues where they can, en masse, identify that a policy is opposed to their preferences and change their votes in response. Where salience is low, too few voters will change their votes. Where complexity is high, criticisms can be subverted through a large number of strategies, including claims of budgetary limitations, designing policies to be symbolic of a position but minimally effective, and deploying expertise to create doubt about factual questions. These make it difficult for journalists and citizens' organisations to communicate failings. Gormley and Culpepper's theories expect increased power of lobbyists and decreased accountability of governments on these issues. In principle, these arguments may apply to individual legislators, too. They may expect threats to re-election based on their positions on highly salient and minimally complex issues, but expect to be able to exercise more freedom on minimally salient and highly complex issues. Economics is fundamentally an issue where public understanding is low, and appeals to expertise and competence are dominant, whereas social issues are more likely to involve strong values judgements, in which resource constraints or practical necessity are unlikely to

convince the public of a policy they disagree with being unavoidable. It may be expected, therefore, that legislators exhibit more strategic behaviour on social issues than economic.

Strategic behaviour may also be responsive to the positions of rival candidates. For example, at the party level Adams and Somer-Topcu (2009) found that parties adjust their platforms leftward or rightward in response to the movements of their rivals (and especially to rivals in their own 'ideological family'). This is taken as evidence that parties adjust their policy platforms in search of a 'Nash equilibrium': a stable position in policy space to capture votes in response to the positions of other parties. Abou-Chadi and Krause (2018) found similarly that mainstream parties shift to the right in response to the entry of radical right parties, suggesting that relative popularity of third parties can affect the positions of mainstream parties. Green parties, however, were not found to have any similar effects on mainstream party positions (Abou-Chadi, 2014). In theory, similar response patterns may be observable at the level of individual legislators. When third parties emerge in their constituencies, they may adjust their platforms towards their new competitors as a way to compete for those voters.

I derive from the existing literature three hypotheses about how legislative behaviour relates to constituency opinion. These hypotheses are not mutually exclusive: it is plausible that all are observable, given that different legislators may have different strategic perspectives, and individual legislators may have mixed patterns of behaviour. I also expect that they may apply differently based on issue topics.

Legislators may act more strategically based on electoral incentives with respect to social issues than economic ones.

The first hypothesis is labelled 'General Responsiveness' (H1). It hypothesises that a legislator is responsive to constituency opinion unconditionally. This may be either due to idealised or normative ideas about their role as a representative, or through the theory that they are primarily accountable to their party organisation, and constituency interests provide an opportunity to dissent, rather than acting as a constraint that forces them to dissent.

H1 (General Responsiveness): MPs adopt positions similar to their constituents' opinions.

The second and third hypotheses are the 'Conditional' hypotheses, and may conform to either of two patterns. Marginality-conditional responsiveness (H2) hypothesises that a legislator will adopt positions closer to their constituents' when their seat is marginal, based on an expectation that acting as a 'better' representative will have electoral payoffs, or based on a 'median voter' theory of vote-winning.

H2 (Marginality-Conditional Responsiveness): MPs are more likely to adopt positions similar to their constituents' if their seat is more marginal.

Party-conditional responsiveness (H3) hypothesises that a legislator will move left or right based on the popularity of left- or right-wing third parties in their constituency, based on the perception that those voters may be more mobile, or taken as evidence of opinion shifts in the constituency given that election results (and vote choice polls) are the most accessible sources of information to a legislator.

H3 (Party-Conditional Responsiveness): When third parties become popular in a constituency, MPs adopt positions in order to win (back) those voters.

Existing Methods

Most classical empirical analyses of parliamentary behaviour are dependent on voting records. Even in parliaments with semi-frequent rebellions like the UK or USA, the best analyses possible with this data are weak. They're limited by not being reproducible in most other countries where backbench rebellions are far rarer. They're also limited by rebellious voting being rare even in the UK, creating sparse meaningful data. Finally, they're limited by a lack of richness to the data which at best harms interpretability and at worst harms validity.

When the only data is whether an MP voted with the party line, against, or abstained, this provides limited information about rebellious behaviour. Abstentions can be due to illness, prior commitments, or simple disinterest in many cases; but in high-stakes bills can be strong, meaningful, and highly damaging acts of rebellion; yet these are treated the same way in voting data analyses. Likewise, votes against the party line are differently substantial depending on the stakes of the vote, and only one form of dissent. Criticism at the committee stage, along with demanding and negotiating amendments, and of course speeches in parliament are all ways to dissent[31]. Rebellious voting is not only more extreme, it occurs in unique circumstances: where an individual MP is too ostracised or powerless to resolve their grievances through

[31] Garner and Letki (2005), analysing why Canada has much lower rates of backbench rebellion despite having a similar institutional setup to the UK, find that Canadian MPs are far less likely to report that they believe they're powerless to influence policy. Their lower rates of rebellion are therefore not necessarily due to higher cohesion or loyalty, but that they dissent through other, more effective, avenues.

changes to the legislation, or where there are deep factional divides which prevent 'peaceful' conflict resolution. It therefore involves a selection bias which only captures dissent if it is organised behind an entrenched faction, or targets a high-salience policy issue. All of this makes patterns in rebellious voting very difficult to interpret, as dissent among party members could be forcing government to give up ground at other (less visible) stages of the legislative process. Only a very specific type of dissent is captured, and patterns could represent anything from increased boldness of MPs or weakness of government, to polarised factionalism, to increased stubbornness of government to accept changes at early stages of the process, to simply changes of the types of topic on the agenda. Votes therefore cannot be appropriately used to measure political positions or partisan cohesion, as they capture just one pattern of dissent.

As well as providing incomplete information about aggregate patterns in dissent, they provide minimal information about individual MPs. Most MPs rarely ever rebel. Figure 3.1 presents data on rebellions from The Public Whip (The Public Whip, 2022) The mean MP rebels 0.9% of the time, and more than half rebel less than 0.8% of the time. In the average 5-year parliament, there are around 1000 divisions, so this means most MPs rebel fewer than 8 times in 5 years. This means having one data point on one policy topic once or twice a year. For any given MP this is insufficient to determine how exactly their policy preferences or political positions deviate from the party line, except in countable specific circumstances where clear patterns can be identified. This means, even where patterns at the aggregate level can be spotted, the mechanisms at the individual-level can't be effectively tested. This is why use of voting data has been valuable in the study of US legislator behaviour based on a fairly robust assumption in US politics specifically that rebellious legislators are centrists (Poole & Rosenthal, 1985; Poole & Rosenthal, 2017). This assumption, however, doesn't hold true in other countries, with prominent rebels in the UK, for example, often being from extreme factions within the main parties, including Eurosceptic factions in the Conservative party and the more radical left in the Labour party (McClean & Spirling, 2007). This can therefore lead to incorrect conclusions when such methods are applied outside of the US context. Methodically determining whether such assumptions are valid before applying a roll call method is also difficult without being able to measure empirically legislators' political positions. This is something my

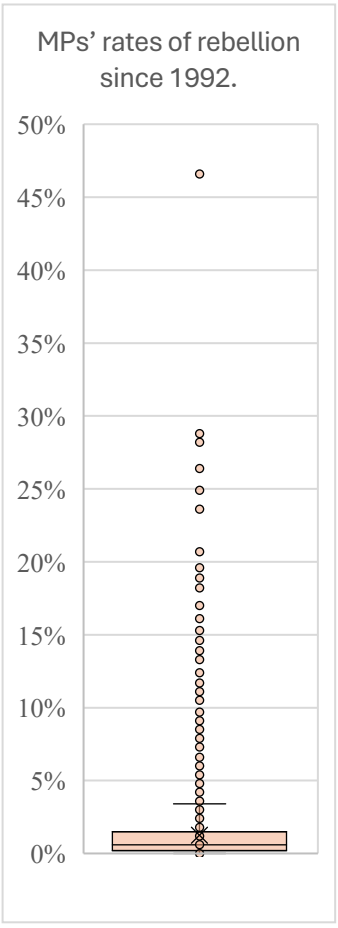

Figure 3.1: Box plot of MPs' rates of rebellion

method will enable researchers to do. My method will identify not only which MPs deviate from the party line, but in which directions their political positions and policy preferences diverge.

Even what findings can be derived from voting patterns are limited to a handful of more rebellious parliaments, including the UK and USA. Cross-country comparative analyses on rates of rebellion are, in part for this reason, rare. Comparisons are also difficult because there is large variation in what is being measured. Not all parliaments have public records of how legislators vote: sometimes votes are anonymised, or only accessible by requesting specific votes individually (Rich, 2014). Furthermore, voting procedures can differ in other important ways. Parliaments have different ways of determining whether a roll call or division of the house should be held. Most hold votes by voice, hand-raising, or declarations by party whips, and if there's a 'clear' result from that, determine outcomes of votes without recording individual members' votes. There may therefore be rebellious actions in these votes which go unrecorded. This is more systematically overlooked in systems which are less likely to go to roll call votes or divisions. Likelihood to go to roll call votes or divisions can have correlates: governments with narrower majorities are less likely to win clear votes from voice voting or hand-raising, meaning more recorded divisions (Hug, 2009). Furthermore, recorded divisions and roll calls are correlated with disunity under a substantial majority, but correlated with unity under a narrow majority. These complex selection effects can harm validity and interpretability even in within-country analyses, but especially create a barrier to cross-country comparisons.

Language models are a relatively recent entry into quantitative studies on legislatures. The most basic models which handle speeches as data simply look at participation: whether an MP spoke in a debate, and nothing about the content of their speech (Proksch & Slapin, 2015). Robust developments on this include analyses based on non-context-specific features of the speeches, such as Slapin and Kirkland's (2020) analysis of speech length, language simplicity, and occurrence of first-person pronouns. These features are clearly interpretable, simple to measure, and have clear links to their outcome variable of rebellious voting. As such they find evidence that speaking frequently in the first-person, with simpler language and longer speeches: emphasising (collective) independence and activeness, communicating their stance clearly with a wide intended audience, and justifying their stance at length. These findings are strong and plausible, but have limited use. Voting records can't be substituted with these variables to measure rebellious behaviour. This is because the mechanism implies they are only expected if they coincide with a subsequent rebellious vote: as a person speaks in this way in order to talk about their decision to vote against the party. These variables also can't be used to estimate political positions or policy preferences.

The majority of attempts to analyse content of parliamentary speeches are based on word frequency models. These models have limited power. They also suffer from being either too general, or on the other extreme from problems of endogeneity. A pre-built model called 'wordfish' (Slapin & Proksch, 2008) is the most common in the literature (Lauderdale & Herzog, 2016; Schwarz, et al., 2017). It has limited power because it is a unigram model, and suffers from endogeneity because it effectively uses the test data as training data. A similar unigram model called 'wordscores'

(Laver, et al., 2003) avoids the latter problem only in part by using training data, but introduces problems of bias from the selection and quality of that training data.

As unigram models, wordfish and wordscores only look at words in isolation of context, whereas 'n-gram' models look at *n* words either side of a word to appreciate words in context. By way of example, it might be the case that 'reduce taxes' is common in right-wing texts, whereas 'increase taxes' is common in left-wing texts. A unigram is unlikely to associate any of the words 'reduce', 'increase', or 'taxes' with any political leaning in particular[32]. A bi-gram model, however, can make use of 'reduce taxes' as an indicator of right-wing politics and 'increase taxes' as an indicator of left-wing politics. Wider context windows allow for more nuanced interpretation of context (with use of negatives and double-negatives of particular concern), but at the expense of introducing noise. In general, so long as there is enough data to support the training of such a model, a larger n-gram window is beneficial to the quality of the estimates it outputs. A lack of quality training data is the key reason why unigram models are the most popular approach still in analysing political texts.

Wordfish also does not use training data. It works by assuming a unidimensional statistical distribution in the frequencies of words across the texts, and that the basis of that distribution is political position. This creates problems of endogeneity. It can identify trends within the sample, but it can't be used to compare beyond the sample, because each sample the model is run on measures frequencies of different words, and

[32] And if it does, this becomes an even bigger problem.

weights those words differently. The scale of the distribution is also local to the sample, which makes magnitudes difficult to interpret: polar opposites in political position estimates may only represent very small real shifts in political position. Furthermore, it is hard to know what is influencing distributions, and much care must be taken to analyse the words being used (and not used). Factors such as Southern-ness or Northern-ness of dialect, education, ethnicity, and class background could all affect language use on a unidimensional scale. These, and plenty of other variables, can all be correlates of political lexicons, and therefore by confounding the relationship between politics and word choice, distort the wordfish model. It's therefore difficult to interpret, to validate, and to generalise and compare.

Wordscores, though it does use training data, still suffers from the above problem, because it treats the 'reference texts' in the same way that wordfish treats its test data. Once the 'reference texts' are provided, the wordscores method assumes that there is a unidimensional statistical distribution in the frequencies of words between the reference texts, and that the basis of that distribution is political position. It then assigns values to the words based on their relative frequency in texts at different positions on that scale based on the scores the researcher has given to the 'reference texts'. Those scores are then applied to the test data, and aggregated to assign scores to each text in the test data. These scores can therefore be affected by unmeasured variables other than political position if they are correlates of the political division that the researcher intends from the reference data. Given the training data for wordscores is by definition more limited than the training data for wordfish, it is also even more likely that these confounding variables are present, as it is even more likely that there are systematic unobserved differences between the documents. For

example, if one Labour and one Conservative manifesto are used as reference texts, it is more likely that they happen to differ in the general writing style (e.g. formality, and many other style-markers of authorship), the gender, or the class background of the author (or their intended audience). Furthermore, whereas it is possible that these variables are correlates of political position, and therefore that in practice capturing them in how the model is trained to recognise political position coincidentally improves the model, with a smaller training set, there is a smaller sample size of authors, which increases the probability that it captures variables not correlated with political position. Even with larger training sets, any such real correlations will not be perfect, and will lead to inaccurate results.

Lauderdale and Herzog (2016) use a model they call 'wordshoal'. This model is an implementation of wordfish, and suffers from most of the same problems. Their deviation from a simple implementation of wordfish is that they appropriately apply wordfish only within each debate (given each debate in theory has its own political axis), but then solve the lack of cross-comparability they would otherwise have using a Bayesian factor analysis. Importantly, by doing this they find that the same words occurring in different debates are differently left- or right-wing contingent on the context of what debate they are said in, which is an important finding on both the weakness of unigram models, and the importance of treating individual topics appropriately, rather than estimating a unidimensional divide directly. They also show that time affects how left- or right-wing words are: noting how talk of preexisting medical conditions is very left-wing in 2009-10 at the start of Obama's first term, but becomes less partisan after the Affordable Care Act passed and consensus developed

over the issue. This reveals how the treatment of words by a model must be sensitive not just to the issue context, but also the temporal context of speech.

The 'sentiment approach' developed by Proksch et al. (2019) resolves a lot of wordfish's problems, but for the specific application of measuring polarisation and consensus in individual debates. They state that by using other approaches, the level of analysis is 'policy dimension, rather than the policy issue being debated in the legislature.' For my analysis, however, precisely what I intend to measure is positions on policy dimensions overall, rather than agreement with specific pieces of legislation being debated. Pure sentiment analysis is therefore not appropriate for my purposes. By using a larger context window than unigram models, however, my model does leverage *both* the semantic content of what a speaker talks about, and the sentiment with which they talk about it. This makes it capable of the best of both approaches for this analysis, as it captures not just what a speaker decides to talk about, but what the way they talk about it says about their position on that policy.

The unsupervised methods discussed also yield general left-right positions, and cannot be used to identify specific axes other than left-right. For the analysis of partisan loyalty, a measurement of adherence to the party line is needed. Within each party, this can be correlated with left-right position, but especially if the party leadership's position is towards the centre of the party's left-right divide, it may not be. An approach tailored to the measurement task at hand, and with training data which ensures the correct concept is being measured, is therefore vital.

Methods

Using a large language model (LLM) as a coder, I create a coded dataset of Hansard speeches which is used to construct the dependent variable: an MP's annual average political position on left-right and libertarian-authoritarian dimensions. Below, I outline how this is constructed, and I explain and analyse the methods involved. I then provide a brief description of the Hansard dataset, before outlining the selection, construction, and sources of the independent and control variables. I conclude with the model specifications, and a summary of expected results with respect to the hypotheses.

Speech Dataset and Dependent Variables

I create a new dataset of statements in the House of Commons from 2015 to 2024, coded by their left-right position and the topic of the debate. This dataset is used to measure MPs' annual left-right position on an economic and a social dimension. Below, I first outline how this dataset is created by using an LLM in the role of a human coder. I justify why this method is possible and robust. I also introduce chain-of-thought methods, which have been absent from existing social science applications of LLMs, and greatly improve the accuracy of their outputs. I then briefly present the dataset created by this process. This is followed by an explanation of how the main dependent variables are derived from the dataset, along with how the independent variables are selected, sourced, and constructed. I then outline the model specifications for the hypothesis tests, and the expected results from the models based on each hypothesis.

I leverage an LLM as a substitute for expert human coders in order to measure the political positions of legislative speech. I develop a two-stage implementation of Llama 3 8B (run locally via Ollama) to determine the political position of sentences spoken in parliament. In the first stage, the LLM is used to generate a long-form explanation of the political position of the given speech in terms of how left- or right-wing it is. In the second stage, the LLM is given the explanation generated in the first stage, and used to summarise the text's conclusion using one of seven codes, ranging from extremely left-wing to extremely right-wing.

Comparisons of LLMs and human coders have found that LLMs are highly accurate, nearly matching the performance quality of expert coders (Heseltine & von Hohenberg, 2024). Ziems et al. (2024) find that LLMs can 'out-perform' humans in 'free-form' generation, explanation, and summarisation tasks[33]. Le Mens and Gallego (2023) find that zero-shot classification of political documents aligns with benchmarks based on expert codes, crowd-workers' codes, and positions derived from roll-call votes. Their analysis spans US, UK, and European contexts, as well as text formats including tweets, manifesto documents, and policy speeches. They develop a method based on averaging the estimates of multiple models, which they find is more accurate than methods based on supervised models.

I implement a method based on leveraging the strength of LLMs in generative tasks and text summarisation. Generative AI computes by evaluating the quality of its own

[33] This was done by having field experts evaluate the quality of annotations (and their explanations) from both humans and LLMs, so that the 'gold standard' annotations are also evaluated on the same metric as the LLM outputs.

responses. Longer responses therefore provide more data by which to evaluate its own quality. Complex tasks, therefore, have been found to be completed with greater accuracy when LLMs output more text, and this observation has led to the development of 'chain-of-thought' methods in the application of LLMs (Nye, et al., 2021; Wei, et al., 2022; Kojima, et al., 2022). I therefore do not prompt the LLM to simply output codes to classify the texts' political leanings, as existing applications of LLMs to social science coding exercises do (Ziems, et al., 2024). I instead prompt it to both classify the political leaning, and explain why. I then extract the 1- to 3-word codes by providing that explanation in a second prompt requesting a summary of its conclusion selected from seven possible answers, ranging from extremely left-wing, to centrist, to extremely right-wing.

This method allows the LLM to correctly identify political leanings based on interpretation of complex semantics and subtext, including leveraging contextual information, similar to the role expected of an expert coder. This is a substantial development over word frequency, n-gram models, which can struggle even due to the introduction of a negative. For example, in an n-gram model a speech for and a speech against 'levelling up', if both mention the scheme by name, risk being both categorised as right-wing due to the use of a phrase that is primarily part of the Conservative platform and lexicon. This becomes even more hazardous when criticism of the policy doesn't take the form of a simple negation: speeches against 'levelling up' can often appear to support the scheme.

Robert Halfon, Conservative
Year: 2022
Debate: Levelling-up and Regeneration Bill
Statement: "Harlow, as she knows, has a fair bit of deprivation, but under this Government it has been levelling up for the past 10 years: an advanced manufacturing centre, millions of pounds; an enterprise zone, millions of pounds; a new hospital coming, hundreds of millions of pounds; a new road junction on the M11 just about to open up, many millions of pounds; infra-structure improvements; a technical school opened up; and a £23 million town fund."

Explanation:
"Based on this speech, I would describe it as leaning towards being right-wing in a 2022 UK context. Here's why:
* The focus is on economic growth and investment, rather than social welfare or public services.
* The emphasis on "levelling up" is often associated with the Conservative-"
Summary:
"Strongly Right-Wing"

Example 1

It is also a potential improvement over BERT-class and non-LLM transformers, as it can leverage the large database of training data not just to identify statistical correlations between encoded tokens, but by generating text, to explicitly draw links between pieces of information. It can therefore connect praise for a scheme such as 'levelling up', to Conservative Party promises, and therefore to right-wing political

positioning[34]. This type of pseudo-logical drawing of links is unique to large language models, and makes them appealing for complex annotation tasks.

I select some speeches on this topic and how they were coded by the LLM as illustrative examples below.

The first example is a speech by Robert Halfon from 2022. The LLM is prompted with the following, followed by the text of the statement:
"Provide a brief description of how left- or right-wing this parliamentary speech would be in a [Year] UK context."
The LLM therefore provided with the year, the context that it is a parliamentary speech in the UK, but no information about the speaker or their partisanship. This prompt yields the explanation text, which is limited to 40 tokens for processing efficiency, and therefore sometimes cuts off mid-sentence[35]. The explanation generated is then provided to the LLM after the following prompt:
"Summarise the conclusion of this text with one of the following statements, and no explanation: Strongly Left-Wing, Centre-Left, Centre, Centre-Right, or Strongly Right-Wing"
This yields the 1- to 3-word codes that are used to quantify the political leaning of the

[34] A potential hazard of this method is that LLMs may have access to Hansard transcripts as part of their training data, and so may 'know' who gave a speech and what their political positions are. However, given that LLMs draw these connections by generating sentences from one word to the next, if it ever used this kind of 'reasoning', it would be evident in the text it outputs. No such reasoning was evident in the dataset produced using my method. The LLM would also only succeed from this kind of process with particularly high-profile MPs, where their positions are well-known, and have been discussed online, publicly, and explicitly with reference to their left-right positioning. This does not apply for a majority of MPs. It would also be only effective at the MP-level, and unlikely to be sensitive to change over time. The fixed effects model design (detailed below) would therefore be resistant to any such biases, as within-MP change over time is what is analysed.
[35] The model also has a temperature of 0.8.

text. The output is limited to 5 tokens to allow for instances of full natural language repsonses such as “Yes, it is strongly right wing.”

In the first example, Halfon is responding to criticism of the ‘levelling up’ scheme. The LLM is able to identify that levelling up is a Conservative partisan phrase. It also makes use of what goes unsaid in the speech – identifying the lack of references to social welfare and the focus instead on investment in industry as evidence of a right-wing perspective. The second example is from Louie French – also Conservative – and despite being possibly interpretable as critical of the policy by emphasising a potential omission from the policy, it correctly identifies the reference to the policy as indicative of general support. By contrast, the third example also references the policy, but is disparaging of it, and the LLM correctly identifies the critical tone.

Louie French, Conservative Year: 2022 Debate: Levelling-up and Regeneration Bill Statement: “Levelling up the country must not forget areas in the south-east such as Bexley, which does not have the infrastructure of inner London but is seeing its population dramatically increase and never gets its fair share of funding, whether grant funding or health spending per head” Explanation: “Based on this speech, I would describe it as centre-right to right-wing. Here's why: * The speaker is using the term "levelling up", which was a key phrase coined by Boris Johnson during his 2019 Conservative Party leadership campaign and has been associated with Conservative policies since then.” Summary: “Strongly Right-Wing”	Lisa Nandy, Labour Year: *2022* Debate: Levelling-up and Regeneration Bill Statement: “In clause 5 is a measure that allows the Government to tear up those missions on a whim: their entire levelling up agenda, the promise made to the people of Britain and on which they won the last general election, presumably when they fail to deliver every single one.” Explanation “Based on this speech snippet, I would estimate it to be from a left-wing or opposition MP in a 2022 UK context. Here's why: * The tone is critical and accusing, suggesting that the government has made empty promises (the "levelling up agenda") and is now trying to” Summary “Strongly Left-Wing”

Examples 2 & 3

I separate speeches into individual sentences. This is firstly because a long speech could articulate multiple differing political positions. It is also because LLMs perform better in response to shorter texts. When presented with longer texts, they can bias the information at the start and end of the document, neglecting information in the middle, leading to inaccurate outputs (Wu, et al., 2023). Separating speeches into individual sentences, therefore, helps to preserve all data provided by the speeches.

The LLM coding method yields a dataset of 2,625,240 statements in Parliament from 2015 to 2024, coded by the political leaning of the statement and categorised by the topic of the debate they were speaking on. Figure 3.2 shows how MPs of each party are distributed on the economic and social dimensions in each of the three Parliaments covered in the period. The two main parliamentary parties are strongly divided in both dimensions, forming two distinct clusters. These clusters persist throughout the three parliaments studied. This division is deeper on the economic axis than the social axis: there is very little overlap between the economic positions of Labour and Conservative MPs, but a sizeable overlap on social positions. SNP and Liberal Democrat MPs are not included in the analysis of this chapter, but have very similar distributions to the Labour Party.

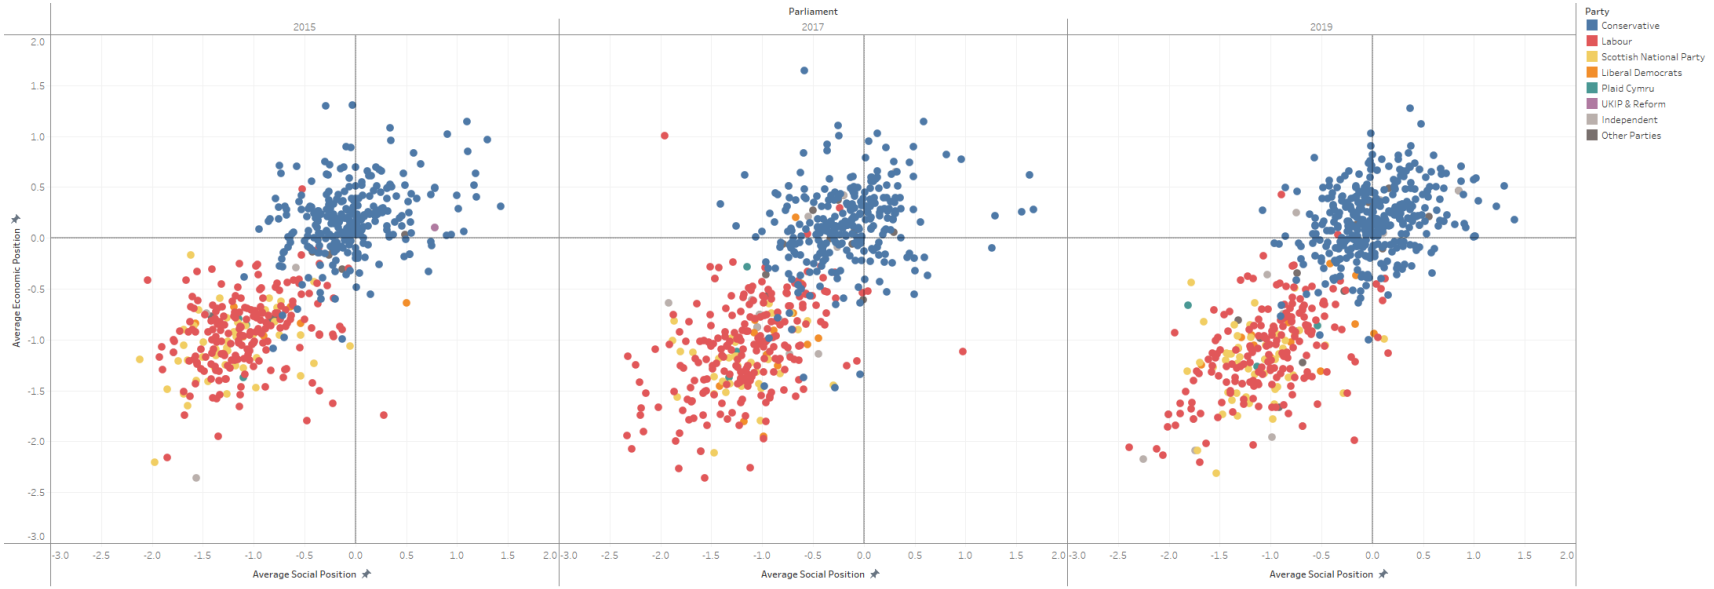

Figure 3.2: Estimated Social and Economic Positions of MPs across 3 Parliaments from 2015 to 2024

Once the speeches, acquired from Hansard transcripts, are coded by the two-stage LLM process, this provides a 7-point numerical scale for all speeches which I use as the dependent variable. For modelling purposes, these are aggregated into annual estimates for each MP. Independent variables are sourced from the British Election Study (Fieldhouse, et al., 2024) for constituency-level public opinion data, as well as demographic data from the UK Census. Control variables are sourced from the Public Whip (2024) for data on MPs' party memberships, and ministerial roles.

Whilst no gold standard exists for this corpus of parliamentary speeches – the new LLM method has enabled it to be extensively coded for the first time – the resulting average positions of MPs can be compared to existing coders' ratings of MPs individually. For example, Hanretty and Lazarov (2025) measured left-right positions of MPs using a survey of councillors. Treating the ranking of MPs who are coded in both datasets (the 2019-2024 parliament) as an ordinal variable, the Krippendorff's alpha (Krippendorff, 2022) (a measurement of inter-coder reliability) for my parliamentary speech ratings and Hanretty and Lazarov's ratings is 0.78, just below the 0.8 threshold for 'reliable' agreement, but high given the radically different methods for rating MPs. Another source for left-right coding of MPs comes from Gaughan (2024), who validates their method (based on network analysis of Twitter followers) by comparison to a small expert survey rating the positions of 30 MPs. The Krippendorff's alpha between my rankings and Gaughan's Twitter follower-based rankings is 0.81, suggesting reliable agreement. Meanwhile, the Krippendorff's alpha between my rankings and the 30 selected MPs rated by expert coders is 0.90, suggesting a very high reliability on these high-profile figures. This suggests that our models may have less agreement when rating MPs who are lower profile. This could

be due to them having fewer speeches, reducing the reliability of my method, but could also be due to poorer accuracy of councillors when rating more obscure MPs in Hanretty and Lazarov's method, or Gaughan's method losing accuracy at lower follower counts. Importantly, however, even if other methods for rating MPs are more reliable, they are less sensitive to change over time. Councillors will not collectively observe changes over time in MPs' positions, especially for more obscure personalities. Twitter followers are unlikely to undergo large-scale replacement, unless there are major headlines about a specific legislator. My method is therefore not just reliable for the coding of MP's positions, but can capture changes over time, enabling the panel model I conduct in this chapter.

I also use Llama to categorise speeches into topics[36] based on the title of the debate, grouping debates into four topics: economic, social, constitutional, or international[37]. The latter three topics are condensed into a single 'social' axis given they all play some role in the conventional 'authoritarian-libertarian' divide. This distinction is used to determine MP positions on each issue dimension individually. I categorise topics using debates rather than individual speeches. This is because the context is usually more important than the statement. For example, emphasising social concerns during a debate on an economic topic is usually indicative of a (left-wing) economic position, as arguments in favour of redistribution often take the form of pointing to

[36] The prompt used for this is *"Is this debate about: Economic Policy, Social Issues, The Constitution, International Relations, or Unknown? Answer in no more than two words with no explanation"* at 55 tokens with a temperature of 0.8. The model is then given the title of the debate as an input. If the 'unknown', then it is then instead given both the debate title and the opening speech.
[37] The latter three topics are all captured by the 'authoritarian-libertarian' axis, but are easier to comprehend for the language model when treated individually.

the needs and conditions of public service users or welfare recipients, or to moral imperatives, regardless of economic costs. Raising these concerns in the context of an economic debate makes them a statement of the legislator's views on economic policy, even though a topic model would be likely to class it, in isolation, as a social speech. Conversely, emphasising costs during a debate on a bill to do with social issues can be indicative of a (right-wing) social position, as arguments against civil rights, humanitarian, or constitutional proposals often take the form of pointing to the costs or practicalities of implementing the proposals. Raising these concerns in the context of a social debate makes them a statement of the legislator's views on social policy, even though a topic model would likely class it, in isolation, as an economic speech. Without categorising at the debate level rather than at the speech level, this interaction between social and economic issue dimensions would cause a selection bias in the categorisation of speeches, where regardless of the topic of discussion, left-wing speeches get categorised as social, and right-wing speeches get categorised as economic.

Independent Variables

I measure the left-right and libertarian-authoritarian positions of constituencies by applying MRP methods to BES data. Constituency-level demographic data on gender, age, education, and ethnicity are sourced from the 2021 UK Census, and used as the post-stratification frame for the MRP model.

I measure the local popularity of UKIP, the Brexit Party, or the Reform Party, and of the Liberal Democrats, using their highest-yet popularity according to polling companies' MRP estimates of constituency-level popularities, and election results since 2015. The sources used include the 2015 General Election results, YouGov

polling from 2017, the 2017 General Election Results, YouGov polling from 2019, the 2019 General Election Results, FocalData polling from 2020[38], Find Out Now/Electoral Calculus polling from 2021 and 2022, and Survation polling from 2023. I also use these MRPs to provide annual estimates of expected seat marginality. I include opinion polls so that between-election fluctuations can be accounted for where available, and to avoid issues caused by the Brexit Party deciding not to contest Conservative seats in 2019[39]. I use data since 2015 due to constituency boundary changes implemented for the 2015 Election which make previous election and polling data incompatible with post-2015 seats. Polls are selected for each year where available, but availability is limited in earlier periods, as constituency-level estimates by major polling sources were not common. The increased availability of MRPs is a key reason why the period from 2015 was selected for this analysis.

Model Specification

The dependent variable, therefore, is a measure of annual average *MP left-right speech position* or *MP lib-auth speech position*, each constructed by aggregating their speeches on economic and social topics respectively. For independent variables, there are measures of *constituency left-right position* and *constituency lib-auth position* based on an MRP model of BES data, as well as constituency-level percentage-point *margin of victory*, *Liberal Democrat highest-yet popularity*, and *UKIP/Brexit/Reform*

[38] Polling outlets vary because YouGov did not produce MRPs in 2020-2023, but other companies did.
[39] This decision was only made on the 11th November 2019, after Parliament had dissolved for the December Election, and therefore would not be expected to affect behaviour leading up to the 2019 election. It would (rightly) not be seen as likely the Brexit Party would repeat this decision at the next election, so it would not be expected to affect behaviour after the 2019 election either. The expected popularity of the Brexit Party in 2019 based on polling data, therefore, is important for measuring the most recent estimates MPs would act on when considering the threat the Brexit Party may pose in their constituency.

Party highest-yet popularity[40] based on either the last election or latest included MRP poll, whichever is most recent. I also interact constituency positions with the local margin of victory. The control variables are whether the MP has a *frontbench position* (which has an interaction effect with their current *party leader*[41]), who the *current Prime Minister* is, and who the *current Opposition Leader* is. Fixed effects are included for *MP* and the *year*, and I include a *lagged dependent variable*[42].

The models used are therefore panel models that can be described with the following equation:

$$y_{i,t} = \alpha y_{i,t-1} + \beta \tilde{x}_{it} + \gamma \mathrm{z}_{it} + \delta_1 \mu_i + \delta_2 \mu_t + \delta_3 \varepsilon_{i,t}$$

Where $y_{i,t}$ represents the dependent variable of speech position for a given MP (i) in a given year (t). Meanwhile $y_{i,t-1}$ is a lagged dependent variable which takes coefficient α. $\tilde{x}_{it}$ is a vector of independent variables[43] which takes a vector of coefficients β, and z_{it} is a vector of control variables[44] which takes a vector of

---

[40] Highest-yet popularity is used because of expected dynamics in third-party threat assessment. If MPs are responsive to the threat, and the vote share of the third party decreases, it is likely to be interpreted as a success of the existing strategy. There is therefore not an expectation (being analysed here) that MPs would move in the inverse direction in response to subsiding threat: only that they adopt new strategies in response to a rising threat. Measuring highest-yet popularity avoids the model being distorted by the potential (absence of) reverse effects when threats subside.

[41] As, for example, the effect of being on Jeremy Corbyn's frontbench is expected to differ from the effect of being on Keir Starmer's frontbench.

[42] For robustness given the biases that could be introduced by a lagged dependent variable, I provide in Appendix 3 an Arellano-Bond Estimator version of the model. The main findings are unchanged.

[43] These are: constituency left-right position, constituency lib-auth position, margin of victory, constituency left-right position × margin of victory, and constituency lib-auth position × margin of victory.
For the models testing for third party threat effects, these also include: Liberal Democrat highest-yet popularity, and UKIP/Brexit/Reform Party highest-yet popularity

[44] These are: *frontbench position*, *current Prime Minister*, and *current Opposition Leader*.
For the models using Labour MPs, these also include *frontbench position* × *current Opposition Leader*.
For the models using Conservative MPs, these also include *frontbench position* × *current Prime Minister*.

coefficients γ. μ_i and μ_t are fixed effects for MP and for year respectively. $\varepsilon_{i,t}$ is the error term.

I use interaction effects to identify a moderating effect of seat marginality. By comparing the direct effects of public opinion to the interaction effects of seat marginality combined with public opinion, it can be determined the extent to which MPs' responsiveness to constituency interests is conditional on re-election threats. Furthermore, the effect of third-party vote shares can show the degree to which strategic responsiveness is conditional on opinions being mobilised behind specific electoral threats. I apply the models to the two major parties individually. This is because there would be reasons to expect some effects, such as marginality, to cause polarisation or convergence effects. If present, they would be visible as negative effects for one party and positive for the other. It is also because strategic responses are likely to be learnt through party institutions or cultural norms, which may have developed in response to real differences in winning strategies for each party.

To test the hypotheses outlined above, I use eight fixed effects models with lagged dependent variables: an economic dimension model and a social dimension model for each of the two main parties in the UK parliament, all of which are run firstly just for the public opinion variables, and secondly also incorporating variables for party popularity. In Table 3.1, I outline the expected results based on each of the hypotheses.

Hypothesis	Constituency	Marginality Interaction
H1: General Responsiveness	MPs' speech positions correlated with constituency position.	
H2: Marginality-Conditional Responsiveness		MPs' speech positions' correlation with constituency positions proportional to the marginality of the seat.
H3: Party-Conditional Responsiveness	MPs' speech positions correlated with third party positions proportional to third party vote shares.	

Table 3.1: Hypotheses about responsiveness to constituency opinion, voting, and marginality, and their empirical expectations.

Results

Constituency Opinion Effects

The predictions of the fixed effects models testing H1 and H2 are presented in Figure 3.3. The full tables with coefficients and significance levels for all variables in the models are presented in Appendix 1. These models test for the effect of constituency opinion and seat marginality on English-constituency MP positions. The *y* axis shows the estimated MP speech positions on the 7-point scale. The *x* axes show the full actual range of MRP-estimated constituency positions on left-right and libertarian-authoritarian topics, with the distribution of constituency positions for MPs of each party displayed via a histogram. The lines show the predicted positions of MPs' speeches based on the position of their constituents. The darker line shows the predicted positions of MPs on the lower quartile of victory margins (those whose seats are marginal); and the paler line shows the predicted positions of MPs in the upper quartile (those whose seats are safe).

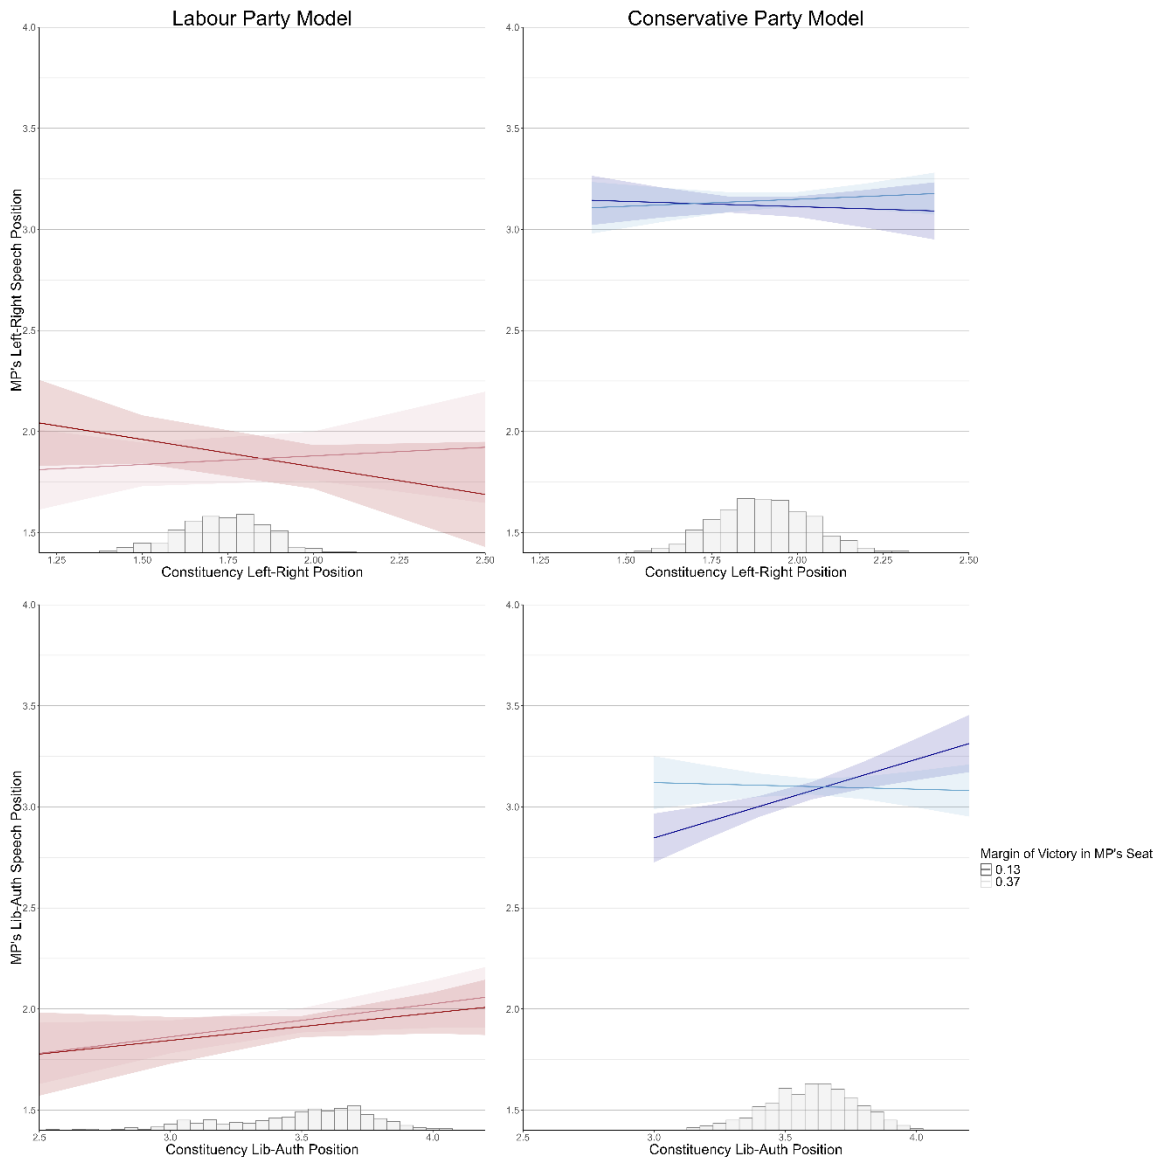

Figure 3.3: Predicted MP Positions Conditional on Constituency Opinion and Seat Marginality.

There are no significant effects of constituency left-right or libertarian-authoritarian positions on MPs' positions spoken in parliament, with the exception of Conservative MPs in more marginal seats, who are responsive to libertarian-authoritarian positions in their constituency.

Also visible in Figure 3.3 is the degree of difference between the two parties. The predicted positions of Labour and Conservative MPs are far apart, and this distance

puts into context the scale of difference that is associated with constituency positions, especially given the large overlap in the positions of constituencies between the two parties.

Overall, MPs are not particularly responsive to differences in constituency opinions, even when marginal. There is no evidence for general responsiveness (H1), and marginality-conditional responsiveness (H2) is found only for Conservative MPs on non-economic issues.

Constituencies do have varied political positions on economic and social issues, but these only affect what goes on in parliament insofar as they result in different parties winning at elections. Whilst there is a rightward skew in the positions of Conservative MPs' constituencies, this is small and leaves a very large overlap, with many Labour MPs representing constituencies that are overall to the right of many Conservative-voting constituencies. Despite often representing people with preferences that diverge from their party's position, MPs are unaffected by this and remain strongly aligned with co-partisans.

Third Party Threat Effects

Figure 3.4 shows the predictions from the fixed effects models for the effect of the highest-yet popularity of UKIP, the Brexit Party, and the Reform Party according to constituency-level opinion poll estimates and election results. The full tables with coefficients and significance levels for all variables in the models are presented in Appendix 2. These models test for the effects of third party threats from parties on the right wing.

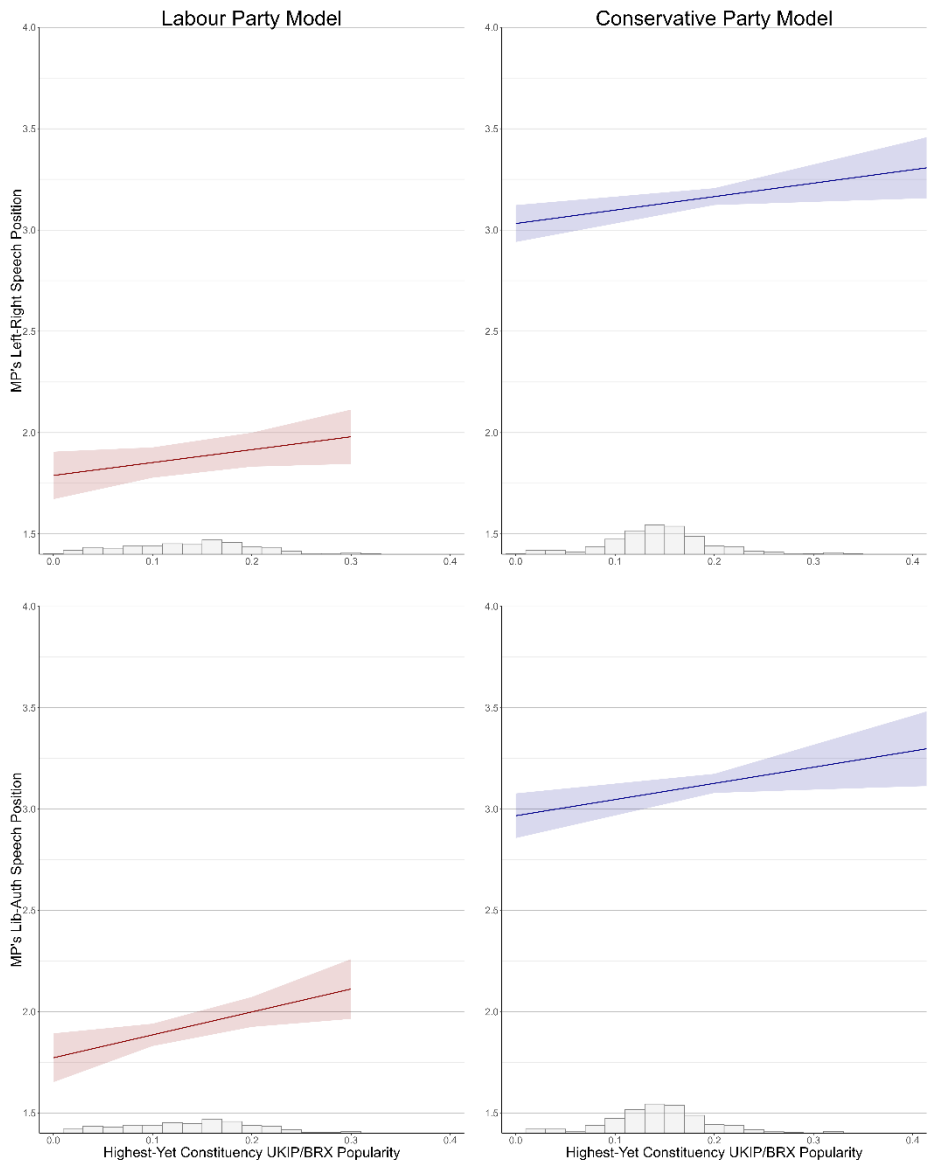

Figure 3.4: Predicted MP Positions Conditional on Constituency-Level UKIP, Brexit Party, and Reform Party Highest-Yet Popularity.

Both Labour and Conservative MPs are consistently responsive to UKIP, Brexit Party, and Reform Party popularity in their constituencies. Economic positions are significantly different only for Conservative MPs, and only at the extremes of radical right third party popularity levels, such that they are affected only by a small amount, and only if radical right third parties become effectively the main opposition in their seat. Non-economic positions, meanwhile, are significantly different in response to

radical right third party popularity levels throughout the distribution for both Labour and Conservative MPs.

This suggests that, when constituency interests are organised behind a third party threat, MPs are inclined to respond to those interests. This confirms the party-conditional responsiveness hypothesis (H3) with respect to the radical right third parties. This may be because voters for third parties – especially relatively newly popular breakthrough parties like UKIP, the Brexit Party, and the Reform Party – are seen as more mobile in updating their vote choices, and therefore key strategic targets. It may also be that the popularity of third parties operates as an indicator of constituency opinion that is more visible to MPs than constituents' average left-right views.

Figure 3.5 shows the predictions of the fixed effects model for the effect of the highest-yet popularity of the Liberal Democrats according to constituency-level opinion poll estimates and election results. The full tables with coefficients and significance levels for all variables in the models are presented in Appendix 2. These models test for the effects of third party threats from the centre on economic issues, and the left on social issues[45].

[45] Whilst similar on social issues to the Labour party for much of this period, relative to the default of competing for Conservative voters, competing for Liberal Democrat voters would involve more left-wing positions.

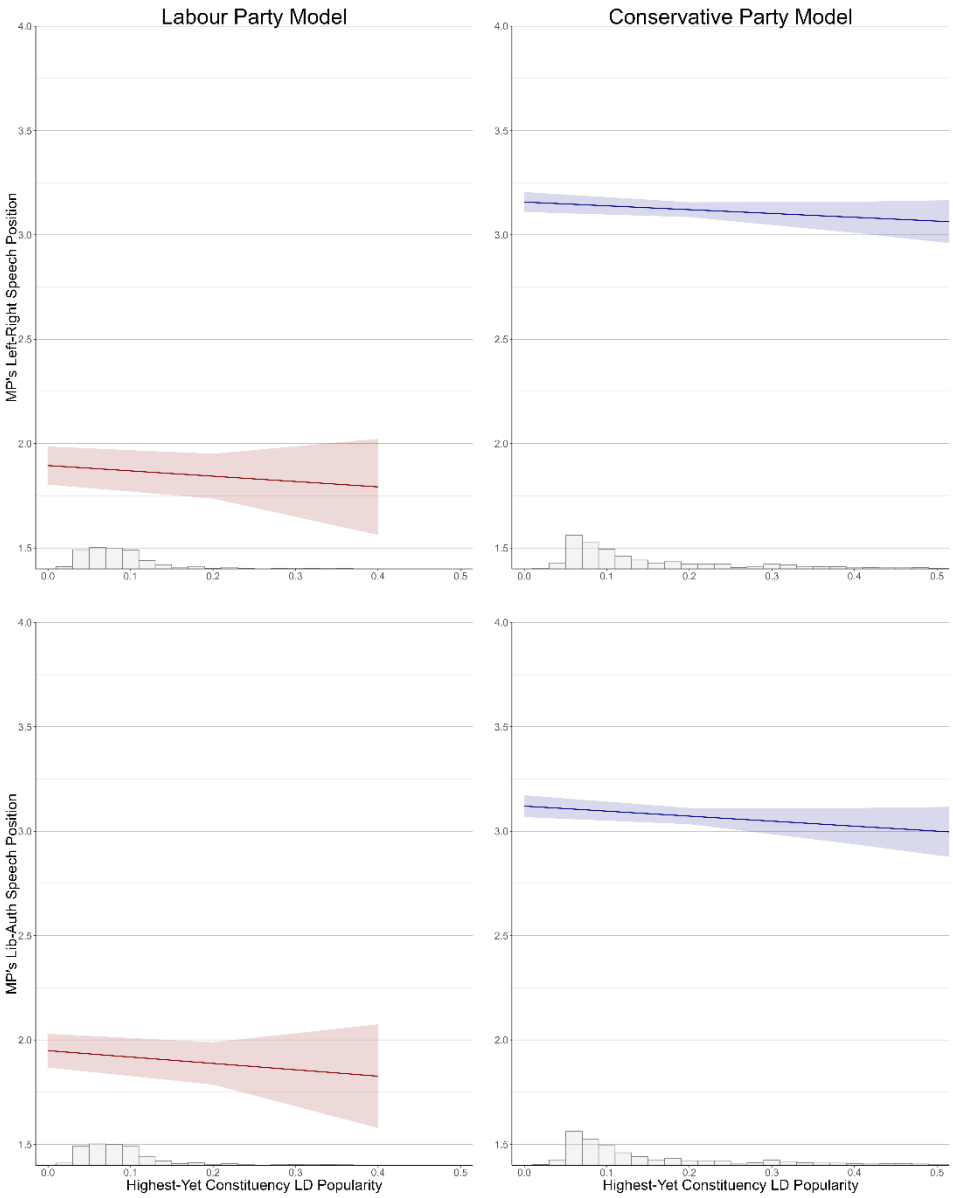

Figure 3.5: Predicted MP Positions Conditional on Constituency-Level Liberal Democrat Highest-Yet Popularity.

There are no significant effects of Liberal Democrat popularity. There are consistently negative coefficients – with MPs estimated to be more left-wing with increases in Liberal Democrat popularity – but none of these are significant.

These results therefore reject party-conditional responsiveness (H3) with respect to the Liberal Democrats. This could be based on MPs' diverging assessments of different parties' prospects. Radical right party breakthrough is an emerging

international trend throughout the period being studied, whereas similar trends did not exist to support the potential for a large and growing threat from the centre. It could also be due to asymmetries in the parties' positions. The EU Referendum created a democratic mandate, which may be a source of asymmetric pressure, with adoption of positions more towards Euroscepticism, anti-immigration, and social conservatism being legitimised, whilst positions towards 'Remain', pro-immigration, multiculturalism, and progressivism being less justifiable.

Conclusions

Partisanship has a huge effect on the positions MPs advocate in parliament, with the two main parties in the House of Commons being highly differentiated in their MPs' positions on both axes, with minimal overlap, despite the political preferences of the set of constituencies each party represents being similar. Where MPs vary within parties, their positions in parliament are consistently responsive to the local popularity of radical right third parties, but not to the local popularity of the Liberal Democrats, or to their constituents' overall left-right or libertarian-authoritarian political positions – except for Conservative MPs in marginal seats (who are responsive to libertarian-authoritarian positions). This is despite a broad range of political positions between different constituencies, with many Labour MPs representing seats whose views are to the right of the average Conservative MP's seat, and vice versa. This pattern is consistent with the theory that legislators primarily act strategically in order to win votes from the supporters of third parties, but includes an asymmetry in how they respond to the right-wing parties versus how they respond to the Liberal Democrats. MPs do not adjust their political positioning in response to changes in their constituents' opinions, even under marginal conditions, but when opinions are organised behind a credible electoral threat in the form of a third party, MPs may adopt positions that would appeal to the voters of that party. In all, the results reject the hypotheses of general responsiveness (H1), confirm marginality-conditional responsiveness (H2) only among Conservative MPs on non-economic issues, and confirm the hypothesis of party-conditional responsiveness (H3) only with respect to the right-wing third parties.

I find that MPs' positions are responsive to the popularity of radical right third parties in their constituencies. When there is evidence from constituency-level polling and election results that UKIP, the Brexit Party, or the Reform Party is popular in their constituency, MPs adopt more right-wing positions when speaking in parliament. This effect is larger for non-economic issues than for economic, and no comparable effects are found in response to the popularity of the Liberal Democrats.

These results are relatively robust to confounders, given the panel model design. National-level effects such as the Covid-19 pandemic or the EU Referendum would cause national-level trends: and so would not explain why it is specifically MPs in constituencies where radical right third parties have gained popularity that move to the right. It is possible that the findings are the result of an interaction between national-level shocks and local-level threats, but this is not a threat to the theory overall.

Why MPs are responsive to radical right third parties but not to the Liberal Democrats is not clear from this evidence alone, but there are a number of possible explanations. These results are similar to those found by Abou-Chadi (2014) with respect to how mainstream parties' platforms move in response to rising popularities of radical right parties versus green parties. They showed that parties shifted their policies in response to radical right parties, but not in response to green parties. Whilst they explained this result with findings that parties de-emphasise environmental issues as their main strategic response. A similar explanation for how parties are responding to the Liberal Democrats, however, is not readily available. As a centre or centre-left party, they combat the mainstream parties on many of their same issues, rather than as an issue

entrepreneur. During much of this period, if there was an issue the Liberal Democrats 'owned' it was maintained opposition to Brexit and support for a second referendum. Brexit and the EU were, however, still salient issues for the mainstream parties, and belong to the same issue dimension that the radical right parties competed on. It may therefore be something unique about radical right parties that parties or legislators respond to their entry with policy shifts, rather than something unique about other parties that makes them less responsive.

It may, for example, be that the rise in radical right-wing voting was a highly anticipated and worrying trend for both parties after the EU Referendum in 2016. In the context of radical right breakthroughs taking place in other European countries, MPs can likely imagine such breakthrough taking place in the UK if they are not careful. There are templates for imagining that, whatever the current level of radical right support, it could continue to increase given precedents from other countries. Meanwhile, there were no such international trends for the breakthrough of centre or centre-left parties, and the Liberal Democrats had issues restoring their 2010 levels of popularity after their participation in the 2010-2015 Coalition. The causes for this likely include those associated with being a junior partner in government, such as voters being disappointed with concessions and compromises to the party's platform (Hjermitslev, 2020), and those associated with 'bloc changes' for centre parties in coalition politics (Arndt & Christiansen, 2022), given that the Liberal Democrats had historically been known for 'Lib-Lab' pacts, and were therefore expected to align with the left if possible. In sum, there were popular reasons for MPs to be sceptical about the potential for a Liberal Democrat breakthrough at upcoming elections throughout the 2015-2024 period, and popular reasons to be fearful about the potential

for a UKIP or Reform Party breakthrough. This may explain the asymmetry in responsiveness to these parties as based on strategic assessment of how serious a threat they are. Future research may consider calculating measures of 'threat' for small parties which incorporate international trends as well as local vote shares to yield estimates of what a party or MP would predict the small party's upper limit of expected success could be.

Alternatively, MPs may be responding strategically based on the differences in the parties' platforms. Given that UKIP and the Brexit Party have a platform representing a radical pole of social political positions, movements MPs make in that direction are more likely to get noticed, identified, and rewarded as clear signals that they can advocate for those positions instead of the alternative party. As a centre party, however, the Liberal Democrats' positions are less distinctive, so the rewards for adopting their positions may be smaller, less noticed, and therefore not worthwhile. The main distinctive position of the Liberal Democrats in this period is their maintained 'Remain' position on Brexit, and support for a second referendum. This, however, was a highly salient topic which became central to both parties' platforms, and Remain positions were seen as toxic for electability, as they risked being seen as undemocratic in the wake of the referendum, and had both anticipated and real costs in their chances for both re-election by the public, and re-selection by their party (Cox & Shapiro, 2025). There were therefore significant barriers during this period and context to adopting the distinctive pro-European position of the Liberal Democrats, whereas if anything pressure towards adopting the Eurosceptic position of UKIP or the Reform Party.

The lack of responsiveness in general to constituent opinions is unlikely to be surprising. Legislators have imperfect information about their constituents' preferences outside of poll results and electoral outcomes. This could explain why the one significant finding of an effect of constituency opinions is for Conservative MPs on social issues: libertarian-authoritarian positions are highly correlated with Brexit positions (Wilson, 2024). Where different constituencies stood on Brexit and how that would affect divisions within the parliamentary Conservative Party were salient issues at the time. Given that most Conservative MPs before the referendum supported Remain, perhaps the effect being captured is that of Remain MPs in Leave constituencies being more likely to shift their positions towards Euroscepticism.

These new findings on legislative behaviour were possible due to new methods of coding large corpora of texts using large language models. Whilst previous implementations of LLMs for this task in social science have had mixed success, I have introduced chain-of-thought methods, allowing the LLM coder to make more complex conclusions based on the texts provided, and improving the accuracy of the results.

A key limitation is that the scales used to measure MP positions and those used for public opinion cannot be made to capture exactly the same set of beliefs on exactly the same scale. Whilst we can nonetheless identify effects between the two scales, a piece of context is missing. It cannot be identified exactly where a given constituency's ideal point would be if an MP represented its views perfectly. It also cannot be identified whether the extremes of constituency opinion represent a similar range of beliefs to the extremes of MP opinion. It also cannot be identified whether

MPs are, on the whole, to the left or to the right of constituencies. These more descriptive questions about how representative MP positions are of their constituencies are therefore still unanswered, and can likely only be answered using like data: either submitting equivalent survey questions to both MPs and voters (such as candidate surveys used for the study of May's Laws (Norris, 1995; van Holsteyn, den Ridder, & Koole, 2017)), or collecting natural language data from voters about their political beliefs (perhaps with open-ended survey questions). This task remains for future research.

Through these methods, I have identified that individual legislators are responsive to constituency-level changes in party popularity. Whilst it's been recognised, through the use of datasets such as the Comparative Manifesto Project and Chapel Hill Expert Survey (Meijers, 2017), that these dynamics exist at the party level, such evidence at the level of individual legislators on this scale is novel. I have also yielded new evidence on the strength of partisanship in the positions legislators adopt in the UK parliament. Even when capturing a measure of left-right position rather than explicit partisanship, there is very little overlap between the two main parties' MPs. Overall, this highlights the importance of voting in determining legislative behaviour. MPs' positions are generally not significantly correlated with public opinion, but a change to the winning party in a constituency can radically change how that constituency is represented in parliament, and changes in third party popularity levels can have effects on incumbent MPs' behaviour.

Appendices

Appendix 1: Regression Tables, Opinion Model

	Dependent variable:			
	MP Left-Right Position (0-10)		MP Auth-Lib Position (0-10)	
	Labour	Conservative	Labour	Conservative
Intercept	2.49**	2.94**	1.16*	-0.86
	(0.45)	(0.37)	(0.48)	(3.03)
MP Left-Right Position in Previous Year	0.17**	0.16**		
	(0.03)	(0.02)		
MP Lib-Auth Position in Previous Year			0.20**	0.15**
			(0.03)	(0.02)
Constituency Left-Right Position (0-10)	-0.47	-0.12		
	(0.25)	(0.19)		
Margin of Victory × Constituency Left-Right Position	1.49	0.51		
	(0.79)	(0.62)		
Constituency Auth-Lib Position (0-10)			0.12	0.62**
			(0.13)	(0.16)
Margin of Victory × Constituency Auth-Lib Position			0.11	-1.76**
			(0.39)	(0.53)
Margin of Victory in Constituency	-2.75*	-0.87	-0.27	6.41**
	(1.38)	(1.19)	(1.34)	(1.91)
MP on Backbenches	-0.08	-0.10*	-0.08	1.35
	(0.04)	(0.05)	(0.05)	(2.98)
MP is on Corbyn Frontbench	-0.23**		-0.21**	
	(0.06)		(0.06)	
MP is on Cameron Frontbench		0.01		1.42
		(0.12)		(2.98)

MP is on May Frontbench		0.01		1.39
		(0.07)		(2.98)
MP is on Johnson Frontbench		-0.03		1.39
		(0.07)		(2.98)
MP is on Truss Frontbench		0.04		0.67
		(0.33)		(2.99)
MP is on Sunak Frontbench				1.31
				(2.98)
Prime Minister: Cameron	0.02	0.20	0.09	-0.33*
	(0.33)	(0.14)	(0.16)	(0.14)
Prime Minister: May	-0.24	0.21	-0.05	-0.13
	(0.33)	(0.13)	(0.15)	(0.13)
Prime Minister: Johnson	-0.09	-0.002	-0.01	-0.09*
	(0.11)	(0.05)	(0.05)	(0.05)
Prime Minister: Truss	-0.33	-0.26	-0.28	-0.18
	(0.74)	(0.30)	(0.36)	(0.30)
Opposition Leader: Corbyn	0.26	-0.27*	0.08	0.01
	(0.29)	(0.12)	(0.14)	(0.12)
Marginal R^2	0.06	0.04	0.06	0.04
Conditional R^2	0.28	0.12	0.28	0.18
Observations	1,444	2,585	1,452	2,600
Note:	*p<0.05; **p<0.01			

Appendix 2: Regression Tables, Party Model

	Dependent variable:			
	MP Left-Right Position (0-10)		MP Lib-Auth Position (0-10)	
	Labour	Conservative	Labour	Conservative
Intercept	2.25**	2.66**	1.6**	-0.53
	(0.48)	(0.38)	(0.53)	(3.03)

MP Left-Right Position in Previous Year	0.17**	0.16**		
	(0.03)	(0.02)		
MP Lib-Auth Position in Previous Year			0.19**	0.14**
			(0.03)	(0.02)
Constituency Left-Right Position (0-10)	-0.37	0.01		
	(0.25)	(0.20)		
Margin of Victory × Constituency Left-Right Position	1.32	0.37		
	(0.80)	(0.62)		
Constituency Lib-Auth Position (0-10)			-0.06	0.47**
			(0.15)	(0.16)
Margin of Victory × Constituency Lib-Auth Position			0.18	-1.55**
			(0.39)	(0.54)
Margin of Victory in Constituency	-2.36	-0.70	-0.43	5.60**
	(1.40)	(1.19)	(1.35)	(1.92)
Highest-Yet UKIP/Brexit Party Popularity in Constituency	0.64	0.66*	1.13**	0.80*
	(0.35)	(0.29)	(0.42)	(0.35)
Highest-Yet Liberal Democrat Popularity in Constituency	-0.26	-0.18	-0.31	-0.24
	(0.35)	(0.13)	(0.39)	(0.15)
MP on Backbenches	-0.08	-0.12*	-0.09	1.49
	(0.04)	(0.05)	(0.05)	(2.98)
MP is on Corbyn Frontbench	-0.23**		-0.22**	
	(0.06)		(0.06)	
MP is on Cameron Frontbench		-0.001		1.56
		(0.12)		(2.98)
MP is on May Frontbench		-0.002		1.53
		(0.07)		(2.98)
MP is on Johnson Frontbench		-0.04		1.54
		(0.07)		(2.98)

MP is on Truss Frontbench		0.01		0.79
		(0.33)		(2.99)
MP is on Sunak Frontbench				1.47
				(2.98)
Prime Minister: Cameron	0.02	0.18	0.09	-0.34*
	(0.33)	(0.15)	(0.16)	(0.14)
Prime Minister: May	-0.26	0.19	-0.07	-0.15
	(0.33)	(0.14)	(0.15)	(0.13)
Prime Minister: Johnson	-0.09	-0.002	-0.01	-0.09*
	(0.11)	(0.05)	(0.05)	(0.05)
Prime Minister: Truss	-0.33	-0.25	-0.28	-0.17
	(0.73)	(0.31)	(0.35)	(0.30)
Opposition Leader: Corbyn	0.26	-0.25*	0.09	0.03
	(0.29)	(0.12)	(0.14)	(0.12)
Marginal R^2	0.07	0.05	0.07	0.05
Conditional R^2	0.28	0.13	0.29	0.19
Observations	1,444	2,585	1,452	2,600
Note:	*p<0.05; **p<0.01			

Appendix 3: Regression Tables, Arellano-Bond Robustness Check

	Dependent variable:			
	MP Left-Right Position (0-10)		MP Lib-Auth Position (0-10)	
	Labour	Conservative	Labour	Conservative
Intercept	2.25**	2.66**	1.68**	-0.53
	(0.48)	(0.38)	(0.53)	(3.03)
MP Left-Right Position in Previous Year	0.17**	0.16**		
	(0.03)	(0.02)		
MP Lib-Auth Position in Previous Year			0.19**	0.14**

			(0.03)	(0.02)
Constituency Left-Right Position (0-10)	-0.37	0.01		
	(0.25)	(0.20)		
Margin of Victory × Constituency Left-Right Position	1.32	0.37		
	(0.80)	(0.62)		
Constituency Lib-Auth Position (0-10)			-0.06	0.47**
			(0.15)	(0.16)
Margin of Victory × Constituency Lib-Auth Position			0.18	-1.55**
			(0.39)	(0.54)
Margin of Victory in Constituency	-2.36	-0.70	-0.43	5.60**
	(1.40)	(1.19)	(1.35)	(1.92)
Highest-Yet UKIP/Brexit Party Popularity in Constituency	0.64	0.66*	1.13**	0.80*
	(0.35)	(0.29)	(0.42)	(0.35)
Highest-Yet Liberal Democrat Popularity in Constituency	-0.26	-0.18	-0.31	-0.24
	(0.35)	(0.13)	(0.39)	(0.15)
MP on Backbenches	-0.08	-0.12*	-0.09	1.49
	(0.04)	(0.05)	(0.05)	(2.98)
MP is on Corbyn Frontbench	-0.23**		-0.22**	
	(0.06)		(0.06)	
MP is on Cameron Frontbench		-0.001		1.56
		(0.12)		(2.98)
MP is on May Frontbench		-0.002		1.53
		(0.07)		(2.98)
MP is on Johnson Frontbench		-0.04		1.54
		(0.07)		(2.98)
MP is on Truss Frontbench		0.01		0.79
		(0.33)		(2.99)
MP is on Sunak Frontbench				1.47

				(2.98)
Prime Minister: Cameron	0.02	0.18	0.09	-0.34*
	(0.33)	(0.15)	(0.16)	(0.14)
Prime Minister: May	-0.26	0.19	-0.07	-0.15
	(0.33)	(0.14)	(0.15)	(0.13)
Prime Minister: Johnson	-0.09	-0.002	-0.01	-0.09*
	(0.11)	(0.05)	(0.05)	(0.05)
Prime Minister: Truss	-0.33	-0.25	-0.28	-0.17
	(0.73)	(0.31)	(0.35)	(0.30)
Opposition Leader: Corbyn	0.26	-0.25*	0.09	0.03
	(0.29)	(0.12)	(0.14)	(0.12)
Marginal R^2	0.07	0.05	0.07	0.05
Conditional R^2	0.28	0.13	0.29	0.19
Observations	1,444	2,585	1,452	2,600
Note:	*p<0.05; **p<0.01			

Chapter 4: The Political Positions of Female and Ethnic Minority Legislators in the UK

Abstract

Descriptive representation is an important mechanism for the representation of marginalised groups' interests in the UK. To what extent, though, do the political positions of female and ethnic minority legislators differ from other MPs? And do they differ in ways that represent differences in the population? I use new methods in coding text data using large language models (LLMs) to determine the left-right positions expressed in over 2 million speeches in the UK Parliament from 2015 to 2024. I also apply these methods to code over 20,000 debates as on topics directly relevant to ethnic minorities' interests, directly relevant to women's interests, or not, so that MPs' positions on these topics can be determined separately from their positions in general. I find that ethnic minority and female legislators behave differently to their white and male peers. Female MPs are significantly more left-wing than male MPs of their own party, especially when discussing issues related to gender. Ethnic minority MPs are only more left-wing than white MPs within the Labour Party. Conservative ethnic minority MPs are more right-wing than white Conservative MPs, unless they are discussing issues related to ethnicity, and/or representing an especially left-wing constituency.

Introduction

The problem of representing minority and marginalised interests in democratic systems is a long-standing paradox, especially in majoritarian systems. It demands that representatives serve the interests of a minority, within a system intended to only

represent the interests of the majority. This problem has often been couched in terms of 'descriptive' and 'substantive' representation (Pitkin 1967), where the presence of representatives who, themselves, belong to minority or marginalised groups (descriptive representation) is theorised to not only constitute representing them, but to lead to substantive activities that represent those groups' interests (substantive representation). This, however, comes into conflict with traditional notions of representation, if it means that those legislators' activities are less representative of the majority of the constituents they represent. It is, however, often the only way for such interests to be represented substantively at all[46]. Institutional factors have complicated descriptive representation, however, as parties' role in candidate selection allows them to 'gatekeep' candidacies and select minority and marginalised group members that fit specific ideological profiles, to ensure they are not representative of those groups' interests or views (Dancygier 2017). Whilst there are good theoretical and qualitative analyses of these questions, more quantitative evidence is needed to determine the quality of descriptive representation in different countries, and how it varies between parties.

Recent advances in methods of analysis for large text data now enable high quality empirical tests of legislative behaviour at the level of individual legislators, over time. Speeches provide millions of data points across a decade, which all provide information about the political positions the speaker chooses to represent. These have so far been analysed little, because coding each speech individual would be an

[46] In the absence of a representative number of constituencies that have a majority population of the marginalised group.

inordinately expensive and time-consuming task. Large language models (LLMs), however, have been approaching (and, for some use cases, surpassing) the quality of human coders in text annotation tasks (Heseltine and von Hohenberg 2024, Ziems, et al. 2024, Le Mens and Gallego 2023). By using an LLM to the code political positions of statements made in the Westminster Parliament from 2015 to 2024, I have created a dataset of the average positions represented by MPs in each year. By also applying these methods to the coding of debate topics as directly relevant or not to the interests of ethnic minorities or of women, I also determine the positions taken by MPs on ethnic and gender issues. LLM coding of parliamentary transcripts, therefore, enables individual-level, multi-dimensional, time-series data on MPs' positions in Parliament. I create this data and use it to yield new empirical findings on the representation of ethnic minority and female interests in the UK.

There are important theoretical implications from the work that has been done so far on the causes of underrepresentation of women and ethnic minorities. The next step of empirically studying the quality of representation of women and ethnic minorities, however, has been reliant on sparse data, such as elite surveys and voting records. I create an extensive dataset on positions taken in parliamentary speeches in the UK, providing a new way to study the platforms and behaviours of female and ethnic minority MPs, and of all representatives on topics relating to the rights and interests of women and ethnic minorities.

Literature Review and Theory

Findings so far have been conflicting on the effects of descriptive representation. Homola (2017) found in Western Europe at the aggregate level that party platforms are more responsive to men's interests than women's, and that this is unaffected by the level of descriptive representation of women among the party's legislators. Acquarone and di Landro (2024) found, similarly, with elite survey data on individual legislators in Latin America, that ethnic minority politicians did not differ from their ethnic majority peers, and that nor does their prevalence within a parliamentary party affect the overall party platform. Norris and Lovenduski (1994) summarise existing literature at the time as inconclusive, with most attempts to identify an effect of social background on policy attitudes yielding little evidence. They suggest that, though it can't be assumed that issues that divide men from women or ethnic groups from each other in mass politics will be carried forward to divide elites, there are strong theoretical reasons to believe that socialisation into groups anyone belongs to their whole life will affect their attitudes[47]. Meanwhile Tate (2002) found in the US at the individual level that the ethnicity of congresspeople had a noticeable effect on their voting records on bills that related to Black interests. At the party level, Greene and O'Brien (2016) found across 20 countries that inclusion of greater numbers of female legislators in a party was associated with a greater diversity of issues in that party's manifesto.

[47] They reference the changing class composition of the parliamentary Labour party as a site where it may be observed, in aggregate, how the demographics and social backgrounds of legislators may affect their policy positions.

Contradictory implications for the behaviour of ethnic minority and female legislators are found in the literature on voter behaviour, which may be theoretically inferred to have impacts on legislators' strategies. There have been findings that the public evaluates candidates based on traits such as ethnicity and gender. Voters have been found to assume women and Black candidates in the US are more left-wing (McDermott, 1998; Koch, 2002; Jacobsmeier, 2015). Zingher and Farrer (2014) contribute some further nuance in the UK and Australian contexts, showing in the UK that Labour ethnic minority candidates win additional support from Labour party supporters and from ethnic minorities. Meanwhile, Conservative ethnic minority candidates have no significant effect on ethnic minority voting patterns, but can reduce their support from white Conservative party supporters. This may be one explanation for Fisher et al.'s (2015) finding that ethnic minority candidates suffer a penalty in vote shares, which is primarily caused by the responsiveness of anti-immigrant voters[48]. Existing literature suggests, therefore, that ethnic minority candidates are primarily punished at the ballot box due to their (mis-)perception as being more progressive or less anti-immigration, with Jacobsmeier (2015) finding that the direct effect of prejudice is outweighed by the indirect effect of (assumed) candidate ideology, and Martin and Blinder (2020) finding that 'pro-minority positions' are punished more by voters than personal ethnicity. When seeking to appeal strategically, therefore, ethnic minority MPs could be expected to compensate

[48] They explain this as being a result of anti-immigrant opinion being strongly correlated with islamophobia, which they do not measure. A mechanism based on simple islamophobia, or a mechanism based on perception of Muslim candidates as pro-immigration would both be consistent with this finding.

for constituents' perceptions that they are more left-wing by constructing a platform that emphasises or creates more right-wing positions: especially for Conservative ethnic minority MPs who are more likely to seek to appeal based on their right-wing credentials. Whereas these public perceptions can be inferred to have an effect on candidate (and MP) strategies, evidence that they do is sparse.

Other theories have focused on selection effects that may cause ethnic minority and female legislators not to be representative of their peers in the general public. These identify a similar strategic incentive, but locate the mechanism at the party level rather than the individual level. Van de Wardt, et al., (2024) for example, hypothesise that parties that shift rightwards (on civil rights issues) use ethnic minority legislators as 'reputational shields'. This allows them to either point to the continued loyalty of ethnic minority legislators as an indicator that their policies are acceptable to ethnic minorities, or use the supportive voices of loyal or aligned ethnic minority legislators as a way to legitimise their policies. This was a dynamic also identified by Clayton, et al. (2019). There are also selection effects implied by the mechanisms theorised as deterrents against ethnic minority participation in elite-level politics. For example, Lajevardi, et al., (2024) found that ethnic minorities in the UK and Sweden fear discrimination and expect to feel unwelcome if they run for elected office – and that this holds true regardless of how politically engaged the individual is. Similarly, Fox and Lawless (2024) found depressed levels of political ambition among 'well-credentialed potential candidates' who were women compared to men. Bird (2005) theorises that visible ethnic minority candidates are likely to be 'entrepreneurs' or 'bridge-builders'. These candidates are likely to gain attention from party institutions and the media for challenging stereotypes about their ethnicity. They cite examples of

progressive Muslim candidates who subvert the expectation that they would be socially conservative, but their theory could equally be applied to van de Wardt, et al.'s (2024) and Clayton, et al.'s (2019) cases of women and ethnic minority candidates' role in legitimising anti-civil rights policies. Furthermore, Farrer and Zingher (2018) found that the nomination of ethnic minority candidates is explained by parties' incentives and strategies. They show that the probability of ethnic minority candidates being nominated is affected by the types of voters the party appeals to, with centre-left parties more likely to field them more broadly as their more cosmopolitan voter base is more favourable to ethnic minority candidates. Meanwhile centre-right parties will be more selective, likely fielding them only in safer seats or in seats with larger ethnic minority populations to counteract the penalty found among anti-immigration voters, who they also seek to appeal to. As well as affecting the number of candidacies, however, party strategies ought also to be expected to affect the types of candidates the parties are willing to field.

Party-level differences in strategies with regard to ethnic minority candidates were analysed in the context of the 2010 General Election in the UK by Sobolewska (2013). They observed that the Labour and Conservative parties both made efforts to increase the number of ethnic minority candidates, but for different reasons and to different effects. Labour intended to retain ethnic minority voters who threatened to switch votes due to the Iraq War and increased media attention on their large number of ethnic minority voters, but small number of ethnic minority MPs. The Conservatives, meanwhile, sought to subvert their reputation as 'the nasty party', or as a racist party hostile to the interests of ethnic minorities. Based on this account, both parties were adopting ethnic minority representation for the purpose of providing 'reputational

shields', but whereas Labour sought to appeal to ethnic minorities themselves, the Conservatives' priority was their overall image as a party. They found that whereas Labour fielded ethnic minority candidates in a broad diversity of seats – with high and low ethnic minority populations, and in safe and marginal seats – the Conservatives' ethnic minority candidates were polarised between safe and unwinnable seats (with the majority in unwinnable seats). This meant that ethnic minority Conservative MPs were more likely to represent conservative strongholds, where selection was more to do with party-level strategic considerations than local ones, with 8 out of 11 elected in 2010 being in safe seats. Meanwhile, Labour MPs were relatively more likely to be selected and elected in marginal seats, where their individual strategy would have to be responsive to the local constituency. Many of the ethnic minority and female MPs elected in this election (including Cameron's 'A list') continued to sit in the 2015-2024 period I study in this chapter, and are vital context as they were the first significant wave of new ethnic minority MPs.

Based on the literature so far, there is a direct contradiction in expectations between the traditional theories of descriptive representation, and newer theories of party and candidate strategy. The newer theories come in two main forms. Firstly, that parties select unrepresentative candidates who belong to marginalised groups to act as 'reputational shields', so that their policy platform can be justified by the continued loyalty of legislators belonging to marginalised groups, even if the platform harms those groups' interests. Secondly, that the candidates themselves may be incentivised to adopt and emphasise positions contrary to the interests of marginalised groups to which they belong, in order to challenge public pre-conceptions. Whereas theories of descriptive representation expect representatives to be more likely to adopt positions

that align with those of marginalised groups to which they belong, strategic theories expect representatives to avoid or even oppose such positions.

If these strategic theories are true, they ought to have effects on how legislators behave in parliament. This is in part because they act as selection effects: as candidates who fit a certain political profile are more likely to be selected by their party, or to be elected by voters. It is also in part because behaviour in parliament is part of how legislators construct a public platform. Speeches create a reputation among parliamentary peers and political commentators in the media, which has indirect effects on public perceptions by affecting how colleagues and the media portray them. Speeches also justify positions taken in votes, and so the content of speeches is restrained to being representative of a legislator's voting behaviour, which is likely to be publicly cited by critics and journalists when discussing their political positions. Strategic incentives that affect how a candidate campaigns, therefore, ought to also affect how an elected legislator behaves in order to protect their re-election chances.

There are therefore two strategic theories which differ in whether they expect the effect to be situated at the level of party strategy, or the level of individual candidates. They therefore yield diverging predictions about how legislators' behaviour would vary in response to the constituency they represent.

The party-level theories suggest that legislators may be unrepresentative of the marginalised groups they belong to because they are selected by their party for political beliefs that are opposed to the most popular and/or most beneficial policies for that group; or simply for their unwavering loyalty to the party line. These are the

theories of 'gatekeeping' and 'reputational shields' (Dancygier, 2017). If so, legislators who belong to marginalised groups would either not differ from other legislators, or be more right-wing than other legislators due to gatekeeping selection. This effect occurs at the party level, and therefore a party's incentive to adopt this practice varies by the strategic costs and benefits from the perspective of the party itself. Right-wing and conservative parties are more incentivised to subvert the left-wing assumptions voters may have about their ethnic minority and female candidates. If the effect is situated at the party level, therefore, the effect ought to be observable more among Conservative MPs than Labour MPs.

The individual-level theories, however, suggest that legislators adopt strategies responsive to perceptions based on their own interests and strategic assessments. This is likely to be based on contextual factors about their own constituency. MPs in more right-wing constituencies are more likely to be punished for assumed left-wing political positions. Whether an MP compensates, therefore, would be conditional on the overall political leaning of the constituency they represent. Likewise, a candidate in a more right-wing constituency who doesn't compensate is more likely to be punished at the ballot box and fail to become an MP.

There are therefore both selection effects and behavioural differences that can exist at the party-level and at the individual-level. At the party-level, selection effects entail parties fielding candidates who fit a specific profile that benefits the party's reputation or strategically compensates for left-wing assumptions, whereas behavioural effects involve MPs adopting strategically compensatory positions as a result of party discipline. At the individual level, selection effects entail candidates' electoral

chances being conditional on strategically compensating for left-wing assumptions, making those who succeed in more right-wing constituencies more likely to be those who compensated for left-wing assumptions. Individual-level behavioural effects entail legislators adopting compensatory positions regardless of whether they are necessary or successful in reality, or even if they are in safe seats.

In this chapter, I analyse MPs and not candidacies. Distinguishing between selection effects and behavioural effects is therefore beyond the scope of this chapter. What I do distinguish between, however, is whether strategies are adopted at the party-level or individual-level. This distinction is important for understanding the role of parties and of voting behaviour in the quality of descriptive representation in the UK.

Based on these implications, I yield four hypotheses about MPs' behaviour. These propose two of the expected effects of descriptive representation (H1 & H2), and two diverging theories of how strategic incentives may reduce or invert the effects of descriptive representation. H3 represents the party-level theories, whilst H4 represents the individual-level theories[49].

H1: Representatives from marginalised groups are more left-wing in general.

H2: Representatives from marginalised groups are more left-wing when discussing issues related to their marginalised group.

H3: H1 & H2 apply conditionally on which party a representative belongs to.

[49] All theories are tested at the individual level, but the party-level theories propose that MPs vary according to which party they belong to, and the individual-level theories propose that MPs vary according to what's beneficial in their constituency.

H4: H1 & H2 apply conditionally on how left- or right-wing a representative's constituency is.

I take left-wing positions to be indicative of substantive representation in these hypotheses. This is because of a relationship between left-wing politics and marginalised interests both theoretically and empirically. Theoretically, redistributive and progressive politics generally favour the interests of marginalised groups by seeking to reduce inequalities. Empirically, members of marginalised groups generally exhibit preferences for redistributive and progressive politics.

Looking first at ethnic minorities, ethnic minorities in the UK consistently vote for left-wing parties (Ipsos 2024). Spatial alignment with Labour was also observed to be a factor[50] in ethnic minorities' preference for left-wing parties in British politics in 2010 (Sanders, et al. 2014). Heath et al. (2013), however, found a conflict in ethnic minority opinion data: they prioritised unemployment as the most important issue in the UK, but exhibit lower preferences for redistribution than their white peers[51]. They hypothesise that this may be due to immigrant backgrounds from countries without strong welfare states or labour movements – which may affect either their political information, or their contentedness with the current welfare state provision. More recent data on ethnic minority preferences has revealed a reversal in this trend, with Stowers finding in 2023 that ethnic minorities held overall more left-wing views than white respondents, with the sole exception of those of Chinese and Indian backgrounds (Stowers 2024), who Heath et al. note to be less economically

[50] Albeit a smaller one than for white voters.
[51] Alongside significantly more progressive views on immigration and affirmative action.

disadvantaged than other groups (Heath, et al. 2013). This change in the overall left-wing preferences of ethnic minorities could be due to increased socialisation into the UK political context, in line with Heath et al.'s predictions, via both cohort and period effects, as immigrant ethnic minorities age, and as new cohorts of 2nd, 3rd, and later generations of ethnic minorities enter the voting population. It could also be explained by improvements in polling methods and weighting practices, which make more disadvantaged and hard-to-reach ethnic minorities better accounted for in more recent polls. Mohyuddin (2024) found, in line with this account, that whereas class voting has eroded among white voters, income levels still have very large and significant effects on ethnic minority voting and economic preferences, suggesting that biases in ethnic minorities being surveyed could have skewed earlier results[52].

Meanwhile, female voters' preference for Labour over the Conservatives is a recent development in UK politics (Campbell and Shorrocks 2023). Studies of redistributive preferences, however, tend to find that women are more left-wing in their policy preferences. A meta-analysis by Bozzano and Scabrosetti (2024) showed that women were generally more favourable to redistribution across European and USA studies. British Social Attitudes data also shows women to have been consistently more left-wing in their values for almost the entire period studied, from 1986 to 2022 (Campbell

[52] Weights can account for this only partially, as there are unobserved variables determining an individual's economic circumstances, such as (parental) wealth, which may mean that even low-income respondents are not representative of low-income ethnic minorities.

and Shorrocks 2023)[53]. Women's historical preference for the Conservatives, therefore, appears to be in spite of their left-wing economic preferences.

I analyse whether speeches are 'on issues related to the interests of women' or 'on issues related to the interests of ethnic minorities' as a way to determine if female and ethnic minority MPs' positions are significantly different on these issues. I avoid employing a specific theory of what it is for a position to be 'feminist' or 'anti-racist'. This is so that the definition of what types of policies can be relevant can be as broad as possible. The theory is simply that, under the expectations of substantive representation, if an issue is likely to (uniquely) affect women or ethnic minorities, female and ethnic minority MPs would be more likely to act (uniquely) in the interests of the group(s) to which they belong.

By testing the hypotheses outlined above, it can be determined to what extent descriptive representation is successful in the UK, and how parties' and individuals' strategies differ. This is key to understanding how better substantive representation for marginalised groups can be achieved.

[53] An experimental study by Ranehill and Weber (2025) also finds that women in a simulated decision-making environment also show strong preferences for redistribution, and that groups with female majorities had more egalitarian policy outcomes

Data & Methods

Empirical tests for the outlined hypotheses require data on MPs' positions in parliament, their ethnicity and gender, and the positions of their constituents. For MPs' positions in parliament, I create an LLM-coded dataset of speeches in parliament from 2015 to 2024, which is aggregated into a yearly dataset on MPs' positions. I find MPs' ethnicities and genders by scraping Wikipedia. I create a dataset on constituency positions using an MRP model of British Election Study (Fieldhouse, et al. 2024) data on left-right economic positions and 'culture war' positions.

I code left-right positions of parliamentary speeches separately from their topic. This is because of LLMs' biases and stereotyping behaviours (Bai, Wang, Sucholutsky, & Griffiths, 2025), which may have harmful effects on the coding of positions, especially if handling concepts such as 'feminist' or 'anti-racist'. Combined with the LLM not being provided information about the MP's own name, demographics, or background, this ought to shield any such biases from affecting the coding of position on the given topics[54]. Instead, the LLM deploys a broad concept of 'left' and 'right', which is affected by all training data from its large dataset which relates to 'left-' and 'right-wing' politics. This will be a broad corpus of academic and conversational references to the concepts, constructing an 'understanding' of them that broadly tracks trends in (online, public) social and academic discussions of them. This may therefore be taken to be a 'descriptive' understanding of what it means for something

[54] It is still possible that the content of the speech being clearly to do with women or ethnic minorities may have effects: but this would in theory affect all such speeches equally, and not cause observable differences between the speeches *of* women and ethnic minorities.

to be left- or right-wing. Its implications for how statements about gender or ethnicity are coded can be seen later in this segment, in Table 4.1.

I code statements in parliament using the LLM Llama 3 8B, run locally via Ollama. This is selected as it can be downloaded and run locally without needing to access an internet service via an API. This maximises reproducibility in case of models no longer being supported, or being altered.

To improve the accuracy of the LLM's coding of statements, I use chain-of-thought methods. The premise of these methods is that LLMs generate responses by testing whether their outputs are valid. Longer outputs provide more data on their own validity, and therefore lead to more accurate outputs. Single-word or single-phrase 'codes' that are more usable as data can then be created by prompting another model to summarise the conclusion of the longer text. Researchers have therefore found that longer outputs improve LLMs' performance when conducting complex reasoning tasks, such as calculations and logic puzzles (Nye, et al. 2021, Wei, et al. 2022, Kojima, et al. 2022). I adapt these methods to be applied instead to the task of coding documents. Although previous successes with this method have been for computational tasks, the improved reasoning abilities are applicable to text coding. Identifying political positions requires identification of unspoken biases and subtext, which LLMs are more likely to do if they output natural language explanations that analyse the text for its leanings. The dataset I produce therefore innovates on existing methods, which usually just prompt an LLM for a code with no explanation (Ziems, et al. 2024). To create my dataset, the LLM is instead first prompted to *"Provide a brief description of how left- or right-wing this parliamentary speech would be in a*

[Year] UK context."[55] A second model is then given the output of the first model, and asked to *"Summarise the conclusion of this text with one of the following statements, and no explanation: Strongly Left-Wing, Centre-Left, Centre, Centre-Right, or Strongly Right-Wing."*[56] This yields the 1-phrase codes that are used to quantify the political leaning of the text. Further details on this method and how this dataset was constructed are available in Chapter 3 of this thesis.

I also apply these LLM coding methods to debate titles and speeches to determine whether they are on topics directly related to the interests of ethnic minorities or women. This is so that MPs' positions on these topics can be measured on a separate axis to their positions generally, as ethnic minority and female MPs' behaviour may vary only (or more) when their own interests are relevant, as predicted in H2. I code both debate titles and speeches so that, if the debate itself is on a bill or topic that is related to these groups' interests, all speeches in that debate are included in the measurement of an MP's positions on the relevant axis. This is important because if, for example, an MP speaks against a bill that protects a group from discrimination by highlighting the costs of the policy, they may only speak around the civil rights issued involved, and so the speech in isolation may not be identified as relevant to the group's interests. It ought, however, to be treated as a right-wing speech on the issue dimension of that group's interests. Individual speeches are also coded for relevance to ethnic minority or female interests. This is so that, if an MP raises the topic of a

[55] For this first output, the model is limited to 40 tokens, based on a trade-off of seeking a sufficient explanation, and optimising compute time to produce the output. The model uses a temperature of 0.8.
[56] For this output, the model is limited to 5 tokens, to allow for responses such as "Yes, it is strongly left-wing". This model also uses a temperature of 0.8.

marginalised group's interests during a more general debate, it is also identified as part of their overall positions on that group's interests. These codes are also generated using chain-of-thought methods.

To code the topics, an LLM is first prompted to answer *"What broad topics in British Politics is this [debate/statement] about? Is it directly related to the interests of women, ethnic minorities, or immigrants?"* For the coding of debates, it is provided with the title of the debate, and the opening statements (which are conventionally a spoken summary of the debate or bill, to introduce it to the house). For the coding of statements, it is simply provided with the statement. Immigrants are included in the prompt, due to initial tests revealing the LLM sometimes did not identify immigration as relevant to ethnic minorities' interests. There are then two separate summary prompts. One is given the broad topics output and asked *"Does this text suggest that the [debate/statement] in question is directly relevant to women's rights? Answer yes or no with no explanation."* The other is asked *"Does this text suggest that the [debate/statement] in question is directly relevant to ethnic minority rights? Answer yes or no with no explanation."* These yes or no prompts are asked separately to simplify the outputs, whilst allowing for ethnic minority and women's interests codes to not be mutually exclusive. Using these codes, every statement is coded as relevant to 'none', 'women', 'ethnic minorities', or 'both': either because it is part of a debate that is relevant, or because it is a stand-alone speech that is relevant.

Combined with the left-right annotations, the coded dataset successfully identifies left-wing and right-wing speeches relating to gender and ethnic topics. Table 4.1

shows illustrative examples of the breadth of issues and topics identified, and left- and right-wing speeches identified on those issues[57].

[57] The gender dimension includes speeches on trans rights, and the left-right coding treats pro-trans positions as left-wing, and anti-trans positions as right-wing. This is in line with evidence that pro-trans views are most common among progressives, and anti-trans views most common among conservatives (Tryl, Surmon, Kimaram, & Burns, 2022). On an interests basis, it is also in line with consistent findings that women are more pro-trans than men (Smith, 2025; Smith, 2020), and of externalities of transphobia for gender non-conforming cis women (Brooks, 2025). It also aligns with the current feminist philosophical literature (Mikkola, 2022).

Gender		Ethnicity	
Left-wing	Right-wing	Left-wing	Right-wing
"Despite the lauded progress that is apparently being made, only 26% of 110 Government posts are occupied by women." - Richard Arkless, Scottish National Party, 2015	*"They are learning from the Government's voluntary approach to women on boards, and I am pleased that Jayne-Anne Gadhia from the finance sector and others in the insurance sector have recently launched voluntary initiatives to ensure that companies publish their own gender pay gap."* - Nicky Morgan, Conservative, 2015	*"The Home Secretary spoke about community cohesion last week, yet her Government's explicit and almost dystopian goal is to create a 'hostile climate', as if we can hermetically seal off the bad migrants, while the rest of the multicultural UK goes about its business as usual."* - Stuart McDonald, Scottish National Party, 2015	*"Stop and search is a vital tool in the police's armoury in keeping people safe on our streets."* - Victoria Atkins, Conservative, 2019
"What has the Home Secretary got to say about the 32% increase in sexual exploitation and sexual offences, which is a really serious matter? Will she tell us what plans she has to involve the perpetrators in the criminal justice system?" - Andrew Gwynne, Labour, 2015	*"Many feminists disagree with that, but increasingly, disagreeing with gender ideology has become a dangerous thing to do, as we heard from the honourable Member for Thurrock (Jackie Doyle-Price)"* - Joanna Cherry, Scottish National Party, 2020	*"When the Prime Minister was forced to publish the review of the risks covid-19 poses to black and minority ethnic groups yesterday, why did he remove reference to the 1,000 responses to the review, many of which cited structural racism and discrimination as root causes of higher risk?"* - Andrew Slaughter, Labour, 2020	*"The standard form of this new logic is this: if 10% of people are characteristic x, then 10% of workers should be x, 10% of every company should be x, 10% of every role should be x and 10% of all chief execs should be x, and where that is not true, it is offered as evidence of discrimination, differences are inequalities, and the logic assumes that what is unequal must be wrong."* - Paul Bristow, Conservative, 2020
"Tragically, of those LGBT people killed in the Americas in 2013-14, 46% were trans women, and more than 2,000 trans gender and gender-diverse people were murdered in 65 countries between 2008 and 2015, according to the trans murder monitoring project." - Dawn Butler, Labour, 2017	*"Does the Minister not agree that perhaps the most inspirational women are those who have lived a life of duty and service, and of honourable- and devotion to their community, and that we should shine a light on our Queen, perhaps the most extraordinary woman of our generation, as an example of what we should aspire to?"* - Jim Shannon, DUP, 2020	*"Many people in Bath, Bristol and the surrounding areas have said that the statue of Edward Colston should have been removed many years ago."* - Wera Hobhouse, Liberal Democrats, 2020	*"Now then, I see that our NHS has published its very own woke alphabet, which includes terms such as 'white fragility' for the letter W."* - Lee Anderson, Conservative, 2021
"Will the Minister acknowledge that sanitary products are not a luxury, but a necessity? My question is: will she accelerate the timetable for removing the tampon tax: yes, or no?" - Karen Lee, Labour, 2018	*"Does my honourable Friend agree that woke activists are of course entitled to their views, and to express them, but that they are not entitled to impose those views as though they were in any way authoritative or unchallengeable?"* - Marco Longhi, Conservative, 2021	*"The legislation will perpetuate the systemic risk that infects our criminal justice system, including by expanding stop and search, which sees black men targeted, and by creating a new trespass offence that criminalises the life of nomadic Gypsy and Traveller communities."* - Caroline Lucas, Green, 2021	*"Before anyone accuses me of stoking a 'culture war' which, as we all know, is the left's new buzzword to try to shut down critical debates about their woke ideas, let me also point out that the official data shows that white people were the most searched ethnic group in this period: 10,000 more over a two-year period."* - Louie French, Conservative, 2024

Table 4.1: Speech samples from gender and ethnic topic categories,and from the 'Strongly Left-Wing' and 'Strongly Right-Wing' categories of each.

With these LLM-coded outputs, a dataset is created for every spoken statement in the UK parliament from the 2015 General Election to the 2024 General Election. This data is aggregated for into a yearly average position for each MP on three axes. The first is their average position on 'none' statements: where there are no ethnic minority or female interests, and so representing their general positions. The second is their average position on 'women' or 'both' statements, representing their positions on women's interests. The third is their average position on 'ethnic minorities' or 'both' statements, representing their positions on ethnic minority interests.

I also create a dataset on MPs' demographics by scraping Wikipedia. Ethnic minority and female MPs are identified using Wikipedia pages which list ethnic minority and female MPs in the UK respectively. MPs' age and education are found by scraping the 'info boxes' on MPs' individual Wikipedia pages. Age and whether the MP went to Oxbridge are to be used as control variables, as they may be correlated with both ethnicity/gender and MPs' political positions.

A descriptive summary of the data on MPs' demographics and political positions on ethnic and gender issues is presented in Figure 4.1. The *y*-axis represents MP's positions on ethnic issues on the left-hand box plots, and their position on gender issues on the right-hand plots. The box plots display the distribution of MPs' positions, grouped by their party, gender, and ethnicity. Some of the trends in the data are already visible. There is a deep partisan divide between Labour and Conservative MPs. Among Labour MPs, there is a clear trend that ethnic minority and female MPs are more left-wing than their white and male peers. These trends are not observable,

however, among Conservative MPs. Whilst female MPs seem to be more left-wing than male MPs, ethnic minority MPs do not appear to be different to white MPs.

Figure 4.1: Box plots of MPs' positions, measuring ethnic positions and gender positions separately. Lower numbers for position represent left-wing positions, and higher numbers are right-wing.

Figure 4.2 presents similar box plots, but with intersectional categories. The y-axis displays the average of MPs' positions on gender and on ethnic issues. Ethnic minority women in the Labour Party are on average even more left-wing than members of just one of either category. Ethnic minority female MPs, however, are similarly distributed to white male MPs among Conservatives. There are, though, very few MPs who fit this category in the parliamentary Conservative Party during this period, making an intersectional analysis of Conservative MPs very difficult.

Figure 4.2: Box plots of MPs' average positions on ethnic and gender issues, grouped by party, gender, and ethnicity. Lower numbers for position represent left-wing positions, and higher numbers are right-wing.

I estimate constituency positions by creating an MRP model on British Election Study opinion data. The model is fit using constituency-level census data on gender, age, ethnicity, and education. Individuals are coded for whether they are left or right of the median position on the set of Likert scale items measuring left-right economic positions. I run an MRP using a logit model to estimate the economic left-wing percentage in each constituency. These are calculated for the 2015, 2017, and 2019 waves (each election year). I do the same process using the British Election Study's set of Likert scale items measuring 'culture war' positions, which are designed to capture positions on contemporary issues related to the interests of women, ethnic minorities, and LGBTQ+ people. These are chosen as they are more directly related to

constituency-level perceptions of women and ethnic minorities, and would therefore be expected to affect what pre-conceptions they hold about female and ethnic minority candidates (van Oosten 2022). These questions were only introduced in 2023, and therefore only one snapshot of constituency-level opinion on these issues is used. Overall constituency-level opinion is measured using the average of these two issue dimensions.

With this dataset on MPs' annual positions, MP demographics, and their constituency profiles, I run two analyses, each of which is run separately for the effects of ethnicity, then for the effects of gender.

Firstly, I test for the traditional expectations of descriptive representation (H1 & H2), and whether they vary at the party level (H3). This is done using linear regression models, taking an MP's position in a given year as the dependent variable, and taking whether an MP is an ethnic minority, and whether they are speaking on a topic relevant to ethnic minorities, as independent variables. I include an interaction effect between the independent variables and the MP's party membership, to capture potential differences in effects between parties. Controls are included for the MP's gender, age, education, and whether they are a backbencher. An equivalent model is constructed for gender, with ethnicity as a control variable. If ethnic minority (or female) MPs are significantly more left-wing than their white (or male) peers on general topics or on topics relating to ethnic minority (or women's) interests, this would confirm H1 and H2 respectively. If they are not, it would reject H1 and H2 respectively. If H1 or H2 are rejected only for MPs in the Conservative party, this would confirm H3.

Secondly, I test for the hypothesis that ethnic minority and female MPs' behaviour varies at the individual level in response to their constituency profile (H4). This is done using linear regression models with the same set of variables, but with an additional interaction effect. The effects of MP ethnicity (or gender) and topic are interacted with not just party membership, but with the constituency left-right profile. If being an ethnic minority (or a woman) has a positive interaction effect with constituency left-right profile, this would mean their left-wing positions observed in the tests for H1 and H2 (or lack thereof) are conditional on average voter position[58]. A positive result would therefore confirm H4: that a strategic incentive at the constituency level mediates the effect of ethnicity on positions in parliament.

[58] Voter positions have previously been linked to propensity to vote against ethnic minority candidates (Fisher, Heath, Sanders, & Sobolewska, 2015)

Results

Descriptive Representation and Party-Level Strategic Hypothesis

The predicted positions of MPs based on their gender or ethnicity, and whether they are speaking on a topic relevant to the interests of women or ethnic minorities respectively, are presented in Figure 4.3. The *y* axis displays the predicted position of MPs on a 0-10 scale, where 10 is the most right-wing, and the *x* axis combines an MP's ethnicity or gender with whether they are speaking on topics relevant to ethnic minority or women's interests. With respect to ethnicity, they confirm H1, H2, and H3: descriptive representation has some expected results, but these are conditional on party. With respect to gender, they confirm H1 and H2, but reject H3, conforming to the expected results of descriptive representation, without party-level variation.

The effects of descriptive representation vary at the party level for ethnic minority MPs. There are observable effects of descriptive representation that are statistically significant, confirming H1 and H2, only among Labour MPs. These are not observed among Conservative MPs. Ethnic minority Conservative MPs are significantly more right-wing in general than their white peers, and on topics relevant to ethnic minority interests are not significantly different to white MPs.

This confirms H3 with respect to ethnic minorities. There is evidence of a party-level strategy to mitigate the effects of ethnic minority candidates which selects or encourages ethnic minority MPs to adopt or emphasise more right-wing platforms. The lack of such party-level evidence among Labour MPs suggests that the left-wing perceptions of candidate positions is not seen as harmful when a party is seeking election on a left-wing platform.

The effects of descriptive representation do not vary at the party level for female MPs. The effects are observable among female MPs of both parties, confirming both H1 and H2. Women are more left wing than men both in general and on issues directly related to women's interests. This may be because the Conservative party does not see assumptions that female candidates will be more left wing as so prevalent or extreme as such assumptions about ethnic minority candidates. It may also be because the larger number of female MPs[59] weakens the power of parties to 'gatekeep', strategically or otherwise.

[59] Achieved through persistent pressure and campaigns, but still not achieving parity (Collignon 2024).

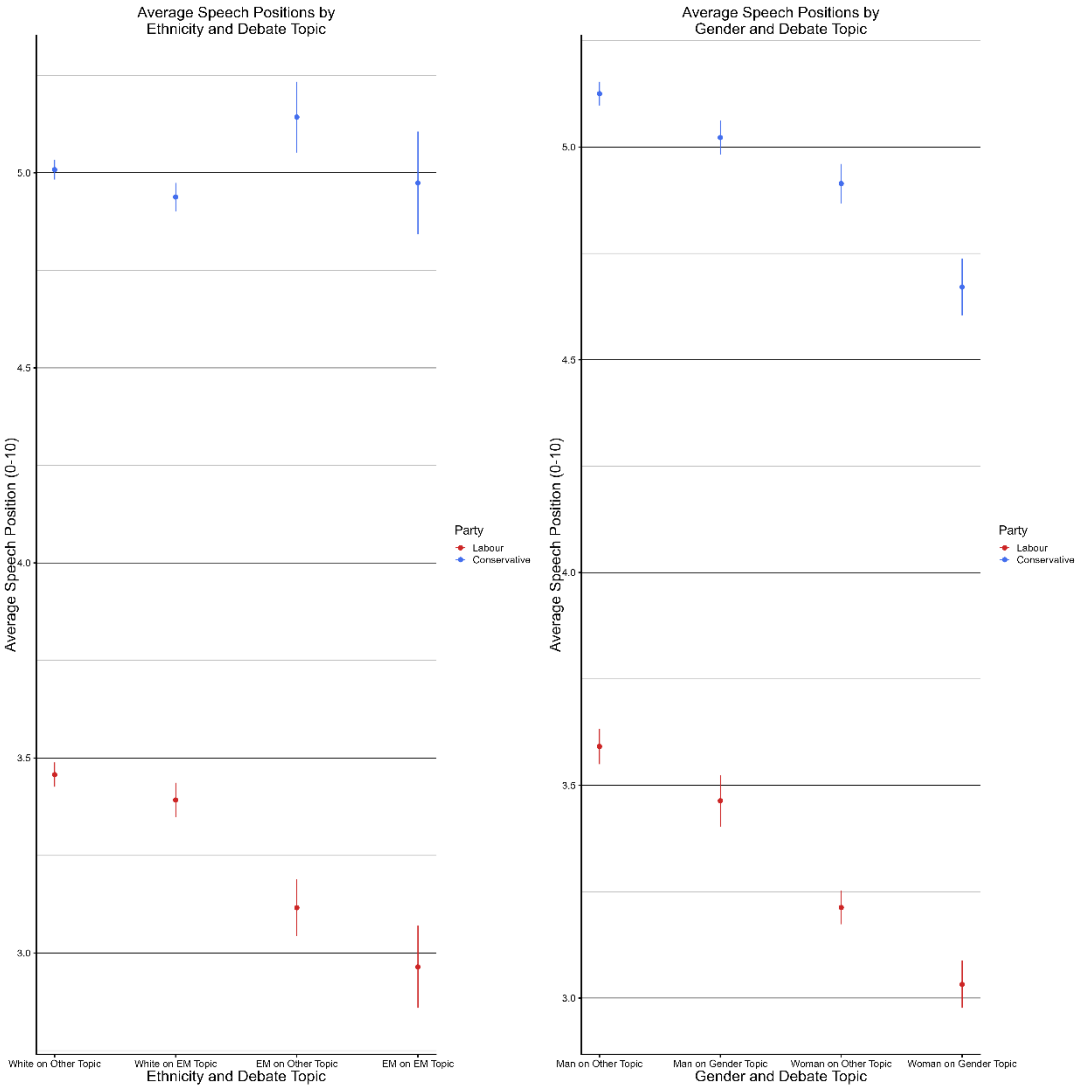

Figure 4.3: Predicted MP positions by ethnicity and whether speaking on a topic relevant to ethnic minority issues; and predicted MP positions by gender and whether speaking on a topic relevant to gender issues. Lower numbers for position represent left-wing positions, and higher numbers are right-wing.

Individual-Level Strategic Hypothesis

The predicted positions of MPs based on their party, demographic, and speech topic, and interacted with their constituency's left-right position are presented in Figure 4.4 for ethnicity, and Figure 4.5 for gender. The *y* axis displays the predicted position of MPs on a 0-10 scale, where 10 is the most right-wing, and the *x* axis is an MP's constituency's left-right placement based on the MRP-estimated prevalence of left-wing economic and cultural views. With respect to ethnicity, H4 is confirmed only

among Conservative MPs. With respect to gender, there are similar, but smaller and more mixed findings.

Ethnic minority Conservative MPs are significantly more sensitive to the left-right political profile of their constituency. This means that the effect observed in Figure 4.3 with respect to H3 is partially conditional on local-level strategic considerations. On average, Conservative MPs represent more right-wing constituencies, meaning that ethnic minorities are more likely to compensate for presumed left-wing positions by emphasising right-wing positions. Where they represent more left-wing constituencies, though, they may be more similar to (or further left than) their white peers. This confirms H4 with respect to ethnic minority MPs: that there are individual-level effects that mean the substantively representative effects of descriptive representation are conditional on local-level considerations about their constituency.

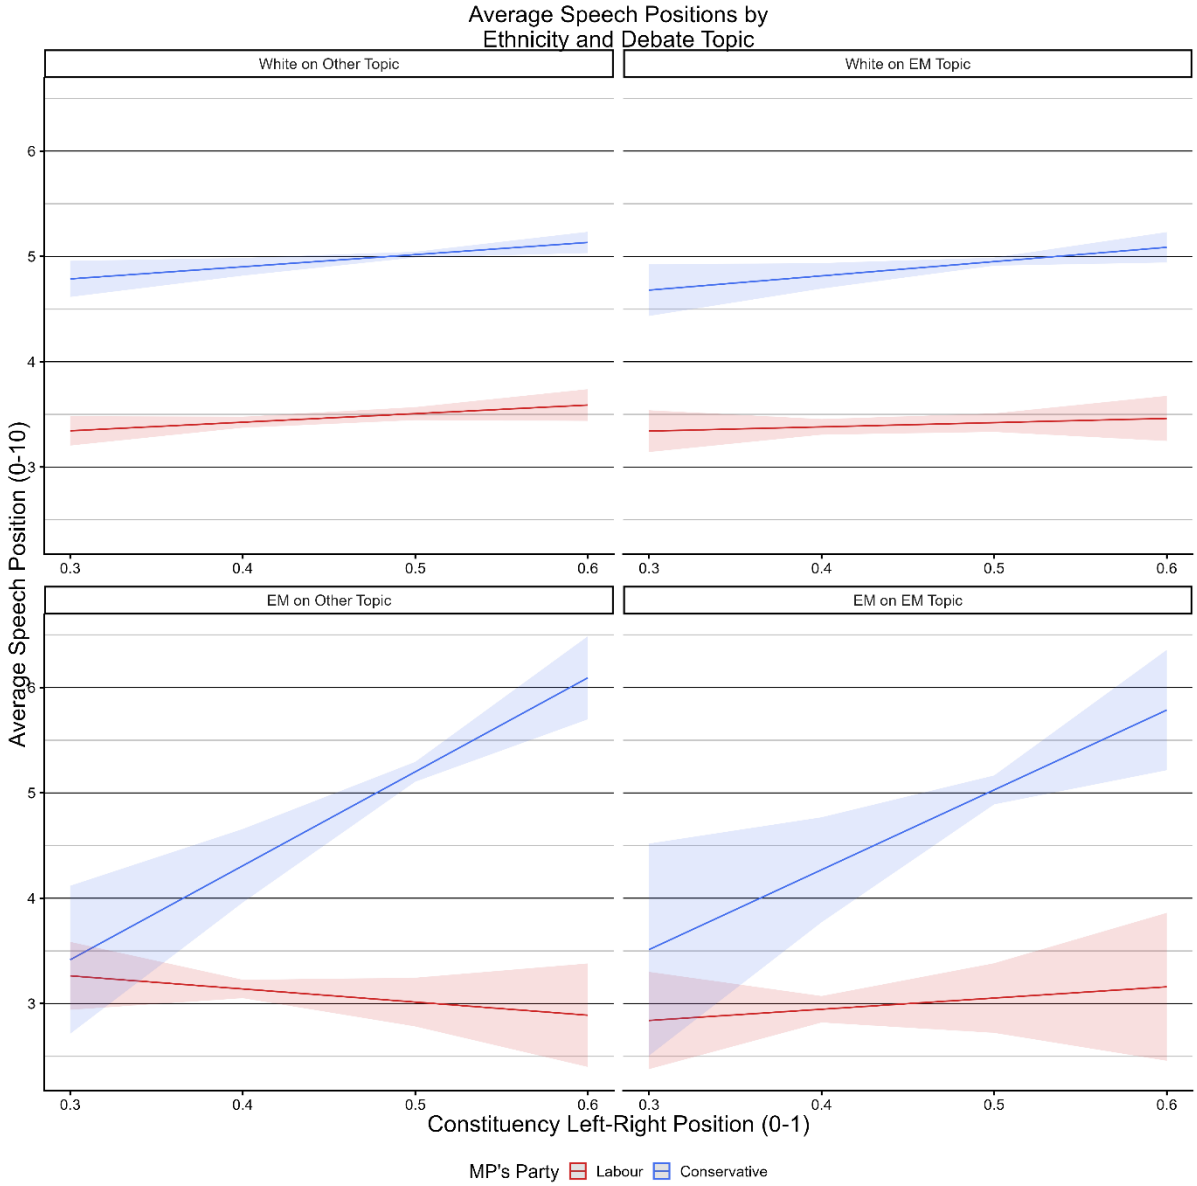

Figure 4.4: Predicted MP positions by ethnicity and whether speaking on a topic relevant to ethnic minority issues, interacted with their constituency's left-right position. Lower numbers for position represent left-wing positions, and higher numbers are right-wing.

Female MPs are found to be more sensitive to overall constituency opinion only on gender issues, and only in the Conservative party. Women speaking on gender issues have the only statistically significant interaction effect with constituency left-right positions, and this effect is only significant among Conservative MPs. This means that H4 applies among female Conservative MPs, but only with respect to gender issues.

This could be due to public pre-conceptions about female candidates being (perceived to be) stronger with respect to gender issues than other issue dimensions, such that compensation for those perceptions is only needed when in an especially right-wing constituency.

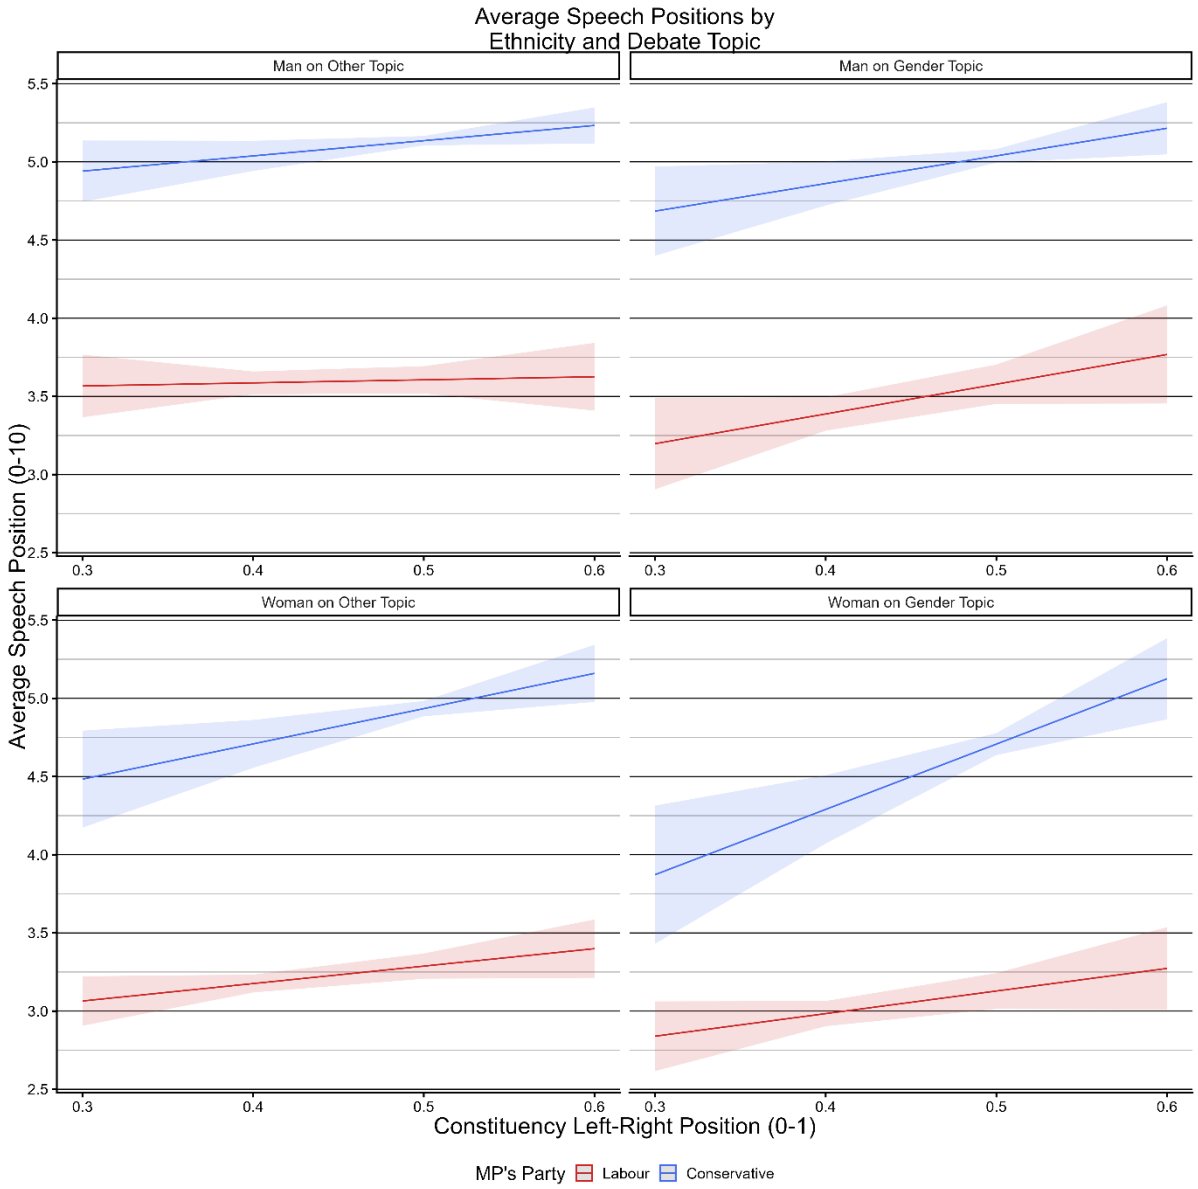

Figure 4.5: Predicted MP positions by gender and whether speaking on a topic relevant to gender issues, interacted with their constituency's left-right position. Lower numbers for position represent left-wing positions, and higher numbers are right-wing.

Conclusions

The relationship between descriptive and substantive representation is complicated by party-level and individual-level strategic considerations. Ethnic minority MPs who belong to right-wing parties have significantly more right-wing positions, especially if they are standing in a more right-wing constituency. The same is not true for female MPs, who are overall more left-wing than male MPs of their own party, especially on gender issues, regardless of party, but with some variation in response to constituency position among Conservatives. This provides important new evidence on the role and quality of descriptive representation in the UK.

The results have shown that descriptive representation has its anticipated effects on MPs' behaviour among all Labour MPs, and among female Conservative MPs. Their speeches reveal more left-wing positions in general, and especially on issues affecting the group to which they belong. This is indicative of MPs positions supportive of the interests of ethnic minorities and women, which would not otherwise be taken by their white or male peers, emphasising the importance of descriptive representation in parliament. Whether these stated positions cause any changes in policy outcomes, however, remains a subject for future work. Though positions are stated in parliament, whether these positions are followed through with (rebellious) voting behaviour, and whether recommendations or criticisms made in those statements are taken seriously by other legislators, are beyond the scope of this chapter. I have, however, analysed one of the links in the chain that ought to connect marginalised groups' interests to policy outcomes: whether the election of members of those groups to parliament yields substantially different representatives, or whether gatekeeping processes by

parties, or strategic incentives among individuals prevent authentic, substantive representation.

Descriptive representation does not, however, have its anticipated effects among ethnic minority Conservative MPs. I find that ethnic minority Conservative MPs are more right-wing than their white peers, and not different to white Conservative MPs on issues directly relevant to ethnic minority interests. This suggests that some of the theories about strategic incentives in response to public stereotypes about ethnic minority candidates are being reflected in real trends in the behaviour of ethnic minority MPs belonging to a right-wing party. This suggests that the strategic considerations are either regarded as more serious, or only regarded as harmful, by the right-wing Conservative Party.

The differing results between ethnic minorities and women may be indicative of differences in the types of barriers, stereotypes, and/or dilemmas they face as representatives. It could be to do with supply-side issues, where the still-small number of ethnic minority representatives (especially in the Conservative party) means that selection effects of party gatekeeping practices have a strong effect on the types of candidates that enter parliament (Norris & Lovenduski, 1994). Women are a far larger group, and the number of female MPs in the Commons is high relative to other marginalised groups (albeit still not representative), reaching 34% (202 MPs) in 2019. Keeping tight control on the ideological profiles of every female candidate would be a much more difficult task than with ethnic minorities, which made up just 10% of the Commons (65 MPs) in 2019 – up from just 6% (41 MPs) in 2015 (Cracknell and Tunnicliffe 2022). The differing results could also be to do with demand-side

issues, which mean that voters' stereotypes about female politicians are weaker or lead fewer voters to attribute left-wing views to female politicians, for example due to stereotypes of female politicians differing from stereotypes about women generally (Schneider and Bos 2014). If so, the demand and supply side are likely deeply linked, as the higher prevalence of female politicians contributes to broadening perceptions about how female politicians behave and what their preferences are. Meanwhile evidence on ethnic minorities has not found inconsistency between stereotypes about politicians and members of the public (van Oosten 2022). Given that both party-level and individual-level differences are observed primarily among ethnic minorities, it is likely that there are both demand- and supply-side effects causing this difference.

The results of this chapter provide an indication of the degree to which descriptive representation is subverted by effects such as party gatekeeping, voters punishing candidates from marginalised groups based on left-wing assumptions about their positions, and parties or candidates strategically compensating for those assumptions. My analysis also distinguishes between mechanisms occurring at the party-level and those occurring at the individual-level. It cannot, however, separate selection effects from behavioural ones. Ethnic minority MPs' behaviour does not align with the expectations of descriptive representation in the Conservative Party, and it is likely this is due either to the effect of how the Conservatives select their candidates, or a strategy that is encouraged of MPs by the party. Their behaviour is, however, also correlated with the profile of the constituency they are in. This suggests that there is also either some individual variance in whether the strategy is adopted, or some constituency-level selection effects causing candidates who do not adopt the strategy not to be elected.

Distinguishing between selection effects and strategic effects will require an analysis that focuses on nominees for party selections, and candidates before elections, to determine the effect of their strategies on selection or election chances. This would need an alternative data source, such as Twitter posts, to measure their political positions. Such an analysis is beyond the scope of this thesis, but would be important to further understanding the mechanisms behind what is observed here.

Appendices

Appendix 1: Regression Table for H1-3

	By Ethnicity	By Gender
Intercept	5.028***	5.063***
	(0.029)	(0.029)
Party: Labour	−1.589***	−1.566***
	(0.021)	(0.024)
White on Ethnic Topic	−0.070**	
	(0.022)	
EM on Other Topic	0.124*	
	(0.048)	
EM on Ethnic Topic	−0.045	
	(0.069)	
White on Ethnic Topic (Labour)	0.006	
	(0.035)	
EM on Other Topic (Labour)	−0.448***	
	(0.063)	
EM on Ethnic Topic (Labour)	−0.430***	
	(0.089)	
Man on Gender Topic		−0.111***
		(0.025)
Woman on Other Topic		−0.184***
		(0.029)
Woman on Gender Topic		−0.422***
		(0.039)
Man on Gender Topic (Labour)		−0.028
		(0.041)
Woman on Other Topic (Labour)		−0.211***
		(0.042)
Woman on Gender Topic (Labour)		−0.155**
		(0.054)
Is a Woman (Control)	−0.292***	
	(0.018)	
Is Ethnic Minority (Control)		−0.150***

	By Ethnicity	By Gender
		(0.026)
Age (Decades)	0.013**	0.012**
	(0.004)	(0.004)
Has Oxbridge Degree	0.105***	0.097***
	(0.018)	(0.017)
Is Backbencher	−0.059**	−0.085***
	(0.018)	(0.018)
Num.Obs.	12612	12612
R2	0.498	0.498
R2 Adj.	0.497	0.498

* p < 0.05, ** p < 0.01, *** p < 0.001

Appendix 2: Regression Table for H4

	By Ethnicity	By Gender
Intercept	4.462***	4.573***
	(0.223)	(0.254)
Party: Labour	−1.341***	−1.141**
	(0.310)	(0.401)
Constituency Left-Right Position	1.153*	0.973
	(0.453)	(0.519)
White on Ethnic Topic	−0.167	
	(0.388)	
EM on Other Topic	−3.702***	
	(0.933)	
EM on Ethnic Topic	−3.203*	
	(1.321)	
Non-EM on Ethnic Topic (Labour)	0.288	
	(0.537)	
EM on Other Topic (Labour)	4.240***	
	(1.114)	
EM on Ethnic Topic (Labour)	2.623	
	(1.566)	
White on Ethnic Topic × Constituency Left-Right Position	0.201	
	(0.790)	
EM on Other Topic × Constituency Left-Right Position	7.772***	

	By Ethnicity	By Gender
	(1.890)	
EM on Ethnic Topic × Constituency Left-Right Position	6.428*	
	(2.678)	
White on Ethnic Topic × Constituency Left-Right Position (Labour)	−0.614	
	(1.148)	
EM on Other Topic × Constituency Left-Right Position (Labour)	−9.833***	
	(2.376)	
EM on Ethnic Topic × Constituency Left-Right Position (Labour)	−6.178	
	(3.343)	
Man on Gender Topic		−0.494
		(0.449)
Woman on Other Topic		−0.842
		(0.477)
Woman on Gender Topic		−2.030**
		(0.628)
Man on Gender Topic (Labour)		−0.386
		(0.704)
Woman on Other Topic (Labour)		0.064
		(0.619)
Woman on Gender Topic (Labour)		0.928
		(0.782)
Man on Gender Topic × Constituency Left-Right Position		0.794
		(0.914)
Woman on Other Topic × Constituency Left-Right Position		1.282
		(0.971)
Woman on Gender Topic × Constituency Left-Right Position		3.203*
		(1.277)
Man on Gender Topic × Constituency Left-Right Position (Labour)		0.911
		(1.523)
Woman on Other Topic × Constituency Left-Right Position (Labour)		−0.363
		(1.317)
Woman on Gender Topic × Constituency Left-Right Position (Labour)		−1.955
		(1.657)
Is a Woman (Control)	−0.311***	

	By Ethnicity	By Gender
	(0.017)	
Is Ethnic Minority (Control)		−0.153***
		(0.026)
Age (Decades)	0.015***	0.015***
	(0.004)	(0.004)
Has Oxbridge Degree	0.101***	0.102***
	(0.018)	(0.018)
Is Backbencher	−0.052**	−0.070***
	(0.018)	(0.018)
Num.Obs.	12445	12445
R2	0.502	0.502
R2 Adj.	0.501	0.501

* p < 0.05, ** p < 0.01, *** p < 0.001

Chapter 5: Conclusion

Through three substantive chapters, I have shown how speech data can be a valuable data source to test theories about legislative behaviour. Whereas a vote for or against a bill in parliament can be for any number of reasons, through speeches legislators tell us those reasons. They justify their votes by framing them from the political positions they choose to support, as well as frequently providing general opinions and critiques on policies, posing questions to ministers and other legislators, and speaking from personal experiences.

Whilst many researchers have identified speech data as key to analysing legislative behaviour, methods have been varied, and findings inconsistent. I provide new improvements on existing LLM methods, introducing chain-of-thought methods as a way to improve accuracy in coding tasks. I also acknowledge that LLMs are not a panacea for all natural language processing tasks. I use clustering methods for the analysis of partisan loyalty: a task which is consistent with the assumptions of n-gram models, as it is indicated through the use of buzzwords and references to party policy agendas. LLMs are a highly accessible and usable tool, and their use in social science currently lacks strong precedents that prevent misuse, ensure reproducibility, and promote interpretability. This is therefore a vital time to be establishing norms of good practice, including the use of local LLMs, as well as methods which maximise accuracy, such as the introduction of chain-of-thought methods.

A common limitation of these methods is often cost and replicability. In part due to the funding constraints of a DPhil project, I've minimised costs. As well as minimising costs (which are themselves often a barrier to replicability), I've used a local LLM (Llama 3 via Ollama). Local LLMs are good practice, as models that can

only be accessed via APIs are liable to being discontinued, making all research that used them non-replicable. More powerful local LLMs are available, however, and may be a direction on which these results can not just be replicated, but improved, in future research with larger budgets. Similarly, validation by direct comparison to a third-party expert-coded sample of speeches from the dataset would be desirable, but was beyond the scope of this project.

I've applied these methods to answer three questions about legislative behaviour. The question of how electoral systems can affect party discipline, of how constituency-level electoral threats can influence legislators' positions in parliament, and of the role of whether descriptive representation does indeed lead to substantive differences in legislators' behaviour.

In each of these, I have yielded new evidence. I have shown that legislators' speeches become more partisan (and that their conscience votes align with partisan divides) after adopting a closed-list proportional mixed member system. I have shown that legislators are not generally responsive to constituency opinion, that their positions in parliament move significantly to the right when radical right parties poll or perform well in their constituency, and that there is no such response for a centre or centre-left party. I have also shown that ethnic minority and female MPs are generally more left-wing, especially on ethnic and gender issues respectively, but that this does not hold true for ethnic minority MPs in the Conservative Party, who are generally more right-wing, and indistinguishable from white Conservative MPs on ethnic issues.

These are important findings that are informative about legislative behaviour. They give support to a number of theories in the literature, and reject some others. An

important limitation, however, is that these chapters only examine one (albeit important) site of legislative behaviour: speeches in the main chamber. This is a good site to analyse, given their direct connection to votes on legislation, and the highly publicised Prime Minister's Questions (and other ministerial question times), which give MPs opportunities to justify their voting intentions, and to raise issues of their own. Activity on select committees, written questions, and the content of bills and amendments, however, are another important site. Whether similar findings are reproducible in these sites of behaviour could be informative about the effects of whipping procedures (which may be tighter in the main chamber, where legislation is at stake and there is greater public attention).

Legislators' activity on social media is also important, as an informal and more directly public-facing site where political positions are expressed. Researchers in this area, however, should be careful with how they handle topics. Some MPs may be cultivating single-issue political profiles, which are easier to maintain on social media where they have total control over what they speak about and when. Parliamentary speech has the advantage of taking place on an externally set agenda. Whilst MPs still have significant freedom about when they do or do not speak, there are more likely to be speeches from every MP on a broad range of topics. In my spatial analyses, I have used at least two axes of political positioning, to at least in part mitigate these effects. This means that, if an MP focuses extensively on the cost of living and austerity, almost all of these speeches will be grouped in the economic dimension, and their views on immigration or military interventionism still play a key role in measuring their positions on the second dimension. This is also more consistent with the meaning of placing someone left or right. A high salience of cost of living and

austerity may bias their position leftwards, but if those are their priorities within economic policy, that is indicative of a left-wing position. If social and international issues were not separated out, however, it would be wrong to assume the focus on these issues is indicative of a left-wing position overall, if they also held positions opposed to gay marriage or immigration. The more issue dimensions that can be constructed, the better the problems of issue salience can be handled. The number of dimensions, however, is limited by data availability on each desired dimension.

My findings have been limited to just two countries: New Zealand and the UK. They do, however, have intuitive implications for other countries – especially those with legislators representing single-member districts – but readers should consider these implications with care. Electoral systems vary, and within them party selections vary, too. For example, although nominally the USA has a SMDP electoral system, it fundamentally differs to the UK through its federalism, elective bicameralism, party primaries, presidentialism, and extensive gerrymandering (Jeong & Shenoy, 2024). All of these factors may have drastic consequences for the expected behaviour of legislators. Federalism may embolden representatives of specific states, and create norms and factionalisms that vary on state lines (Cherny, 2023). Primaries may also incentivise factionalism and 'the personal vote', as ambitious individuals see clear payoffs if they can win a popular public platform (Carey & Shugart, 1995), and factions may gather around future primary candidates as much as current party leaders. Presidentialism affects house majorities, and therefore the role of rebellious voting and behaviour within the legislative system, given legislation regularly has to pass through minority chambers (Bianco & Smyth, 2020). Gerrymandering may systemically increase and decrease the re-election threats of specific legislators

depending on the success of their party in control of boundary re-drawing (Bertelli & Carson, 2011). This may have not just direct effects on behaviour, but contribute to a culture of rebellion, such as frequent appeals to bipartisanship[60]. Generalisability, therefore, cannot be taken for granted in any study of legislative behaviour. Similar analyses in other countries would be important work for future studies where possible[61].

Throughout these chapters, parties have played a core role, but whereas I have ensured that changes in party positions are not a potential confounder, I have not directly analysed changes to party positions. In Chapter 2, party positions were accounted for by comparing speeches to their most recent manifesto documents. In Chapter 3, by controlling for current party leaderships. In Chapter 4, by the models being cross-sectional. The adjustment of MPs to new party platforms and leaderships, however, would be an interesting direction for future study. As party leaderships, and party-level strategic considerations, change, parties often reinvent themselves and their platforms, but do not completely replace their personnel. Which legislators adjust to the new position of their party? Does failing to adjust have identifiable consequences for re-election or re-selection? And do those who do not adjust adopt more extreme views within their party? There are a wealth of theories from the literature on party politics, often using manifesto or expert survey data, concerning

[60] This, especially, could have effects on speech analysis methods. Justifications for rebellious behaviour which reference bipartisanship could be assumed to be moderate, when sometimes interests align at the fringes of each party.
[61] The results of Chapter 2 could only be reproduced after the fact of another electoral system change. It is therefore already a vital source of information for considering options for electoral reform in FPTP countries.

how and why parties change their positions, which provide a basis for deriving implications for the behaviour of individual MPs.

Within its own scope, however, this thesis has provided important new evidence about legislative behaviour in New Zealand and the UK. Legislators' speeches in parliament reflect strategic positions, often following patterns that conform to theoretical expectations, but with important exceptions that had not previously been identified and proven empirically.

Future research should continue to develop the science of analysing speech and natural language data. In the study of social science, the importance of spoken and written communication cannot be stressed enough. Through new tools and practices, new approaches to social science are available to us, with wide-ranging implications beyond the study of legislatures, many of which are already being explored. These include the use of social media text data to estimate vote choice (Cerina & Duch, 2024), the use of LLMs as interviewers to collect data (especially for open-ended responses) (Wuttke, et al., 2025), and measuring changes and trends in news media (Barrie, Ketchley, Siegel, & Bagdouri, 2025).

Even where researchers do not intend to deploy these methods themselves, the creation of new datasets using these methods will make a broader set of possible variables available for all research. For example, if analysing the effect of political engagement on vote choice over time, it may be prudent to include a measure of left-right bias in the media, or the favourability of media sources to particular parties. The effect direction of political engagement may be conditional on the content of media that is accessed by more politically-engaged people. Likewise, if analysing the effects

of constituency-level campaign activities on vote choice, it may be important to control for the position of each voter's local MP relative to their party, as this may have impacts on parties' resource allocations, and on voters' choices. Whilst these kinds of variables have typically been difficult to measure, with increased awareness, continued work on the methods and standards in the handling of text data, and open science practices, they may become standard tools in the arsenal of social scientists.

Bibliography

Abou-Chadi, T. (2014). Niche Party Success and Mainstream Party Policy Shifts – How Green and Radical Right Parties Differ in Their Impact. *British Journal of Political Science, 46*(2), 417-436.

Abou-Chadi, T., & Krause, W. (2018). The Causal Effect of Radical Right Success on Mainstream Parties' Policy Positions: A Regression Discontinuity Approach. *British Journal of Political Science, 50*(3), 829-847.

Acquarone, I. E., & di Landro, G. (2024). Historically marginalized groups and ideological representation in legislatures. *Legislative Studies Quarterly, 7*, 00: 1-8.

Adams, J., & Somer-Topcu, Z. (2009). Policy Adjustment by Parties in Response to Rival Parties' Policy Shifts: Spatial Theory and the Dynamics of Party Competition in Twenty-Five Post-War Democracies. *British Journal of Political Science, 39*(4), 825-846.

AICBT. (2015). *Authorship Attribution with Python*. Retrieved February 25, 2021, from https://aicbt.com/authorship-attribution/

Andersen, R., Tilley, J., & Heath, A. F. (2005). Political Knowledge and Enlightened Preferences: Party Choice Through the Electoral Cycle. *British Journal of Political Science, 35*(2), 285-302.

Ansolabehere, S., & Jones, E. (2010). Constituents' Responses to Congressional Roll-Call Voting. *American Journal of Political Science, 54*(3), 583-835.

Arndt, C., & Christiansen, F. J. (2022). Is disloyalty rewarded? The electoral consequences ofbloc changes of Scandinavian Centre Parties 1977–2019. *Scandinavian Political Studies, 45*(4), 383-575.

Babaki, B. (2020, May 12). *MinSizeKmeans*. Retrieved January 17, 2021, from https://github.com/Behrouz-Babaki/MinSizeKmeans#readme

Bai, X., Wang, A., Sucholutsky, I., & Griffiths, T. L. (2025). Explicitly unbiased large language models still form biased associations. *PNAS, 122*(8), 1-9.

Barker, F., & Lavine, S. (1999). The Individual Parliamentary Member and Institutional Change: the Changing Role of the New Zealand Member of Parliament. *The Journal of Legislative Studies*, 5 (3-4): 105-130.

Barker, F., & McLeay, E. (2000). How Much Change? An Analysis of the Initial Impact of Proportional Representation on the New Zealand Parliamentary Party System. *Party Politics*, 6 (2): 131-154.

Barrie, C., Ketchley, N., Siegel, A., & Bagdouri, M. (2025). Measuring Media Criticism with ALC Word Embeddings. *Political Analysis*, 1-17.

Barrie, C., Palaiologou, E., & Törnberg, P. (2025). Prompt Stability Scoring for Text Annotation with Large Language Models. *arXiv*.

Barrie, C., Palmer, A., & Spirling, A. (2025). Replication for Large Language Models: Problems, Principles, and Best Practices for Political Science. *Unpublished Manuscript*.

BBC News. (2016, October 25). *Zac Goldsmith quits as MP over Heathrow decision*. Retrieved June 1, 2020, from Zac Goldsmith quits as MP over Heathrow decision

BBC2 England. (2015, February 4). *Inside the Commons, Lifting the Lid (00:46:24)*. Retrieved February 25, 2021, from https://learningonscreen.ac.uk/ondemand/index.php/prog/08B00ACF?bcast=115047860

Bertelli, A. M., & Carson, J. L. (2011). Small changes, big results: Legislative voting behavior in the presence of new voters. *Electoral Studies, 30*(1), 201-209.

Bianco, W., & Smyth, R. (2020). The Bicameral Roots of Congressional Deadlock: Analyzing Divided Government Through the Lens of Majority Rule. *Social Science Quarterly, 101*(5), 1712-1727.

Bird, K. (2005). The Political Representation of Visible Minorities in Electoral Democracies: A Comparison of France, Denmark, and Canada. *Nationalism and Ethnic Politics, 11*(4), 425-465.

Bowler, S., Farrell, D. M., & Katz, R. S. (1999). *Party discipline and parliamentary government.* Columbus: Ohio State University Press.

Bozzano, M., & Scabrosetti, S. (2024). Preferences for Redistribution: Two Decades of Gender Gaps and Generational Differences in Europe. *CESifo Economic Studies, 70*(3), 320-345.

Brooks, L. (2025, August 12). *'I've been spat on': gender non-conforming women tell of toilet abuse after UK's supreme court ruling*. Retrieved from The Guardian:

https://www.theguardian.com/world/2025/aug/12/ive-been-spat-on-gender-non-conforming-women-tell-of-toilet-abuse-in-aftermath-of-supreme-court-ruling

Burke, E. (1774). Speech to the Electors of Bristol. In P. B. Kurland, & R. Lerner, *The Founders' Constitution.* Chicago: University of Chicago Press.

Butler, D., & Nickerson, D. (2011). Can Learning Constituency Opinion Affect How Legislators Vote? Results from a Field Experiment. *Quarterly Journal of Political Science, 6*, 55-83.

Cain, B., Ferejohn, J., & Fiorina, M. (1987). *The personal vote: constituency service and electoral independence.* London: Harvard University Press.

Campbell, D. T., & Ross, H. L. (1968). The Connecticut Crackdown on Speeding: Time-Series Data in Quasi-Experimental Analysis. *Law & Society Review*, 3 (1): 33-53.

Campbell, R., & Shorrocks, R. (2023). *British Social Attitudes 40 | The evolution of the gender gap?* London: National Centre for Social Research.

Canes-Wrone, B., Brady, D. W., & Cogan, J. F. (2002). Out of Step, Out of Office: Electoral Accountability and House Members' Voting. *American Political Science Review, 96*(1), 127-140.

Canes-Wrone, B., Minozzi, W., & Reveley, J. B. (2011). Issue Accountability and the Mass Public. *Legislative Studies Quarterly, 36*(1), 5-35.

Carey, J. M. (2007). Competing Principals, Political Institutions, and Party Unity in Legislative Voting. *American Journal of Political Science*, 51 (1): 92-107.

Carey, J. M. (2009). *Legislative Voting and Accountability.* Cambridge: Cambridge University Press.

Carey, J. M., & Shugart, M. (1995). Incentives to Cultivate a Personal Vote: a Rank Ordering of Electoral Formulas. *Electoral Studies*, 14 (4): 417-439.

Carroll, R., & Nalepa, M. (2020). The Personal Vote and Party Cohesion: Modeling the Effects of Electoral Rules on Intraparty Politics. *Journal of Theoretical Politics*, 32 (1): 36-69.

Carson, J. L., Koger, G., Lebo, M. J., & Young, E. (2010). The Electoral Costs of Party Loyalty in Congress. *American Journal of Political Science, 54*(3), 583-835.

Cerina, R., & Duch, R. (2024). The 2024 US Presidential Election PoSSUM Poll. *Political Science & Politics, 58*(2), 286-297.

Ceron, A. (2015). Brave Rebels Stay Home: Assessing the Effect of Intra-Party Ideological Heterogeneity and Party Whip on Roll-Call Votes. *Party Politics*, 21 (2): 246-258.

Cherny, N. (2023). Do Provincial Governors Control National Legislators? Quasi-Experimental Evidence of Argentine Federalism. *Bulletin of Latin American Research, 42*(5), 622-634.

Chiru, M. (2021). Does Electoral Reform Change MPs' Behavior? Evidence from Romania. *Political Studies Review*, 19 (3): 355-375.

Clayton, A., O'Brien, D. Z., & Piscopo, J. M. (2019). All Male Panels? Representation and Democratic Legitimacy. *American Journal of Political Science, 63*(1), 1-264.

Collignon, S. (2024). Addressing Barriers to Women's Representation in Party Candidate Selections. *The Political Quarterly, 95*(2), 363-367.

Cowley, P. (2014). *Almost 70% of people know the name of their MP*. Retrieved March 23, 2025, from https://www.britishelectionstudy.com/bes-findings/almost-70-of-people-know-the-name-of-their-mp-by-professor-phil-cowley-university-of-nottingham/

Cox, C., & Shapiro, I. (2025). Party Realignment and Single-Issue Voters. *PLoS One, 20*(3).

Cox, G. W., & McCubbins, M. D. (1994). Bonding, Structure, and the Stability of Political Parties: Party Government in the House of Representatives. *Legislative Studies Quarterly, 19*(2), 215-231.

Cracknell, R., & Tunnicliffe, R. (2022). *Social Background of MPs 1979-2019.* London: House of Commons Library.

Culpepper, P. D. (2010). *Quiet Politics and Business Power: Corporate Control in Europe and Japan.* Cambridge: Cambridge University Press.

Dancygier, R. M. (2017). *Dilemmas of Inclusion: Muslims in European Politics.* Princeton: Princeton University Press.

Delli Carini, M. X., & Keeter, S. (1991). Stability and Change in the U.S. Public's Knowledge of Poltiics. *Public Opinion Quarterly, 55*(4), 583-612.

Farrer, B. D., & Zingher, J. N. (2018). Explaining the nomination of ethnic minority candidates: how party-level factors and district-level factors interact. *Journal of Elections, Public Opinion and Parties, 28*(4), 467-487.

Fieldhouse, E., Green, J., Evans, G., Mellon, J., Prosser, C., Bailey, J., . . . Perrett, S. (2024). *British Election Study Internet Panel Waves 1-29*.

Fishel, J. (1972). On the Transformation of Ideology in European Political Systems: Candidates for the West German Bundestag. *Comparative Political Studies*, 4 (4): 406-437.

Fisher, S. D., Heath, A. F., Sanders, D., & Sobolewska, M. (2015). Candidate Ethnicity and Vote Choice in Britain. *British Journal of Political Science, 45*(4), 883-905.

Flavin, P., & Law, W. (2022). Ideological proximity, issue importance, and vote choice. *Electoral Studies, 75*.

Fouirnaies, A., & Hall, A. (2022). How Do Electoral Incentives Affect Legislator Behavior? Evidence from U.S. State Legislatures. *American Political Science Review, 116*(2), 662-676.

Fox, R. L., & Lawless, J. L. (2024). The Invincible Gender Gap in Political Ambition. *PS: Political Science and Politics, 57*(2), 231-237.

Garner, C., & Letki, N. (2005). Party Structure and Backbench Dissent in the Canadian and British Parliaments. *Canadian Journal of Political Science, 38*(2), 463-482.

Gaughan, C. (2024). Estimating Ideal Points of British MPs Through Their Social Media Followership. *British Journal of PolitIcal Science*, 1506-1518.

Goffman, E. (1956). *The Presentation of Self in Everyday Life.* Edinburgh: University of Edinburgh, Social Sciences Research Centre.

Gold, H. (1989). *New Zealand Politics in Perspective.* Auckland: Longman Paul.

Gormley, W. T. (1986). Regulatory Issue Networks in a Federal System. *Polity, 18*(4), 595-620.

Greene, Z., & Haber, M. (2014). Leadership Competition and Disagreement at Party National Congresses. *British Journal of Political Science*, 46 (3): 611-632.

Greene, Z., & Haber, M. (2015). The consequences of appearing divided: An analysis of party evaluations and vote choice. *Electoral Studies, 37*, 15-27.

Greene, Z., & O'Brien, D. Z. (2016). Diverse parties, diverse agendas? Female politicians and the parliamentary party's role in platform formation. *European Journal of Political Research, 55*, 435-453.

Hanretty, C., & Lazarov, V. (2025). The left-right positions of British MPs inferred from a survey of local councillors. *Working paper*. Retrieved from https://github.com/chrishanretty/pairwise_mps_2024/blob/main/article/article.pdf

Hanretty, C., Mellon, J., & English, P. (2021). Members of Parliament are Minimally Accountable for Their Issue Stances (and They Know It). *American Political Science Review, 115*(4), 1275-1291.

Hansard. (2019). End of Life Choice Bill — Third Reading. *New Zealand Parliament (Pāremata Aotearoa).* Wellington: Hansard. Retrieved from https://www.parliament.nz/en/pb/hansard-debates/rhr/combined/HansDeb_20191113_20191113_16

Hazan, R. Y., & Rahat, G. (2010). *Democracy within Parties: Candidate Selection Methods and Their Political Consequences.* Oxford: Oxford University Press.

Heath, A. F., Fisher, S. D., Rosenblatt, G., Sanders, D., & Sobolewska, M. (2013). The Ethnic Minority Agenda(s). In A. F. Heath, S. D. Fisher, G. Rosenblatt, D. Sanders, & M. Sobolewska, *The Political Integration of Ethnic Minorities in Britain* (pp. 60-80). Oxford: Oxford University Press.

Heseltine, M., & von Hohenberg, B. C. (2024). Large Language Models as a Substitute for Human Experts in Annotating Political Text. *Research and Politics, 11*(1), 1-10.

Hirschman, A. O. (1970). *Exit, Voice and Loyalty: Responses to Decline in Firms, Organisations, and States.* Cambridge, M.A.: Harvard University Press.

Hix, S. (2004). Electoral Institutions and Legislative Behaviour: Explaining Voting Defection in the European Parliament. *World Politics*, 56 (2): 194-223.

Hjermitslev, I. B. (2020). The electoral cost of coalition participation: Can anyone escape? *Party Politics, 26*(4), 510-520.

Homola, J. (2017). Are Parties Equally Responsive to Women and Men? *British Journal of Political Science, 49*, 957-975.

Huber, J. D., & Powell, G. B. (1994). Congruence Between Citizens and Policymakers in Two Visions of Liberal Democracy. *World Politics*, 291-326.

Hug, S. (2009). Selection Effects in Roll Call Votes. *British Journal of Political Science, 40*, 225–235.

Hughes, M. M. (2013). Diversity in National Legislatures Around the World. *Sociology Compass, 7*(1), 23-33.

Hunt, R. N. (1964). *German Social Democracy.* New Haven: Yale University Press.

Ipsos. (2024). *Political Attitudes By Ethnicity.* Paris: Ipsos.

Jacobsmeier, M. L. (2015). From Black and White to Left and Right: Race, Perceptions of Candidates' Ideologies, and Voting Behavior in U.S. House Elections. *Political Behavior, 37*, 595-621.

Jeong, D., & Shenoy, A. (2024). The Targeting and Impact of Partisan Gerrymandering: Evidence from a Legislative Discontinuity. *The Review of Economics and Statistics, 106*(3), 814-828.

Kam, C. (2009). *Party Discipline and Parliamentary Politics.* Cambridge: Cambridge University Press.

Ko, H. (2015). The Changing Shape of Welfare Policy in Korea: Analysis of Break Point in the Welfare Budget. *Asian Social Work and Policy Review*, 9 (2): 151-162.

Koch, J. (2002). Gender Stereotypes and Citizens' Impressions of House Candidates' Ideological Orientations. *American Journal of Political Science, 46*(2), 453-462.

Kojima, T., Gu, S. S., Reid, M., Matsuo, Y., & Iwasawa, Y. (2022). Large Language Models are Zero-Shot Reasoners. *arXiv.org preprint*.

Krippendorff, K. (2022). *The Reliability of Generating Data.* New York: CRC Press.

Lajevardi, N., Mårtensson, M., & Vernby, K. (2024). Do Minorities Feel Welcome in Politics? A Cross-Cultural Study of the United States and Sweden. *British Journal of Political Science, 54*(4), 1435-1444.

Lauderdale, B. E., & Herzog, A. (2016). Measuring Political Positions from Legislative Speech. *Political Analysis*, 374-394.

Laver, M., Benoit, K., & Garry, J. (2003). Extracting Policy Positions from Political Texts Using Words as Data. *American Political Science Review*, 97 (2): 311-331.

Le Mens, G., & Gallego, A. (2023). Positioning Political Texts with Large Language Models by Asking and Averaging. *arXiv.org preprint*.

Lijphart, A. (1989). Democratic Political Systems: Types, Cases, Courses and Consequences. *Party Politics*, 19 (3): 414-429.

Lindsey, D. (2008). A Brief History of Conscience Voting in New Zealand. *Australian Parliamentary Review*, 23 (1): 144-171.

Lindsey, D. (2011). Conscience Voting in New Zealand. Dissertation, University of Auckland.

Martin, N. S., & Blinder, S. (2020). Biases at the Ballot Box: How Multiple Forms of Voter Discrimination Impede the Descriptive and Substantive Representation of Ethnic Minority Groups. *Political Behaviour, 43*, 1487-1510.

McClean, I., & Spirling, A. (2007). UK OC OK? Interpreting Optimal Classification Scores for the U.K. House of Commons. *Political Analysis, 15*, 85-96.

McDermott, M. L. (1998). Race and Gender Cues in Low-Information Elections. *Political Research Quarterly, 51*(4), 847-1013.

McDermott, M. L. (2016). *Masculinity, Femininity, and American Political Behavior.* New York: Oxford Academic.

Meijers, M. J. (2017). Contagious Euroscepticism: The impact of Eurosceptic support on mainstream party positions on European integration. *Party Politics, 23*(4), 413-423.

Meyerrose, A., & Watson, S. (2024). The Effects of Import Shocks, Electoral Institutions, and Radical Party Competition on Legislator Ideology: Evidence from France. *British Journal of Political Science, 54*, 547-572.

Mikkola, M. (2022). Feminist Perspectives on Sex and Gender. *Stanford Encyclopedia of Philosophy*.

Mohyuddin, Z. (2024). *The myth and reality of ethnic minority support for Labour.* London: UK in a Changing Europe.

Muggeo, V. M. (2020, June 23). Package 'segmented': Regression Models with Break-Points / Change-Points Estimation.

Munis, B. K. (2021). Place, candidate roots, and voter preferences in an age of partisan polarisation: observational and experimental evidence. *Political Geography, 85*.

Musazzi, U. M., Rocco, P., Brunelli, C., Bisaglia, L., Caraceni, A., & Minghetti, P. (2018). Do laws impact opioids consumption? A breakpoint analysis based on Italian sales data. *Journal of Pain Research*, 11: 1665-1672.

mySociety. (2023). *ParlParse: Overview*. Retrieved from Theyworkforyou: http://parser.theyworkforyou.com/

New Zealand Labour Party Incorporated. (2019, November 30). Constitution and Rules 2019. *Party Documents*. Retrieved from https://drive.google.com/file/d/1AW3cu4RtVdGMW67FGhZTh2lDv5TJ4d27/view

New Zealand Parliament. (2001, December 21). Electoral (Integrity) Amendment Act 2001.

New Zealand Parliament Paremata Aotearoa. (2010, October 14). *Chapter 7 Parties and Government*. Retrieved June 15, 2020, from https://www.parliament.nz/en/document/00HOOOCPPNZ_71

New Zealand Parliament Paremata Aotearoa. (2018, October 3). Electoral (Integrity) Amendment Act 2018.

Norris, P. (1995). May's Law of Curvilinear Disparity Revisited: Leaders, Officers, Members and Voters in British Political Parties. *Party Politics, 1*(1), 29-47.

Norris, P., & Lovenduski, J. (1994). *Political Recruitment: Gender, Race and Class in the British Parliament.* Cambridge: Cambridge University Press.

Nye, M., Andreassen, A. J., Gur-Ari, G., Michalewski, H., Austin, J., Bieber, D., . . . Odena, A. (2021). Show Your Work: Scratchpads for Intermediate Computation with Language Models. *arXiv.org preprint*.

Osnabrügge, M. (2021). 30 New Zealand: Parliamentary Speeches under a Mixed-Member Proportional Electoral System. In H. Back, M. Debus, & J. M.

Fernandes, *The Politics of Legislative Debates* (pp. 613-632). Oxford: Oxford University Press.

Osnabrügge, M., Ash, E., & Morelli, M. (2020). Measuring Topics Using Cross-Domain Supervised Learning: Methods and Application to New Zealand Parliament. *Center for Law & Economics Working Paper Series, ETH Zürich*.

O'Sullivan, D. (2018, January 31). *Conscience vote on euthanasia bill exposes democratic weakness of New Zealand's voting system*. (T. Conversation, Editor) Retrieved 05 28, 2020, from The Conversation: https://theconversation.com/conscience-vote-on-euthanasia-bill-exposes-democratic-weakness-of-new-zealands-voting-system-90838

Palmer, M. (1995). Toward an Economics of Comparative Political Organization: Examining Ministerial Responsibility. *Journal of Law, Economics, and Organization, 11*(1), 164-188.

Parsons, A. (2024). *Who is your MP? Factors that lead to greater and decreased knowledge of a person's MP.* London: mySociety.

Patty, J. W. (2008). Equilibrium Party Government. *American Journal of Political Science*, 636-655.

Pitkin, H. F. (1967). *The Concept of Representation.* Oakland: University of California Press.

Poole, K., & Rosenthal, H. (1985). A Spatial Model for Legislative Roll Call Analysis. *American Journal of Political Science*, 29 (2): 357-384.

Poole, K., & Rosenthal, H. (2017). *Ideology and Congress: A Political Economic History of Roll Call Voting.* Abingdon: Routledge.

Proksch, S.-O., & Slapin, J. B. (2015). *The Politics of Parliamentary Debate.* Cambridge: Cambridge University Press.

Proksch, S.-O., Lowe, W., Wäckerle, J., & Soroka, S. (2019). Multilingual Sentiment Analysis: A New Approach to Measuring Conflict in Legislative Speeches. *Legislative Studies Quarterly, 44*(1), 1-198.

Ranehill, E., & Weber, R. A. (2025). Gender preference gaps and voting for redistribution. *Experimental Economics, 25*(3).

Reeve, J. (2020). *Text-Matcher*. Retrieved from https://github.com/JonathanReeve/text-matcher

Řehůřek, R., & Sojka, P. (2010). Software Framework for Topic Modelling with Large Corpora. In R. Řehůřek, & P. Sojka, *Proceedings of the LREC 2010 Workshop on New* (pp. 40-45). Valletta, Malta: ELRA.

Renwick, A. (2018). Electoral System Change. In E. S. Herron, R. J. Pekkanen, & M. S. Shugart, *The Oxford Handbook of Electoral Systems.* Oxford: Oxford University Press.

Rice, S. A. (1925). The Behavior of Legislative Groups: A Method of Measurement. *Political Science Quarterly*, 40 (1): 60-72.

Rich, T. S. (2014). Party Voting Cohesion in Mixed Member Legislative Systems: Evidence from Korea and Taiwan. *Legislative Studies Quarterly*, 39 (1): 113-135.

Rivers, C. (2012). *The Congressional Black Caucus, Minority Voting Rights, and the U. S. Supreme Court.* Ann Arbor: University of Michigan Press.

Roberts, N. S. (2001). New Zealand. In D. Nohlen, F. Grotz, & C. Hartmann, *Elections in Asia and the Pacific : A Data Handbook: Volume II: South East Asia, East Asia, and the South Pacific.* Oxford: Oxford University Press.

Rush, M., & Giddings, P. (2011). Parliamentary careers and ministerial ambitions. In M. Rush, & P. Giddings, *Parliamentary Socialisation: Learning the Ropes or Determining Behaviour?* (pp. 120-121). London: Palgrave Macmillan UK.

Sanders, D., Heath, A., Fisher, S., & Sobolewska, M. (2014). The Calculus of Ethnic Minority Voting in Britain. *Political Studies, 62*, 230-251.

Schneider, M. C., & Bos, A. L. (2014). Measuring Stereotypes of Female Politicians. *Political Psychology, 35*(2), 133-305.

Schwarz, D., Traber, D., & Benoit, K. (2017). Estimating Intra-party Preferences: Comparing Speeches to Votes. *Political Science Research and Methods*, 5 (2): 379-396.

Sieberer, U. (2010). Behavioral consequences of mixed electoral systems: Deviating voting behavior of district and list MPs in the German Bundestag. *Electoral Studies, 29*, 484-496.

Slapin, J. B., & Kirkland, J. H. (2020). The Sound of Rebellion: Voting Dissent and Legislative Speech in the UK House of Commons. *Legislative Studies Quarterly*, 45 (2): 153-176.

Slapin, J. B., & Proksch, S.-O. (2008). A Scaling Model for Estimating Time-Series Party Positions from Texts. *American Journal of Political Science*, 52 (3): 705-722.

Slapin, J. B., Kirkland, J. H., Lazzaro, J. A., Leslie, P. A., & O'Grady, T. (2018). Ideology, Grandstanding, and Strategic Party Disloyalty in the British Parliament. *American Political Science Review*, 112 (1): 15-30.

Smith, M. (2020). *Where does the British public stand on transgender rights in 2020?* London: YouGov.

Smith, M. (2025). *Where does the British public stand on transgender rights in 2024/25?* London: YouGov.

Sobolewska, M. (2013). Party Strategies and the Descriptive Representation of Ethnic Minorities: The 2010 British General Election. *West European Politics, 36*(3), 615-633.

Stephens, G., & Leslie, J. (2011). Parties, organizational capacities and external change: New Zealand's National and Labour parties, candidate selection and the advent of MMP. *Political Science, 63*(2), 205-218.

Stowers, S. (2024). *The Changing Politics of Britain's Ethnic Minority Voters.* London: Political Insight.

Strøm, K. (2003). Parliamentary Democracy and Delegation. In K. Strøm, W. Müller, & T. Bergman, *Delegation and Accountability in Parliamentary Democracies* (pp. 56-105). Oxford: Oxford University Press.

Tate, K. (2002). Legislative Styles and Voting Records. In *Black Faces in the Mirror: African Americans and Their Representatives in the U.S. Congress* (pp. 73-95). Princeton: Princeton University Press.

The Public Whip. (2019, November 6). *Voting Record — Dennis Skinner MP, Bolsover (10544)*. Retrieved June 15, 2020, from https://www.publicwhip.org.uk/mp.php?mpn=Dennis_Skinner

The Public Whip. (2022). *Rebel MPs — All on record*. Retrieved April 11, 2023, from The Public Whip: https://www.publicwhip.org.uk/mps.php?parliament=all&sort=rebellions

The Public Whip. (2024). *Raw Data*. Retrieved October 17, 2024, from https://www.publicwhip.org.uk/project/data.php

Tryl, L., Surmon, T., Kimaram, A., & Burns, C. (2022). *Britons and Gender Identity: Navigating Common Ground and Division.* London: More in Common.

Tsebelis, G. (2002). *Veto Players: How Political Institutions Work.* New York: Russell Sage Foundation.

UK Parliament. (2025, September 9). *Register of All-Party Parliamentary Groups: Registers published in 2025.* Retrieved from UK Parliament: https://www.parliament.uk/mps-lords-and-offices/standards-and-financial-interests/parliamentary-commissioner-for-standards/registers-of-interests/register-of-all-party-party-parliamentary-groups/registers-published-in-2025/

van de Wardt, M., Sobolewska, M., & English, P. (2024). Ethnic minority MPs as reputational shields? How Western European political parties respond to public opinion shifts on immigration policy. *European Political Science Review, 16*(4), 503–520.

van Holsteyn, J. J., den Ridder, J. M., & Koole, R. A. (2017). From May's Law to May's Legacy: On the opinion structure within political parties. *Party Politics, 23*(5), 471-486.

van Oosten, S. (2022). What shapes voter expectations of Muslim politicians' views on homosexuality: stereotyping or projection? *Electoral Studies, 80*.

Ward, L. J. (1998). 'Second-Class MPs'? New Zealand's Adaptation to Mixed-Member Proportional Representation. *Political Science*, 49 (2) 125-152.

Wei, J., Wang, X., Schuurmans, D., Bosma, M., Ichter, B., Xia, F., . . . Zhou, D. (2022). Chain-of-Thought Prompting Elicits Reasoning in Large Language Models. *arXiv.org preprint*.

Williams, B. D. (2012). Institutional Change and Legislative Vote Consensus in New Zealand. *Legislative Studies Quarterly, 37*(4), 559-573.

Wilson, H. (2024). *How Brexit reshaped party politics*. Retrieved April 14, 2025, from https://natcen.ac.uk/how-brexit-reshaped-party-politics

Wu, Y., Iso, H., Pezeshkpour, P., Bhutani, N., & Hruschka, E. (2023). Less is More for Long Document Summary Evaluation by LLMs. *arXiv.org preprint*.

Wuttke, A., Aßenmacher, M., Klamm, C., Lang, M. M., Würschinger, Q., & Kreuter, F. (2025). AI Conversational Interviewing: Transforming Surveys with LLMs as Adaptive Interviewers. *Proceedings of LaTeCH-CLfL*, 179-204.

YouGov. (2019, May 22). *Only 5% of Brits can name a single one of their MEPs*. Retrieved June 2, 2020, from https://yougov.co.uk/topics/politics/articles-reports/2019/05/22/only-5-brits-can-name-single-one-their-meps

Ziems, C., Held, W., Shaikh, O., Chen, J., Zhang, Z., & Yang, D. (2024). Can Large Language Models Transform Computational Social Science? *Computational Linguistics, 50*(1), 237-291.

Zingher, J., & Farrer, B. (2014). The electoral effects of the descriptive representation of ethnic minority groups in Australia and the UK. *Party Politics, 22*(6), 691-704.