

Database of Parliamentary Speeches in Ireland, 1919–2013

Alexander Herzog
School of Computing
Clemson University
Email: aherzog@clemson.edu

Slava J. Mikhaylov
Institute for Analytics and Data Science
Department of Government
University of Essex
Email: s.mikhaylov@essex.ac.uk

Abstract—We present a database of parliamentary debates that contains the complete record of parliamentary speeches from Dáil Éireann, the lower house and principal chamber of the Irish parliament, from 1919 to 2013. In addition, the database contains background information on all TDs (Teachta Dála, members of parliament), such as their party affiliations, constituencies and office positions. The current version of the database includes close to 4.5 million speeches from 1,178 TDs. In addition, we also present three applications in which we make use of the new data.

I. INTRODUCTION

Almost all political decisions and political opinions are, in one way or another, expressed in written or spoken texts. Great leaders in history become famous for their ability to motivate the masses with their speeches; parties publish policy programmes before elections in order to provide information about their policy objectives; parliamentary decisions are discussed and deliberated on the floor in order to exchange opinions; members of the executive in most political systems are legally obliged to provide written or verbal answers to questions from legislators; and citizens express their opinions about political events on internet blogs or in public online chats. Political texts and speeches are everywhere that people express their political opinions and preferences.

It is not until recently that social scientists have discovered the potential of analyzing political texts to test theories of political behavior. One reason is that systematically processing large quantities of textual data to retrieve information is technically challenging. Computational advances in natural language processing have greatly facilitated this task. Adaptation of such techniques in social science – for example, Wordscore [1] or Wordfish [2] – now enable researchers to systematically compare documents with one another and extract relevant information from them. Applied to party manifestos, for which most of these techniques have been developed, these methods can be used to evaluate the similarity or dissimilarity between manifestos, which can then be used to derive estimates about parties' policy preferences and their ideological distance to each other.

One area of research that increasingly makes use of quantitative text methods are studies of legislative behavior and parliaments [3]–[6]. Only a few parliaments in the world use roll-call votes (the recording of each legislator's decision

in a floor vote) that allow for the monitoring of individual members' behavior. In all other cases, contributions to debates are the only outcome that can be observed from individual members. Using such debates for social science research, however, is often limited by data availability. Although most parliaments keep written records of parliamentary debates and often make such records electronically available, they are never published in formats that facilitate social science research. A significant amount of labor is usually required to collect, clean and organize parliamentary records before they can be used for analytical purposes, often requiring technical skills that many social scientists lack.

The purpose of this paper is to present a new database of parliamentary debates to overcome precisely this barrier. Our database contains all debates as well as questions and answers in Dáil Éireann, covering almost a century of political discourse from 1919 to 2013. These debates are organized in a way that allows users to search by date, topics or speaker. More importantly, and lacking in the official records of parliamentary debates, we have identified all speakers and linked their debate contributions to the information on party affiliation and constituencies from the official members database. This enables researchers to retrieve member-specific speeches on particular topics or within a particular timeframe. Furthermore, all data can be retrieved and stored in formats that can be accessed using commonly used statistical software packages.¹

In addition to documenting this database, we also present three applications in which we make use of the new data. In the first study, we analyze budget speeches delivered by all finance ministers from 1922 to 2008 and show how the policy agenda and ministers' policy preferences have changed over time. In the second application we compare contributions that were made on one particular topic: the 2008 budget debate. Here we demonstrate how text analytics can be used to estimate members' policy preferences on a dimension that represents pro- versus anti-government attitudes. Finally, we estimate all contributions from members of the 26th government that formed as a coalition between Fianna Fáil and the Progressive Democrats in 2002. Here we estimate the policy positions of all cabinet ministers on a pro- versus anti-spending dimension

¹The database is made available on the Harvard Dataverse at <http://dx.doi.org/10.7910/DVN/6MZN76>.

and show that positions on this dimension are highly correlated with the actual spending levels of each ministerial department.

II. OVERVIEW OF DATABASE CONTENT

Parliamentary debates in Dáil Éireann are collected by the Oireachtas' Debates Office and published as the Official Record. The Debates Office records and transcribes all debates and then publishes them both in printed as well as in digital form. All debates are then published on Oireachtas' website as single HTML files. At the time of writing, the official debates website contains 549,292 HTML files. We scraped all debate contributions and the names of all speakers, the date, and the topic of each debate. In addition, we collected the information on each speaker in the historical debates.

The final database contains all debates and written answers from the first meeting of the Dáil on 21 January 1919 through to 28 March 2013, covering every Dáil session that has met during this period. In total, the database contains 4,443,713 individual contributions by 1,178 TDs. The data is organized in a way that facilitates analysis for substantive questions of interest to social scientists containing information on the following features: the name and surname of the speaker; the unique ID the speaker; the speaker's party affiliation and constituency; the title of the Dáil debate as recorded in the official records; and the date of the debate.

III. ANALYZING THE CONTENT OF PARLIAMENTARY DEBATES

In the following three sections we demonstrate how the data can be used for social science research using three different applications. First, we analyze the budget speeches of all finance ministers from 1922 to 2008. Budget speeches are delivered by Finance Ministers once a year, with the exception of emergency budgets. Analyzing this data, we show how policy agendas and ministers' fiscal preferences have changed over time. Second, we retrieve all speeches from one particular year and on one particular topic from our database – the 2008 budget debate – to estimate the policy positions of all speakers who contributed to the budget debate and to compare how similar or dissimilar their preferences were. We find that policy positions are clustered into two groups: the government and the opposition; but we also find considerable variation within each group. Finally, we take all contributions made during the term of one government and use the data to estimate the policy positions of all cabinet members on a dimension representing pro- versus anti-spending. We demonstrate the validity of estimated policy positions by comparing them against actual spending levels of each cabinet ministers' department and show that the two measures are almost perfectly correlated with each other.

A. Estimation of Finance Ministers' Policy Positions

Wordfish [2] is a method that combines Item Response Theory with text classification. In this model each document is treated as a separate actor's position and all positions are estimated simultaneously. If a minister maintains a similar

position from one budget speech to the next, this means that words with similar frequencies were used over time. At the same time any movement detected by the model towards a position held by, for example, his predecessor, means that the minister's word choice is now much closer to his predecessor than to his own word usage in the previous budget speech.

Before including documents in the analysis, we have removed all numbers, punctuation marks, and stop words. In addition, we delete words that appear in less than 20% of all speeches [7]. Figure 1 shows the results of estimation, with an overlaid regression line.

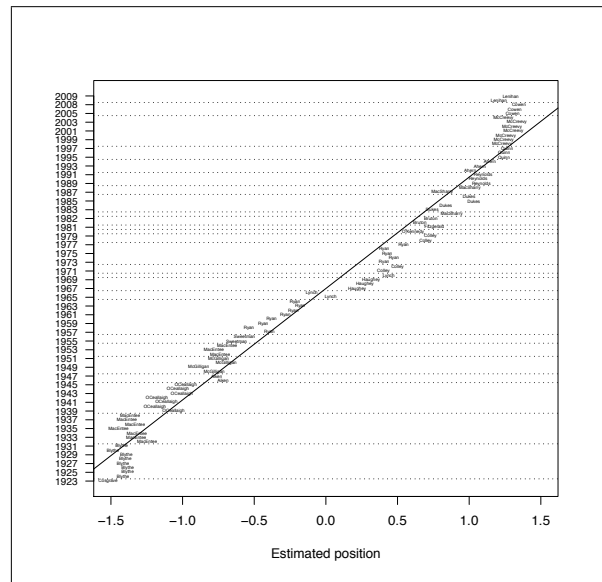


Fig. 1. Finance ministers' policy positions as estimated from all budget speeches (1922–2009) with an overlaid linear regression line.

The results in Figure 1 indicate a concept drift – the gradual change over time of the underlying concept behind the text categorization class [8]. In the political science text scaling literature, this issue is known as agenda shift [7]. In supervised learning models like Wordscore, this problem has typically been dealt with by estimating text models separately for each time period [9], [10], where the definition of the dimensions remains stable through the choice of training documents. However, this approach is not easily transferrable to inductive techniques like Wordfish, where there may be substantively different policy dimensions at different time periods, rendering comparison of positions over time challenging, if not impossible. A clear presence of the concept drift issue in Wordfish estimation should be a cautionary note for using the approach with time series data, even though the original method was specifically designed to deal with time-series data as indicated in the title of the paper [2].

Looking at Figure 1 we can also observe that some ministers have similar preference profiles while others differ significantly. For example, Ahern and Reynolds are very similar in

their profile but differ from a group consisting of Quinn, McCreevy, Cowen, and Lenihan who are very close to each other. There also appears to be a dramatic shift in agenda between the tenures of Lynch and Haughey (and also during Taoiseach Lynch's delivery of the budget speech for the Minister for Finance Charles Haughey in 1970). Overall, it appears that topics covered in budget speeches develop in waves, with clear bands formed by, for example, Lenihan, Cowen, McCreevy and Quin; Ahern and Reynolds; MacSharry, Dukes, Bruton, Fitzgerald, O'Kennedy and Colley; R. Ryan, Colley, Lynch (for Haughey); MacEntee, McGilligan and Aiken; Blythe and MacEntee.

One intuitive interpretation of our Wordfish results is that budget speeches by finance ministers are related to underlying macroeconomic dynamics in the country. We consider the relationship between estimated policy positions of Ministers and three core economic indicators: unemployment, inflation, and per capita GDP growth rates. Figure 2 shows the three economic indicators, inflation (1923–2008), GDP growth (annual %; 1961–2008) and unemployment rate (1956–2008), over time.

Figure 3 show Ministers' estimated positions plotted against the three indicators.

As expected, the results presented here show that the policy positions of some Ministers can be partly explained by the contemporaneous economic situation in the country. However, the fact that some of the Ministers are clear outliers highlights the effect of individual characteristics on policy-making. One of the avenues for research that arises from this exercise is to analyze the determinants of these individual idiosyncrasies, possibly looking at education, class, and previous ministerial career. Such questions can now be easily investigated by researchers using our database.

B. Speakers' Policy Position in the 2008 Budget Debate

In the previous section, we used budget speeches from each year and compared them over time. In this section, we restrict the analysis to a single year but take multiple speeches made on the same topic. More specifically, we estimate the preferences of all speakers who participated in the debate over the 2008 budget. We extract these speeches from the database by selecting all contributions to the topic "Financial Resolution" for the 2008 budget. This leaves us with a total of 22 speakers from all five parties.

To estimate speakers' position we use Wordscore [1] – a version of the Naive Bayes classifier that is deployed for text categorization problems [11]. In a similar application, [1] have already demonstrated that Wordscore can be effectively used to derive estimates of TDs policy positions.

Wordscore uses two documents with well-known positions as reference texts (training set). The positions of all other documents are then estimated by comparing them to these reference documents. The underlying idea is that a document that, in terms of word frequencies, is similar to a reference document was produced by an author with similar preferences. The selection of reference documents furthermore determines

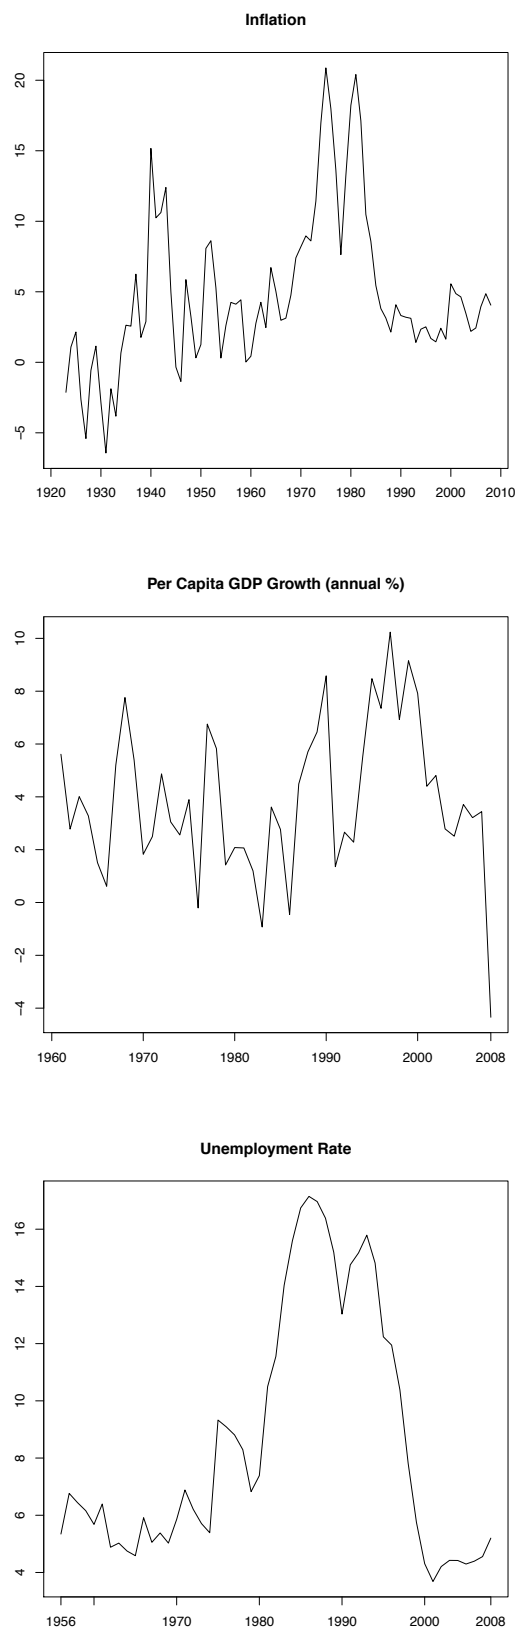


Fig. 2. The Irish economy over time: Inflation (1923–2008), Per Capita GDP growth (annual %; 1961–2008) and unemployment rate (1956–2008).

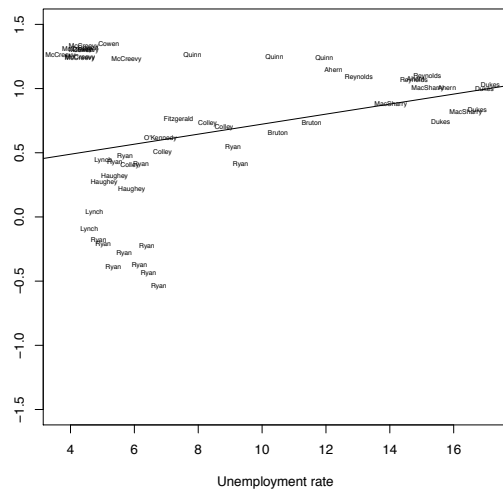
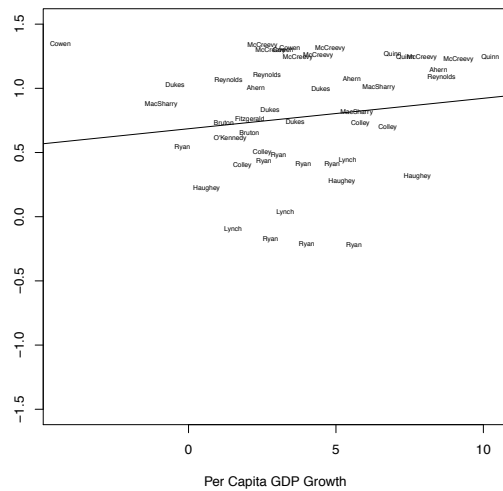
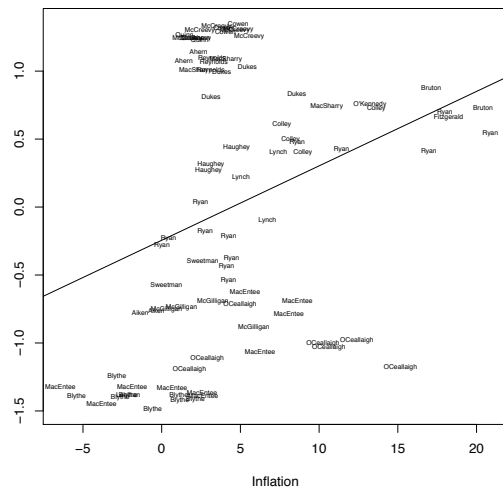


Fig. 3. Estimated finance ministers' positions against inflation (1923–2008), GDP growth (annual %; 1961–2008), and unemployment rate (1956–2008), with an overlaid linear regression line.

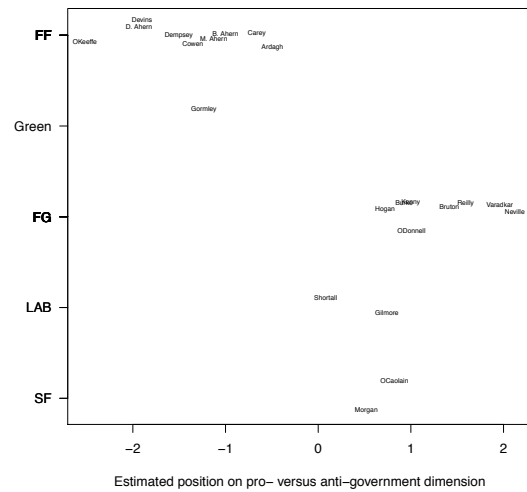


Fig. 4. Estimated positions of all speakers in the 2008 budget debate. Estimated dimension represents pro- versus anti-government positions. Scaling of x-axis is arbitrary. Speeches of Bertie Ahern (FF, Taoiseach) and Enda Kenny (FG party leader) were used as reference texts for being respectively pro- or anti-government.

the (assumed) underlying dimension for which documents' positions are estimated. For example, using two opposing documents on climate change would scale documents on the underlying dimension “climate politics”.

We assume that contributions in budget debates have the underlying dimension of being either *pro* or *contra* the current government. Our interpretation from reading the speeches is that, apart from the budget speech itself, all other speeches largely either attack or defend the incumbent government and to a lesser extent debate the issues of the next budget. We can therefore use contributions during the budget debate as an indicator for how much a speaker is supporting or opposing the current government, here consisting of Fianna Fáil and the Green Party. As our reference texts we therefore chose the speeches of Bertie Ahern (Taoiseach) and Enda Kenny (FG party leader). The former should obviously be strongly supportive of the government while the latter, as party leader of the largest opposition party, should strongly oppose it. Figure 4 shows estimated positions for all speakers grouped by party affiliation.

The estimated positions are clustered into two groups, one representing the government and one the opposition. Within the government cluster, Deputy Batt O’Keeffe (Minister of State at the Department of Environment, Heritage and Local Government) is estimated to be the most supportive speaker for the government, while Deputy Pat Carey (Minister of State at the Department of Community, Rural and Gaeltacht Affairs) and Deputy Sean Ardagh are estimated to be relatively closer to the opposition. Deputy John Gormley, leader of the Green

party and Minister for the Environment, Heritage and Local Government in the FF-Green coalition, is estimated to be in the centre of the government cluster. Among all positions in the opposition cluster, the speech of Róisín Shortall is the closest to the government side, with Neville being the farthest out.

C. Ministers' policy position in the 26th government

The government cabinet in parliamentary democracies is at the core of political decision making, yet it is difficult to model intra-cabinet bargaining as the preferences of most cabinet members are unknown. Cabinet decisions are usually made behind closed doors and the doctrine of joint cabinet responsibility prevents ministers from publicly opposing decisions, even if they disagree with them. Using ministers' speeches and their responses during question times offer a unique opportunity to infer their preferences on policy dimensions of interest. In our final application we estimate policy positions for all cabinet members in the 26th government. The dimension on which positions are estimated represents pro- versus contra-government spending (or spending left-right). We show that estimated positions are highly correlated with departments' actual spending, which means that estimated positions are not only meaningful but can also be used to predict actual policy-making.

The 26th government was formed as a coalition between Fianna Fáil and the Progressive Democrats after the election for the 29th Dáil in 2002. The cabinet was reshuffled on 29 September 2004 and we only include ministers' speeches until that date. To estimate ministers' policy positions, we retrieve the complete record of each minister's contribution in parliament from the first meeting on 6 June 2002 until the date of the reshuffle. On average, each minister made 3,643 contributions with an average number of 587,077 words.

We again use Wordscore to estimate positions as it allows us to define the underlying policy dimension by choosing appropriate reference texts. We estimate positions on a social-economic left-right dimension that reflects pro- versus contra-government spending. We therefore use contributions by Mary Coughlan (Minister for Social and Family Affairs) and Charlie McCreevy (Minister for Finance) as reference texts, assuming that the former is more in favor of spending than the latter. Figure 5 shows the results of estimation grouped by the two parties.

As expected, we find that the two PD members, Mary Harney and Michael McDowell, are at the right side of the dimension. We estimate the most left-wing members to be Éamon Ó Cuív (Minister for Community, Rural and Gaeltacht Affairs), Noel Dempsey (Minister for Education and Science), and Micheál Martin (Minister for Health and Children). The most right-wing members are John O'Donoghue (Minister for Arts, Sport and Tourism), Charlie McCreevy (whose contributions we used as right-wing reference text), and Michael Smith (Minister for Defense).

How valid are these estimated positions? In order to have substantive meaning, our estimates should be able to predict political decisions on the same policy dimension. We therefore

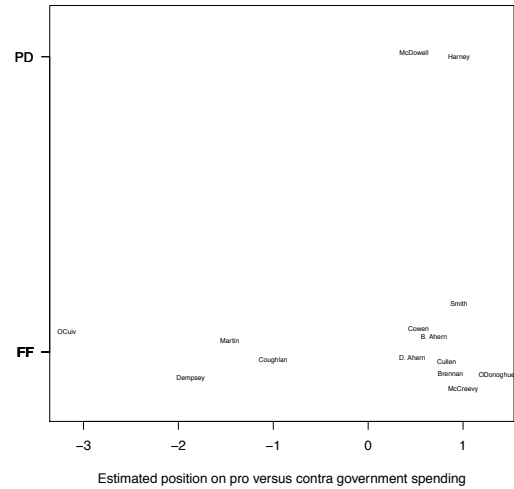


Fig. 5. Estimated positions for all cabinet members in the 26th government (29th Dáil) using Wordscore. Positions are jittered along the y-axis. Estimation is based on each minister's contribution in Dáil Éireann before the cabinet reshuffle on 29 September 2004. Speeches by Mary Coughlan (Minister for Social and Family Affairs) and Charlie McCreevy (Minister for Finance) are used as left and right reference texts, respectively

use ministers' estimated positions to predict their departmental spending level [3]. Our outcome variable is each department's spending as share of the total budget in 2004 modeled as a function of estimated policy positions. We conjecture that more left-wing ministers should have higher spending levels than right-wing ministers, which we test by estimating spending as a function of policy position using a linear regression.

Figure 6 shows the two variables plotted against each other together with the estimated regression line from equation ???. In one analysis shown we include all cabinet members. In the other, we exclude non-spending departments with small budgets, such as the office of the Taoiseach or the Department of Foreign Affairs.²

Figure 6 reveals that there is a negative, albeit weak, relationship between estimated positions and spending, with more left-wing cabinet members having higher spending levels than right-wing members. The correlation between the two variables is -0.53 ($p = 0.0523$) which is not significant at the 0.05 level. However, if we only take members from high-spending departments into account (second pane in Figure 6) we find a significant linear relationship between the two variables with a correlation coefficient of -0.95 ($p = 0.0002$).

²The eight high-spending departments we include are, in decreasing order of budget share, the Department of Health and Children, Department of Education and Science, Department of Social and Family Affairs, Department of the Environment and Local Government, Department of Transport, Department of Enterprise, Trade and Employment, Department of Defence, and the Department of Arts, Sport and Tourism. These eight departments together account for more than 95 per cent of the total budget in 2004.

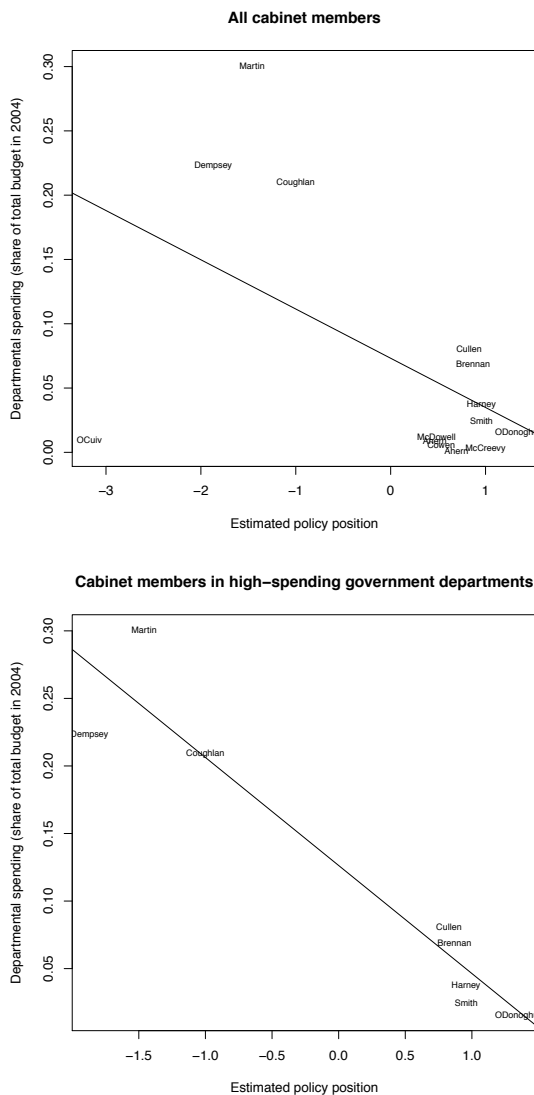


Fig. 6. Cabinet ministers' policy position plotted departmental spending as share of total government budget in 2004. For the analysis of high-spending departments we remove the Office of the Taoiseach or the Department of Foreign Affairs, with the remaining eight departments accounting for more than 95 per cent of the total budget in 2004

This result provides some level of validation for our data and analysis.

These results also open up an intriguing question about the endogeneity of observable policy preferences of ministers. Do higher spending portfolios receive more pro-spending ministers or do ministers adapt their policy preferences after appointment and literally grow into the job? This and related questions are outside the scope of this paper and can be pursued by researchers with the help of our database of parliamentary speeches.

IV. CONCLUSION

Policy preferences of individual politicians (ministers or TDs in general), are inherently unobservable. However, we have abundant data on speeches made by political actors. Application of natural language processing allows us to estimate the policy positions of individual actors from these speeches.

In this paper we present a new database of speeches that was created with the purpose of allowing the estimation of policy preferences of individual politicians. For that reason we created a relational database where speeches are related to the members database and structured in terms of dates, topics of debates, and names of speakers, their constituency and party affiliation. This gives the necessary flexibility to use available text scaling methods in order to estimate the policy positions of actors.

We also present several examples for which this data can be used. We show how to estimate the policy positions of all Irish Ministers for Finance, and highlight how this can lead to interesting research questions in estimating the determinants of their positions. We show that for some ministers the position can be explained by the country's economic performance, while the preferences of other ministers seem to be idiosyncratic. In another example we estimate positions of individual TDs in a budget debate, followed by the estimation of policy positions of cabinet members of the 26th Government.

With the introduction of our database, we aim to make text analysis an easy and accessible tool for social scientists engaged in empirical research on policy-making that requires estimation of policy preferences of political actors.

REFERENCES

- [1] M. Laver, K. Benoit, and J. Garry, "Extracting policy positions from political texts using words as data," *American Political Science Review*, vol. 97, no. 2, pp. 311–331, May 2003.
- [2] J. B. Slapin and S.-O. Proksch, "A scaling model for estimating time-series party positions from texts," *American Journal of Political Science*, vol. 52, no. 3, pp. 705–722, 2008.
- [3] D. Giannetti and M. Laver, "Policy positions and jobs in the government," *European Journal of Political Research*, vol. 44, no. 1, pp. 91–120, 2005.
- [4] S.-O. Proksch and J. B. Slapin, "Position taking in European Parliament speeches," *British Journal of Political Science*, vol. forthcoming, 2009.
- [5] D. Hopkins and G. King, "A method of automated nonparametric content analysis for social science," *American Journal of Political Science*, vol. 54, no. 1, pp. 229–247, 2010.
- [6] J. Grimmer, "A bayesian hierarchical topic model for political texts: Measuring expressed agendas in senate press releases," *Political Analysis*, vol. 18, no. 1, pp. 1–35, 2010.
- [7] S.-O. Proksch and J. B. Slapin, "How to avoid pitfalls in statistical analysis of political texts: The case of Germany," *German Politics*, vol. 18, no. 3, pp. 323–344, 2009.
- [8] C. D. Manning, P. Raghavan, and H. Schütze, *Introduction to Information Retrieval*. Cambridge University Press, 2008.
- [9] A. Baturo and S. Mikhaylov, "Life of brian revisited: Assessing informational and non-informational leadership tools," *Political Science Research and Methods*, vol. 1, no. 01, pp. 139–157, 2013.
- [10] A. Herzog and K. Benoit, "The most unkindest cuts: Speaker selection and expressed government dissent during economic crisis," *Journal of Politics*, vol. 77, no. 4, pp. 1157–1175, 2015.
- [11] K. Benoit and P. Nulty, "Classification methods for scaling latent political traits," 2013, paper prepared for presentation at the Annual Meeting of the Midwest Political Science Association, April 11–14, 2013, Chicago.