Party Power and Legislative Decision-Making in Congress

by

Andrew Ojala Ballard

Department of Political Science
Duke University

Date:_______________________

Approved:

____

John Aldrich, Co-supervisor

____

Michael Munger, Co-supervisor

____

Scott de Marchi

____

David Rohde

Dissertation submitted in partial fulfillment of
the requirements for the degree of Doctor
of Philosophy in the Department of
Political Science in the Graduate School
of Duke University

2018

ABSTRACT

Party Power and Legislative Decision-Making in Congress

by

Andrew Ojala Ballard

Department of Political Science
Duke University

Date:_______________________
Approved:

____
John Aldrich, Co-supervisor

____
Michael Munger, Co-supervisor

____
Scott de Marchi

____
David Rohde

An abstract of a dissertation submitted in partial
fulfillment of the requirements for the degree
of Doctor of Philosophy in the Department of
Political Science in the Graduate School of
Duke University

2018

Abstract

Questions of party power and legislative outcomes are central to our understanding of Congress. And yet, our knowledge of these concepts has many gaps. We know little about exactly how parties negotiate legislative deals with their members—and hold them to their deals—or about how parties exert agenda control before bills reach the floor. We also have limited knowledge of how predictable the outcomes of bills are, particularly from their content. These questions remain unanswered largely because political scientists have traditionally not had data or methods at their disposal to answer them. In my dissertation, I provide some answers to these questions by focusing on modern text analysis techniques.

In chapter one, I examine how parties make deals with recalcitrant members on landmark legislation, and more importantly, how members are held to their deals. Particularly, I argue that public statements are a tool wielded by the party, and when members are convinced to vote for a bill they are provided incentive to make a statement in support of the bill to lock them into their vote. Using a novel data set of all public statements made by members of Congress, and two pieces of landmark healthcare legislation (the Affordable Care Act in 2009-2010 and the American Health Care Act in 2017), I show that members of Congress do make public statements after they make deals to vote with the party, and that these statements are likely for the purpose of precommitting to vote with the party.

In chapter two, I seek to quantify the level of agenda control exerted by the majority party on bills that never reach a final passage vote. To do this, I present the first systematic estimates of how members would have voted on bills based on a characterization of bill content from bills' text, would they have come to a final passage vote. I find that the majority party not only exercises negative agenda control, but also considerable positive agenda control. I also find that the minority party in the House is systematically and consistently shut out of the agenda process, but not the policymaking process.

In chapter three, I investigate how predictable the outcome of bills is and whether bill content has a part to play in the predictability of bill outcomes. I am able to predict where bills end up in the US House of Representatives with high accuracy, and that knowledge about the content of a bill has a sizable effect on how well we can predict bill outcomes.

Contents

List of Tables

List of Figures

Acknowledgements

No dissertation is completed without help. I have been fortunate enough to find myself among some of the brightest and best people while working on mine. My successes are also theirs, for I would not be here today, and this dissertation would not exist, without them. I cannot possibly thank them all in this space, but I would like to take this time to thank a few of these wonderful individuals.

First, I would like to thank my co-chairs John Aldrich and Mike Munger, and my committee members Dave Rohde and Scott de Marchi. In addition to dealing with my various shortcomings over six years and helping turn them into a tenure track job, each of you is an excellent professional role model. You brought me into projects early, which has been invaluable to my professional life. I've learned a great deal about mentoring graduate students, particularly, that I will pass on to my own students.

But no role models loom larger in my life than my parents, Charles Ballard and Carolyn Ojala. You taught me how to value hard work, to search for interesting things in the world, and to be a good person. There are other lessons, to be sure, like putting the cap on the toothpaste and not biting my nails (still working on that one), but those on how to improve myself and the world are what stick with me most. I love you both so much. Thank you. As is the nature with relationships between parents and children, I won't ever be able to adequately thank you for all the things you've done for me, but I can sure try.

My friends have also made graduate school worthwhile. Thanks to Kat, for late night chats over movies and snacks, often when we need it the most. To Spencer, for sharing our little kingdom in Gross, and for (Too many? Not enough?) breaks for laughter and gum. To Clare, for helping me through the toughest times in graduate school, even when you didn't have any reason to. To Aspen, for being the best person I know, and making me better just by being around. And to Brennan, Chris, and Devin, for agreeing to live with me, and for crossing over from being just roommates to friends.

I also want to thank the staff of the Duke Political Science Department. Without your help, graduate school would have been hell. Thanks to Kathy Ivanov and Fonda Anthony for their help with forms, payments, and really, with everything. Thanks to Susan Emery for being the person you go to if you need to get things done, and the person on whose bad side you least want to be. Thanks to Shaun King, Vilay Nidiffer, and Jim Wallace for excellent IT and media support. Thanks to Suzanne Pierce for helping me set up my first course, an experience I'll never forget. Thanks to Rita McCoy and Doris Cross for their support with reimbursements. And thanks to Steffani Shouse and Amanda Shipp for keeping the boat upright.

Finally, I want to thank the folks who sent me on this path. When I finished college in the height of the recession, a BA in Psychology did not open many doors. I decided to go to graduate school, but did not know what I wanted to study. So thank you to Jorge Delva and Liz Gerber for teaching me how to do research and for pointing me down this path. None of this would have happened without you.

Introduction

The current political era is one of heightened partisan competition. We may actually be living in the most contentious period of party politics in American history (Lee 2009, 2016; Jochim and Jones 2012; Grossmann and Hopkins 2016). Although the link between polarizing competition in Congress and a polarized mass public is less certain (see Abramowitz and Saunders 2008; Abramowitz 2010; and Fiorina and Abrams 2008 for opposing views), most scholars agree that today's parties are competing as fiercely as ever in Congress.

And yet, some patterns in recent research do not fit with our understanding of how legislative outcomes should manifest when parties are polarized. For instance, the majority party is relatively more powerful in the House than the Senate, since most decisions can be made in the House with a simple majority and this is not true in the Senate. If parties are polarized, we would expect the minority party to be shut out of the legislative process, but the vast majority of bills that pass in the House do not do so over the objections of the minority party (Curry and Lee 2018).

What, then, are some of the implications of hyperpartisanship in Congress for party power and legislative outcomes? What can parties actually do? What do they control? How does this lead to legislative outcomes? In my dissertation, I investigate these questions with a novel approach and new sources of data.

The study of Congress has long been dominated by the study of legislative voting behavior. Votes are important because they determine policy outcomes, but researchers have also focused on studying voting records because science tends to gravitate toward what is easy to measure or what has already been measured and for which we have data. Voting records are easily observable and meticulously recorded.[1]

However, members of Congress spend most of their time with behaviors other than voting, such as various public statements signaling their intentions, crafting bill content, raising donations, etc. Many have argued that not only is most of a member's time spent doing things other than voting, but that these activities are more influential to policy and members' career prospects (*e.g.*, Woon 2009; Hall 1996; Bauer, Sola Pool, and Dexter 1972).

The output of many non-voting behaviors—such as the text of bills and public statements like floor speeches, committee hearings, and press releases—have been recorded by various sources for decades, but they have only recently been used in quantitative studies. Past scholars were unable to analyze such large sources of text data for practical reasons, but modern text analysis techniques break down these methodological roadblocks. Scholars have recently begun to analyze congressional texts for purposes such as categorizing press releases (Grimmer 2009; Grimmer et al. 2012)

---

[1] This is not to say that scholars have until recently only studied votes in Congress. There are a number of other less easily recorded, but still observable, behaviors that scholars have focused on, such as time spent in committee (*e.g.*, Hall and Wayman 1990), time spent meeting with lobbyists (McCubbins and Schwartz 1984; Langbein 1986), and qualitative work on "Home Styles" and presentation of self (Fenno 1978; Gulati 2004).

and floor speeches (Quinn et al. 2010) by topical content, measuring (Sim et al. 2013) and predicting (Diermeier et al. 2012; Yu et al. 2008) the ideological positions of politicians in the US, and measuring whether members of Congress support or oppose proposed legislation (Thomas et al. 2006). Building on the work of these scholars, I examine questions of party power and legislative outcomes using text analysis of bill text and public statements.

In chapter one, I examine how parties make deals with recalcitrant members on landmark legislation in a highly polarized environment when the majority has no support from the minority party for its legislation, and cannot allow too many of its own members to vote against a bill. Importantly, I examine how members are held to their deals; I argue that public statements are a tool wielded by the party, and when members are convinced to vote for a bill they are provided incentive to make a statement in support of the bill to lock them into their vote. Using a novel data set of all public statements made by members of Congress, and two pieces of landmark healthcare legislation (the Affordable Care Act in 2009-2010 and the American Health Care Act in 2017) I show that members of Congress do make public statements after they make deals to vote with the party, and that these statements are likely made for the purpose of precommitting to vote with the party.

In chapter two, I seek to quantify the level of agenda control exerted by the majority party on bills that never reach a final passage vote between 1995 and 2011. To do this, I first develop a characterization of each bill's content via a shallow neural network method called doc2vec (Mikolov et al. 2013a, 2013b). Extending a recent model to predict votes in Congress (de Marchi, Dorsey, and Ensley 2017), I present the first systematic estimates of how members would have voted on bills based on a characterization of bill content from bills' text, would they have come to a final passage vote in the House. I find that the majority party not only exercises negative agenda control, but also considerable positive agenda control. I also find that the minority party in the House is systematically and consistently shut out of the agenda process, and that the degree of majority party control of the legislative process I observe did not change substantially throughout the study period. But at the same time, the proportion of bills that received a final passage vote (and those that passed) which were introduced by the minority party increased nearly four-fold, and about three-quarters of bills pass with Yea votes from a majority of the minority party. While the minority party is shut out of the agenda process, they are not shut out of the policymaking process. The minority party may even be more important for policymaking than they were 20 years ago.

In chapter three, I investigate how predictable bill outcomes are, and whether bill content has a part to play in the predictability of bill outcomes. Using the same characterization of bill content and sample of bills as in chapter two, I find that I am able to predict where bills end up in the US House of Representatives with high accuracy, and that knowledge about the content of a bill has a sizable effect on how well I am able to predict bill outcomes. Predictions are better for determining whether a bill will make it out of committee, versus what happens to a bill after it leaves committee. Again, like in chapter two, the results in chapter three do not change over time. During a period where the degree of partisan polarization is thought to have increased, the stability of my results leads us to rethink the implications of party polarization for legislative processes and outcomes in Congress.

1. Public Statements in Congress as Precommitment Devices

"A political party is a team of men [and women] seeking to control the governing apparatus by gaining office in a duly constituted election." Anthony Downs (1957, 25)

"Nothing about this effort was secret — it was public from the very first meeting that happened at the governor's mansion in January. It was a broadly supported delegation effort from the beginning. And it was never a condition of my support for the bill. There should be some concerns about specific arrangements that were made, or for specific promises of support. This was not one of them. And the record will show that."
- Sen. Mary Landrieu (D, LA)[2]

1.1 Introduction

An overarching theme of American politics is that parties are teams, with collective goals of influencing policy (Downs, 1957), but which also advance members' individual electoral and legislative goals (*e.g.*, Aldrich 2011; Rohde 1991; Cox and McCubbins 2005, 2007; Fenno 1978; Mayhew 1974; Grynaviski 2010; Sniderman and Stiglitz 2012; Snyder and Ting 2002). Party government also provides an organizing infrastructure for legislatures (Cox and McCubbins 2005, 2007; Weingast and Marshall, 1988), while affording opportunities for creating brand recognition (Aldrich 2011; Grynaviski 2010; Carty 2004; Snyder and Ting 2002) that helps individual legislators win elections and pass laws. In exchange, parties demand loyalty from members, protecting the party's reputation and advancing its policy agenda.

---

[2] Senator Landrieu explained why federal projects worth $300 million benefiting her state of Louisiana were included in the Affordable Care Act bill she ultimately voted for, after initially expressing doubts. http://thehill.com/homenews/senate/79823-sen-landrieu-hits-back-over-louisiana-purchase

Consider the inherent collective action problem. On one hand, members of Congress need the party's brand to be effective, and they need access to parties' other electoral and legislative resources. But given the heterogeneity in district composition and member goals, MCs also want to be able to advance their own policy agenda or the desires of their constituents. When the interests of parties and their member diverge, parties must develop and use a variety of "carrot-and-stick" measures to induce recalcitrant members to cooperate in supporting partisan positions on many issues.

The tension between group and individual goals is typical of a broad range of collective action problems in political science and economics, such as the "franchising" problem faced by restaurants, car dealerships, and a variety of other retail enterprises. Individual business operators need the franchise's brand name recognition and the economies of scale in national advertising benefits, but can profit more by shirking on the conditions the franchise agreement imposes. There has been a fruitful application of this logic to political parties, showing that associating with the party's brand is advantageous to members, but that members also want to act on their own (*e.g.*, Carty 2004; Grynaviski 2010). Because parties often need support for bills from members who would rather not vote with the party, a deeper understanding of how and when parties are able to control their members will illuminate the process determining whether legislation is passed or fails to pass.

I present here an account of policymaking that explains how parties attempt to extract “Yea” votes necessary to pass bills from members who want to vote “Nay” and how they attempt to ensure those members follow through and vote with the party. Public statements—in committee hearings, on the floor of the House or Senate, and in press releases—act as a signal of the member’s intent to vote “Yea” that allow the party to solve the problem of holding members to their intended vote and at the same time enticing other members to vote with the party.

Using a novel dataset of text documents produced by members of Congress, I detail this process with two examples from recent health care debates: the Affordable Care Act in 2009-10 and America’s Health Care Act in 2017. When party leaders “turned” reluctant members to vote with the party on a bill, the party pushed those members to make a public statement in favor of the bill. Further, I show that these statements were indeed precommitment devices, separating this purpose from credit claiming (*e.g.*, Mayhew 1974; Kingdon 1981).

1.2 Precommitment and Passing Legislation

The relationship between parties and their members has often been likened to corporate franchises or financial sureties (*e.g.*, Snyder and Ting 2002; Grynaviski 2010), whereby much of the value that a party provides to its members is related to how the party’s brand—or its set of policies, credibility in name recognition, and overall quality—is perceived by the public. Parties invest in consistent public stances on issues, therefore an affiliation signals how a candidate might act in office (Aldrich 2011). Parties also communicate to the public that a party member is a quality candidate:

> It is reasonable for people to interpret pledges by franchises, political parties, and the like as a credible signal of candidate or product quality if they believe that the parent organization suffers from long-term consequences if it makes a sufficient number of bad pledges. Thus, the reputation of McDonald's suffers if too many of its local stores offer cold French fries and poor customer service. (Grynaviski 2010, 51)

Members of Congress will benefit if their party membership sends a strong, positive signal to the public. In order to produce a consistent brand name connoting such a signal, parties need to be able to control member behavior. The party can induce loyal behavior from recalcitrant members—for example by granting prestigious committee assignments (Rohde and Shepsle 1973; Ballard et al. 2018) or supporting a member's legislative agenda (Cox and McCubbins 2007). But while parties need members to protect the brand with their behavior, at times the best decision for an individual member may be to deviate from the party. The party's promises mean little to its members if they are voted out of office, which becomes more likely if members ignore the preferences of their constituents (Canes-Wrone 2015; Canes-Wrone, Brady, and Cogan 2002; Nyhan et al. 2012). This potential divergence between party and member interests often creates tension in Congress.

The tension between member and party goals is particularly important while passing legislation. It is often the case that some party members who are necessary to the party's winning coalition have incentive to vote against a bill the party favors. To enact their policy goals and protect the party brand, parties must convince these members to vote with the party through a combination of threats and inducements. Then, equally importantly, the party must ensure the votes of its recalcitrant members, for instance by compelling holdouts to make a statement in favor of the bill at hand.

To illustrate this process, consider the following example. Suppose that the majority party in a legislature wants to pass a bill, needing a simple majority to do so. At some point before the bill is brought to the floor for a vote, the majority party surveys its members, finding that enough members would prefer to vote "Nay," either for electoral reasons (their constituents largely oppose the measure) or personal reasons (the bill violates their own beliefs), that the party is short of a majority.[3] If the bill is to pass, the party must take action.

Just as United Airlines can ask for "volunteers" to give up their seats, the party can attempt to solve the collective action problem by offering inducements or threats to some combination of the holdout members. While the recalcitrant members would prefer to vote against the bill, they have a stake in the success of the party as they want to benefit from its brand. Thus, each would be more willing to vote "Yea" if they could be assured that their vote would contribute to final passage.

---

[3] The leadership is assumed to know, as members are assumed privately to reveal, their sincere preferences and the preferences of their constituents.

When a single member needs to be convinced to vote with the party, there is no collective action problem. However, parties must often convince multiple members, each of whom would be more likely to vote "Yea" if the bill would pass, but none of whom can pass the bill on their own. There are two goals for the party in such situations. When convincing multiple members, the party needs to ensure that other members know a deal was made and that it will be upheld. The party also wants other members to know that a member has been "turned" so that each remaining holdout will be more amenable to voting with the party, as the bill is one step closer to passing. Members making public statements in favor of the bill are a possible solution to the collective action problem that accomplishes both of the party's goals.[4] Further, making a public statement can be beneficial for both parties and members. For parties, such statements are a credible commitment to future behavior that helps solve their collective action problem.[5] If a member makes a public statement in favor of a bill, it is more difficult politically—but not impossible—to then vote against that bill. Each additional "Yea" also makes the bill closer to passing, therefore easing the path for convincing other recalcitrants. For members, these statements can serve as both a precommitment device

[4] Public statements are not the only way for parties to make sure that members follow through on their deals and vote with the party. For instance, parties may stipulate that they will only "pay" members if the bill passes. Rather, I argue that there is strategic value of political speech to policymaking; public statements can be (and often are) used by parties as precommitment devices.
[5] Studying precommitments is common in the economics (Kreps and Scheinkman 1983; Frydman et al. 2000), psychology (Ariely and Wertenbroch 2002; Kivetz and Simonson 2002), law (Levmore 1996), and political theory literatures (Elster 2000; Holmes 1988), but less so in empirical political science.

and as a way for the member to claim credit for their legislative accomplishments to their constituents.

1.3 A Tale of Two Health Care Debates

I will illustrate the important of statements with two important health care bills from the last decade as case studies. In 2009 and 2010, Democrats were able to whip enough votes to pass the Patient Protection and Affordable Care Act (ACA) by a razor-thin margin in each chamber. Republicans came close to repealing the ACA in the spring and summer of 2017 by narrowly failing to pass the American Health Care Act of 2017 (AHCA) through both chambers. Both the Democrats' success in 2009-10, and the Republicans' success and failure in 2017, can be explained partly in terms of how successful parties were at convincing their members to vote with the party and then pushing them to make public statements in favor of a bill as precommitment devices. Democrats in 2009-10 were able to provide enough incentive for their members—and to hold them to their deals—to pass a bill, but Republicans in 2017 were not.

1.3.1 Data and Method

In my analysis of the context and purpose of statements made by members of Congress, I use three sources of text data: congressional floor speeches, committee hearings, and press releases. I refer to these collectively as 'public statements.' The data come from LexisNexis, and each document has tags for the issue subjects to which it pertains.[6] I use only those documents relevant to the ACA and AHCA, according both to subject tags provided by LexisNexis and the date of each document.[7]

The dependent variable in the following analysis is a binary indicator for whether members made a statement in support of their party's bill, between the introduction of the bill and the last vote on the bill in each chamber. This includes four groups of members: (1) Senate Democrats who made a statement in support of the ACA between October 8 (bill introduced) and December 24, 2009 (bill passed), (2) House Democrats who made a statement in support of the ACA between December 24, 2009 (bill received from Senate) and March 21, 2010 (bill passed), (3) House Republicans who made a statement in support of the AHCA between March 8 (bill introduced) to May 4, 2017 (bill passed), and (4) Senate Republicans who made a statement in support of the AHCA between May 4 (bill received from House) to July 28, 2017 (final failed vote).[8] Members are considered to have made a statement in favor of a bill if they made any statement—on the floor, in committee, or in a press release—that could signal an intent to vote "Yea" on the bill. The total number of each type of public statement, by party/year and chamber, can be found in Table 1.

---

[6] LexisNexis uses a proprietary machine learning algorithm that tags documents with subjects, along with a percentage indicating the predicted probability that a document deals with a particular issue. Documents are only tagged with an issue subject if the predicted probability is at least 50%.
[7] Documents relating to the ACA and AHCA have tags "Health Care Reform", "Obama Health Care Reform", and/or "Trump Health Care Reform." There are no document tags for specific bills, such as an "Affordable Care Act" tag.
[8] The author coded these by hand.

Table 1: Frequency of public statements of each type on health care reform bills, made by members of the majority party.

	Democrats (2009-10)		Republicans (2017)	
	House	Senate	House	Senate
Press Releases	83	35	72	40
Committee Hearings	20	15	3	2
Floor Speeches	132	78	141	67
Total	235	128	216	109

I also use data on which members were convinced to vote with the party, which members received policy concessions, public opinion on each bill, and additional demographic and institutional variables: number of terms served in Congress, member race and gender, and dummy variables for whether the member is in the party leadership or sits on a prestige committee.[9] Information on which members were convinced to vote with the party was hand-coded from news accounts (Democrats in 2009-10) and whip counts (Republicans in 2017). It was necessary to use news accounts for Democrats in 2009-10 because detailed whip counts were not available for those votes. These members were only coded as having made a statement in favor of the bill if the statement was made *after* they changed their intended vote.[10] Public opinion estimates come from

[9] Prestige Committees in the House are Appropriations, Budget, Rules, and Ways and Means. Prestige Committees in the Senate are Appropriations, Commerce, and Foreign Relations. Leadership positions in the House are the Speaker of the House, Majority Floor Leader, and Majority Whip. Leadership positions in the Senate are the Majority Floor Leader, Majority Whip, and President Pro Tempore.

[10] For Democrats in 2009-10, there were few holdouts in the House or Senate; all deals went through the party leadership and were widely publicized. The same is not true of Republicans in 2017. Because it is not immediately apparent how to code which Republicans received policy concessions on the bills, I run models with different coding schemes. In this paper, I present models where members of the Freedom Caucus are coded as having received policy concessions on the AHCA, as well as those members who proposed

Warshaw and Broockman (2017), which are state-level multilevel regression and poststratification (MRP) estimates for the proportion of citizens who support, oppose, or are not sure whether they support/oppose the AHCA.[11] For each bill, I assess whether members who were convinced to vote with the party were more likely to make a statement in favor of the party's bill and whether this was affected by having gained policy concessions and/or the opinions of their constituents.

1.3.2 The Affordable Care Act (2009-2010)

On October 8, 2009, the ACA was introduced in the Senate. Because of Republican vows to filibuster and vote against any Democrat-led health care reform effort, Democrats needed the approval of all 58 Senate Democrats, plus both Independents (who caucused with Democrats). Most Senate Democrats needed little convincing to vote "Yea", but seven members held out: Sens. Evan Bayh (D-IN), Kent Conrad (D-ND), Mary Landrieu (D-LA), Carl Levin (D-MI), Joe Lieberman (I-CT), Ben Nelson (D-NE), and Bernie Sanders (I-VT). The Democratic leadership would have to

amendments to various bills that made it to a vote (*e.g.*, Rep. MacArthur and Sen. Ted Cruz). In Appendix A, I also use membership in the Tuesday Group as the measure of who received policy concessions. The substantive results remain the same.

[11] I use the difference between estimated support and opposition (support minus oppose) for each state. I do estimate a model where the MRP estimate is merely the estimated support for each state, but with no changes in the results. These results may be found in Appendix A. With MRP, issue-specific estimates are only generally stable at the state level—rather than the district level—over such short time periods, due to small sample sizes.

convince each to vote with the party and coordinate among them to hold each to their agreement.

Each member eventually agreed to vote “Yea,” and all gained substantial policy concessions to do so. For example, Landrieu was given $300 million in Medicaid funding for Louisiana. Lieberman had the public option removed from the bill. Nelson received language allowing states to decide whether to cover abortion in their insurance exchanges and an increase in the Medicare reimbursement for Nebraska. Between agreeing to vote “Yea” and the final vote, each member made a statement in support of the bill. The ACA passed the Senate, 60-39.

As in the Senate, House Democrats found no bipartisan support for the ACA. Of their 253 seats in the House, Democrats had to find 216 “Yea” votes; the party could allow no more than 37 defectors.[12] Working to pass the bill in the House met little opposition. However, a group of pro-life Democrats led by Bart Stupak (D-MI) opposed the bill because it did not explicitly outlaw the use of federal funds for abortions.[13] In response, President Obama signed an executive order reaffirming that federal funding could not be used for abortions, and Stupak’s group agreed to vote “Yea.” The group then issued a joint press release, and the ACA passed the House, 219-212.

---

[12] At the time, only 431 of the 435 seats in the House were filled, so 216 votes were required to pass the ACA rather than the usual 218.

[13] These included Reps. Jim Cooper (D-TN), Kathy Dahlkemper (D-PA), Joe Donnelly (D-IN), Steve Driehaus (D-OH), Marcy Kaptur (D-OH), Paul Kanjorski (D-PA), Alan Mollohan (D-WV), and Nick Rahall (D-WV).

The evidence here supports that the statements made by Democrats after they agreed to vote “Yea” on the ACA were precommitment devices. Each Democrat who made a deal with the party leadership then made a statement in support of the bill, compared with only 75% of other Democrats in the Senate and 29% of other Democrats in the House (37% of all other members). In each chamber, this difference in proportions is significant at $\alpha = .01$. Table 2 shows the proportion of Democrats who made a statement in favor of the ACA—by chamber and whether the member was convinced to vote with the party.

Further statistical modeling is unnecessary for the case of Democrats and the ACA. Democrats always made a statement in support of the ACA after being convinced to vote with the party, regardless of how their constituents felt about the ACA or whether members received policy concessions. At the same time, each Democrat who was convinced to vote with the party also received policy concessions, so separating the effect of these variables is not possible. There is also some danger of selecting on the dependent variable when analyzing public statements made during the Democrats’ push to pass the ACA in 2009 and 2010. It may be that some Democrats made a deal with the party leadership that was not made public. A public statement in support of the bill may be less likely to follow such deals, but there is no way to know with the publicly available data included here.

Table 2: Majority party caucus statements in favor of health care reform. Asterisks indicate a significant difference in means between convinced members and all others at α = .01.

2009-10 (Democrats)						
	Senate[14]		House		Combined	
	Convinced	All Others	Convinced	All Others	Convinced	All Others
Made Statement	7	40	9	70	16	110
Total	7	53	9	244	16	297
Percentage	100.0%*	75.4%	100%*	28.6%	100.0%*	37.0%
2017 (Republicans)						
	Senate		House		Combined	
	Switched	All Others	Switched	All Others	Switched	All Others
Made Statement	10	19	23	50	33	69
Total	10	42	34	206	44	248
Percentage	100.0%*	45.2%	67.6%*	24.2%	75.0%*	27.8%

1.3.3 American Health Care Act (2017)

Attempts to pass the AHCA by Republicans in 2017 provide a better test than the case of Democrats with the ACA for whether parties push members to make public statements in favor of a bill as precommitment devices. Here, my measure of whether Republicans were convinced to vote with the party is not whether they made a public

[14] Includes Bernie Sanders (I-VT) and Joe Lieberman (I-CT).

deal with the party. Instead, I have data for two (intended) votes in the House[15] and three in the Senate. I code members as having been convinced to vote with the party if they switched their vote from "Nay" to "Yay" between any of the votes within their chamber. All vote changes are recorded, meaning there is no danger of selecting on the dependent variable. Further, to my knowledge, at least some Republicans who switched their vote did not receive policy concessions on the AHCA.

Republicans held 238 seats in the House and needed 216 votes to pass the AHCA. Without support from Democrats, there could be no more than 22 Republican "Nay" votes. After being pulled from the floor in March, The bill was revived in April by Rep. Tom MacArthur (R-NJ), who proposed an amendment allowing states to choose whether to waive the ACA's essential health benefits. Though MacArthur was chairman of the more moderate Tuesday Group, his amendment was meant to entice the conservative Freedom Caucus. This appeal to conservatives paid off; of the 34 members who ended up voting "Yea" on the AHCA after declaring their intent to vote "Nay" in March, 31 were members of the Freedom Caucus (or were hard-line conservatives). Of these 34 members, 23 then made a public statement in favor of the bill. The House passed the AHCA, 217-213.

---

[15] Speaker Paul Ryan (R-WI) pulled the AHCA from a vote in March 2017, after learning that the party did not have enough votes to pass the bill. I have a whip count of each member's intended vote in March, and their actual vote when the bill passed in May: https://www.washingtonpost.com/graphics/politics/ahca-house-vote/

Republicans held 52 seats in the Senate. Vice President Mike Pence would cast any tie-breaking vote, so Republicans could allow no more than two defectors and win. The Senate produced three bills, but none passed: the Better Care Reconciliation Act (BCRA), the Obamacare Repeal Reconciliation Act (ORRA), and the Health Care Freedom Act (HCFA). In all cases, zero Democrats voted "Yea."[16]

The Senate leadership did not generally try to make deals with specific members. Instead, appeals made to specific members were often in the form of threats.[17] Some of these threats may have worked. For example, Sen. Dean Heller (R-NV) voted "Yay" on the HCFA after President Trump threatened him over his "Nay" votes on the BCRA and the ORRA.[18] Sen. Shelley Capito (R-WV) was chastised for her "Nay" vote on the ORRA and then voted "Yea" on the HCFA.[19] In all, 10 Senators changed their vote from "Nay" to "Yea"—either between the BCRA and the ORRA or between the ORRA and the

---

[16] https://www.washingtonpost.com/graphics/2017/politics/ahca-senate-whip-count/

[17] Rep. Buddy Carter (R-GA) suggested Sens. Collins (R-ME) and Murkowski (R-AK) deserved a physical reprimand after not voting in favor of a motion to proceed with debate on the HCFA (http://time.com/4875702/buddy-carter-snatch-a-knot-in-their-ass-lisa-murkowski/). Interior Secretary Ryan Zinke also attacked Murkowski, telling her and Sen. Dan Sullivan (R-AK) that their requests for Department of Interior funds could be in jeopardy if they did not vote with the party (https://www.washingtonpost.com/news/powerpost/paloma/daily-202/2017/07/28/daily-202-trump-s-hardball-tactics-backfire-as-skinny-repeal-goes-down/597a7cf630fb045fdaef0fd5/).

[18] President Trump made electoral threats to Heller, and pro-Trump PAC launched an ad campaign against him, over his "Nay" vote on the ORRA (http://www.cnn.com/2017/07/19/politics/dean-heller-trump/index.html).

[19] Rep. Blake Farenthold (R-TX) challenged "female Senators from the Northeast" to a hypothetical duel over opposition to Republican health care bills. Sens. Susan Collins (R-ME), Shelley Moore Capito (R-WV), and Lisa Murkowski (R-AK) were key holdouts. (https://www.washingtonpost.com/news/morning-mix/wp/2017/07/25/texas-rep-farenthold-says-he-would-challenge-female-gop-senators-to-a-duel-if-they-were-south-texas-men/).

HCFA—and each made a statement in favor of the bill on which they decided to vote "Yea."[20]

In both the House and the Senate, Republicans who switched their vote to vote "Yea" with the party were more likely to then make a statement in support of the AHCA than all other members (see Table 2). However, the Republican leadership did not do enough to assuage the concerns of Sens. Murkowski, Collins, and John McCain (R-AZ), all of whom voted against the final bill (the HCFA), leading to its defeat 49-51.

1.4 Evidence of Statements as Precommitment Devices

If members' public statements are indeed precommitments to voting with the party, three things should be true. (1) Members who are convinced to vote with the party should be more likely to make a public statement in favor of a bill than other members. The party does not need members who were always going to vote "Yea" to make a public statement in favor of the bill to help ensure their votes. However, members who are convinced to vote with the party may make a statement for the purpose of precommitment, claiming credit, or both. If members make public statements as a precommitment to vote with the party, rather than to claim credit, then the likelihood of making a statement in favor of the bill should not be affected by (2) whether members gained any policy concessions on the bill or (3) public opinion on the

[20] Sens. Bob Corker (R-TN), Tom Cotton (R-AR), Lindsey Graham (R-SC), Mike Lee (R-UT), Jerry Moran (R-KS), and Rand Paul (R-KY) voted "Nay" on the BRCA but "Yea" on the ORRA and the HCFA. Sens. Lamar Alexander (R-TN), Shelley Capito (R-WV), Dean Heller (R-NV), and Rob Portman (R-OH) voted "Nay" on the ORRA but "Yea" on the HCFA.

bill. Members who gained policy concessions have a legislative accomplishment to claim credit for, and MCs are more likely to claim credit when they have something to say that they think their constituents will like to hear. Thus, if the probability of making a statement is unchanged by whether members gained policy concessions or how much a member's constituents like the bill, there is evidence that these statements are precommitment devices.

I used logistic regression and Lasso regression (Tibshirani 1996) to investigate whether public opinion and having gained policy concessions affected the probability that Republicans made a statement in favor of the AHCA. For each model, the dependent variable is whether a member of the majority party made a statement of support for their party's AHCA bill (Yes or No). The main independent variables are: (1) a dummy variable for whether a member switched his or her vote from "Nay" to "Yea" (switched), (2) a dummy variable for whether a member gained policy concessions from the party in exchange for his or her "Yea" vote (concessions), and (3) the level of support for the bill from a member's constituents (MRP). Each model also includes the number of terms served in Congress (tenure), and dummy variables for race (white), gender (male), whether the member holds a party leadership positions (leader), and whether the member serves on a prestige committee (prestige). Logistic regression model results are displayed in Table 3. The models in columns 1, 2, and 3 include only one of the main independent variables each, plus covariates. The full model with all variables is in column 4.

Across all models, members who switched to vote with the party were more likely to make a statement in favor of their party's health care reform bill, and this effect was highly significant. Furthermore, this variable leads to substantial increases in the pseudo R^2 (as seen in columns 3 and 4). Much more of the variance in whether members make public statements in favor of a bill was explained by switching to vote with the party than by public opinion and policy concessions.

In model 2, where the only main independent variable is whether a member received policy concessions on the bill, members who received policy concessions were more likely to make a statement in favor of their party's bill. However, when I also considered whether those members switched to vote with the party (the full model in column 4), this significant effect of policy concessions goes went, even changing sign. Public opinion did not significantly affect the likelihood of members making a statement in favor of the AHCA in any model.

Table 3: Logistic regression results predicting whether a Republican made a statement in favor of the AHCA. Note: * = $p < 0.05$, ** = $p < 0.001$

	(1)	**(2)**	**(3)**	**(4)**
Switched			2.19**	2.75**
			(0.39)	(0.57)
Concessions		0.95*		-1.02
		(0.41)		(0.65)
MRP	0.02			0.004
	(0.02)			(0.02)
Tenure	-0.01	-0.01	-0.01	-0.01
	(0.02)	(0.02)	(0.02)	(0.02)
Male	-0.36	-0.49	-0.72	-0.68
	(0.43)	(0.44)	(0.45)	(0.45)
White	0.75	0.73	1.05	1.14
	(0.68)	(0.68)	(0.76)	(0.81)
Leader	0.19	0.28	0.53	0.52
	(0.91)	(0.91)	(0.91)	(0.91)
Prestige	0.52*	0.61*	0.59*	0.53
	(0.26)	(0.26)	(0.27)	(0.28)
Intercept	-0.83	-1.12	-1.49	-1.47
	(0.81)	(0.78)	(0.86)	(0.94)
Observations	290	290	290	290
Nagelkerke Pseudo R^2	0.03	0.06	0.20	0.21

The models in Table 3 are consistent with the hypothesis that members made public statements as precommitment devices and not with the hypothesis that they were claiming credit to their constituents. However, more can still be done to describe the relative importance and effect of each variable in the full model. I used Lasso regression to further pull apart the effects of the independent variables. Lasso imposes a penalty (lambda) on the regression model, such that the absolute values of the coefficients cannot exceed some threshold. As the penalty increases, the coefficients for each variable trend toward zero. Variables whose coefficients reach zero (or "leave" the model) first have a smaller effect on the dependent variable. By varying lambda from the value where the first variable enters the model to an un-penalized model, we can see which variables are more important for predicting the dependent variable. Lasso estimates of coefficients for the full logistic regression model are displayed in Figure 1. These vary over the range of the logged regression penalty.[21]

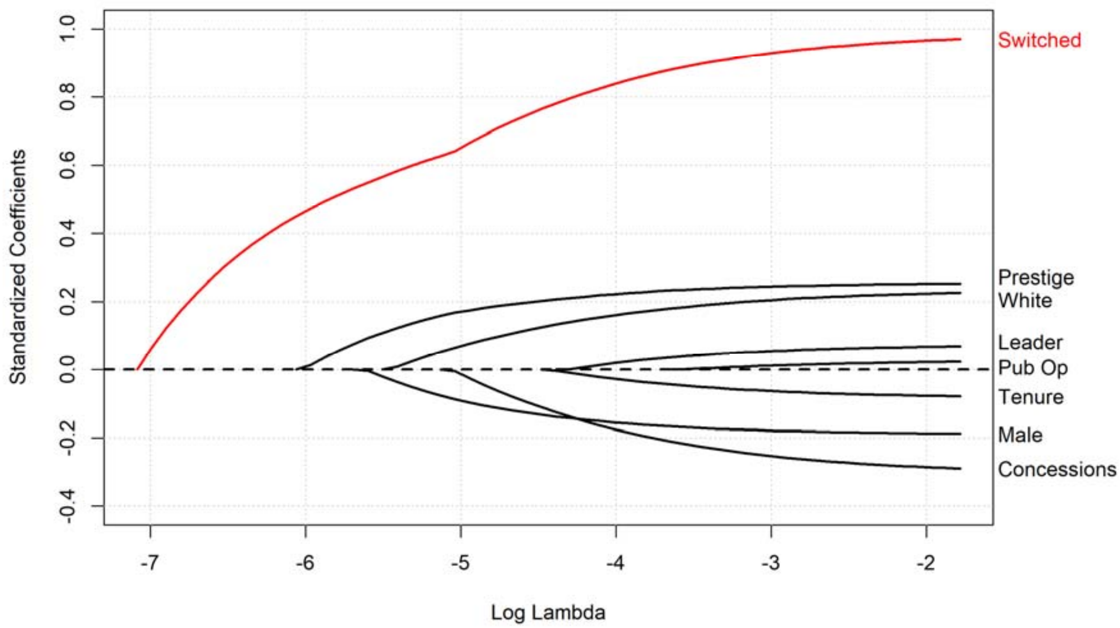

Figure 1: Lasso regression coefficient estimates for the full model, over the range of the logged regression penalty. The penalty decreases from left to right.

---

[21] This model was estimated using the 'glmnet' package in R, which automatically chooses a range for lambda according to the coordinate descent algorithm proposed by Friedman, Hastie, and Tibshirani (2010).

The first variable to enter the model was whether members switched to vote with the party. The coefficient for this variable was always the largest and was more than twice as large as the next highest in the un-penalized model.[22] As indicated by the logistic regression results above, the driving force behind whether Republicans made statements in favor of the AHCA is switching to vote "Yea" with the party and not more favorable public opinion or having gained policy concessions. To estimate how much more likely those who switched to vote with the party were to make a public statement than those who did not, I computed the predicted probability of making a statement in favor of the party's bill for members who switched their vote to vote with the party, and for those who did not, based on the full model shown in column 4 of Table 3.[23] These predicted probabilities are displayed in Figure 2. This model predicts that members who switched to vote with the party were 58 percentage points more likely to have made a statement in favor of the party's bill than members who were not convinced to vote with the party (82.0% versus 24.1%).[24]

---

[22] The variables in this model were standardized, so the coefficients are directly comparable.

[23] These probabilities are based on 1000 simulations of the dependent variable from the full model, based on the coefficients and variance-covariance matrix. All variables except "Convinced" (which was varied) were held at their median values.

[24] Even when taking the model uncertainty into account, there is a predicted difference of 35 percentage points between the middle 95% of the distribution (69.4% versus 34.8%).

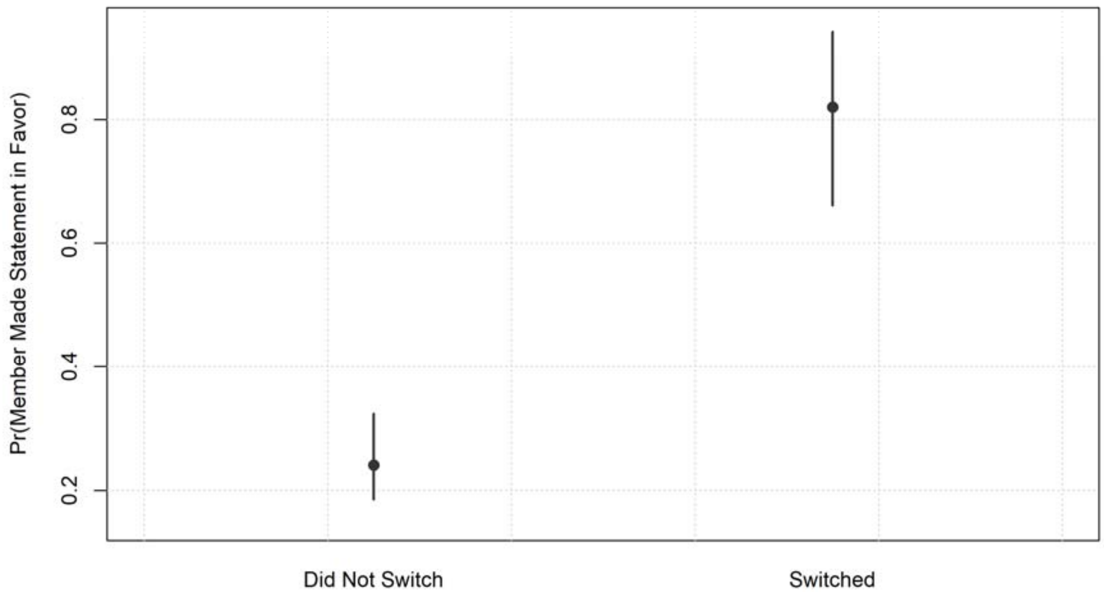

Figure 2: Predicted probability of Republicans making a statement in favor of the AHCA, based on whether the member was convinced to vote with the party. Bars represent the middle 95% of the distribution of simulated responses.

These predictions come from the full model, which was run on all of the data. Therefore, these are in-sample predictions. To better characterize the predictive accuracy of the model, and to test whether the model is over fit, I performed 10-fold cross-validation on the Lasso regression of the full model. In this way, not only could I test the predictive accuracy of the model, but I could also show how that accuracy changes over the range of the regression penalty. The distribution of out-of-sample predictive accuracy for all 10 folds, across all levels of lambda, can be found in Figure 3.

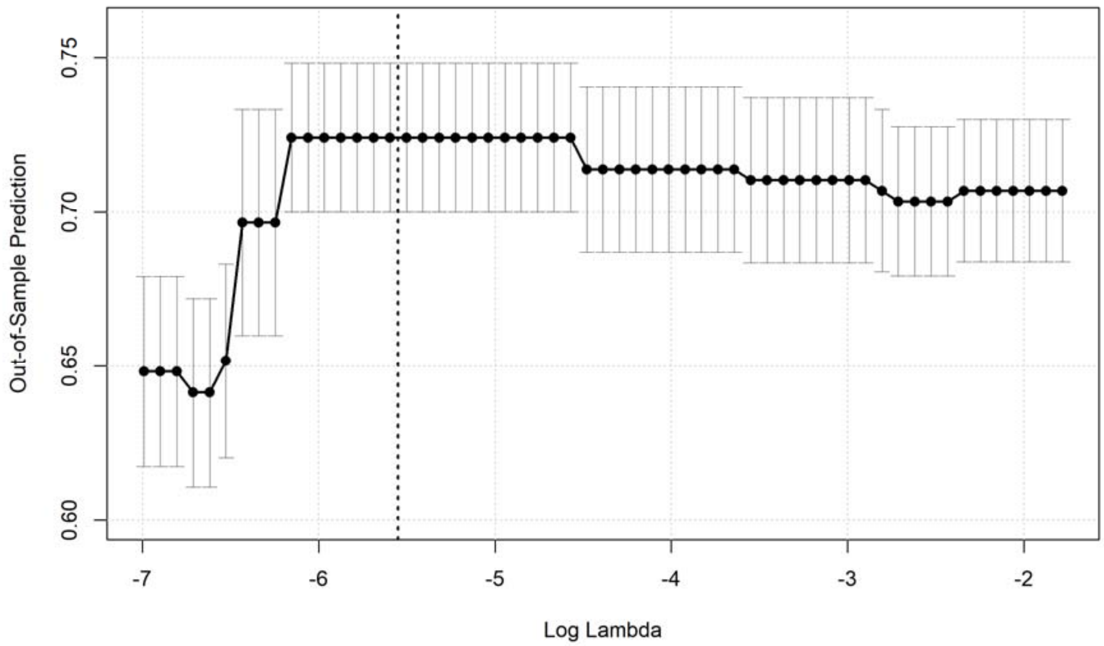

Figure 3: Out-of-sample predictive accuracy from 10-fold cross-validation for Lasso regression of the full model, varying over the range of the regression penalty, which decreases from left to right. Points are the mean out-of-sample accuracy, and bars denote the middle 95% of the distribution of cross-validated accuracy, across all 10 folds. The dotted vertical line indicates the point at which a second variable enters the model.

Interestingly, the out-of-sample prediction rate was maximized when the variable for whether members switched to vote with the party was the only predictor (before a second variable enters the model). A model predicting whether members made a statement in favor of the AHCA using only whether the member switched to vote with the party correctly predicted 72.4% ± 3.9% of the cases.[25] The accuracy of the model did not increase as more variables enter the model.

Taken together, these analyses offer strong evidence that Republicans made statements in support of the AHCA as precommitment devices. After Republicans switched to vote "Yea" with the party, most made a public statement in favor of the AHCA. Importantly, electing to make a statement was not mediated by public opinion or whether members gained policy concessions on the bill. These other variables did not increase predictive accuracy, were not statistically significant in the full model, and entered the Lasso model long after whether members switched to vote with the party. Even when members had good reason to use public statements to claim credit for legislative accomplishments, their behavior was only consistent with the use of statements as precommitment devices.

1.5 Discussion

Parties and their members often have competing incentives that lead to collective action problems. Members benefit from the party's brand, and for the party, protecting its brand means getting its members to vote for the party's preferred policies. However, members have their own individual interests that sometimes conflict with those of the party. At times, parties must secure votes from members whose personal or electoral

[25] Other cross-validation methods yield similar out-of-sample predictive accuracy, such as repeated 10-fold cross-validation (67.4%), leave-one-out cross-validation (65.8%), and bootstrap resampling (68.2%).

incentives give them cause to vote against the party, even though voting with the party would protect the party's brand, which benefits the members. Further, these recalcitrant members may only be willing to vote with the party if they can be assured the bill will pass. Parties must therefore not only convince members to vote the party line: they must assure members that they will not stick out their necks for a bill that fails, and they must hold members to their deals. When parties convince recalcitrant members to vote with the party, those members also often public statements in favor of a bill. While such statements might be credit claiming, I show that members still make statements even when they have little reason to claim credit for their votes. In these situations, members' statements serve instead as precommitment devices, effectively locking members into voting with the party.

This study is one of the first to show that statements made by members of Congress have strategic value to lawmaking. It is often argued that legislators' public statements amount to little more than cheap talk, having scant strategic value and no discernible signal that is separable from voting decisions (*e.g.*, Farrell and Rabin 1996). I break from this proposition and show that public statements have a strategic purpose for signaling precommitment. As such, they are separable from voting behavior. Further work should be done to confirm these results and extend them to more bills and situations to answer questions such as: Are parties are more likely to push members to make precommitment statements in certain issue areas, or on more salient issues? When the party can find bipartisan support, do they push for both in- and out-party members to publicly commit their support for the bill? How can we better assess whether it is in fact the party pushing members to make these statements?

The political environment surrounding the bills I investigate can also provide additional context for the results. For instance, both bills were controversial. Voting in favor of the ACA appears to have cost Democratic incumbents six percentage points in the 2010 midterm elections (Nyhan et al. 2012), and the AHCA was the "most unpopular piece of major legislation that Congress ha[d] considered in decades" (Warshaw and Broockman 2017). As such, it could be argued that claiming credit for having won policy concessions would soften the potential electoral blow of voting "Yea" on an unpopular bill. And yet, regardless of public opinion or having gained policy concessions, most Republicans and all Democrats made a statement in favor of their party's bill, once they agreed to vote with the party. Subsequent research could apply this work to bills with more variance in public opinion in order to separate statements made for the purpose of precommitment from statements made to soften the electoral blow of an unpopular decision to one's constituents.

The bills I studied are also indicative of the contentious party environment in Congress today. Democrats initially looked to pass the ACA with bipartisan support but found none. Republicans did not even look for bipartisan support for the AHCA. Other high-profile bills have passed largely along party lines in recent years,[26] and the nature of policy bargaining in Congress changes when neither party is willing to allow any members to support the other party's policies. If the parties in Congress are well-sorted and/or polarized, and the majority party believes they are unlikely to receive support from the minority party, then the majority party does not need to consult the minority party when trying to pass legislation. At the same time, the majority also needs support from their own members more often. When one party controls both chambers in Congress or has a fully unified government, the minority party may have even less of a role in policy-making. In such a system, the policies that Congress considers would be increasingly undesirable to the minority party, perhaps without being less likely to pass. While some voters prefer their representatives try to reach bipartisan solutions (Harbridge and Malhotra 2011), this preference is becoming less common among the electorate and such preferences do not outweigh other considerations, like party ties (Harbridge, Malhotra, and Harrison 2014). The ways in which Congress creates policy, and the reactions of public opinion, may intensify party polarization. Studying when and how parties convince their members to vote with the party, and how members respond, can illuminate the ability of parties to pass legislation in polarized conditions.

---

[26] For example, "Cap and Trade" passed the House in 2009 by a vote of 219-212, with just 8 Republicans voting "Yea." The Dodd-Frank Wall Street Reform and Consumer Protection Act passed through the House and Senate in 2009 with just 6 Republican "Yea" votes. A budget resolution reducing federal spending on health care passed the House and Senate in 2005 with just a single Democratic "Yea."

This discussion may appear to paint a bleak picture for the future of bipartisan politics in the United States. However, there is much more that we can learn about how polarization and strategic political speech play a role in the policymaking process, particularly when collective incentives of parties conflict with the individual incentives of their members. Paying close attention to what members of Congress say can tell us not only about their relationships with their constituents but with their party as well.

2. Bill Text and Agenda Control in the U.S. House of Representatives

2.1 Introduction

A common finding in American legislative studies is the ability of the majority party to exert power over the agenda and in doing so bias policy outcomes toward the preferences of the majority. Schattschneider (1960) called this "mobilization of bias" the core function of parties in democracies. The very notion of bias in this context is compelling because it implies a comparison to a "true" or honest representation, presumably of the median voter (Krehbiel, 1998). If the bias notion is correct, parties are doing more than simply competing over who represents the median voter. Instead, party composition and action directly affects what gets done, and even what gets said, in the legislature (for reviews, see Aldrich 2011, and Masket 2009).

Understanding how and when parties matter is fundamental to our knowledge of legislative politics. But to date studies of party influence and agenda power have been limited by a lack of reliable measurements of the use of party power in shaping the agenda, and through it, legislation. How can we tell when the party has exerted positive or negative influence on a bill or other legislative vehicle? And compared to what? There is no available counterfactual of when party "bias" had not been mobilized. Many studies have found evidence of majority party agenda control in both the House and Senate (*e.g.*, Carson, Madonna, and Owens 2016; Cox and McCubbins 2005; Den Hartog and Monroe 2011; Gailmard and Jenkins 2007; Jenkins and Monroe 2012, 2016; Monroe and Robinson 2008) as well as in state legislatures (Anzia and Jackman 2012; Jackman 2013). However, these have almost exclusively dealt with negative agenda power exercised on the floor of Congress, where institutions have combined with partisan control to block legislation. There has been little systematic empirical evidence of positive agenda power, where parties secure legislative language or the passage of laws that would not otherwise have passed simply based on overall chamber majorities (but see Cox and McCubbins 2005, Chapter 10). Furthermore, the focus on votes and passage of legislation has come at the expense of measuring negative agenda power in processes that take place before committee or floor votes in Congress. Because majority party bias at the level of formulation and direction of legislation at the early stages are at the heart of what is claimed by the theories of partisan control (Aldrich and Rohde, 2001; Cox and McCubbins. 2005), studying agenda power before votes is necessary to truly understand how the majority exercises power.

Cox and McCubbins (2005) provided one of the first means of comparing the success and failure of majority party power, through introducing the concept of situations where the party gets rolled.[27] Rolls are a measure of negative agenda control—the majority party is rolled when a bill passes that a majority of the majority party opposes. A strong party should experience few rolls, since much of its power comes

[27] In parliamentary systems, of course, if the government gets rolled, or loses a vote on legislation it has supported, this means an election will be called. Consequently, party discipline is much more explicit and powerful in such systems.

from keeping bills off the floor that a majority of its members dislike.[28] Since Cox and McCubbins proposed the measure, many studies have used roll rates as their measure of agenda control (*e.g.*, Anzia and Jackman 2013; Cox, Kousser, and McCubbins 2010; Den Hartog and Monroe 2011; Gailmard and Jenkins 2007; Gamm and Kousser 2010; Jackman 2013; Jenkins and Monroe 2012), but others have raised objections to their entrenchment as the dominant measure of agenda power in the literature. For instance, rolls may be explained by non-institutional factors, such as the size and cohesion of the majority party (Krehbiel 2006), and how valuable the party leadership perceives a bill to be (Krehbiel and Woon 2005). Rolls are also based on final passage votes, which the majority does not often expect to lose, particularly in contemporary Congresses. There has also been disagreement whether it is appropriate to use counts or rates when discussing rolls (see Gailmard and Jenkins 2007, Jenkins and Monroe 2016 for examples). As such, rolls may only measure the upper bound of party power, rather than a measure of how often and effectively parties actually *use* their institutional powers to control the agenda. Although they have been used as a measure of negative agenda power, rolls more closely resemble a measure of the majority's ability to control its members' voting decisions. Rolls do not deal with how well majority parties are able to keep policies

[28] Such an occurrence requires not only for the Speaker to schedule a vote on legislation that is unpopular with most of the members of their party, but for that vote to pass. Speakers since the mid-1990s have informally required support from a "majority of the majority" in order to schedule a final passage vote on a piece of legislation, known as the Hastert Rule. Even so, each Speaker since Newt Gingrich "violated" the rule at least a handful of times.

disliked by most of their membership from being considered but instead with how well they are able to keep such policies from becoming law.

As such, there has recently been a resurgence of interest in moving beyond roll rates (*e.g.*, Jenkins and Monroe 2016), in particular emphasizing instead the importance of pre-floor agenda control (rolls only ever happen on final passage votes) in Congress and a need for measures of positive agenda control.[29] Such measures, however, have not actually been developed yet, in part because they also require a method to infer how members would have voted on bills that never reached the floor. Here, I introduce such a method. Using the text of the bills, I train a neural network to develop a vector representation of each bill, with which I can infer how a given member of Congress would have voted on a given bill, had that bill come to a vote on the floor of the House. I then use these inferences to present some of the first measures of pre-floor party agenda control, both positive and negative, in Congress.

With these estimates, I quantify and compare both pre-floor and floor agenda power. I specifically examine bills after they are 1) *introduced*, 2) *referred* to committee, 3) *reported* out of committee, and as they are 4) *voted* on. I find ample evidence of both positive and negative agenda power exercised by the majority party throughout the legislative process, but most of the agenda power is used when bills are still in

---

[29] As in Jenkins and Monroe (2016), I use both agenda power and agenda control throughout, but differentiate between the two. By agenda control, I mean the theoretical ability of different actors to control the agenda. By agenda power, I mean the exercise of that control.

committee. Furthermore, I show that the minority party is systematically and consistently shut out of the agenda process, but that most minority party members still vote Yea on most final passage votes, suggesting both that the agenda process and the process of influencing policy are distinct, and that the minority party is not shut out of the policymaking process, despite their inability to shape the agenda.

2.2 Background

Work on agenda power is part of the broader study of political parties. This research can be broken into three main questions: Why do parties exist? How are they organized? What do they do? Researchers generally agree that parties exist to solve the legislative (*e.g.*, Aldrich 2011; Smith and Gamm 2001; Smith 2007) and electoral (*e.g.*, Cox and McCubbins 2007; Evans and Oleszek 2001) collective action problems that confront legislators. In order to do so, parties organize by delegating decision-making authority to central agents to reduce transaction costs (*e.g.*, Rohde 1991; Sinclair 1983, 1999; Gamm and Smith 2002) and distribute proposal and veto power throughout the party to various senior and/or loyal members (Cox and McCubbins 2005). In practice, (majority) parties in Congress work as voting (Rohde 1991; Aldrich 2011; Aldrich and Rohde 2001) and procedural (Cox and McCubbins 1991, 2005) coalitions, controlling both the voting behavior of their own members and the agenda considered by the chamber.

2.2.1 Measuring Agenda Power

There is some disagreement about the relative prevalence of various interests and sources of power— *e.g.*, primacy of the power to control members' votes versus control of the agenda—but theories of parties in Congress are generally complementary. In particular, all the dominant perspectives on parties (Cox and McCubbins 2005, 2007; Aldrich 2011; Rohde 1991; Aldrich and Rohde 2001) share the position that the majority party's goals when exercising its agenda power are to further their members' interests. First, the majority party seeks to block bills that most of their members oppose (negative agenda control). Second, the majority party seeks to turn into laws bills which most of its members support (positive agenda control).

This implies a 2x2 typology of agenda control, based on the majority party's preferences toward a given measure (support or oppose) and the eventual outcome (pass or fail). Jenkins and Monroe (2016) proposed just such a typology, which can be found in Table 4. While rolls have been extensively studied, we must at minimum measure the other portions of the agenda power typology to be able to understand how the majority manipulates the agenda in Congress. However, observing the majority's full influence on disappointments and blocks in Congress still eludes scholars, in part because agenda control in Congress is often a pre-floor activity—accomplished in the House by majority-controlled committees (particularly Rules) or by the Speaker—fewer behavioral measures of which are typically recorded.[30] Some disappointments and blocks may happen on the floor and are recorded. But from the sheer number of bills that die in committee (after being referred; see Table 5), we are surely missing a substantial proportion of blocks and disappointments by studying only floor behavior.

Table 4: Typology of agenda power—majority position versus bill outcome.[31]

	Pass	***Fail***
Support	Success	Disappointment
Oppose	Roll	Block

There are several legislative chokepoints to consider when examining (pre-)floor agenda power in the House. Many bills die after being referred to committee. Some make it out of committee, but receive no vote on the floor. Others are voted on but do not pass. Still more are passed by one chamber, but not taken up by the other. A small portion of bills are introduced but never referred to committee. At times, a bill can even bypass the committee system entirely.[32] I categorize bills according to which of three pre-floor stages along the legislative processes a bill reached—Introduced, Referred, or Reported—as well as whether a bill received a final passage vote. In other words, the

[30] For this reason, Jenkins and Monroe do not analyze Congress in their 2016 paper. Instead they examine 85 state legislatures, the Mexican Chamber of Deputies, and the Canadian House of Commons, as these have a more observable mixture of pre-floor and floor agenda power.
[31] Table 4 comes from Jenkins and Monroe (2016).
[32] As Speaker, Newt Gingrich would often assign a task force to mark-up bills, a duty traditionally performed in committee.

final action before a bill "died"—though this moniker does not apply to "Voted" bills that then pass and become public laws.[33] For instance, a bill categorized as Referred died in committee, after being referred. And a Reported bill died after being reported out of committee, but before a vote.

Table 5: Bills reaching each legislative stage in the House of Representatives, by Congress.[34]

	Introduced	*Referred*	*Reported*	*Voted*	*Total*
104th	55	3,975	441	213	*4,684*
105th	67	4,427	464	242	*5,300*
106th	52	5,195	534	337	*6,118*
107th	27	5,383	391	260	*6,061*
108th	46	5,124	392	298	*5,860*
109th	53	6,140	343	271	*6,807*
110th	44	6,705	444	426	*7,619*
111th	49	6,032	335	399	*6,815*
Total	*393*	*43,081*	*3,344*	*2,446*	*49,264*

The number of bills that finish their journey at each of these stages of the legislative process, can be found in Table 5 for all the Congresses for which I have data (104th-111th, corresponding to 1995-2010). As we can see, the vast majority of bills that are

[33] However, because I only consider the House, and bills must also pass the Senate, passage in this sense does not mean becoming law.

[34] These tallies are computed from the text versions that are available for different bills (which come from the Government Publishing Office) and from data collected by political scientists at PIPC (Crespin and Rohde 2018) and the Congressional Bills Project (Adler and Wilkerson 1995-2011).

introduced in the House die in committee (87.4%), while very few die immediately after being introduced (0.8%). More of the bills which make it out of committee die without a vote (6.8 versus 5.0% of all bills). Categorizing these bills according to the agenda control typology in Table 4 will illuminate the degree of agenda control in the House.

There are two forms of legislative behavior of particular interest when considering pre-floor votes which go unexplained by the current typology, and they are predicated on the expected outcome of an eventual floor vote (even if such a vote never occurs). First, if the majority successfully keeps a bill they oppose from reaching a final passage vote, how we interpret that execution of agenda power differs if the bill would have passed or failed. In the case of the former, the majority has successfully blocked a potential roll—this is a stronger use of agenda power than blocking a bill that would not have had enough support to pass anyway. I call these *dodges*, where the majority has blocked a bill that they oppose, but would have passed. The inverse case is also notable. Bills that are disappointments (when the majority party supports a measure that fails) could have passed or failed when put to a full floor vote. When the majority fails to shepherd a bill they support through the legislative process, even though it would have had the votes to pass if the measure had been put to a vote, I call *missed opportunities* for the majority.

With the addition of dodges and missed opportunities, we have a new typology classifying the use of agenda power in pre-floor legislative stages in Congress (Table 6) to supplement the existing topology for agenda power in floor legislative stages (Table 4). For different stages along the legislative process—both before the floor and on the floor—there are different expectations for the prevalence of each outcome of this typology. In each case, expectations are based on which actors have veto and/or proposal power in each stage of the legislative process, and how their interests compare with those of the majority party. Bills reaching each stage in the legislative process encounter different agenda setters with different preferences and relationships to the majority party. As such, we can expect a different proportion of bill types[35] in each pre-floor stage.[36]

Table 6: Typology of pre-floor agenda power activities—majority position versus bill outcomes.

	Predicted Pass	**Predicted Fail**
Support	Missed Opportunities	Disappointments
Oppose	Dodges	Blocks

Very few bills are introduced but not referred to a committee. Indeed, bill referral to committee is near-automatic in the House (done by the House parliamentarian in conjunction with the majority party leadership). Those that never make it to committee

[35] By "bill types" I mean outcomes in my typology of agenda power (Table 6) and that of Jenkins and Monroe (2016; Table 4 of this dissertation).
[36] Because successes and rolls require a bill to pass, there will be no successes or rolls on bills that never reach a final passage vote, which necessarily precludes all bills that die before reaching the floor from being successes or rolls.

are more likely to bypass the committee process than die just after being introduced. Because referral to committee is usually an automatic process, and the majority party can bypass the committee system, bills that the majority supports are incredibly unlikely to die after being introduced, but with no further legislative action. Because of this, we can expect that bills which die after introduction are unlikely to be those which the majority party favors (disappointments or missed opportunities), as the majority can easily take up a bill if it wants. More likely than not, we would expect such bills to be blocks or dodges—those which the majority opposes. Furthermore, because blocks also do not have enough support to pass, there should be more blocks than dodges that die after being introduced.

As much of the legislative action in Congress traditionally happens in committees so too does much of the agenda control. As we can see from Table 5, more bills die in committee than any other stop along the legislative process. Like those bills that die after introduction, bills that the majority supports are not likely to die in committee. For one, the majority party controls more seats and the chair position in every committee and subcommittee in both chambers. Further, committee chairs are usually selected in part due to their loyalty to the party (particularly among Republicans, but also among Democrats), rather than by the once ironclad norm of seniority from the pre-reform era (Rohde 1991). We should find few missed opportunities or disappointments among those bills that die in committee.

Bills that die in committee are similar to bills that die after introduction in that they are likely to be opposed by the majority party. These are bills that the majority is blocking or dodging by keeping them away from a passage vote. Bills that the majority wants to bury, because they would be divisive to their membership or because they do not fit with the party's policy preferences are most likely to die in committee. Indeed, many more bills die in committee than at any other stage, per Table 5. Most bills die in committee because passing a committee vote is the largest pre-vote hurdle that a bill faces, and the majority's control of the committee system gives them the most power to stop unwanted bills relative to other stages. As such, I expect relatively more bills dying committee will be dodges than in any other legislative stage. The majority wants to avoid being rolled if at all possible, because it is costly to their legislative agenda and to the power of the caucus. Therefore, the majority party, having some knowledge about which bills could potentially roll them, is more likely to use their most concentrated agenda power to stop their least-preferred outcome.

Agenda-setters in the House majority party also have a great deal of influence on bills that are reported out of committee but have not yet reached a final passage vote. The Rules Committee can substantially shape what is allowed for debate on a measure using special rules, which have become more common in recent decades (Curry 2018). A bill is voted on in the House after the designated debate period for the bill has ended, unless a motion to recommit the bill to a committee is passed. Such motions are usually proposed by opponents of the bill being considered as an attempt to change some provision. The House majority leadership can also pull a bill if they do not think they have enough Yea votes to pass it. In either event, the majority is not likely to need to dodge bills by killing them after being reported, but before being voted on. Few bills in the House die at this stage, and those that do will least likely be disappointments or missed opportunities, relative to other stages. A vote will automatically happen once a bill has been reported out of committee without action, so for a bill to die in this stage it will most likely be pulled by the majority.

There is variance in the degree of control the majority party has over each stage of the pre-voting legislative process in the House, but the majority does still enjoy strong institutional control of the various roadblocks a bill must surpass in order to come to a final passage vote. As such, I hypothesize that I will find a great deal of agenda control exercised by the House majority party. In particular, I expect three trends. First, I anticipate finding strong evidence of negative agenda control. This hypothesis will be supported if I find that bills which a majority of the majority party opposes are more likely to die before reaching a final passage vote. Second, I also expect to find a great deal of positive agenda control. Positive agenda control should manifest in bills supported by the majority being more likely to reach the floor for a vote, and for such bills to be a small proportion of the bills dying in each legislative stage. Third, I expect that the committee process is doing the bulk of the yeoman's work of agenda control in the House. I will find support for this hypothesis if a higher proportion of bills that the minority party favors and/or the majority party opposes die in committee, relative to other stages.

The method I present below allows me to measure agenda control exercised by the House majority party throughout the legislative process, including on pre-floor legislative actions. This is made possible for a new method I have developed for predicting the outcome, had the bill come to a vote, of bills that did not reach a final passage vote. With these predicted outcomes, I can quantify successes, rolls, disappointments, and blocks plus dodges and missed opportunities for the first time. I expect to find evidence of a great deal of both positive and negative agenda power in the House. Furthermore, there should be systematic variance in terms of which types of bills will likely be affected in each stage along the legislative process. When the majority caucus has more power over the legislative process, as in during the committee stage, they will use it to block unwanted bills. This execution of agenda control allows the majority to protect its members from bad votes by dodging bills which might roll the party and to make sure that bills supported by the majority pass safely through the chamber.

2.3 Methods

2.3.1 Data

In the analysis that follows, I first predict members' votes in the House. Then, I use the model to predict votes to infer how members would have voted on bills, had they come to the floor for a vote. From these estimates, I compute the level of support for each bill from the majority and minority parties, and the overall chamber. In order to accomplish these tasks, I use summary data on congressional bills,[37] roll call votes, members of Congress, and votes cast in Congress.[38] I combine these with congressional election returns data for both primary and general elections,[39] and all text versions of each bill before Congress.[40] I employ these data for the 104th-111th Congresses.[41]

2.3.2 Bill Text Versions and Legislative Stages

Each bill, regardless of where its journey through Congress ends, has at least one version of its full text entered into the Congressional Record. For instance, as soon as a bill is introduced, a text version is entered into the record based on whether the bill originated in the House (in this case, the text version is labeled 'ih' for 'Introduced in House'). As a bill moves through the legislative process, additional versions of the full

---

[37] Data on Congressional bills come from the Congressional Bills Project (Adler and Wilkerson 1995-2011), and from PIPC (Crespin and Rohde 2018).

[38] Data on roll calls, MCs, and votes come from Voteview (Lewis et al. 2017).

[39] Election returns data come from Primary Timing Project (Boatright, Moscardelli, and Vickrey 2017).

[40] Bill text was scraped from the U.S. Government Publishing Office (GPO).

[41] These Congresses correspond to the years 1995-2013. GPO only keeps data in an easily accessibly format back to 1995, and version recording is incomplete after the 112th Congress. Further, matching between the Congressional Bills Project, PIPC, and Voteview datasets is currently not possible past the 111th Congress.

text are entered into the record. These later versions often denote a point along the legislative process (*e.g.*, 'rfh' for 'Referred to Committee in House').

Using these bill versions, I classify the point along the legislative process that each bill reached in the House. I specify four different stages: *introduced*, *referred* to committee, *reported* by committee, and *voted* on. All bills that are introduced in Congress have an Introduced text version entered into the Congressional Record. Almost all of the bills that are introduced are also referred to at least one committee that has jurisdiction over the issues present in the bills. Even if a bill is referred to more than one committee, the text versions referred to each committee are identical. Referral marks the beginning of the committee process in Congress, where most of the mark-up and amendments to bills happen.[42] Fewer bills are then reported out by a committee. Most but not all that are reported out of committee are voted on.

Note that these stages are usually sequential: bills classified as "referred" were also introduced, but not reported or voted on.[43] I classified bills as reaching a certain legislative stage by taking the text version for each bill that was most recently entered into the Congressional Record, and matching it to a stage by its bill text version code.

[42] Although there has been a decline of regular order, and particular in usage of the committee system, in Congress since the 1980s (Curry and Lee 2018), most bills still go through the committee stage.
[43] Of course, sometimes bills that are referred but not reported can be voted on, in the case of discharge petitions in the House. There has also been a documented decline in regular order in Congress, where certain stages in the legislative process are more likely bypassed today than in the past (Curry and Lee 2018). Still, though, these stages are sequential for most bills.

The version codes that correspond to the four legislative stages I identify can be found in Table 7.

Table 7: Text versions by legislative stage.

	Version Code	Meaning
Introduced	ih	Introduced in House
Referred	rfh	Referred in House (Received from Senate)
	rih	Referred to House Committee with Instructions
	rth	Referred to Committee in House
Reported	cdh	Committee Discharged House
	rh	Reported in House
Voted	cph	Considered and Passed House
	eah	Engrossed Amendment House
	eh	Engrossed in House
	eph	Engrossed and Deemed Passed by the House
	enr	Enrolled Bill
	fph	Failed Passage House[44]
	fah	Failed Amendment House[45]

[44] This version only applies to bills and joint resolutions.
[45] This version only applies to engrossed (replacement) amendments.

2.3.3 A Model to Predict Congressional Voting

In order to infer how members would have voted on a piece of legislation, had it come to the floor, I must first be able to classify members' voting decisions with a high degree of accuracy. The framework I use for predicting members' votes comes from de Marchi, Dorsey, and Ensley (2017; hereinafter DDE), who built a predictive model of votes in the House based on bill text that yielded the highest out-of-sample predictive accuracy to date (see, for comparison, Gerrish and Blei 2012).

I extend the framework in a number of ways. First, I analyze additional Congresses (104^{th}-111^{th}), compared to DDE (111^{th}-113^{th}). Second, I consider bills authored by members of both the majority and minority parties, rather than just the majority. Third, and most importantly, while research predicting votes has to date only considered final passage votes, I also characterize the policy content of bills which never received a final passage vote, in order to infer how members would have voted on those bills. Because it is not reasonable to apply the topic model employed by DDE for characterizing bills to those bills which did not reach a final passage vote, I use a different method to characterize the content of bills (doc2vec instead of the Structural Topic Model) and to estimate my predictive regression model (logistic regression with stochastic gradient descent instead of a Lasso penalty). The specifics of the model of the model are explained in greater detail below.

2.3.4 The Plausibility of Predicting Votes which Never Happened

There are potential difficulties to overcome when predicting votes on bills that never came to the floor for a vote. For instance, one could ask whether bills that don't come to a vote, and bills that do, are part of the same data-generating process. Final passage votes on bills are a unique combination of salience, importance, and bargaining. Unlike other roll call votes, such as procedural votes and those on amendments to bills, final passage votes have the potential to create laws, therefore they are, on average, more important to members of Congress than other types of votes. This idea is at the heart of one of the long-standing criticisms of the commonly employed DW-NOMINATE scores (Poole and Rosenthal 1985, 1997). DW-NOMINATE scores are computed using *all* roll call votes in Congress, including final passage votes but also including amendments, procedural votes, confirmation hearings, etc. Because these votes may not all be part of the same process, using a scaling method that includes all roll call votes has been shown to conceal some of the intricacies of the legislative process (Aldrich et al 2014; Crespin and Rohde 2010; Roberts et al 2016).

Bills that reach a final passage vote have also often undergone a complex bargaining process between members. In order to shore up support for a bill, its sponsors (often through the majority leadership) will whip votes and make deals with members in exchange for their vote. Often, such deals are made so that the policy content of the bill changes, or in the form of logrolling. When members cast their votes on the final passage of a bill, then, their decision-making is potentially affected by this process (see chapter 1 of this dissertation).

While these considerations certainly suggest that votes on final passage are qualitatively different from other types of roll calls, do they also mean that bills which reach a final passage vote are qualitatively different than bills which do not? Consider the predictive model put forth by DDE: their model (and my extension using doc2vec) is highly predictive of members' voting decisions using only characteristics about the member's voting history and tenure, their district, the party breakdown by Congress, and the topical content of bills. None of these variables captures the bargaining process in Congress, for instance. Instead, the model assesses, "Given what we know about a member and their district, how will they vote on the bill at hand, based on its content?" Importantly, the model implies that members cast their voting decisions, in large part, based on the content of bills at hand.

On bills which never made it to the House, all of the DDE variables exist, including the policy content of each bill. Though bills that never reached a vote do not successfully navigate the bargaining and agenda-setting hurdles required, they do still have policy content. If members' votes can be determined by the policy content of bills, then simulating how members would have voted on a bill, based on the last text version of that bill, is reasonable. One can correspondingly think of the process of predicting votes on bills which never received a vote as asking the question, “How would each member vote if *this* version of a bill got a vote?”

2.3.5 Considering Majority- and Minority-Sponsored Bills

Like many applications for text analysis in political science, DDE approximate the topical content of bills from their text using a topic model, in this case an extension of Latent Dirichlet Allocation (LDA; Blei, Ng, and Jordan 2003) called the structural topic model (STM; Roberts et al. 2014). Topic models such as LDA/STM work by generating topic weights based on the frequency of tokens (usually words, but sometimes short phrases) within and between documents. However, the existence and frequency of a word in documents within a corpus does not tell us anything about the meaning of that word or the context in which it appears. For example, the word “abortion” appearing a similar number of times in two bills, one authored by a Republican and the other by a Democrat, would likely take on a different meaning in context in each bill. However, a topic model would not pick up on this difference and would instead give each document similar weights for topics containing the word abortion, since each deals with the topic at a similar frequency.[46] As I explain below, doc2vec preserves the context in which words appear and characterizes their meaning by doing so. This approach

[46] Assuming that each document also contains a relatively similar proportion of other words in topics that include the word “abortion”.

ameliorates the issue of unclear usage of words if we assume that a word's usage in context can help discern between different meanings of a word.

Finally, if there were a qualitative difference between the usage of words between bills authored by the majority and minority parties, we would expect a sizable difference in the model's predictive power on majority- and minority-introduced bills. Voting Yea on a bill is often akin to agreeing with the majority party's position, as most bills are introduced by the majority. It would be particularly troubling if the predictive accuracy were higher for bills sponsored by the majority than for bills sponsored by the minority, as it is generally more difficult to predict the votes of members of the minority party. However, this is not at all what I find (see Table 9); the predictive accuracy for bills introduced by the majority is nearly identical to the predictive accuracy of bills introduced by the minority.

2.3.6 Characterizing Policy Content with doc2vec

To model a semantic understanding of the policy positions in each bill, I turn to doc2vec, a machine learning technique used in text analysis shown to preserve information about the semantic context of language within documents (Le and Mikolov 2014). Doc2vec is an extension of word2vec (Mikolov et al. 2013a, 2013b), which creates word embeddings—or vector representations of words—via a shallow neural network.[47]

---

[47] See Bengio et al. (2006), Collobert and Weston (2008), Turian et al. (2010), and others for more on word embedding techniques.

In order to understand the intuition behind doc2vec, a primer on word2vec is useful.[48] In word2vec, words in document are transformed into a vector of real numbers, while still maintaining information about context and meaning. This is done by training a neural network to predict the words that fall into some window around a randomly selected target word.[49] The network is trained to do this by assessing pairs of words found in the training documents, consisting of the target word and all words within the window size of that word. Consider the following sentence, with a window size of 2:

"How quickly daft jumping zebras vex"

All the training pairs in the example sentence can be found in Table 8.[50] For the target word "How", the words "quickly" and "daft" are within the window size, so we would feed the network the training samples (how, quickly) and (how, daft). For the target word "quickly", the words "How", "daft", and "jumping" are within the window size, so we would feed the network the training samples (quickly, how), (quickly, daft), and (quickly, jumping). As the network cycles through all the words in the document(s),

---

[48] A more technical explanation of word2vec and doc2vec can be found in Appendix D.

[49] This parameter is called "window size" and is supplied by the researcher. In the literature on word embedding, window sizes of 5, 10, and 20 are most often used. I use a window size of 10 in my analysis. Smaller window sizes run the risk of not considering enough of the context in which words appear, so that context and meaning are not well-preserved. Larger window sizes run the risk of rendering the context in which a word appears less meaningful. Researchers often vary their analysis over a range of window sizes and choosing that which provides the best results for the task at hand (often some classification task, as I do with predicting bill stages).

[50] Note that the text must still be pre-processed. Namely, I removed the preamble of each bill, then removed extra white space, converted all the text to lower-case, and padded punctuation marks with a space on either side. This decision to not remove punctuation entirely stems from the idea that there may be some contextual value in punctuation, which is common in word2vec literature. I did not stem or lemmatize the text, which is also common in word2vec, but would be necessary with topic modeling techniques.

targeting each in turn, it learns based on how often each pairing of words occurs. Based on these statistics, the network will create an output probability that each word in the document occurs within the window size of the target word, via multinomial logistic regression. Because contexts are preserved, this means that once the model is trained, using an input of "United" will result in a higher probability of "States" or "America" than of "Brazil".

The trick is that we are not actually interested in the task for which we have trained the network. Once the network is trained, an input word (in the form of a one-hot vector)[51] will yield an output vector of probabilities for each word in the vocabulary—which is a list of all unique words in the training documents—being within the window size of the input word. However, in order to compute these probabilities, the vector passes through a hidden layer of the network, which is a matrix of weights where each row represents a unique word in the vocabulary. These weights are actually the word embeddings in which we are interested.[52]

---

[51] In a one-hot vector, all elements are 0 except for a single 1. This is one way to represent words as numbers, with a 0 for all elements except the one corresponding to the input word, which is a 1. The length of each one-hot vector is equal to the number of unique words in our training documents, the collection of which is called the vocabulary.

[52] The researcher must also select the number of "features", which sets the length of the word vectors given by the weights of the hidden layer. The original word2vec paper used 300 features, which has become standard in the literature (Mikolov et al. 2013a). I use two concatenated vectors, each of length 300, to represent each document. Each vector comes from a doc2vec model trained with a different iterative algorithm.

Table 8: Target words, nearby words, and training pairs for an example sentence. In the "Nearby words" column, target words are italicized, while words within the window size are in bold. The window is designated with brackets.

Target word	Nearby words	Training pairs
how	[*How* **quickly daft**] jumping zebras vex	(how, quickly) (how, daft)
quickly	[**How** *quickly* **daft jumping**] zebras vex	(quickly, how) (quickly, daft) (quickly, jumping)
daft	[**How quickly** *daft* **jumping zebras**] vex	(daft, how) (daft, quickly) (daft, jumping) (daft, zebras)
jumping	How [**quickly daft** *jumping* **zebras vex**]	(jumping, quickly) (jumping, daft) (jumping, zebras) (jumping, vex)
zebras	How quickly [**daft jumping** *zebras* **vex**]	(zebras, daft) (zebras, jumping) (zebras, vex)
vex	How quickly daft [**jumping zebras** *vex*]	(vex, jumping) (vex, zebras)

Doc2vec extends this model beyond words, so that sentences, paragraphs, and/or documents can be characterized by vector representations. In doc2vec, not only does the network create a vector representation of each word in the vocabulary, but every document (or sentence, or paragraph) is also represented by a vector. When training the network, this document can be thought of as another word, which "acts as a memory that remembers what is missing from the current context" (Le and Mikolov 2014, 3), but does so across *all documents*. Once the model is trained, the vector for the document can be used as a description of the features of the document, based on the words the document contains. I use the resulting vector representation of each bill as a characterization of the policy positions in that bill. Because these vectors are comprised of real numbers, they can easily be adapted for use in predictive regression models.

There is also the question of which bills and versions to include when training the model. In my analyses, I train a model using all bills from each Congress, as well as all text versions of each bill, simultaneously. While this would be incorrect to do with a method such as LDA, as multiple versions of the same bill might bias the topic weights within and between bills, neural networks generally perform better with as much data as possible. The only way that using multiple text versions from the same bill could bias the document vectors is if there were a unique context for a common set of words that is used in a bill that has enough versions to bias the results for other documents.

However, this is unlikely for two reasons. First, the word would also have to be common in other documents in order for the context to carry over from the original document. If this were true, the anomalous context of the word in the document would by definition be less anomalous. Second, there would have to be an incredibly large number of bill versions of a single bill with such an anomaly. There are 49,264 bills totaling 83,732 unique text versions in my sample. No bill has more than 10 text versions, and 99% of bills (48,769) have five or fewer text versions.

Following DDE, I estimate members' votes with a logistic regression model, using both a characterization of bill texts and district/member- and bill-level covariates. Specifically, I estimate the vote of each member on each bill, based on characteristics about member and their district and about the bill itself, which includes the doc2vec neural weights. The member-level predictors I use are 1) member's DW-NOMINATE scores (both first and second dimensions), 2) the member's margin of victory in the most recent general election, 3) the district's partisan voting index (PVI) value, 4) a dummy variable for whether the member is in the majority party, and 5) a dummy variable for whether the member is the bill's sponsor. The bill-level predictors I use are 1) the chamber of origin, 2) the sponsor's party, 3) the number of cosponsors, 4) dummy variables for committees to which the bill was referred, and 5) the doc2vec vector for bill each bill. Each model is estimated separately by party and Congress.

To estimate my model, I use a logistic regression estimated via stochastic gradient descent (SGD). SGD is an algorithm that can be applied to many models, which iteratively optimizes the loss function of the given model. SGD aids in classification and prediction for tasks with large amounts of data (Bottou 2010; Zhang 2004), and there are

1,075,498 votes in my sample. I assess the out-of-sample predictive accuracy with 10-fold cross-validation.

2.4 Results

2.4.1 Validating the Model

I begin by validating out-of-sample predictive accuracy of the logistic regression model I use to predict the votes of members of the House; these results can be found in Table 9. I include the rates of correct predictions, false Yeas (predicted Yea, actual Nay) and false Nays (predicted Nay, actual Yea), broken down by members of the majority party, the minority party, and in total, for all bills that went up for a final passage vote.

The out-of-sample predictive accuracy of the model is high, indicating that the model is not overfit. While the overall accuracy is high, it is not surprising that there is a substantial difference between the accuracy in predicting the votes of majority party members versus minority party members. First, bills that receive a final passage vote are much more likely to be introduced by majority members, so it is rare that a bill receiving a vote will not have majority support. Further, since more bills are authored by majority party members, and the leadership has tight control over many legislative activities, a Yea vote is often "tantamount to supporting the position of the majority party given agenda control" (de Marchi, Dorsey, and Ensley 2017, 15). Therefore we observe a higher baseline rate of Yea voting for majority members (95.0%) than minority members (73.7%).

Table 9: Out-of-sample predictive accuracy for final passage votes in the House for the 104th-111th Congresses.[53]

All Bills	Correct	False Yea	False Nay
Majority Votes	95.1%	3.1%	1.8%
Minority Votes	87.3%	5.3%	7.4%
Total	91.5%	4.1%	4.4%

Introduced by Majority Party	Correct	False Yea	False Nay
Majority Votes	95.0%	3.2%	1.8%
Minority Votes	87.4%	5.3%	7.3%
Total	91.5%	4.1%	4.4%

Introduced by Minority Party	Correct	False Yea	False Nay
Majority Votes	96.0%	2.6%	1.4%
Minority Votes	86.5%	5.2%	8.3%
Total	91.6%	3.8%	4.6%

[53] The predictive accuracy is quite consistent across Congresses. The highest predictive accuracy is found in the 108th Congress (94.4%) and the lowest rate is found in the 110th (86.0%). However, the 110th is the only Congress in which the overall predictive accuracy is below 90%. More information can be found in Appendix B.

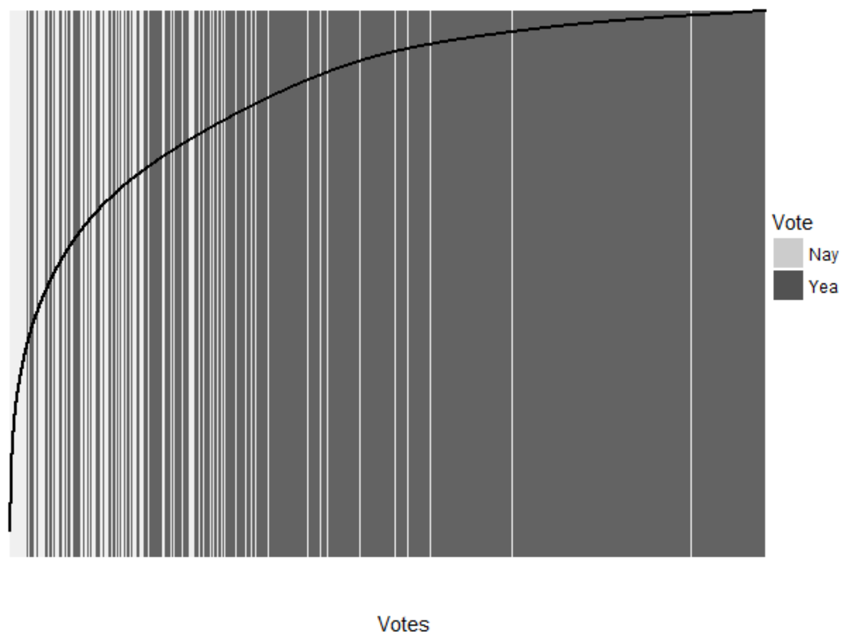

Figure 4: Separation plot of the cross-validated model. The x-axis is comprised of individual observations, ordered according to increasing predicted probability of voting Yea. Actual Nay votes are light gray vertical lines, actual Yea votes are dark gray vertical lines, and the black curve is the predicted probability of a Yea vote for each observation.

While the percent correctly predicted is a common and useful method for assessing a model's predictive power, additional metrics provide a fuller picture of how well the model is able to predict votes in Congress. Figs. 4 and 5 contain graphical representations of the model's predictive power. Figure 4 is a separation plot of the cross-validated model, which details "the models' ability to consistently match high-probability predictions to actual occurrences of the event of interest, and low-probability predictions to nonoccurrences of the event of interest" (Greenhill et al. 2011, 991). As we can see, most of the actual Nay votes are associated with lower predicted probabilities, and most of the actual Yeas are associated with higher predicted probabilities. This indicates that the model is adept at discerning between Yea and Nay votes.

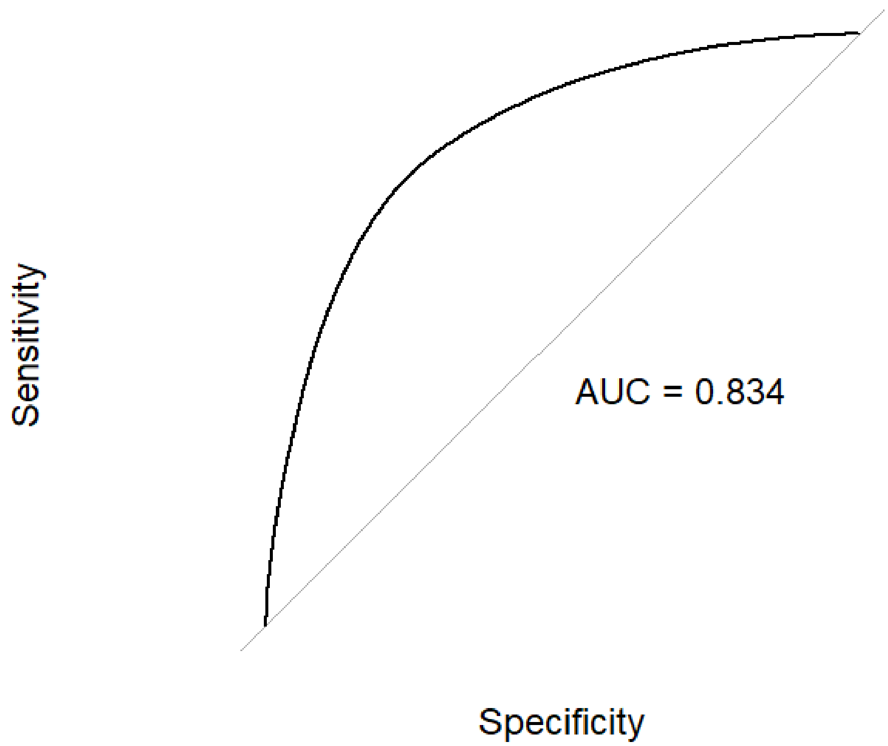

Figure 5: ROC curve plot of the cross-validated model.

Figure 5 is a plot of the receiver operating characteristic (ROC) curve for the model, which shows how well the model is able to discern between true positives (sensitivity) and false positives (one minus specificity/true negatives) at all possible thresholds between 0 and 1. I assess the ROC curves by looking at the area under the curve, or AUC. The AUC corresponds to the overall probability that the model will assign a higher predicted probability to a case that is a “success” than one that is a “failure” when these are selected at random (Berrar and Flach 2012). Models with an AUC of greater than 0.70 are thought to be reasonable predictive tests (Hastie, Tibshirani, and Friedman 2009). This model’s AUC of 0.83 suggests that it is a highly predictive test.

Together, the separation plot, ROC curve, and percent correctly predicted statistics show that the model is able to predict final passage votes in Congress based on the text of bills. Since I am interested in extending the model to predict the outcome of bills that never made it to the floor, I must also assess how well outcome of final passage votes is predicted, in addition to individual members’ decisions. The predicted versus actual bill outcomes for bills in my sample can be found in Table 10. Of the 2,457 bills in my sample, the model correctly predicts the outcome of 2,441 (99.3%) bills, including 99.7% of bills that passed, and 92.7% of bills that failed. The model only incorrectly predicted the outcome of 16 bills between 1995 and 2011, a rate of one final passage vote per year.

Table 10: Predicted versus actual bill outcomes.

	Actual Pass	Actual Fail
Predicted Pass	2,339	8
Predicted Fail	8	102
Total (Percent Correct)	2,347 (99.7%)	110 (92.7%)

2.4.2 Predicting Votes on Non-Voted Bills

Having demonstrated that this model is highly predictive of members' voting decisions and vote outcomes although it uses a different text analysis technique, more data, and a larger sample of bills than DDE, I can now apply the model output to predicting how members would have voted on bills which never came to a vote. The prediction process is a relatively simple one. Each pre-vote bill has values for all the covariates in the DDE model including contextualized text data for the last version of the bill included in the Congressional Record. I then multiply the matrix of independent variables from by the coefficients from the logistic regression described in the section above.

I can then compute three quantities of interest from these predictions: 1) whether the bill is predicted to pass, 2) what proportion of the total membership of the chamber is predicted to vote in favor of the bill, and 3) what proportion of the majority party's membership is predicted to vote in favor of the bill. With these, I can assess how the majority party protects its members and advances its own legislative agenda along each step of the legislative process.

2.4.3 Agenda Power along the Legislative Process

Table 11 contains the count of bills ending in each stage (Introduced, Referred, Reported, Voted) according to each of the six agenda control outcomes outlined in Tables 4 and 6. Two main patterns are immediately clear. First, most bills which the majority supports are successful.[54] There are 2,454 bills in my sample that came to a final passage vote, almost all of which (95.5%) passed. Of the 2430 bills that a majority of the majority party supported, 97.8% made it to the floor for a vote and 95.2% passed. Just 30 bills that passed were rolls, passed over the objections of most of the majority party. While much of the literature on agenda control has focused on negative agenda control, this pattern is also evidence of strong *positive* agenda control exercised by the majority party in the House. A majority of the majority party was only predicted to support 53 bills that never made it to a vote on the floor. Missed opportunities were even rarer, as just 16 of the 53 bills for which there is support from a majority of the majority party that never made it to a vote would have passed had they reached a final passage vote. So not only is the majority party usually successful in shepherding its preferred policies through the legislative process, they are discerning in limiting their embarrassments by not bringing bills to a vote that may fail, instead killing them before they reach the floor.

Second, the vast majority of bills that did not reach a vote would not have passed if they were considered. This is not surprising—bills that have enough support to pass are more likely to move through the various hurdles along the legislative process. Of those predicted to fail, almost all (> 99%) in each stage are blocks. Dodges, which the majority also opposes but were predicted to pass, are the second-most common outcome

[54] Though these are predictions, I use language that classifies each outcome according to the model's prediction, for clarity.

for bills which do not come up for a vote. The majority party is remarkably adept at keeping bills which most of their members oppose from ever coming to the floor. This result confirms patterns in existing research of strong negative agenda control exercised by the majority party (*e.g.*, Cox and McCubbins 2005). In the case of blocks, this is not especially remarkable. Blocks would not pass even if they did come to a vote, so the majority faces less of a challenge in keeping those bills from coming to a vote. However, keeping dodges from coming to a vote is theoretically a more difficult enterprise, as they ostensibly have the support to pass.

The results in Table 11 are also remarkably consistent across Congresses. No more than 15% of the missed opportunities, disappointments, and dodges dying in each stage happened in a single Congress. There is a similar pattern of blocks as well.

Table 11: The frequency of each agenda control type for bills dying in each pre-vote legislative stage. Missed opportunities have support from a majority of the majority, and are predicted to have passed. Disappointments have support from a majority of the majority, and are predicted to fail. Dodges do not have support from a majority of the majority party, and are predicted to pass. Blocks do not have support from a majority of the majority party, and are predicted to fail.

	Missed Opportunities	Disappointments	Dodges	Blocks
Introduced	3	1	21	393
Referred	26	11	190	43,081
Reported	0	0	16	3,344

	Successes	Disappointments	Rolls	Blocks
Voted	2,313	58	31	44

Having established evidence of both positive and negative agenda control exercised by the majority party, I can more closely examine particular types of bills. The two as yet unmeasured bill types in the pre-vote agenda control typology (Table 6) are missed opportunities and dodges. In both cases, a bill fails to reach a final passage vote even though it had the support of enough members in the House for it to pass. Figure 6 shows the percentage of bills in each stage that are predicted to pass, but did not—split by whether they are missed opportunities or dodges.

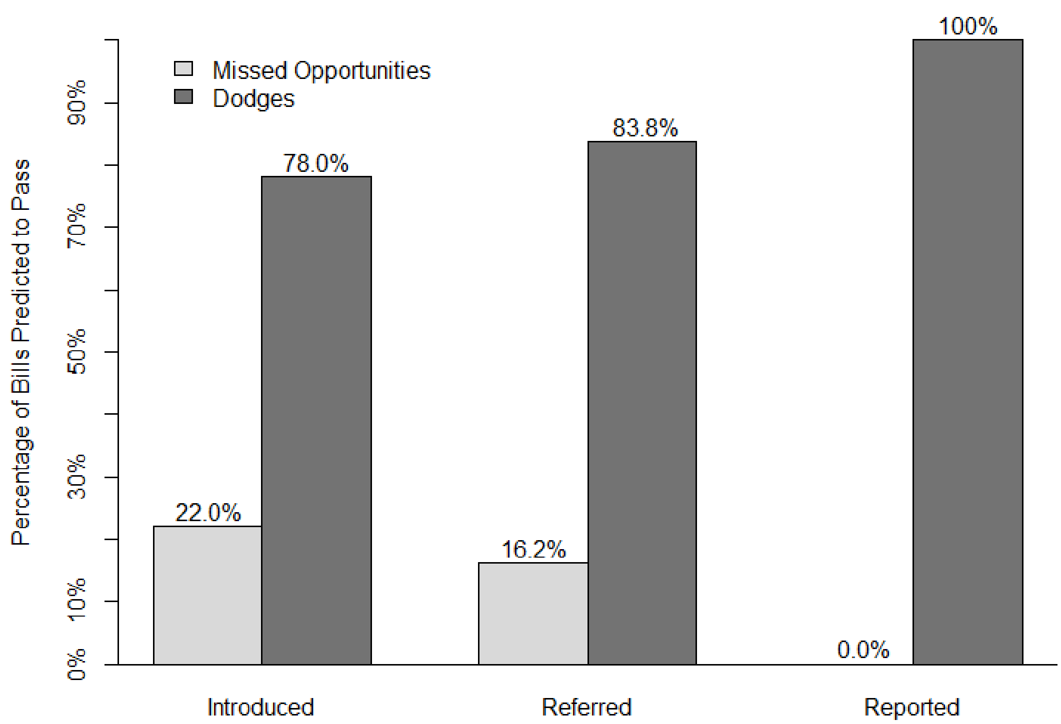

Figure 6: Percentage of bills predicted to pass that are missed opportunities or dodges, among bills dying in each pre-vote legislative stage.

Figure 6 largely confirms the expectations of negative agenda control based on the gatekeeping power and incentives of the majority power in each stage. The majority party has the most control over whether a bill proceeds while in committee, so the majority is more able to kill dangerous bills in committee than otherwise. While a larger proportion of bills with the support to pass are dodges among Reported bills, there are only 13 such bills, while 186 dodges died in committee—more than in any other stage. The bills displayed in Figure 6 also continue the trend of consistent application of agenda power over the study period. In every Congress, more bills dying in committee that were predicted to pass were dodges than missed opportunities—there were no such bills in five of the eight Congresses I study that died after a committee report.

Because the model predicts that almost all bills that do not make it to a vote on the House floor are blocks, it is worth exploring these votes in more depth. Particularly, while I have computed whether bills have support from a majority of the majority party (and a majority of the chamber), I have not considered the level of minority party support for bills, or the role of the minority party in the agenda process.

Figure 7 shows the proportion of blocks in each stage that have support from more than half of the minority party. By definition, blocks do not have support from a majority of the majority party, but most blocks also do not have support from a majority of the minority party. In no stage does the proportion of blocks supported by most of the minority party exceed 25%, and the average level of support from the minority party on bills across all three pre-vote stages is low (14.3%). This suggests that most blocks have a small coalition of support and have little chance to pass before even considering the institutional hurdles each bill must overcome. Further, these findings give credence to those from Figure 6. A higher proportion of bills dying in committee have support from most of the minority party than those dying in any other stage, and this is true in every Congress. The majority party is best able to bury bills that do not fit with its goals in committee, and the majority uses that power to keep many of the minority party's preferred policies off the floor.

However, most members of the minority party did support most bills that came to a vote (72.5%). This suggests that there is some difference between setting the agenda and influencing policy. While the minority party has little ability to determine which bills are considered, the fact that most minority party members vote Yea on nearly three-quarters of the bills coming up for a vote in the House suggests that the policies the chamber votes on largely fit with the preferences of many in both parties.

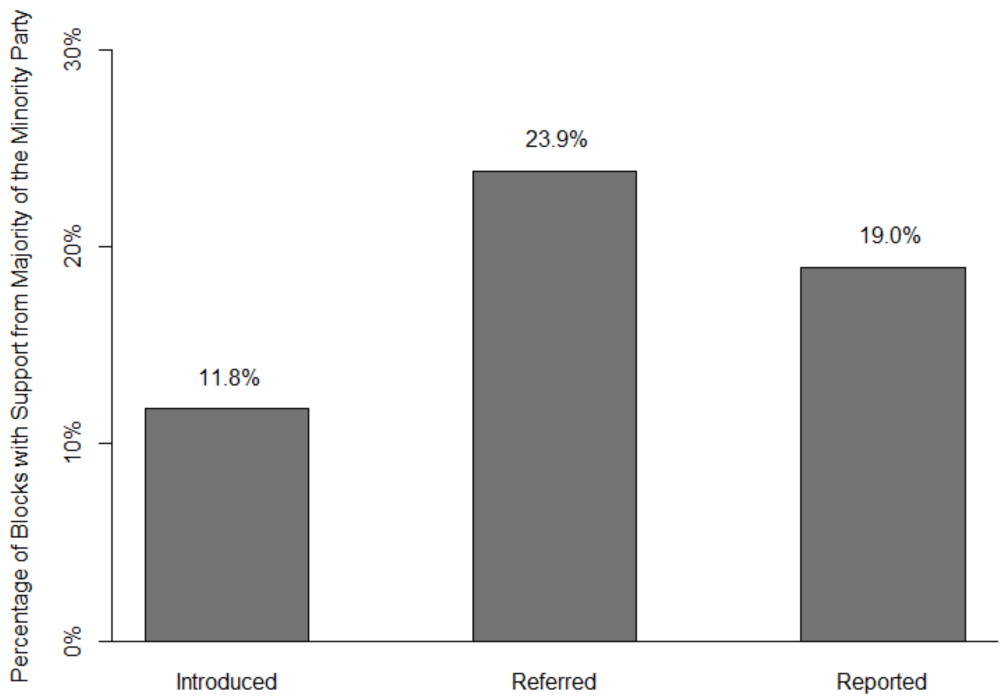

Figure 7: Percentage of blocks that have support from a majority of the minority party, for bills dying in each pre-vote legislative stage.

Another way to examine the minority party's role in the legislative process is to look at how well bills fare that are introduced by members of the minority party. In Figure 8, I show the proportion of all bills that were introduced by a member of the minority party in each Congress, according to the legislative stage in which each bill died. Because minority-introduced bills are less likely to pass, we would expect that when members of the minority party introduce bills, they are more likely to be done so in order to grandstand and/or take a position. Such bills are more likely to die earlier in the legislative process, which implies that the proportion of bills introduced by members of the minority party should decrease with each successive legislative stage. On average, this is exactly what is shown in Figure 8.

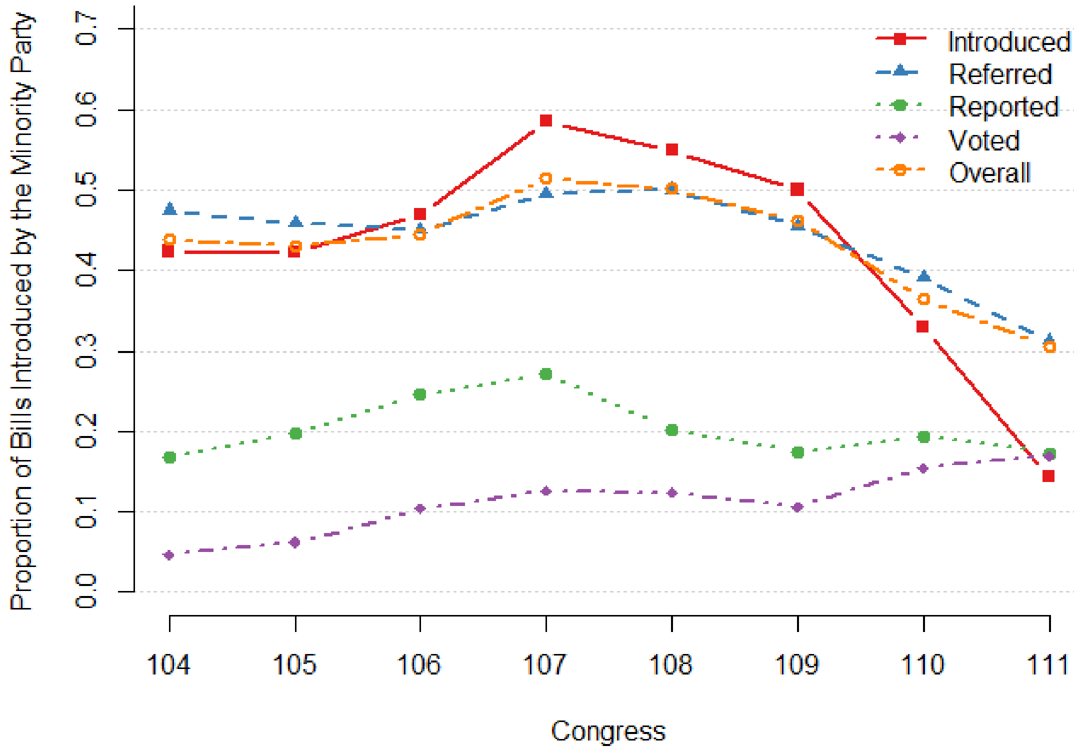

Figure 8: Percentage of bills dying in each pre-vote legislative stage that were introduced by a member of the minority party in each Congress.

As shown in Figure 8, across all the Congresses in my sample, more than 40% of those dying after introduction (46.9%) and those dying in committee (43.7%), were introduced by a member of the minority party, and compared to 20.5% of those dying after a committee report and just 11.9% of those receiving a final passage vote. Along the legislative process, the proportion of bills that were introduced by a member of the minority party steadily decreases. However, the change in these proportions over time complicates matters. The proportion of bills dying in committee (and the total proportion of bills introduced by members of the minority party) consistently decreased between the 104th and 111th Congresses. At the same time, the proportion of bills that received a final passage vote and were introduced by a member of the majority party increased from 4.7% in the 104th Congress to 17.0% in the 111th Congress. At a time when party polarization is thought to have increased and when the level of positive and negative agenda control exerted by the majority party remained relatively consistent by most measures, the proportion of bills introduced by the majority party reaching a final passage vote increased nearly four-fold.

2.5 Discussion

In this paper I have presented the first assessment of agenda power in the House both before the floor and on final passage votes. I extended a model to predict the voting decisions of members of Congress based on the content of each bill, as characterized via a neural network trained on the text of bills. I used the model to predict votes on bills which never came to a vote, to assess how the majority party in the House exercises its agenda power before bills come to a vote as well as on the floor.

Consistent with much of the research studying the House to date (*e.g.*, Jenkins and Monroe 2016, 2012; Cox and McCubbins 2005, 2007), I find evidence of substantial negative agenda control exercised by the majority party. Bills that do not have the support from a majority of the majority party are incredibly unlikely to make it to a vote on the floor, regardless of whether the bill in question has support from enough of the chamber to pass a final passage vote. However, I also find evidence that the majority party employs strong positive agenda control. The vast majority of bills which have support from a majority of the majority party do in fact come to a vote, and pass—the least common type of bill is one which has support from most of the majority party, but dies in a pre-vote legislative stage. I also find that the minority party's policy agenda is largely shut out of the legislative process. Almost all of the bills which have support from a majority of the minority party die before reaching a final passage vote. Nevertheless, most bills that do come to a vote have the support of most members of the minority party, so while the minority may have little control over the agenda, policies that are voted on are likely to be amenable to their members. Lastly, when looking to bury a bill, the majority party most often does so in the committee process than any other legislative stage. A higher proportion of bills that die in committee either have support from a majority of the minority, and/or enough support to roll the majority, than those dying in any other stage.

Using the model to predict how members would vote on bills that never reached a final passage vote is possible because the independent variables in the original model are available for all bills, regardless of whether they received a floor vote. One potential concern is that the models' predictions are not actually driven by the text of each bill and thus would not apply to bills which did not receive a floor vote. On final passage votes, the model over-predicts Yea votes—the overall predicted probability of a Yea on an actual vote is 94.9%—but the predicted probability of a Yea vote across all bills that died before reaching a final passage vote is just 9.4%. Given that the only variables in the model specific to each particular bill are 1) the bill's vector of neural weights, and 2) the number of cosponsors on the bill, much of this change can be attributed to the content of each bill.[55]

This paper continues a long tradition of studying agenda power in Congress, but is the first to provide direct measures of agenda control for bills that never reached a final passage vote in the House. As such, my approach opens many doors for future research. First and foremost, my method allows for a future comparison between the House and Senate. Most of the existing research on agenda power in Congress has focused on the House, because the mechanisms of agenda power are more easily identified and formal in the House than in the Senate. In fact, there is only one direct comparison between the levels of agenda control in the House and Senate (Gailmard and Jenkins 2007) and it only uses roll rates as the means of comparison. Comparing the predicted votes of members on pre-vote bills will allow a more direct and robust

---

[55] I also estimated models that did not include the number of cosponsors to each bills, and the predictive results are not different to the point of statistical significance. This suggests that members' voting decisions do in fact respond to the content of bills. These results can be found in Appendix B.

comparison between the levels of negative *and* positive agenda control between the House and Senate.

Of course, the model does not perfectly predict all votes and bill outcomes. Future research would do well to look into these particular votes and bills to glean why the model predicted incorrectly. Doing so would serve two purposes. The first purpose would be to tune the model and improve predictions in the future. The second purpose would be to assess whether the bills and votes which the model incorrectly predicts are interesting in their own right. In a similar vein, an in-depth examination of the history and outcomes of non-blocks (missed opportunities, disappointments, dodges) at each stage would likely improve our understanding of Congressional politics. Regardless, this research opens the door to many lines of inquiry to further our understanding of party politics and the legislative process in Congress.

More broadly, we should rethink the relationship between party power, agenda control, and party polarization. During the period I study, parties in Congress are thought to have become more polarized (Theriault 2008; Grossmann and Hopkins 2016), and this increase in polarization is often associated with an increase in the use of party power, particularly from the majority party (Rohde 1991; Aldrich 2011; Aldrich and Rohde 2000). In many ways, I find that the majority party completely controls the agenda. The majority party's preferred bills are likely to come to a vote, and the minority party's preferred bills are unlikely to come to a vote. Bills are most likely to die in committee, exactly what we would expect in a system where the majority party has strong agenda control. The minority party is effectively shut out of determining which bills are considered by the chamber and are voted on. And yet, during this time period, the proportion of bills receiving votes that were introduced by a member of the minority party increased, and most bills that were voted on had support from a majority of the minority in every Congress I study. This suggests a disconnect between the majority party's power to influence the agenda and the majority party's power to influence the policy content of the bills that the House passes. In addition, this evidence suggests that there may be more compromise and agreement on policy between the two parties than many Congressional scholars have thought, which does fit with trends in recent research (*e.g.*, Curry and Lee 2018). In light of this evidence, we need to rethink the relationship between polarization and party power. We also need to think about the differences between controlling the agenda and controlling policy content. Find when the majority party can and cannot reliably exert their will over the minority in the House will greatly aid our understanding of Congressional politics.

3. Bill Content and Legislative Success in the U.S. House of Representatives

3.1 Introduction

Because the votes of members of Congress on final passage bills ultimately determine federal law, there has been no shortage of interest in members' voting decisions (*e.g.*, Enelow and Koehler 1980; Cox and McCubbins 1991; Kingdon 1981; Poole and Rosenthal 1985; McCarty, Poole, and Rosenthal 2016; Krehbiel and Rivers 1990; Coates and Munger 1995; Aldrich 2011; Rohde 1991; de Marchi, Dorsey, and Ensley 2017). Research suggests that members' votes on bills are determined by how their preferences interact with the policy content of the bill and the political context of the vote (*e.g.*, Crespin and Rohde 2010). Members usually know how they will vote on a bill well in advance, and they base these decisions on what is actually in the bill (de Marchi, Dorsey, and Ensley 2017; Gerrish and Blei 2012).[56] But how well does this relationship between bill content and members' decisions aggregate? That is, how well can we predict a bill's success in the House of Representatives, and how large a role does the bill's content play in its predictability?

Using machine learning methods to produce a semantic understanding of the policy content in a bill, I build a statistical model to predict how successful bills will be at navigating the legislative process in the House. First, I create a representation of each

[56] Specifically, by the interaction between the bill's policy content and the preferences of various actors (the member, their party, their constituents, their donors, etc.).

bill before the House. Then, I use the resulting output to predict the legislative success of each bill—whether it dies in committee, whether it is reported by a committee, and whether it passes a final passage vote. I find that where a bill finishes in the House is highly predictable. Furthermore, when considering only the bills that make it out of committee, the models can discern with great accuracy which bills will ultimately pass. Inclusion of bill content accounts for ~20% of the improvement in my models' predictive performance over the baseline rate where all bills are predicted to die in committee—by far the most common fate. However, which text version of a bill I use to predict where the bill will end up does not seem to have an effect.

3.2 Background

Scholars have studied many aspects of bills in Congress. Much of the focus has been on members' voting behavior, where measures like roll call votes have helped produce a deluge of scholarly output (*e.g.*, Poole 2007; Iyengar and Kinger 2010; Cox and McCubbins 2005; McCarty, Poole and Rosenthal 2016; Epstein and O'Halloran 1999; Jones and Baumgartner 2005; Ansolabehere et al. 2003; Bartels 2000, 2016; Ansolabehere, Snyder, and Stewart 2001). It is often argued that each member's ideological stance can be recovered from how they vote on bills and procedural measures, most notably resulting in Poole and Rosenthal's NOMINATE scores (Poole and Rosenthal 1985; McCarty, Poole, and Rosenthal 2006; Lewis et al. 2017).

Researchers have also noted the importance of pre-floor behaviors in policy-making. While final passage votes ultimately determine which bills pass, most of the actual policy-making occurs long before the final vote. By the time members vote on a bill, its content has been set through extended deliberations in committee and on the floor. Members can have their greatest impact on policy in the early stages of policymaking, rather than by voting on final passage, by, "gathering information, drafting bills, building coalitions, and keeping pace with the actions of various interests" (Woon 2009, 29). Indeed, those who participate in the committee system tend to have more influence on policy and the workings of Congress (Hall 1996), so members have every incentive to focus their attention on crafting policy rather than mulling over voting decisions.

Following this logic, there have been a number of studies of members' non-voting behaviors in Congress, particularly those behaviors before a bill reaches the floor (*e.g.*, Matthews 1960; Hall 1987, 1996; Schiller 1995; Wawro 2000; Woon 2009). Scholars have investigated such pre-voting phenomena as bill (co)sponsorship—in terms of issue attention (Woon 2009), legislative influence (Fowler 2006), how sponsorship signals support for a bill (Wilson and Young 1997), and the purpose of cosponsorship for individual members (Kessler and Krehbiel 1996)—party rewards for entrepreneurial efforts (Wawro 2000), and the (sub)committee system in Congress (*e.g.*, Fenno 1978; Hall 1987, 1996; Schiller 1995; Groseclose and King 2001).

A common thread through these studies is that they seek to understand how members of Congress interact with bills. More importantly, these studies suggest that not only do members create and shape the policy positions in bills (*e.g.*, Fenno 1973; Cox and McCubbins 2007), but that members also react to and choose their actions based on the content in bills (*e.g.*, de Marchi, Dorsey, and Ensley 2017; Gerrish and Blei 2012). Different legislators with different preferences and roles within the chamber and party have power over a bill at various points throughout the legislative process. A bill's success, then, is a function of its content, and of the preferences of gatekeepers that have power over its progress.

Throughout the process of passing legislation, there are many points where members' preferences can be expressed. Consider: a member may want to enact legislation that imposes stricter regulations for emissions in automobiles, so she puts a bill in the hopper. The bill deals with commerce and the environment, so it would likely be referred to the House Committee on Energy and Commerce, and perhaps also to the House Committee on Natural Resources. The preferences of the members of the committee, likely in concert with the (majority) party leadership, will then lead to a marked up bill, and determine the content of those changes. This process continues until the bill either dies, or passes a final passage vote and continues to the Senate or the President, if the Senate had already passed the bill. At each stage, the preferences of the actors who have control over whether the bill progresses, plus the rules at each stage of the legislative process, determine whether a bill will continue its journey through the House. In my analyses, I assume that it is the interaction of the preferences of members with the rules of the House lead to choices about bills in specific contexts.

This assumption leads to three hypotheses for predicting where bills will end up in the House. The first is, simply, that the content of bills will matter. If the preferences of members on bills determine what happens to a bill, then knowing something about the content of bills, which is thought to impact preferences, will make it easier to predict what is decided about the bill. This means that models including a characterization of bills' policy content should be more accurate in predicting where bills end up in the House than models that do not include this information.

Second, assuming that members' preferences over bills leads to decisions about those bills implies that the version of a bill used to characterize the bill's policy content should matter. In particular, using a text version of a bill that was entered during a later stage of the legislative process should better predict where the bill ends up in the House. As a bill moves through the legislative process, it often goes through a number of changes to its policy content in committee and on the floor. These changes would make it more difficult to predict outcomes based on an earlier version of the bill. In terms of predictive accuracy, this hypothesis means I expect that models including a characterization of bills coming from more recent text versions will more accurately predict where bills end up than models including a characterization of bills from earlier text versions.

Third, if members' preferences lead to decisions made on bills, then decisions should be more predictable when the preferences of the gatekeepers at a particular stage are more homogeneous. Because the majority has strong control of the committee process (*e.g.*, Cox and McCubbins 2005; Jenkins and Monroe 2012, 2016; Gailmard and Jenkins 2017; Chapter 2 of this dissertation), the opinions of the majority party are most important for predicting whether a bill will make it out of committee, relative to other stages. For contrast, it is difficult for the majority to pass bills without minority party support even under the most favorable circumstances for the majority; it is rare that a bill passes over the opposition of the minority (Curry and Lee 2018; Chapter 2 of this dissertation). While the concerns of the majority still usually loom larger, broad support is more important at the voting stage than at any other. Because the minority party has more gatekeeping power on the floor relative to other stages, the preferences of the gatekeepers are least homogeneous when voting on a bill, relative to the committee stage. These differences between the overall preferences of gatekeepers at the committee and voting stages, coupled with the assumption that members' preferences over bills lead to decisions on those bills, implies that I should be better able to predict whether bills will make it out of committee than I am able to predict what happens to bills after they leave committee.

3.3 Methods

3.3.1 Data

To test my hypotheses, I predict where bills end their journey in the House. For this purpose I use data on congressional bills, including institutional covariates like data on committee referrals[57] combined with information about bill sponsors—such as their committee positions—and all versions of the full text of each bill before Congress,[58] for the 104th-111th Congresses.[59]

3.3.2 Modeling a Bill's Journey through the House

The first step to modeling the journey of bills through the House is to create a numeric representation of each bill that can be used in predictive statistical models. For this task, I use doc2vec, a machine learning method for creating vector representations of complex text corpora (Mikolov et al. 2013a, 2013b). For a detailed description of doc2vec, see Chapter 2 and Appendix D of this dissertation.

In my analyses, I predict which stage along the legislative process (referred, reported, or passed) each bill reaches. Independent variables are the doc2vec vector for each bill, and both bill-level and Congress-level predictors for the 104th-111th Congresses. The bill-level predictors are 1) a dummy variable for whether the bill had already passed

---

[57] Data on Congressional bills come from the Congressional Bills Project (Adler and Wilkerson 1995-2015) and PIPC (Crespin and Rohde 2018).
[58] Bill text was scraped from the U.S. Government Publishing Office (GPO).
[59] These Congresses correspond to the years 1995-2011. GPO only keeps data in an easily accessibly format back to 1995, and other sources are not complete for the 112th Congress.

the Senate, 2) a dummy variable for whether the bill's sponsor is a member of the majority party, 3) a dummy variables for whether the bill's sponsor is a member of any committee to which the bill was referred, 4) a dummy variable for whether the bill's sponsor is either the ranking member or committee chair of any committee to which the bill was referred 5) the number of cosponsors, and 6) dummy variables for each committee to which the bill was referred. Each of the bill-level predictors is a proxy for support of an important gatekeeping power, be it the majority party, the entire chamber or Congress, or committee leaders. The Congress-level predictor is a dummy variable for whether the House majority party is also the Senate majority party.

While it might be assumed that the legislative stages are necessarily sequential—as they would be under regular order—there has been a documented decline in regular order in recent decades (Curry 2017; Curry and Lee 2018). To deal with this concern, I estimated multinomial logistic regression models, rather than ordinal. Both methods would enable me to assess how well I can predict, from its content and characteristics about the bill, where a bill will end up in the House, but the multinomial model does not necessitate that the stages are entirely sequential. Of the three stages I analyze, I am most interested in whether a given bill will pass the House, because those bills are one step closer to becoming law. As such, I also estimate models where the dependent variable is binary, to see if it is possible to predict whether a bill will pass a final passage vote based on its content and context.

There is a question of which text versions should be used to predict a bill's outcome. I use two sets of text versions, in each case analyzing a single text version per bill: 1) the earliest text version of each bill (its introduction version), and 3) the last version of each bill. By last version, I mean the version entered into the Congressional Record *before* each bill reached the stage at which it died. For bills dying in committee, this means the introduction version. For bills that die after being reported out of committee, this is the version which was referred to committee. For bills that pass, this is the version that was reported out of committee. This approach allows me to test the proposition that changes in bill content which presumably reflect member preferences are relevant for predicting where a bill ends its journey through the House.

To maximize predictive accuracy, I also estimate models separately by Congress and whether each bill was introduced by a member of the majority or minority party, as doing so has been shown to increase out-of-sample accuracy (de Marchi, Dorsey, and Ensley 2017).[60] Lastly, I estimate each model with and without the neural weights characterizing each bill's content as predictors, to get a better idea of how much bill content contributes to the model's predictive accuracy.[61]

---

[60] However, note that I trained the neural network on all bills and text versions simultaneously. There may be differences in the context of how specific words or phrases are used between Congresses, but the increased predictive accuracy from including more documents when training the network has been shown to outweigh concerns of changing contexts (Gurciullo and Mikhaylov 2017).

[61] The models that do not include neural weights are identical for each set of bill texts.

To estimate my models, I use multinomial and simple logistic regression with a ridge penalty (Hoerl and Kennard 1970). Ridge regression adds a term to the parameter optimization, minimizing both the sum of squared residuals and the sum of the squared coefficients. Ridge regression, like other penalized/regularized regression methods such as lasso regression (Tibshirani 1996), can help in model selection to balance the bias-variance trade-off. However, unlike the lasso, the coefficients in a ridge regression model trend toward 0 much more slowly; in practice, predictors are not removed entirely from the model. This approach is useful in situations where the researcher has many predictor variables and wants to reduce variance in the coefficient estimates (though increasing bias), but also believes that each predictor has some value to the model. Because there is no way to map specific elements from the embedded vectors produced by doc2vec onto specific parts of each document—and therefore also no way to say that the same vector element has a similar meaning to multiple documents—the large number of elements in each vector created by doc2vec make regularized regression necessary. Because each element must be considered by the model in order to capture the full content of each bill, I have chosen to use ridge regression, performed with the glmnet package in R (Friedman, Hastie, Simon, and Tibshirani 2017). For each model, the results presented are based on 10-fold cross validation.

3.4 Results

I will first assess the out-of-sample predictive results for the multinomial and binomial models in turn. Then I will assess the effect of bill content and text sample used on the predictive accuracy of my models.

3.4.1 Predicting Where Bills Finish in the House

The first question to address is how well the best model (including neural weights and using the last text version of each bill) can predict in which stage—referred, reported, or passed—a bill finishes. The out-of-sample predictive results from this multinomial model are in Table 12. For each stage, I compare the model's predictive accuracy to a baseline where all bills are classified as dying in committee ("Referred") because more bills die in that stage (83.8%) than any other. Consequently, the baseline percent correct for Reported and Passed bills is one minus the proportion of bills ending in each stage (*i.e.*, 13.2% of bills pass, so the baseline percent correct for Passed bills is 1 – 0.132 = 0.868 or 86.8%.)

Table 12: Out-of-sample predictive results for the multinomial model. The Baseline Percent Correct is the percentage of cases correctly predicted if all observations were classified as Referred. The Percent Correct is the percentage of observations that the model correctly predicts. The Change from Baseline is the percent correct minus the baseline percent correct. False Positives and False Negatives are classifications of the types of errors produced by the model. [62]

	Frequency	*Baseline Percent Correct*	*Percent Correct*	*Change from Baseline*	*False Positives*	*False Negatives*
Referred	41,278	83.8%	98.6%	14.8	0.3%	1.1%
Reported	1,488	97.0%	14.0%	-83.0	6.4%	79.6%
Passed	6,487	86.8%	54.0%	-32.8	21.7%	24.3%
Total	49,253	83.8%	90.2%	6.4	3.3%	6.5%

There are a few trends especially worth noting in Table 12. First, the model correctly predicts the ending stage of 90.2% of bills out of sample, an increase of 6.4 percentage points over the baseline. Second, all of these gains come from increased accuracy in predicting when bills die in committee, where the model correctly predicts 98.6% of cases, an increase of 14.8 percentage points over the baseline. Admittedly, while the multinomial model is incredibly adept at separating out which bills die in committee, the percent correctly predicted is low for Passed and Reported bills. The model only correctly categorizes 14% of bills which die after a committee report and 54% of bills which pass, decreases of 83 and 32.8 percentage points from the baseline,

[62] To calculate the Total row I include the neural weights and use the last version of each bill's text. As I show below, these are the conditions in which the predictive accuracy is highest.

respectively. These trends are also quite consistent between congresses, as we can see from Figure 9. Even though the model does not do a good job accurately predicting Reported or Passed bills, enough bills die in committee that the overall predictive rate, relative to the baseline, is positive in each Congress.

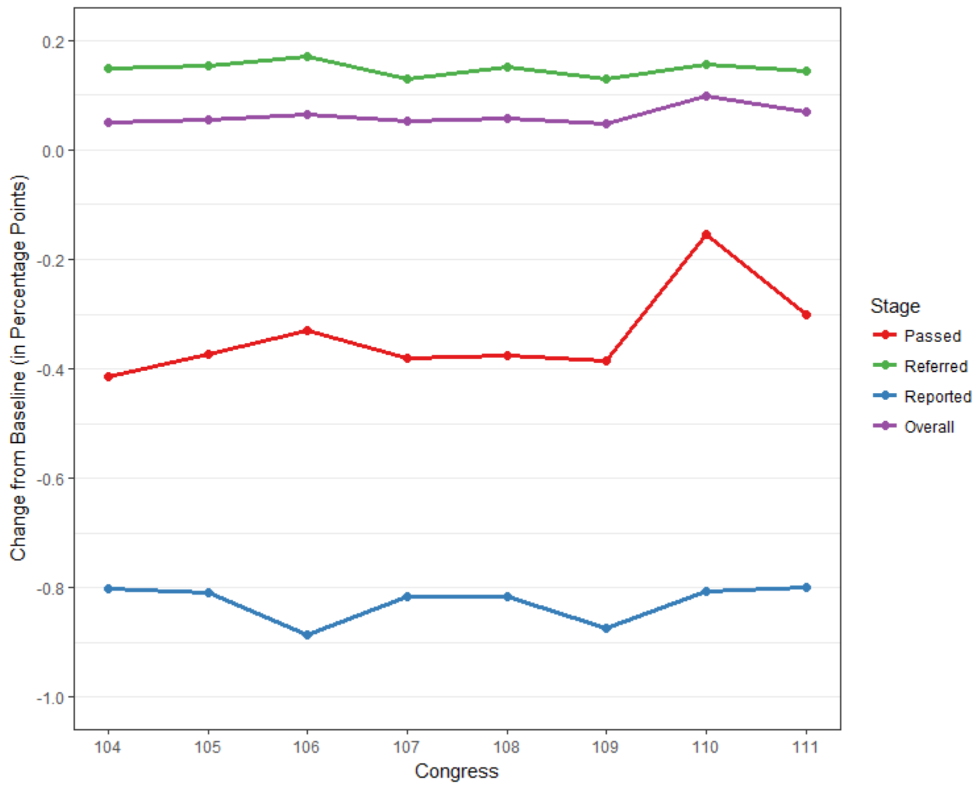

Figure 9: Changes in the multinomial model's predictive accuracy relative to the baseline for each stage over time.

Lastly, most of the model's errors are false negatives; the model often classifies 'Reported' or 'Passed' bills as 'Referred' (at rates of 79.6% and 24.3%, respectively). No stage has more false positives than false negatives, and there are nearly twice as many false negatives than false positives overall (6.5% versus 3.3%). Further evidence of this trend can be found in Figure 10, which displays separation plots for bills ending in each stage. While the model does over-predict bills as Referred, the separation plot results indicate we should be more encouraged than Table 12 alone would lead us to be. All Reported bills are concentrated among the top (right) quartile of predicted probabilities for whether a bill died after being Reported, and more than 90% of Passed bills are in the top quintile of predicted probabilities for whether a bill Passed in the House. This pattern suggests that the predictive accuracy for bills in these categories may be sensitive to the categorization threshold from Table 12, which is one weakness of using percent correctly predicted in assessing predictive accuracy. To assess whether different thresholds would lead to better results, I show plots of the ROC curve for Passed and Reported bills according to the multinomial model in Figure 11.

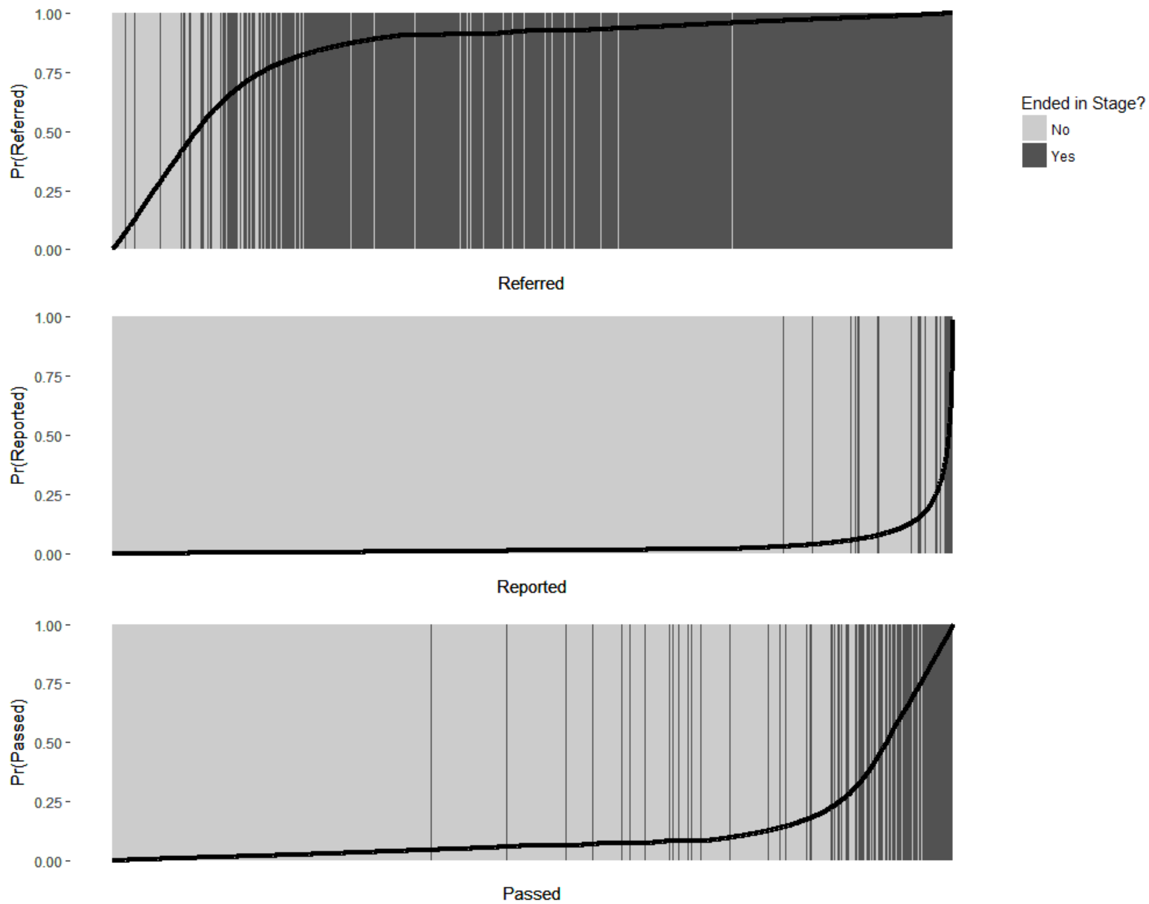

Figure 10: Separation plots of the multinomial model's predictions for each stage. The x-axis is comprised of individual observations, ordered by increasing predicted probability of being categorized according to each stage. Light gray vertical lines are bills that did not end up in each stage, dark gray vertical lines are bills that did end up in each stage, and black curves are the predicted probability of being categorized as each stage for each observation.

As the ROC curve plots in Figure 11 show, the model predicts whether bills will end up at each of the three stages when a wide range of possible predictive thresholds are taken into account.[63] The AUC for each ROC curve is above .925, indicating the model is an excellent test.[64] The AUC for Referred bills is even slightly lower than that of Reported or Passed bills.

The evidence I have presented so far suggests that the model is better at predicting when a bill will die in committee, versus what happens *after* a bill leaves the committee stage. Even then, the over-prediction of bills as dying in committee leads to high rates of false negatives for Reported and Passed bills. This is evidence in favor of my hypotheses that, due to more homogeneity in the preferences of gatekeepers at the committee stage, it will be easier to predict whether bills make it out of committee than what happens to a bill once it leaves committee.

[63] For each, I used the pROC package in R to generate and plot the ROC curves. For each of the ROC curves in Figure 11, the algorithm predicted the accuracy of the model to predict each stage at 42,033 different thresholds.

[64] ROC curves are most commonly applied to binomial tests, but it is possible to approximate a multiclass ROC via pairwise analysis (Landgrebe and Duin 2007). So to generate the ROC curve plots in Figure 11, I compared the predicted probability of each category against those of the other two categories together.

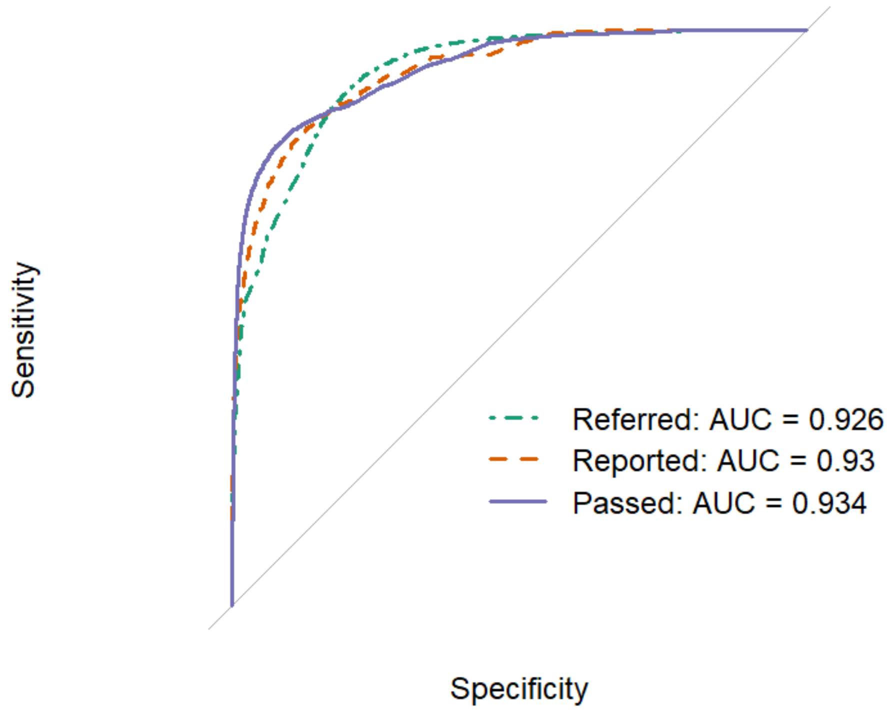

Figure 11: ROC curve plots with AUC for each stage in the multinomial model.

3.4.2 Predicting Whether Bills Pass

While the multinomial model's predictive accuracy for whether bills pass is supported by the ROC curve and separation plots, the low percent correctly predicted rate relative to the baseline is still troubling. Furthermore, distribution of incorrect predictions in the multinomial model makes it difficult to know which bills are likely to pass. If a bill is predicted to fail (to die in committee or after a committee report), there is just a 75.7% chance that the bill actually failed. If a bill is predicted to pass, there is just a 78.3% chance that the bill actually passed. This is far worse than the rate of correct positive predictions for Referred (99.7%) or Reported (93.6%) bills. Some of this finding may be an artifact of the estimation strategy, as it can be difficult to predict low-probability events with a multinomial framework, such as Reported or Passed bills. To further investigate how well I can predict the outcome of bills, I turn to the binary logistic regression model. As described above, this model uses the same input variables and estimation algorithm as the multinomial model, but with a binomial version of the multinomial dependent variable in which successes are if a bill passes in the House.[65] The multinomial model performs well at predicting which bills will die in committee, so I estimate the binomial model both on the entire sample of bills and on just those bills which made it out of committee.

The out-of-sample predictive results for the two binomial models can be found in Table 13. Each is an improvement over the baseline rate of predicting all bills not to pass (in the case of the model using all bills) or predicting all bills to pass (in the case of the model only using bills which make it out of committee). The model with all bills shows a larger improvement over the baseline (5.7 versus 2.6 percentage points), but similar to the multinomial model, this is likely driven by the high frequency of bills which died in committee. Consistent with this notion is the fact that, like the multinomial model, almost all the errors in the binomial model with all bills are false negatives—cases where a bill that actually passed is predicted to fail. In fact, with this model, fewer than half of bills that actually passed are predicted to pass (3,206 / 6,487 = 49.4%).

[65] All Passed bills remain coded "Passed" but bills ending in all other stages are recoded to "Not Passed".

Table 13: Out-of-sample predictive results for the binomial models.

	Bills	*Bills Passed*	*Baseline*	*Percent Correct*	*Improvement over Baseline*	*False Positives*	*False Negatives*
All	49,253	6,487	86.8%	92.5%	5.7	0.9%	6.6%
After Cmte	7,975	6,487	81.3%	83.9%	2.6	15.1%	1.0%

The second binomial model, considering only those bills which made it out of committee, does a much better job correctly categorizing bills that actually passed. The model correctly predicted 98.7% of bills that actually passed, and, with a rate of false negatives of just 1.0%, it is the only model that has a higher rate of false positives than false negatives. However, this second model predicted a lower percent correct (83.9% versus 92.5%) and had a smaller improvement over the baseline than the model with all bills. True, the false positive rate of 15.1% is much higher than in the binomial model with all bills (0.9%), but it is still lower than the rate of false positives for Passed bills in the multinomial model (21.7%).

As with the multinomial model, the percent correctly predicted does not give a full picture of how well the binomial models perform. The separation plots in Figure 12 suggest that the model with all bills was better at separating bills that pass from bills that do not than the model only considering bills after committee. 90% of successes in the all bills model are above the 86th quantile of the distribution of predicted probabilities, whereas 90% of the failures in the model with bills after committee are below the 31th quantile of the distribution of predicted probabilities. In other words, the delineation between the predicted passes and failures is about half as clear in the model with bills after committee as it is in the model including all bills. The model with all bills is also better at identifying which bills pass according to the ROC curves plotted in Figure 13. While each model has a high AUC, the AUC for the model with all bills is .043 higher, indicating slightly better predictive performance.

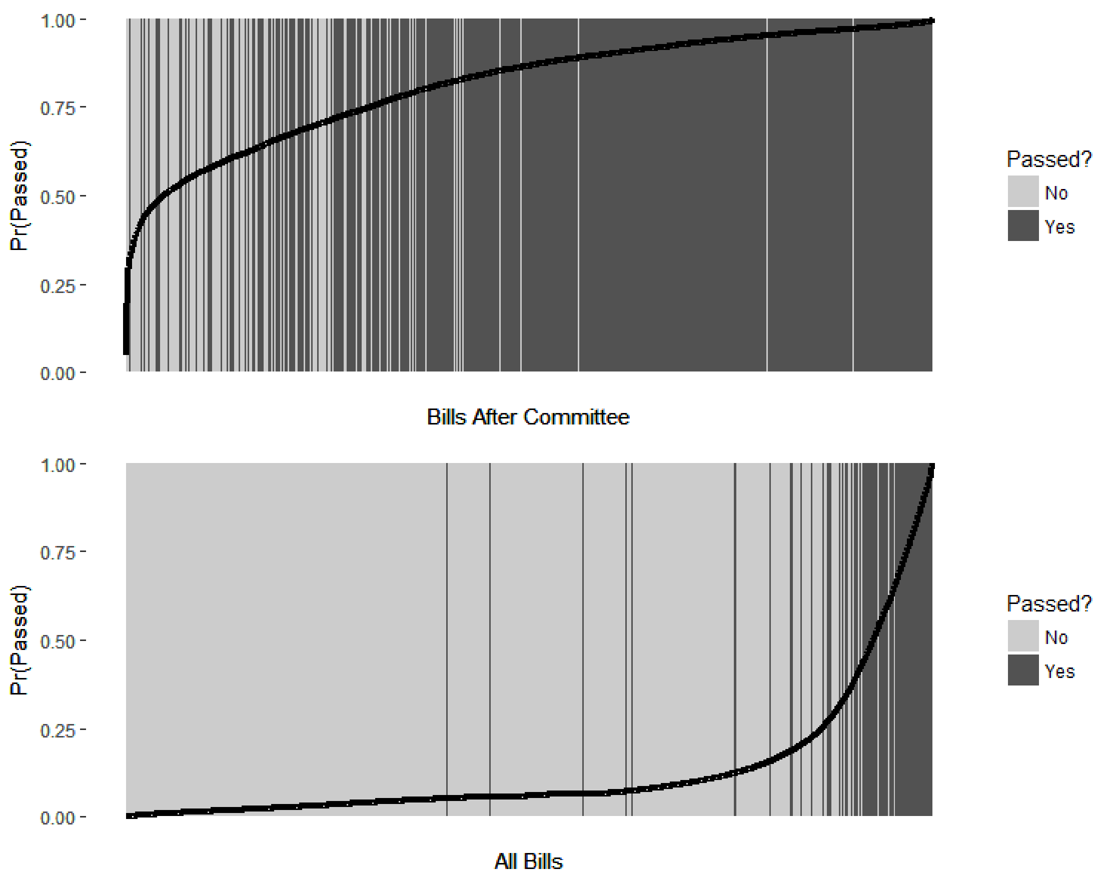

Figure 12: Separation plots of each binomial model. The y-axis is the probability of passage. The x-axis is observations, ordered by increasing predicted probability of passage. The black lines are the predicted probability of passage for each observation.

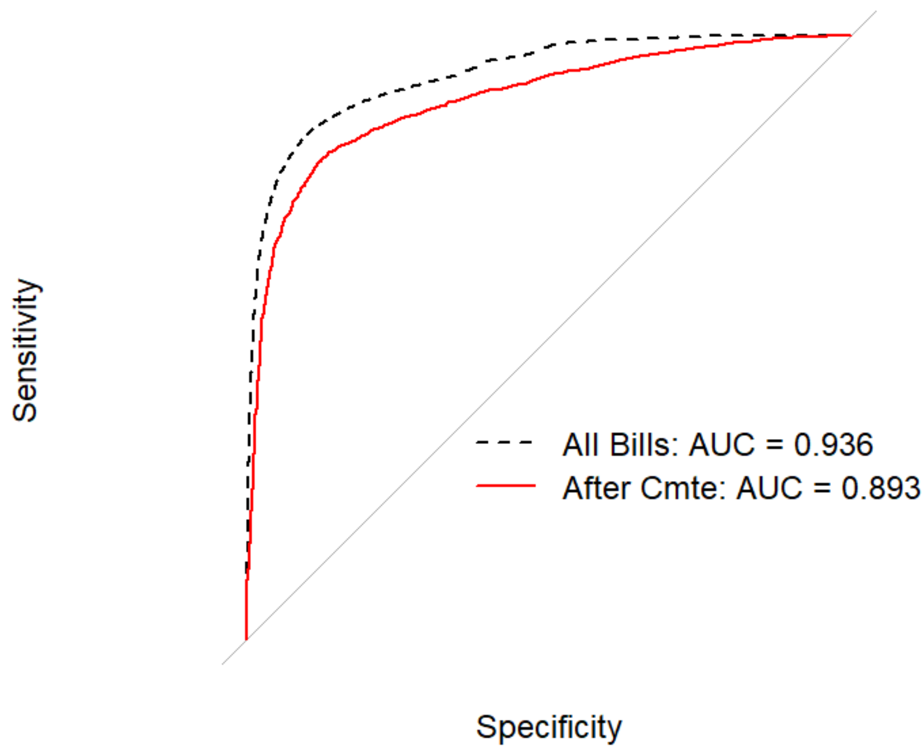

Figure 13: ROC curve plots for each binomial model.

3.4.3 The Effect of Bill Sample and Neural Weights

The full models I consider, both the multinomial and the binomial, include the neural weights for each bill as predictors and use the last text version of each bill. To test my hypotheses for how these choices affect the predictive power of my models, I ran two other models, one where I do not include the neural weights as predictors and another where I use the neural weights for the earliest text version of each bill, rather than the last version.[66] In Table 14, I display the predictive results for the full binomial model using all bills and these two adjusted models.

Table 14: Predictive results for the binomial model using 1) all variables, 2) no neural weights, and 3) the earliest text version of each bill.[67]

	Bills Passed	*Baseline Percent Correct*	*Percent Correct*	*Improvement over Baseline*	*False Positives*	*False Negatives*
1) Full Model	6,487	86.8%	92.5%	5.7	0.9%	6.6%
2) No Weights	6,487	86.8%	91.5%	4.7	1.1%	7.4%
3) Earliest Version	6,487	86.8%	92.4%	5.6	0.8%	6.8%

In terms of percent correctly predicted, there is a small effect of using the neural weights. The full model has a predictive accuracy rate that is 1.0 percentage points greater than the model without neural weights, which is 21.3% of the increase in predictive accuracy of the full model over the baseline accuracy rate. There is even less of an effect of which text version is used in prediction. The increase in predictive accuracy for using the last version is 0.1 percentage points over using the earliest version, a 1.7% increase in improvement over the baseline accuracy rate. For all models, the distribution of error types is similar. The range of false positive rates is 0.3 percentage points and the range of false negative rates is 0.8 percentage points

[66] I cannot run a model with no neural weights and the earliest version of each bill because the choice of which text version to use implies using the neural weights as predictors.
[67] I also computed these results for the multinomial model and found no substantial differences. These results can be found in Appendix C.

To assess the overall predictive accuracy of the models, I again use ROC and separation plots to compare the full model to those that use the earliest text version of each bill or do not include neural weights. Figure 14 contains plots of the ROC curves for each of the models in Table 14, with each of the alternate models are compared to the full model, and it shows there is little difference in the predictive accuracy of any of the models. The full model has the highest AUC (0.936) and it is 0.013 higher than the lowest AUC—the model without the neural weights at 0.923—a similar difference as that between the percent correctly predicted of the two models.

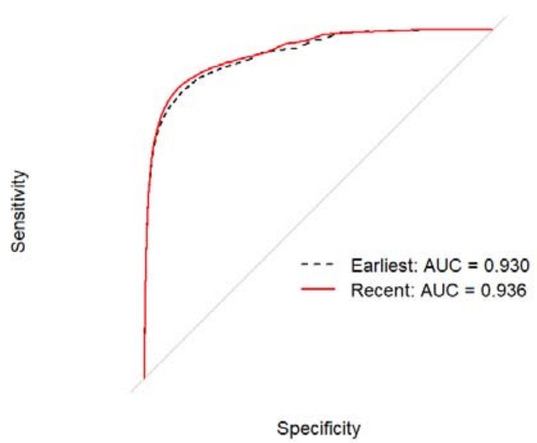

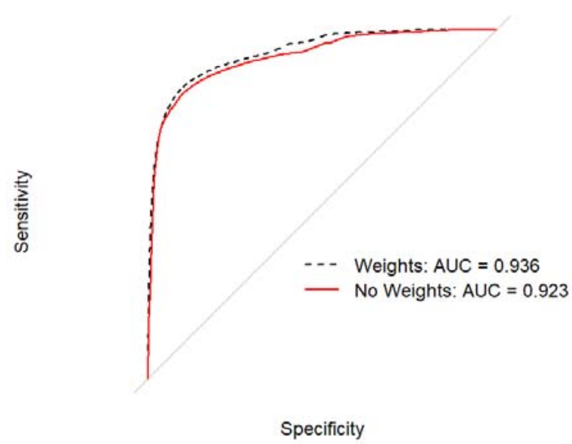

Figure 14: ROC curve plots of binomial models, varying whether the models included neural weights or used the earliest text version of each bill.

Figure 15 contains separation plots of each of the three models in Table 14. While the models each separated bills that actually pass from those that fail to a high degree, bills that actually pass are more tightly grouped toward the top of the distribution of predicted probabilities for the full model than for either the model without the neural weights or the model that uses the earliest text version of each bill. These results support one of my hypotheses—that including bill content will increase the predictive ability of the models, but does not support the hypothesis that using more recent text versions would also increase predictive accuracy.

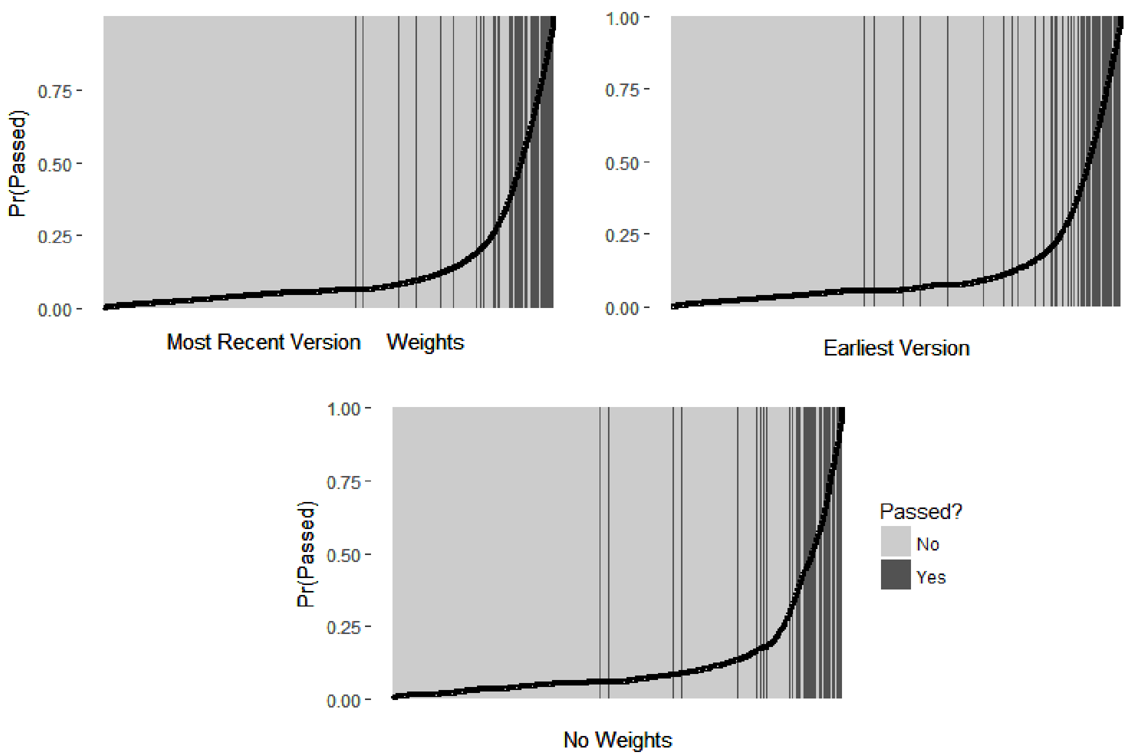

Figure 15: Separation plots of binomial models, varying whether the models included neural weights or used the earliest text version of each bill. The y-axis is the probability of passage. The x-axis is observations, ordered by increasing predicted probability of passage. The black lines are the predicted probability of passage for each observation.

3.5 Discussion

In this paper I interrogated whether it is possible to determine the fate of a bill in the House—whether it dies in committee, after being reported out of committee, or passes a final passage vote. When predicting between these three categories, I was able to correctly predict where 90.2% of bills end up in Congress (out of sample). Separation plots and ROC plots also show that the model is quite adept at predicting where a bill will end its journey in the House. For instance, the AUC of a ROC curve for each category is above 0.92. However, the model over-predicts bills as dying in committee, which means that the model is actually relatively inaccurate at predicting whether bills get a committee report but do not pass a vote, or if they pass a final passage vote in the House.

Because we most want to know whether a bill will pass, or not, I also fit binomial models for each case, considering both all bills and only the sample of bills which have already made it out of committee. In terms of percent correctly predicted, the model using only bills that made it out of committee tends to overestimate which bills will pass, perhaps suggesting that none of my models is adept at discerning between Passed and Reported bills. In this way, the models are exceptional at not missing bills which bills. Like the multinomial model, though, each of the binomial models performs better when considering the ROC plot. On the whole, I was able to predict where a bill would end up in the House with a high degree of accuracy.

In addition to testing how well one could predict the outcome of bills in the House, I was also interested in how much of an effect the *content* of a bill had on the predictability of its journey through the House. I found that the content of bills (measured by the neural weights generated with doc2vec) do not account for a large portion of the predictive accuracy of my models to predict where bills end up in the House. A model predicting whether a bill passes that contains neural weights as predictors does 1.0 percentage points better than a model without neural weights, in terms of out-of-sample predictive accuracy. However, the baseline rate for predicting whether a bill will pass is high (86.8%); since the full model including the neural weights as predictors correctly predicts 92.5% of bills correctly, a difference of 1.0 percentage points is 21.3% of the increase in the model's predictive accuracy relative to the baseline. While the content of bills has a meaningful but small effect when trying to predict where bills end up in Congress, the same cannot be said for *which* text versions are used. There is no discernible difference between how well I am able to predict votes based on the earliest version of a bill's text and the last version of a bill's text.

These results support some, but also negate some, of our intuitions about what drives bill passage. Most importantly, bill content *does* matter for predicting where a bill will end up in the House. Further, this is to my knowledge the first paper to provide predictions of bill outcomes in Congress.

While I expected the models to best predict whether bills make it out of committee, this pattern manifests primarily in the percent correctly predicted rather than the separation and ROC plots. The conflicting evidence between different model arrangements and assessments indicates that this hypothesis will require further investigation. Predicting whether a bill will make it out of committee remains an important goal, since committees are where most bills die. Future work should consider the many particulars of the committee process, including committee assignment, subcommittee handling, and chair features, can impact which bills progress.

Lastly, there is no evidence that the choice of text version had an effect on the model's predictive performance. We know that most bills go through changes in committee and on the floor via amendments, so what does it mean that a bill's content can predict where a bill will end up but the version of the bill we use does not? While this phenomenon is good for prediction—if it does not matter which text version of the bill is used in my model to predict bill outcomes, then this model can be applied to predict the outcome of a bill as soon as it is referred to committee—the question still remains as to *why* the bill text version does not matter. It could be that because most bills have 1 or 2 text versions, that on aggregate, there are not enough changes between the earliest and last text versions of a bill to make a difference when predicting the outcome of all bills. Alternatively, even if they do have many text versions, it may be bills do not go through large enough changes between versions to make enough of a difference in the predictive model. Finally, it could be an artifact of the methodology. While doc2vec is good at picking up the overall context, structure, and semantic meaning of documents, it may not be as sensitive to small changes in bills between text versions. Distinguishing between these drivers is an area for future research as the implications for interpreting bill passage differ depending on whether the phenomenon is the result of practical features of model parameterization or of congressional behavior.

Here I was able to predict where bills would end their journey in the House, but passing the House does not mean a bill becomes law. The next logical step is to try and predict which bills become law and which do not. This is a more complicated process, and as such I would expect the predictive accuracy to be lower. Often, bills will pass the House that have no chance of passing in the Senate, as the House is commonly seen as a place for ideological grandstanding, where bills are introduced with no chance of passage (Sinclair 1998). Then again, this is also often true of bills that come to the House from the Senate.[68] Furthermore, the legislative processes in the House and Senate are quite different, which may lead to difficulty in accurately predicting which bills will become law.

There are also a number of ways in which the predictive model I present here could potentially be improved even for predicting votes purely in the House. One such

[68] Of the 49,253 bills in my sample, 1,840 were passed by the Senate before reaching the House and more than half of these (59.4%) eventually became public laws. Of the 6,487 bills in my sample that passed the House, 5,394 of these were passed first by the House and about half of those (49.4%) became public laws.

improvement could come in the form of ensemble learning, which is a paradigm for prediction that has become popular in the machine learning community. There are many flavors of ensemble learning, but all involve training multiple predictive models and systematically joining their algorithms to yield a single prediction from the ensemble for each observation. Ensemble methods often lead to better predictions than any single model in the ensemble (*e.g.,*Wolpert 1992; Breiman 1996; Monteith et al. 2011). Future work could apply ensemble methods to predicting legislative behavior and outcomes in Congress to obtain better results.

We should generally continue to try and better understand how the content of bills shapes the decisions made about them. Others (de Marchi, Dorsey, and Ensley 2017; Gerrish and Blei 2012) and I (see Chapter 2 of this dissertation) have shown that the content of bills is highly predictive of the voting decisions made by members of Congress. In this case, bill content is less predictive of the eventual success of bills in the House, but still has a part to play. We should strive to know more about the complex process by which members of Congress create and mold the content of bills, and in turn to learn more about how content affects members' decision-making. When a bill is introduced, it is referred to at least one committee in the House, based on the issues the bill considers. Sometimes a bill is referred to more than one committee, but this is a choice; no bill need be referred to more than one committee. By looking at bill content, we could for instance assess which bills are more likely to be referred to more than one committee, and what effect this has on voting decisions both in and out of committee. We could also investigate which members are more likely to change the content of bills in specific ways over the course of a bill's journey, or if bills of a certain type are more likely to bypass the committee process and come to a vote without a committee report. Because most members do not actually work on any specific bill, not only do members of Congress shape bills, but bills can shape members of Congress. A deeper understanding of these dynamics will enrich our knowledge of the legislative system in the United States.

Conclusion

Across three papers, I have come to some complementary conclusions about party power and legislative outcomes in Congress. In chapter one, I assessed some of the upper limits of party power to whip votes and ensure members follow through with their agreements. When the majority party is trying to pass a contentious bill and cannot garner any support from the minority party, the majority must get as many of their members to toe the party line as possible. This means making deals with members, and holding them to their deals. When members agree to vote with the party on a divisive bill, they make statements in support of that bill, even when there is no other obvious incentive to make a statement.

In chapter two, I investigated majority party control of the agenda process in the House. I found that the majority party has almost complete control of the legislative process. There is strong evidence of both positive and negative agenda control, and the minority party is less likely to support or have introduced bills as they progress through the legislative process. Interestingly, this trend did not change appreciably between 1995 and 2011, even though party polarization in Congress increased over that time by many measures. As we would expect an increase in party polarization to lead to less reaching across the aisle and more single-party control of the legislative process, finding similar levels of majority party control highlights the need to rethink how polarization could affect agenda control. Further, there was an increase in the proportion of bills introduced by members of the minority party that received a final passage vote. These findings suggest we should not only rethink the relationship between polarization and agenda, but the relationship between agenda control and control of policy.

In chapter three, I continue my analysis using the content of bills from chapter two, but use it to predict the outcome of bills rather than members' votes. I find that bill outcomes are highly predictable, and that including a doc2vec characterization of each bill substantially improves models relative to a baseline rate. However, there is no discernible effect of using different text versions of bills.

The approaches I bring to bear on my research questions open up a number of interesting lines of research. First and foremost, while I mostly consider the House, there are two comparisons that should be made to the Senate. Namely, how do the levels of majority party agenda control and predictability of bill outcomes compare between the House and Senate? Relative changes and differences between these patterns in the House and Senate over time can help us understand how institutions shape legislative behavior, how party polarization manifests in each chamber, how different theories of party and congressional operations map onto differences between the chambers, and more.

Second, my research highlights the importance of the committee system. In chapter two, I confirm the results of many others by finding that most bills dying in the House do so in committee. In chapter three, I found that I could more accurately predict which bills would end up dying in committee than I could what happens after committee. At the same time, the models in chapter three over-predicted bills as dying in committee. There have been many accounts of the committee system, but no congressional theory adequately explains many of the facets of the behavior and makeup of congressional committees (Groseclose and King 2001). A more thorough examination of the types of bills which my models are able to correctly and incorrectly predict could further aid in our understanding of the committee system, potentially leading to theoretical as well as empirical insights.

My research also exposes the need for a combination of new and old methods in order to truly understand political phenomena. The methods I use enable me to characterize large sets of text in ways that scholars simply could not until recently, but my work also raises questions that would be better solved by careful qualitative content analysis. For instance, why does the model predict some bill outcomes or votes more accurately than others? Because of the high accuracy of my models, these incorrect predictions provide a small enough sample of bills and votes that they could be analyzed on a case by case basis for patterns that would improve future predictive modeling, as well as aiding in how we understand the factors that dominate the minds of members of Congress as they make various legislative decisions.

Lastly, my findings call for more research on party power and polarization. Particularly, it is implausible that parties became more polarized if polarization is related to both control of the agenda and control of the chamber's policy output. While the majority party retained similar levels of agenda control throughout the period I study, the minority party's impact on the policymaking process actually increased. We need to think more about how agenda control and policy control are related. They may be related concepts, and control of the agenda may lead to control of the chamber's policy output, but clearly, lacking control of the agenda does mean that a party is shut out of the policymaking process.

Appendix A— Public Statements in Congress as Precommitment Devices: Supplement

In this appendix, I show alternative model specifications for the main model results in Chapter 1, which can be found in Table 3. These can be found in Table 15. The first column is the full model from the text (Table 3 column 4). The second column includes results where the variable for whether members received concessions is whether they were members of the Tuesday Group, rather than the Freedom Caucus. The third column uses the positive support for the AHCA in each state, rather than the difference, which is reported in the text.

As we can see, the main results are unchanged. In each alternate model, whether members switched to vote with the party is the only significant predictor of whether members made a public statement in favor of the bill at the 0.01 level. In column 2, where members of the Tuesday Group are considered to have received concessions rather than members of the Freedom Caucus, whether members received concessions is significant at the 0.05 level, though the coefficient is also substantially smaller (-0.82 versus -1.02). Further, in this model, membership on a prestige committee was slightly more highly correlated with members' public statements (coefficient change from 0.53 to 0.61 with an identical standard error. In the model using MRP support rather than the different between estimated support and opposition, the coefficients are virtually unchanged.

Table 15: Alternative model specifications for the results presented in Table 3. Column 1 is the full model reported in the text. Column 2 changes which members are thought to have gained policy concessions—here, those in the Tuesday Group, in the text, those in the Freedom Caucus. Column 3 uses an alternate measure of public opinion, namely the support in each state as estimated via MRP, rather than the difference between support and opposition. Note: * = *p* < 0.05, ** = *p* < 0.001

	In-text	**Tuesday Group**	**MRP Support**
Switched	2.75**	2.19**	2.76**
	(0.57)	(0.40)	(0.57)
Concessions	-1.02	-0.82*	-1.01
	(0.65)	(0.42)	(0.65)
MRP	0.004	-0.01	-0.00
	(0.02)	(0.02)	(0.04)
Tenure	-0.01	-0.01	-0.01
	(0.02)	(0.02)	(0.02)
Male	-0.68	-0.78	-0.69
	(0.45)	(0.45)	(0.45)
White	1.14	0.88	1.12
	(0.81)	(0.74)	(0.81)
Leader	0.52	0.41	0.52
	(0.91)	(0.91)	(0.91)
Prestige	0.53	0.61*	0.54
	(0.28)	(0.28)	(0.28)
Intercept	-1.47	-1.28	-1.52
	(0.94)	(0.89)	(1.75)
Observations	290	290	290
Nagelkerke Pseudo R^2	0.21	0.21	0.21

Appendix B—Bill Text and Agenda Control in the U.S. House of Representatives: Supplement

This is the general appendix for chapter 2. Here, I provide information on supplementary analyses and alternative model specifications. First, I show that the predictive accuracy of the model is relatively consistent across Congresses. This information is displayed in Figure 16.

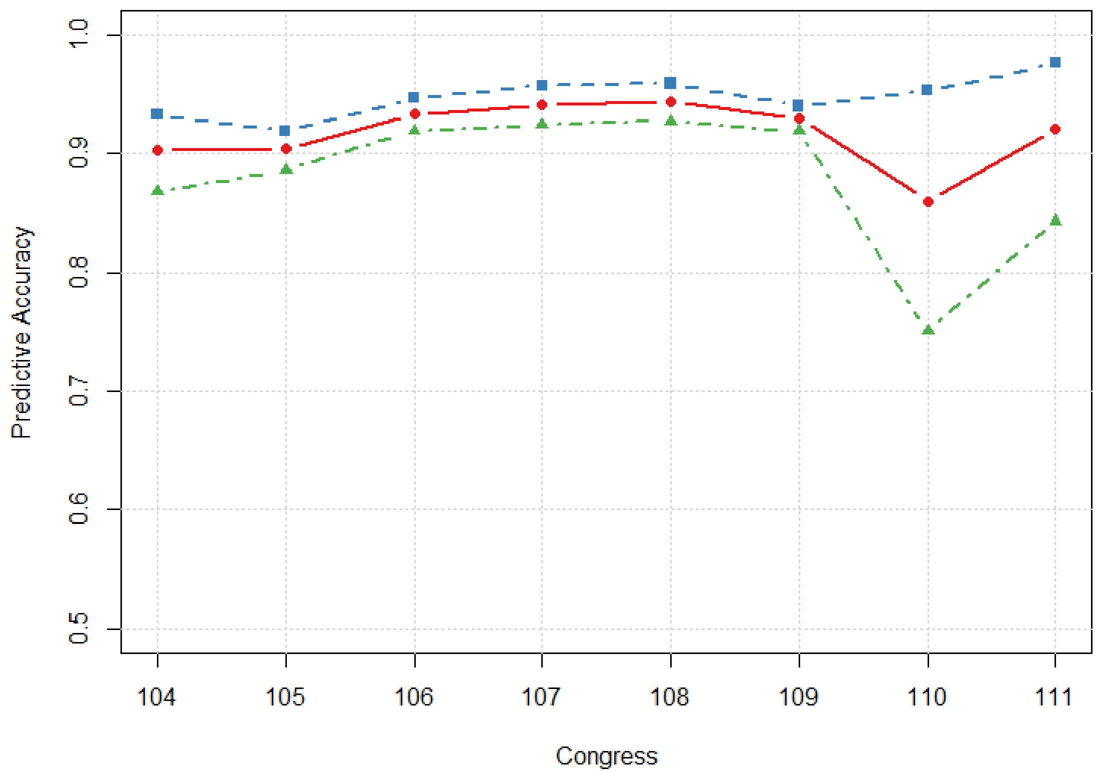

Figure 16: Out-of-sample predictive accuracy for members' votes in each Congress. The lines represent votes for all members, those in the majority party, and those in the minority party.

As shown in Figure 16, the predictive accuracy of the model is relatively consistent over time. There is a similar pattern of high predictive accuracy in the 104th-109th Congresses, with the 110th the largest aberration, though the accuracy in the 111th does not quite return to pre-110th levels. The overall predictive accuracy is greater than 90% in each Congress except the 110th. Further, as demonstrated in the text, the patterns of agenda control are not substantially different between the 104th-109th and 110th-111th Congresses.

Second, I include an alternative model specification that does not include the number of cosponsors as a predictor, in order to isolate the predictive effect of the doc2vec bill vectors. The neural weights and the number of cosponsors are the only independent variables I use at the bill level, so I want to assess whether removing the number of cosponsors from the model decreases the predictive accuracy enough to worry that the neural weights add little to the model. These results can be found in Figure 17.

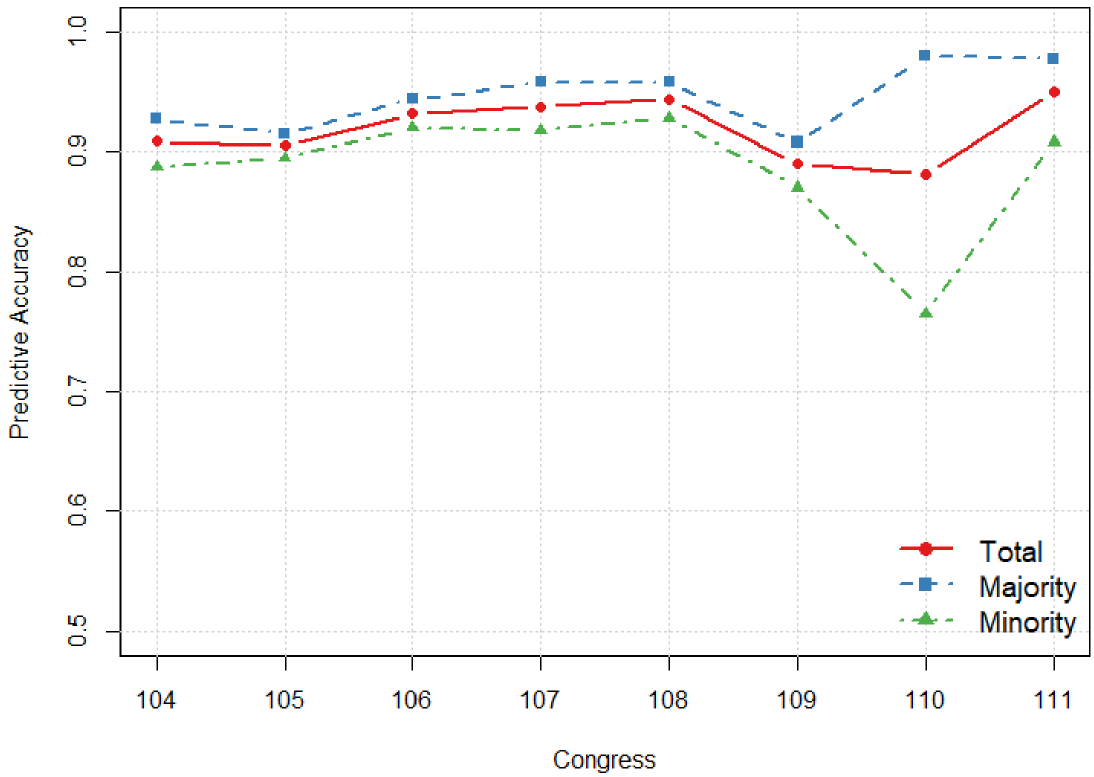

Figure 17: Out-of-sample predictive accuracy for members' votes in each Congress from a model that does not include the number of cosponsors for each bill. The lines represent votes for all members, those in the majority party, and those in the minority party.

Comparing Figures 16 and 17, we can see that the predictive accuracy is nearly identical for the full model and a model that does not include the number of cosponsors for each bill. The pattern of high predictive accuracy, with a dip for the minority party in the 110th, is exactly the same in each model. Because the bill text and the number of cosponsors are the only bill-level predictors in the model, we can say that the model's high predictive accuracy is not an artifact of using the number of cosponsors as a proxy for bill support.

Appendix C—Bill Content and Legislative Success in the U.S. House of Representatives: Supplement

This is the general appendix for chapter 3. Here, I provide information on an alternative model specification. Table 16 include results for the model comparison in Table 14, but for the multinomial model. In all cases, the alternative specifications do not change the thrust of the substantive results.

Table 16: Predictive results for the multinomial model using 1) all variables, 2) no neural weights, and 3) the earliest text version of each bill.

	Bills Passed	*Baseline Percent Correct*	*Percent Correct*	*Improvement over Baseline*	*False Positives*	*False Negatives*
1) Full Model	6,487	83.8%	90.2%	6.4	3.3%	6.5%
2) No Weights	6,487	83.8%	89.8%	6.0	3.7%	6.5%
3) Earliest Version	6,487	83.8%	90.1%	6.3	3.6%	6.3%

As we can see from Table 16, there is functionally no difference between between the accuracy of the full model (using the most recent version) and the full model using the earliest text version of each bill. The full model is a 6.4 percentage point improvement over the baseline rate of 83.8% when each bill is classified as dying in committee, and the model using the earliest text version of each bill is a 6.3 percentage point improvement over the baseline. A difference of 0.1 is just 2% of the full model's improvement over the baseline Furthermore, the distribution of false positives and negatives is nearly identical between the two models.

In the binomial context, more of the full model’s improvement over the baseline predictive accuracy rate can be attributed to the inclusion of the neural weights of the bill texts than in the multinomial context. The full multinomial model is a 6.4 percentage point improvement over the baseline, and a multinomial model that does not include the neural weights is a 6.0 percentage point improvement over the baseline. This accounts for 7% of the full model’s improvement over the baseline rate. While this proportion is larger than for the earliest text version of each bill, it is smaller than the proportion of the full model’s improvement over the baseline in the binomial model (20%).

Appendix D—Doc2vec Technical Appendix

D.1 The word2vec Framework

Word2vec (Mikolov et al. 2013a, 2013b), and by extension doc2vec, is not a single model, but a group of related models that represent words as real number vectors. As explained in Chapter 2, word2vec models create vector embeddings—or vector representations—of words via a neural network by predicting the context in which words appear. In word2vec models, an input is passed through a hidden layer, which transforms the input into a method that can be used in the output. The inputs and outputs vary by model, but the hidden layer is a n x k matrix, where n is the number of "nodes" which determine the length of each vector, k is the number of unique words in the corpus, and each column is the vector embedding for each word.

The two main algorithms for learning the word embeddings are continuous bag of words (CBOW) and skip gram, each predicts words in context, though differently. In CBOW, the model input is the context around the word defined by the window size. For a window size of 2, the input would be the two words preceding and following a target word. The output is a predicted probability for each unique word in the corpus being the word surrounded by the input context. The skip gram model does the opposite, predicting the context around a word, given the word. So the input in a skip gram model is a single word, and the output is a predicted probability for each word in the corpus to appear within the window size of the input word.

A visual representation of one possible framework for learning word vectors via word2vec can be found in Figure 18, which comes from Le and Mikolov (2014). In this framework, the matrix W is our $n \times k$ matrix with words represented as columns of one-hot vectors. Given some set of context words, the goal is to maximize the average probability of choosing the correct target word (in this case "on"), given the set of context words, summed across all words in the corpus being used as the target word.

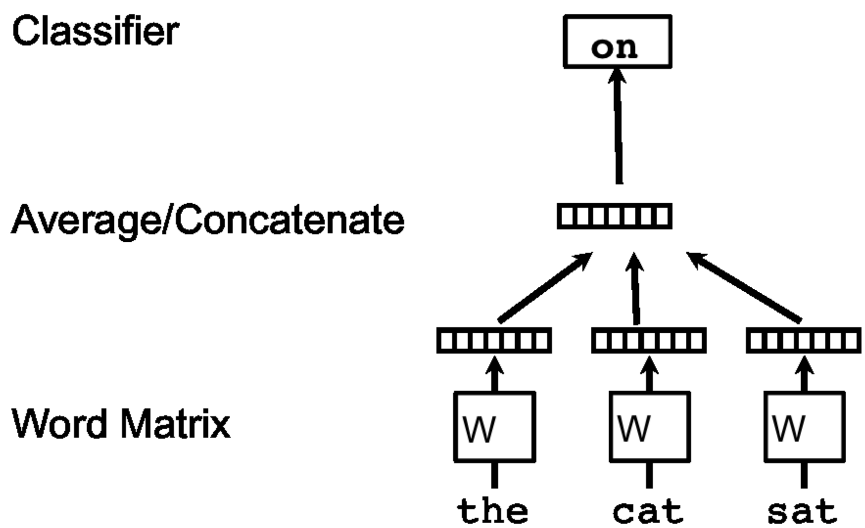

Figure 18: A word2vec framework for learning word vectors, with three words ("the", "cat", and "sat") that represent context being used to predict the fourth word ("on"). Input words are represented by columns in the matrix *W*.

The classifier is generally a multinomial logistic regression model (more commonly referred to as softmax regression in the machine learning literature). The word vectors themselves are then optimized using stochastic gradient descent, with the gradient computed via backpropagation. Backpropagation is a technique for calculating the gradient of the error function as it relates to the model's neural weights which does so backward through the neural network, in our case starting with the output layer and then moving backward to the hidden layer. In word2vec, this means that the model updates the weights of the hidden layer (which are our embedded word vectors) based on the gradient of the error function of the output layer. With each iteration of gradient descent the model updates the hidden layer weights such that they reflect higher predictive accuracy in the output layer. When the model converges, "words with similar meaning are mapped to a similar position in the vector space." (Le and Mikolov 2014, 2).

D.1.1 Negative sampling

The original implementation of word2vec models was trained via multinomial regression, but this method is impractical for large datasets. For each training sample in a multinomial regression framework of word2vec, the algorithm slightly adjusts the neural weights of *every* word in the *output layer* order to predict the sample more accurately. As mentioned in Chapter 2 of the text, we are not even interested in the output layer weights, and rather in the hidden layer weights. However, updating the output layer weights is necessary for learning the hidden layer word embeddings since word vectors are updated via backpropagation.

Consider the case of my data, with more than 80,000 documents, about 10,000 unique words, a window size of 10, and 600 neurons in the hidden layer. For each training sample (a two word pair consisting of the target word and one of the words within the window size of the target word), the model will update the output layer weights for each of the approximately 10,000 * 600 = 6,000,000 weights in the model. There are also billions of training samples (each word in each document used as a target word, paired with each other word within the window size), so this process quickly becomes computationally inefficient. So Mikolov et al. (2013b) developed a new way to train the model called negative sampling. Under negative sampling, the algorithm will only update the output layer weights for a researcher-defined number of "negative" words (for which the model predicts a 0) and one "positive" word (the target word). Mikolov et al. showed that using negative sampling rates between 2 and 20 cut down on computational time without affecting predictive accuracy, which in the case of my data means only updating 0.02 to 0.24% of the 6M weights in the output layer for each training sample.[69]

In the text, I describe a skip gram model, and I train each of my doc2vec models with negative sampling. While skip gram is slower than CBOW, Mikolov et al. (2013b) showed that the skip gram model did a better job at representing rare words than the CBOW model.

D.2 Moving from word2vec to doc2vec

Consider the framework from Figure 18. Because the word vectors are learned by averaging across the error function of the prediction task for each word, each word in the context surrounding the target word has a role to play in learning the word vector

[69] In the hidden layer, only the neural weights for the target word are updated with each training sample.

for the target word. One way to think of this role is as a form of memory. Each word surrounding the target word, and the vector representations of those context words, represents a remembered piece of the context in which the target word appears. This idea, that each word is a representation of the memory of the context surrounding the target word, and that each piece of the context contributes to a prediction task about the target word, is the inspiration for the expansion to doc2vec from word2vec. In doc2vec, each paragraph (or document) can be represented as its own vector. This vector is a column in a matrix D of paragraph vectors of dimensions $n \times l$—where n is the number of nodes in the neural network and l is the number of paragraphs—and contributes to the prediction task for each target word.[70] This paragraph vector is used to broaden the context in which the target word exists, acting "as a memory that remembers what is missing from the current context—or the topic of the paragraph" (Le and Mikolov 2014). This framework is displayed in Figure 19, which like Figure 18 comes from Le and Mikolov (2014).

In such a model, the vector for each paragraph is used for all of the training samples and target words within that paragraph, but not across paragraphs in the corpus. However, the word matrix W is used for all contexts in the corpus. In addition to learning each target word vector using the words that surround it and the paragraph

[70] Le and Mikolov (2014) use paragraphs as the initial example for expanding word2vec to doc2vec. Following their example, I do the same. However, in the text of this dissertation I use document vectors rather than paragraph vectors. The unit at which the broader vectors are estimated is trivial to change and does not require any alterations to the model.

vector, an additional step must be performed in order to learn the paragraph vectors. Like for the word vectors, this is done by gradient descent via backpropagation, though in this case we hold the other model parameters (word vectors and multinomial logistic regression weights) constant.

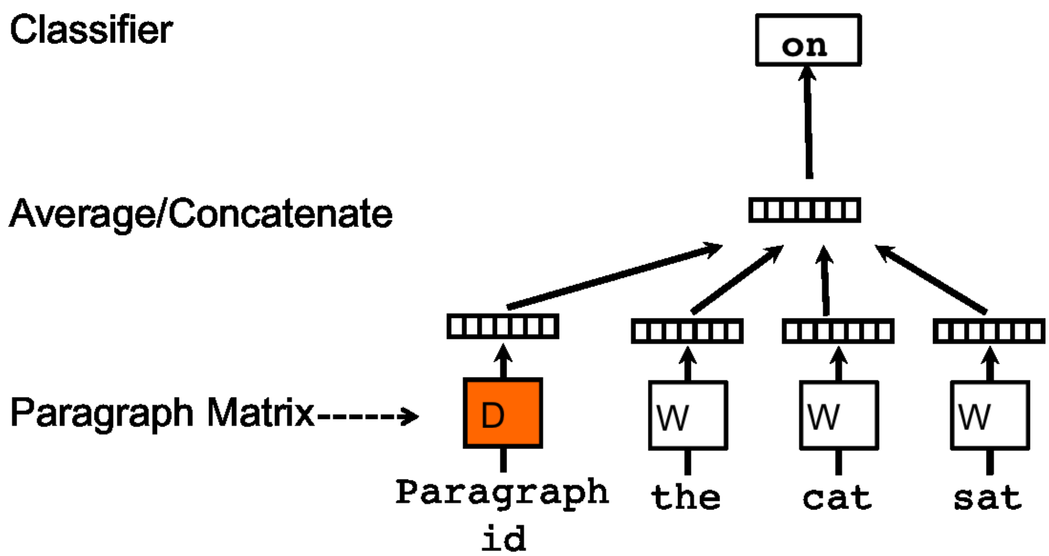

Figure 19: A doc2vec framework for learning paragraph vectors. The only change from Figure 16 is that in this model, the vectors from each context word and the paragraph vector are used to predict the target word.

Just like with more traditional word2vec models, there are multiple frameworks commonly used in doc2vec. The model described above is known as the Distributed Memory version of Paragraph Vector (PV-DM), since each paragraph vector fills in the broader context of the paragraph when predicting each target word. Another popular method of learning paragraph vectors is called the Distributed Bag of Words version of Paragraph Vector (PV-DBOW). In this case, the model ignores context words in the input, and predicts words randomly sampled from within the paragraph in the output.

In other words, a text window from within the paragraph is sampled, and then each word within the text window is predicted, given the paragraph vector, which is updated with each iteration via stochastic gradient descent.

In the text of this dissertation, I estimate document vectors with both PW-DM and PV-DBOW, and then concatenate the vectors from each model to be used as the document vectors when I predict votes (Chapter 2) and bill outcomes (Chapter 3). Concatenated vectors have been shown to be more accurate in predictive tasks (Rathore 2018).

References

Abramowitz, A. (2010). *The Disappearing Center: Engaged Citizens, Polarization, and American Democracy*. New Haven, CT: Yale University Press.

Abramowitz, A. I., & Saunders, K. L. (2008). Is Polarization a Myth?. *The Journal of Politics*, *70*(2): 542-555.

Adler, E. Scott and John Wilkerson, *Congressional Bills Project: 1995-2011*. NSF 00880066 and 00880061.

Aldrich, J. H. (2011). *Why Parties?: a Second Look*. Chciago: University of Chicago Press.

Aldrich, J. H., Montgomery, J. M., & Sparks, D. B. (2014). Polarization and Ideology: Partisan Sources of Low Dimensionality in Scaled Roll Call Analyses. *Political Analysis 22*(4): 435-456.

Aldrich, J. H., & Rohde, D. W. (2001). *The Logic of Conditional Party Government: Revisiting the Electoral Connection*. PIPC.

Ansolabehere, S., De Figueiredo, J. M., & Snyder Jr, J. M. (2003). Why is There So Little Money in US Politics?. *Journal of Economic Perspectives*, *17*(1): 105-130.

Ansolabehere, S., Snyder Jr, J. M., & Stewart III, C. (2001). Candidate Positioning in US House Elections. *American Journal of Political Science*: 136-159.

Anzia, S. F., & Jackman, M. C. (2012). Legislative Organization and the Second Face of Power: Evidence from US State Legislatures. *The Journal of Politics*, *75*(1): 210-224.

Ariely, D., and Wertenbroch, K. (2002). Procrastination, Deadlines, and Performance: Self-Control by Precommitment. *Psychological science*, *13*(3): 219-224.

Ballard, A., Lerner, J., and Minhas, S. (2018). Mapping Agendas: Text Analysis of Party Loyalty through Congressional Rhetoric. Unpublished Manuscript.

Bartels, L. M. (2016). *Unequal Democracy: The Political Economy of the New Gilded Age*. Princeton: Princeton University Press.

Bartels, L. M. (2000). Partisanship and Voting Behavior, 1952-1996. *American Journal of Political Science*: 35-50.

Bauer, Raymond A., Ithiel de Sola Pool, and Lewis Anthony Dexter (1972). *American Business and Public Policy: The Politics of Foreign Trade*, 2d ed. Chicago: Aldine.

Bengio, Y., Lamblin, P., Popovici, D., & Larochelle, H. (2006). Greedy Layer-Wise Training of Deep Networks. In *Advances in Neural Information Processing Systems:*153-160.

Berrar, D., & Flach, P. (2011). Caveats and Pitfalls of ROC Analysis in Clinical Microarray Research (and How to Avoid Them). *Briefings in Bioinformatics*, *13*(1): 83-97.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet Allocation. *Journal of Machine Learning Research*, *3*(1): 993-1022.

Boatright, Robert G., Vincent G. Moscardelli, and Clifford Vickrey (2017). Congressional Primary Elections Data. On line, https://wordpress.clarku.edu/primarytiming.

Bottou, L. (2010). Large-Scale Machine Learning with Stochastic Gradient Descent. In *Proceedings of COMPSTAT'2010*: 177-186.

Breiman, L. (1996). Bagging Predictors. *Machine Learning*, *24*(2): 123-140.

Canes-Wrone, B. (2015). From Mass Preferences to Policy. *Annual Review of Political Science 18*: 147-165.

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

Carson, J. L., Madonna, A. J., & Owens, M. E. (2016). Regulating the Floor: Tabling Motions in the US Senate, 1865-1946. *American Politics Research*, *44*(1): 56-80.

Carty, R., K. (2004). Parties as Franchise Systems: The Stratarchical Organizational Imperative. *Party Politics*, 10(1): 5-24.

Coates, D., & Munger, M. (1995). Legislative Voting and the Economic Theory of Politics. *Southern Economic Journal*: 861-872.

Collobert, R., & Weston, J. (2008). A Unified Architecture for Natural Language Processing: Deep Neural Networks with Multitask Learning. In *Proceedings of the 25th International Conference on Machine learning:* 160-167.

Cox, G. W., Kousser, T., & McCubbins, M. D. (2010). Party Power or Preferences? Quasi-Experimental Evidence from American State legislatures. *The Journal of Politics*, *72*(3): 799-811.

Cox, G. W., & McCubbins, M. D. (2007). *Legislative Leviathan: Party Government in the House*. Cambridge: Cambridge University Press.

Cox, G. W., & McCubbins, M. D. (2005). *Setting the Agenda: Responsible Party Government in the US House of Representatives*. Cambridge: Cambridge University Press.

Cox, G. W., & McCubbins, M. D. (1991). On the Decline of Party Voting in Congress. *Legislative Studies Quarterly*: 547-570.

Crespin, Michael H. and David Rohde. (2018). Political Institutions and Public Choice Roll-Call Database. Retrieved from http://cacexplore.org/pipcvotes/

Crespin, M. H. and Rohde, D. W. (2010). Dimensions, Issues, and Bills: Appropriations Voting on the House Floor. *The Journal of Politics 72*(4): 976-989.

Curry, J. (2018) Knowledge, Expertise, and Committee Power in the Contemporary Congress. Unpublished manuscript.

Curry, J. and Frances E. Lee (2018). What is Regular Order Worth? Partisan Lawmaking and Congressional Processes. Unpublished manuscript.

de Marchi, S., Dorsey, S., and Michael Ensley (2017). Predicting Roll Call Votes in the House of Representatives. Presented at the *Annual Meeting of the Midwest Political Science Association*. Chicago, IL.

Den Hartog, C., & Monroe, N. W. (2011). *Agenda Setting in the US Senate: Costly Consideration and Majority Party Advantage*. Cambridge: Cambridge University Press.

Diermeier, D., Godbout, J. F., Yu, B., & Kaufmann, S. (2012). Language and Ideology in Congress. *British Journal of Political Science*, *42*(1): 31-55.

Downs, A. (1957). An E*conomic Theory of Democracy.* New York: Harper and Row.

Elster, J. (2000). *Ulysses Unbound: Studies in Rationality, Precommitment, and Constraints*. Cambridge: Cambridge University Press.

Enelow, J. M., & Koehler, D. H. (1980). The Amendment in Legislative Strategy: Sophisticated Voting in the US Congress. *The Journal of Politics*, *42*(2): 396-413.

Epstein, D., & O'Halloran, S. (1999). *Delegating powers: A Transaction Cost Politics Approach to Policy Making under Separate Powers*. Cambridge: Cambridge University Press.

Evans, C. L., & Oleszek, W. J. (2001). Message Politics and Senate Procedure. *The Contentious Senate: Partisanship, Ideology, and the Myth of Cool Judgment*: 107-127.

Farrell, J., and Rabin, M. (1996). Cheap Talk. *The Journal of Economic Perspectives*, *10*(3): 103-118.

Fenno, R. F. (1978). *Home Style: House Members in Their Districts*. New York: Harper-Collins.

Fenno, R. F. (1973). *Congressmen in Committees*. Boston: Little, Brown and Company.

Fiorina, M. P., & Abrams, S. J. (2008). Political Polarization in the American Public: Misconeptions and Misreadings. *The Journal of Politics 70*(2): 556-560.

Fowler, J. H. (2006). Legislative Cosponsorship Networks in the US House and Senate. *Social Networks 28*(4): 454-465.

Friedman, J., Hastie, T., Simon, N., & Tibshirani, R. (2017). glmnet: Lasso and Elastic-net Regularized Generalized Linear Models. *R Package Version*, *1*(4).

Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization Paths for Generalized Linear Models via Coordinate Descent. *Journal of Statistical Software*, *33*(1): 1-22.

Frydman, R., Gray, C., Hessel, M., and Rapaczynski, A. (2000). The Limits of Discipline: Ownership and Hard Budget Constraints in the Transition Economies. *Economics of Transition*, *8*(3): 577-601.

Gailmard, S., & Jenkins, J. A. (2007). Negative Agenda Control in the Senate and house: Fingerprints of Majority Party Power. *The Journal of Politics*, *69*(3): 689-700.

Gamm, G., & Kousser, T. (2010). Broad Bills or Particularistic policy? Historical Patterns in American State Legislatures. *American Political Science Review*, *104*(1): 151-170.

Gamm, G., & Smith, S. S. (2002). Emergence of Senate Party Leadership. *US Senate Exceptionalism*: 393-413.

Gerrish, S., & Blei, D. M. (2012). How They Vote: Issue-Adjusted Models of Legislative Behavior. In *Advances in Neural Information Processing Systems:* 2753-2761.

Greenhill, B., Ward, M. D., & Sacks, A. (2011). The Separation Plot: A New Visual Method for Evaluating the Fit of Binary Models. *American Journal of Political Science*, *55*(4): 991-1002.

Grimmer, J., Messing, S., & Westwood, S. J. (2012). How Words and Money Cultivate a Personal Vote: The Effect of Legislator Credit Claiming on Constituent Credit Allocation. *American Political Science Review*, *106*(4): 703-719.

Grimmer, J. (2009). A Bayesian Hierarchical Topic Model for Political Texts: Measuring Expressed Agendas in Senate Press Releases. *Political Analysis*, *18*(1): 1-35.

Groseclose, T., & King, D. C. (2001). Committee Theories Reconsidered. In *Congress Reconsidered, 7th Ed.*: 191-216). Washington, DC: CQ Press.

Grossmann, M., & Hopkins, D. A. (2016). *Asymmetric Politics: Ideological Republicans and Group Interest Democrats*. Oxford: Oxford University Press.

Grynaviski, J. D. (2010). *Partisan Bonds: Political Reputations and Legislative Accountability*. Cambridge: Cambridge University Press.

Gulati, G. J. (2004). First Impressions: Congressional Homepages and Presentation of Self on the WWW. *Harvard International Journal of Press Politics, 9*(4).

Gurciullo, S., & Mikhaylov, S. (2017). Detecting Policy Preferences and Dynamics in the UN General Debate with Neural Word Embeddings. *arXiv:1707.03490*.

Hall, R. L. (1996). *Participation in Congress*. New Haven, CT: Yale University Press.

Hall, R. L. (1987). Participation and Purpose in Committee Decision Making. *American Political Science Review*, *81*(1): 105-127.

Hall, R. L., & Wayman, F. W. (1990). Buying Time: Moneyed Interests and the Mobilization of Bias in Congressional Committees. *American Political Science Review*, *84*(3): 797-820.

Harbridge, L., Malhotra, N., and Harrison, B. F. (2014). Public Preferences for Bipartisanship in the Policymaking Process. *Legislative Studies Quarterly*, *39*(3): 327-355.

Harbridge, L., and Malhotra, N. (2011). Electoral Incentives and Partisan Conflict in Congress: Evidence from Survey Experiments. *American Journal of Political Science*, *55*(3): 494-510.

Hastie, T., Tibshirani, R., & Friedman, J. (2009). Unsupervised Learning. In *The Elements of Statistical Learning*: 485-585. New York, NY: Springer.

Hoerl, A. E., & Kennard, R. W. (1970). Ridge Regression: Biased Estimation for Nonorthogonal Problems. *Technometrics*, *12*(1): 55-67.

Holmes, S. (1988). Precommitment and the Paradox of Democracy. *Constitutionalism and Democracy*, *195*(1): 199-221.

Iyengar, S., & Kinder, D. R. (2010). *News That Matters: Television and American Opinion*. Chicago: University of Chicago Press.

Jackman, M. C. (2013). Parties, Median Legislators, and Agenda Setting: How Legislative Institutions Matter. *The Journal of Politics*, *76*(1): 259-272.

Jenkins, J. A., & Monroe, N. W. (2016). On Measuring Legislative Agenda-Setting Power. *American Journal of Political Science*, *60*(1): 158-174.

Jenkins, J. A., & Monroe, N. W. (2012). Buying Negative Agenda Control in the US House. *American Journal of Political Science*, *56*(4): 897-912.

Jochim, A. E., & Jones, B. D. (2013). Issue Politics in a Polarized Congress. *Political Research Quarterly*, *66*(2): 352-369.

Jones, B. D., & Baumgartner, F. R. (2005). *The Politics of Attention: How Government Prioritizes Problems*. Chicago: University of Chicago Press.

Kessler, D., & Krehbiel, K. (1996). Dynamics of Cosponsorship. *American Political Science Review*, *90*(3): 555-566.

Kingdon, John. 1981. *Congressmen's Voting Decisions*. New York: Harper and Row.

Kivetz, R., and Simonson, I. (2002). Self-Control for the Righteous: Toward a Theory of Precommitment to Indulgence. *Journal of Consumer Research*, *29*(2): 199-217.

Krehbiel, K. (2006). Partisan Roll Rates in a Nonpartisan Legislature. *The Journal of Law, Economics, & Organization*, *23*(1): 1-23.

Krehbiel, K. (1998). *Pivotal Politics: A Theory of U.S. Lawmaking*. Chicago, IL: University of Illinois Press.

Krehbiel, K., & Rivers, D. (1990). Sophisticated Voting in Congress: a Reconsideration. *The Journal of Politics*, *52*(2): 548-578.

Krehbiel, K., & Woon, J. (2005). *Selection Criteria for Roll Call Votes.* Presented at the *Annual Meeting of the American Political Science Association*, Washington, D.C.

Kreps, D. M., and Scheinkman, J. A. (1983). Quantity precommitment and Bertrand Competition Yield Cournot Outcomes. *The Bell Journal of Economics*: 326-337.

Landgrebe, T. C., & Duin, R. P. (2007). Approximating the Multiclass ROC by Pairwise Analysis. *Pattern Recognition Letters*, *28*(13): 1747-1758.

Langbein, L. I. (1986). Money and Access: Some Empirical Evidence. *The Journal of Politics*, *48*(4): 1052-1062.

Lee, F. E. (2016). *Insecure Majorities: Congress and the Perpetual Campaign*. Chicago: University of Chicago Press.

Lee, F. E. (2009). *Beyond Ideology: Politics, Principles, and Partisanship in the US Senate*. Chicago: University of Chicago Press.

Le, Q., & Mikolov, T. (2014). Distributed Representations of Sentences and Documents. In *International Conference on Machine Learning*: 1188-1196.

Levmore, S. (1996). Precommitment Politics. *Virginia Law Review*: 567-627.

Lewis, Jeffrey B., Keith Poole, Howard Rosenthal, Adam Boche, Aaron Rudkin, and Luke Sonnet (2017). *Voteview: Congressional Roll-Call Votes Database*. https://voteview.com/

Matthews, D. R. (1960). *United States Senators and Their World*. Chapel Hill, NC: University of North Carolina Press.

Masket, Seth. (2009). *No Middle Ground: How Informal Party Organizations Control Nominations and Polarize Legislatures*. Ann Arbor, MI: University of Michigan Press.

Mayhew, D. R. (1974). *Congress: The Electoral Connection*. New Haven, CT: Yale University Press.

McCarty, N., Poole, K. T., & Rosenthal, H. (2016). *Polarized America: The Dance of Ideology and Unequal Riches*. Cambridge, MA: MIT Press.

McCormick, C. (2016). *Word2Vec Tutorial - The Skip-Gram Model*. Retrieved from http://www.mccormickml.com

McCubbins, M. D., & Schwartz, T. (1984). Congressional Oversight Overlooked: Police Patrols Versus Fire Alarms. *American Journal of Political Science*, *28*(1) 165-179.

Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013a). Efficient Estimation of Word Representations in Vector Space. *arXiv:1301.3781*.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013b). Distributed Representations of Words and Phrases and Their Compositionality. In *Advances in Neural Information Processing Systems:* 3111-3119.

Monroe, N. W., & Robinson, G. (2008). Do Restrictive Rules Produce Nonmedian Outcomes? A Theory with Evidence from the 101st– 108th Congresses. *The Journal of Politics*, *70*(1): 217-231.

Monteith, K., Carroll, J. L., Seppi, K., & Martinez, T. (2011). Turning Bayesian Model Averaging into Bayesian Model Combination. In *The 2011 International Joint Conference on Neural Networks*: 2657-2663.

Nyhan, B., McGhee, E., Sides, J., Masket, S., and Greene, S. (2012). One Vote Out of Step? The Effects of Salient Roll Call Votes in the 2010 Election. *American Politics Research*, *40*(5): 844-879.

Quinn, K. M., Monroe, B. L., Colaresi, M., Crespin, M. H., & Radev, D. R. (2010). How to Analyze Political Attention with Minimal Assumptions and Costs. *American Journal of Political Science*, *54*(1): 209-228.

Poole, K. T., & Rosenthal, H. (1997). *Congress: A Political-Economic History of Roll Call Voting*. Oxford: Oxford University Press.

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

Poole, K. T. (2007). Changing Minds? Not in Congress!. *Public Choice*, *131*(3-4), 435-451.

Rathore, M. (2018). gensim doc2vec & IMDB Sentiment Dataset. Retrieved from https://markroxor.github.io/gensim/static/notebooks/doc2vec-IMDB.html

Roberts, J. M., Smith, S. S., & Haptonstahl, S. R. (2016). The Dimensionality of Congressional Voting Reconsidered. *American Politics Research*, *44*(5): 794-815.

Roberts, M. E., Stewart, B. M., Tingley, D., Lucas, C., Leder-Luis, J., Gadarian, S. K., ... & Rand, D. G. (2014). Structural Topic Models for Open-Ended Survey Responses. *American Journal of Political Science*, *58*(4): 1064-1082.

Rohde, D. W. (1991). *Parties and Leaders in the Postreform House*. Chicago: University of Chicago Press.

Rohde, D. W., and Shepsle, K. A. (1973). Democratic Committee Assignments in the House of Representatives: Strategic Aspects of a Social Choice Process. *American Political Science Review*, *67*(3): 889-905.

Schattschneider, E.E. 1960. *The Semisovereign People: A Realist's View of Democracy in America*. Holt, Rinehart and Winston.

Schiller, W. J. (1995). Senators as Political Entrepreneurs: Using Bill Sponsorship to Shape Legislative Agendas. *American Journal of Political Science*: 186-203.

Sim, Y., Acree, B. D., Gross, J. H., & Smith, N. A. (2013). Measuring Ideological Proportions in Political Speeches. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*: 91-101.

Sinclair, B. (1999). Transformational Leader or Faithful Agent? Principal-agent Theory and House Majority Party Leadership. *Legislative Studies Quarterly*, 421-449.

Sinclair, B. (1998). *Legislators, Leaders, and Lawmaking: The US House of Representatives in the Postreform Era*. Baltimore, MD: Johns Hopkins University Press.

Sinclair, B. (1983). Purposive Behavior in the US Congress: A Review Essay. *Legislative Studies Quarterly*, *8*(1): 117-131.

Smith, S. S. (2007). *Party Influence in Congress*. Cambridge: Cambridge University Press.

Smith, S. S., & Gamm, G. (2001). The Dynamics of Party Government in Congress. In *Congress Reconsidered, 7th ed.:* 245-269. Washington, DC: Congressional Quarterly Press.

Sniderman, P. M., and Stiglitz, E. H. (2012). *The Reputational Premium: A Theory of Party Identification and Policy Reasoning*. Princeton: Princeton University Press.

Snyder Jr, J. M., and Ting, M. M. (2002). An Informational Rationale for Political Parties. *American Journal of Political Science*: 90-110.

Theriault, S. M. (2008). *Party Polarization in Congress*. Cambridge: Cambridge University Press.

Thomas, M., Pang, B., & Lee, L. (2006). Get Out the Vote: Determining Support or Opposition from Congressional Floor-Debate Transcripts. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*: 327-335.

Tibshirani, R. (1996). Regression Shrinkage and Selection via the Lasso. *Journal of the Royal Statistical Society B*: 267-288.

Turian, J., Ratinov, L., & Bengio, Y. (2010, July). Word Representations: a Simple and General Method for Semi-Supervised Learning. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*: 384-394.

Warshaw, C., and Broockman, D. (2017, June). G.O.P. Senators Might Not Realize it, but Not One State Supports the Republican Health Bill. *The New York Times*. Retrieved from https://www.nytimes.com/2017/06/14/upshot/gop-senators-might-not-realize-it-but-not-one-state-supports-the-ahca.html?

Wawro, G. (2000). *Legislative Entrepreneurship in the US House of Representatives*. Ann Arbor: University of Michigan Press.

Weingast, B. R., and Marshall, W. J. (1988). The Industrial Organization of Congress; Or, Why Legislatures, Like Firms, Are Not Organized as Markets. *Journal of Political Economy*, *96*(1): 132-163.

Wilson, R. K., & Young, C. D. (1997). Cosponsorship in the US Congress. *Legislative Studies Quarterly*, 25-43.

Wolpert, D. H. (1992). Stacked Generalization. *Neural Networks*, *5*(2): 241-259.

Woon, J. (2009). Issue Attention and Legislative Proposals in the US Senate. *Legislative Studies Quarterly*, *34*(1): 29-54.

Yu, B., Kaufmann, S., & Diermeier, D. (2008). Classifying Party Affiliation from Political Speech. *Journal of Information Technology & Politics*, *5*(1): 33-48.

Zhang, T. (2004). Solving Large Scale Linear Prediction Problems Using Stochastic Gradient Descent Algorithms. In *Proceedings of the Twenty-First International Conference on Machine Learning.*

Biography

Andrew Ojala Ballard was born on May 23, 1988 in Lansing, Michigan to Charles Ballard and Carolyn Ojala. He graduated from East Lansing High School in 2006 and earned a Bachelor of Arts degree in Psychology from the University of Michigan in 2010. Andrew began his Ph.D. research in fall 2012, and he received a Master of Arts degree in Political Science from Duke University in 2015. Andrew has published three papers:

Aldrich, J., Andrew Ballard, Joshua Lerner, and David Rohde (2017). Does the Gift Keep on Giving?: House Leadership PAC Donations Before and After Majority Status. *The Journal of Politics, 79*(4): 1449-1453.

Johnston, C., and Andrew Ballard (2016). Economists and Public Opinion: Expert Consensus and Economic Policy Judgments. *The Journal of Politics, 78*(2): 443-456.

Ballard, A., D. Sunshine Hillygus, and Tobias Konitzer (2016). Campaigning Online: Web Display Ads in the 2012 Presidential Campaign. *PS: Political Science and Politics, 49*(3): 414-419.

Andrew has received a number of honors and awards from Duke University, including two Dissertation Support Awards from the Program for the Study of Democracy, Institutions, and Political Economy (DIPE), a Research Award from DIPE, three Research Grants from the Department of Political Science, and a Summer Research Fellowship, Conference Travel Award, and Bass Instructional Fellowship from the Graduate School. He also received a Travel Award from the Southern Political Science Association and was selected for the Congressional Fellowship Program by the American Political Science Association. Beginning August 2018 he will serve as an Assistant Professor in the Department of Government at American University.