2022

A House's Speech Divided: Novel Applications Of Text-As-Data For The Study Of Elite Polarization In The U.s. House Of Representatives (1983-2016)

Jacob Pearl
University of Pennsylvania

A House's Speech Divided: Novel Applications Of Text-As-Data For The Study Of Elite Polarization In The U.s. House Of Representatives (1983-2016)

Abstract

Current models of elite polarization imply that the behaviors and ideologies of Democrats and Republicans have become increasingly distinct. The congressional roll-call voting record is the most relied-on indicator of congressional polarization, however, voting behavior is limited in its scope, ability to provide deeper insights into the nature of elite polarization, and can be affected by external non-ideological factors. This dissertation leverages the richness of the congressional record and introduces a flexible computational method, the dynamic topic model, to study three unique but related indicators of political polarization across three decades of debate from the floor of the House of Representatives (1983-2016). Using the output of the dynamic topic mode – and through the lens of political communication – this dissertation reveals patterns of increasing polarization in not only what Democrats and Republicans talk about, but also how political issues are discussed. Furthermore, this dissertation interrogates elite ideologies through belief network analysis and finds that the networks of political beliefs held by Democrats and Republicans have not significantly diverged since 1983. This dissertation introduces a novel approach to the study of political polarization in Congress and provides three applied use-cases for studying political polarization through text-as-data and relevant quantities to political communication.

Degree Type
Dissertation

Degree Name
Doctor of Philosophy (PhD)

Graduate Group
Communication

First Advisor
Yphtach Lelkes

Keywords
Computational Methods, Congress, Political Communication, Text-as-data

Subject Categories
Communication | Political Science

A HOUSE'S SPEECH DIVIDED: NOVEL APPLICATIONS OF TEXT-AS-DATA FOR THE STUDY

OF ELITE POLARIZATION IN THE U.S. HOUSE OF REPRESENTATIVES (1983-2016)

Jacob Maurice Pearl

A DISSERTATION

in

Communication

Presented to the Faculties of the University of Pennsylvania

in

Partial Fulfillment of the Requirements for the

Degree of Doctor of Philosophy

2022

Supervisor of Dissertation

Yphtach Lelkes, Associate Professor of Communication

Graduate Group Chairperson

John B. Jemmott III, Kenneth B. Clark Professor of Communication and Psychiatry

Dissertation Committee

Emily Falk, Professor of Communication, Psychology, and Marketing

David Lydon-Staley, Associate Professor of Communication

Daniel Hopkins, Professor of Political Science

ACKNOWLEDGMENT

Through the course of a five-year doctoral program, a student's primary objective is to develop their intellectual and methodological skills to become an expert in their field. Five years is a significant time in a person's life, and it is of course absurd to not also expect that a great deal of development occurs, and a great number of milestones are reached in other areas of life during that time. I am incredibly fortunate to have experienced and learned so many wonderful things in all aspects of my life during my time in grad school, and the growth that I've experienced would undoubtedly not have occurred without the support, guidance, kindness, and love of so many people.

To my loving wife, Julia. Thank you for the endless support you have shown me throughout this rollercoaster ride. You should make your millions writing a handbook for the partners of doctoral students; it can be called "Buckle Up: a guide to living with and managing a doctoral student." Your patience and understanding through my struggles, self-doubt, successes, and over-enthusiasm was invaluable to my completion of this program. You are my rock and my best friend, and I don't know what I'd do without you.

To my son, Wren. Although you are many years from being able to read this (and I doubt you ever will), I feel it is important to acknowledge the role you had in the completion of this dissertation. Your arrival was the kick in the rear I needed to finish this project, and your smile was the carrot at the end of every day to work hard and efficiently to see it through.

To my advisors old and new, Emily Falk and Yphach Lelkes. Thank you both for your support and guidance throughout all my time at Annenberg. Emily, thank you for taking me on as a wide-eyed cognitive neuroscience student, and allowing me to explore the many methods and ideas that excited me, and thank you for your generous advice and support as I transitioned away from neuroscience and into political communication research. Yph, thank you for taking me on as a fourth-year student with limited political communication research experience and helping me make that transition. Thank you for all the knowledge you have given me, and for not just making

political communication exciting, but also fun. I'm so thankful for your trust in my work, and your support of my research goals; every meeting with you fills me with optimism and faith that what I'm doing is worth my time and a contribution to our field.

To my colleague, teacher, confidant, and friend, Matthew Brook O'Donnell. The methods and advances made in this dissertation are in large part thanks to your innumerable contributions. Your availability to discuss all manner of methods and help me get over the many roadblocks on my way to this final product was invaluable. You are a brilliant scientist and researcher, and I am so grateful to have developed such a close working and personal relationship with you during my time at ASC.

To the late Emile Bruneau, whose initial trust in my research ideas meant so much to me in my first year at ASC. Thank you for opening my eyes to how science can have a positive impact in the world and specifically in our politically divided country. Thank you, Emile, for taking me on as a student in the first years of grad school, thank you for opening worlds up to me, and thank you for showing me and everyone around you what it means to care deeply about others whether they are known to us or not. Although we only worked together for a short time, I feel connected to you and your life's work in a profound way. I think of you often.

Finally, to all Annenberg staff and faculty. Thank you for all the kindness and support you show to all students. We are all so lucky to have our time at this school and are fortunate to have the opportunity to spend our days lost in heady conversations about information systems or staring at datasets that we hope hold the answers to our questions. The freedom to explore communication phenomena, however, we may define them, is something I am immensely thankful for. I don't know of any other program that offers as much flexibility as this one, and that is a true gift for all burgeoning scholars. Thank you all so so much.

ABSTRACT

A HOUSE'S SPEECH DIVIDED: NOVEL APPLICATIONS OF TEXT-AS-DATA FOR THE STUDY

OF ELITE POLARIZATION IN THE U.S. HOUSE OF REPRESENTATIVES (1983-2016

Jacob Maurice Pearl

Ypthach Lelkes

Current models of elite polarization imply that the behaviors and ideologies of Democrats and Republicans have become increasingly distinct. The congressional roll-call voting record is the most relied-on indicator of congressional polarization, however, voting behavior is limited in its scope, ability to provide deeper insights into the nature of elite polarization, and can be affected by external non-ideological factors. This dissertation leverages the richness of the congressional record and introduces a flexible computational method, the dynamic topic model, to study three unique but related indicators of political polarization across three decades of debate from the floor of the House of Representatives (1983-2016). Using the output of the dynamic topic mode – and through the lens of political communication – this dissertation reveals patterns of increasing polarization in not only what Democrats and Republicans talk about, but also how political issues are discussed. Furthermore, this dissertation interrogates elite ideologies through belief network analysis and finds that the networks of political beliefs held by Democrats and Republicans have not significantly diverged since 1983. This dissertation introduces a novel approach to the study of political polarization in Congress and provides three applied use-cases for studying political polarization through text-as-data and relevant quantities to political communication.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF ILLUSTRATIONS

INTRODUCTION

"Today, on this January day, my whole soul is in this: Bringing America together. Uniting our people. And uniting our nation. I ask every American to join me in this cause. Uniting to fight the common foes we face: Anger, resentment, hatred. Extremism, lawlessness, violence. [. . .] With unity we can do great things. [. . .] We can right wrongs." – Biden, 2021

These solemn words from President Joseph R. Biden Jr.'s inaugural address speak to the significant concern expressed among academics, politicians, and the American people alike that the U.S. is more deeply, and perhaps dangerously divided along partisan lines than it has been since the Civil War. These concerns are not simply political rhetoric or based on subjective observations. Rather, they are borne out of empirical evidence that Democrat and Republican politicians have become increasingly ideologically polarized and averse to inter-party cooperation (McCarty, Poole, & Rosenthal 2016; Sulkin & Schmitt 2014) - a trend that has contributed to significant legislative gridlock and the erosion of public trust in congress (Theriault, 2008). Indeed, congressional polarization is understood as a significant determinant of presidential success, use of cloture votes to end debate, use of the filibuster in the Senate (Bond & Messing, 2015). Furthermore, congressional gridlock leads to a substantial decrease in the public's trust in the legislative branch and; as a consequence, their willingness to behave in accordance with laws (Jones, 2015). Whether the public has followed their elected officials towards their ideological extremes is still debated (Fiorina, 2016; Abramowitz & Saunders, 1998), however the polarizing influence of political elites, their rhetoric, and partisan media messaging is substantial and in today's media environment seemingly unavoidable (Rodriguez et al., 2017; Peterson, Goel, & Iyengar, 2018; Bakker, Lelkes, & Malka, 2020; Levendusky, 2010).

Multiple lines of evidence indicate that since the 1960's congressional Democrats and Republicans have grown increasingly polarized (Coleman, 1997; Stonecash, Brewer, & Mariani, 2018; Theriault, 2008; Poole & Rosenthal, 1985). Perhaps the most relied on indicator of congressional polarization is the foundational work of Poole and Rosenthal and their model of legislative voting, DW-NOMINATE (Poole & Rosenthal, 1985), which compares the spatial

1

positions of legislators in a low dimensional projection of their voting behavior via

multidimensional scaling. Modeling the roll-call vote record of Senators and House

Representatives across the entire history of the U.S. congress, DW-NOMINATE finds that

Democrats and Republicans have not only moved further apart from one another in their voting

behavior, but have also become increasingly internally consistent (McCarty, Poole, & Rosenthal

2016; McCarty, 2016).
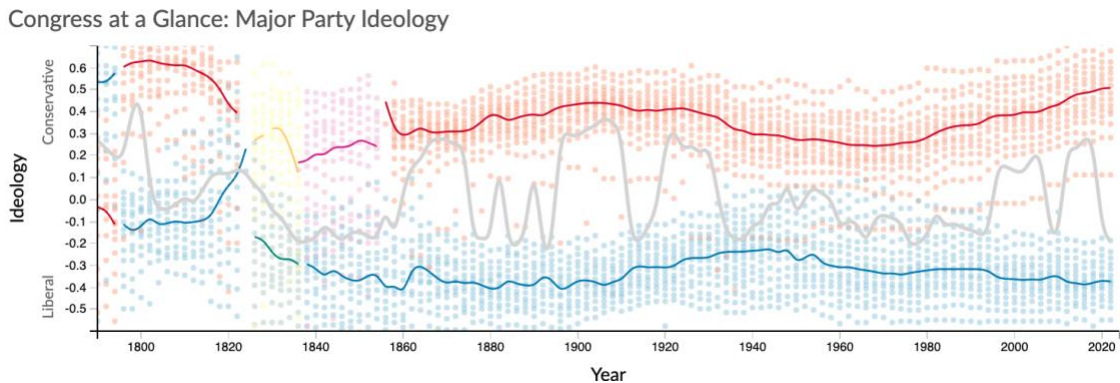
Figure 1: DW-NOMINATE Ideology scores 1800's to present day. (Voteview, 2022).

The division represented in Figure 1 are widely accepted as evidence of ideological polarization.

Roll call votes, however, provide only a limited perspective on polarization within congress.

Beyond the potentially non-ideological drivers of voting behavior – like logrolling or leadership

pressure (Sulkin & Schmitt, 2014; Bateman & Lapinski, 2016; Lee, 2016) – there exists the fact

that only a small fraction of bills ever even make it to the floor for a recorded vote. Indeed, of the

140,707 bills introduced since 2001, only 4.8% made it out of committee to receive a vote

(GovTrack 2021), As such, alternative resources ought to be turned to for richer and more

comprehensive information regarding the political thinking and behavior of the political parties that

contribute to polarization.

There is perhaps no source of information more abundant in today's world than written

and spoken language. In the narrower context of congressional politics, this is also true. Indeed,

written into the U.S. Constitution is a mandate that "each House shall keep a journal of its

proceedings", and since 1873 every word spoken in congress has been recorded, verbatim, in the *U.S. Congressional Record.* Within the Congressional Record, there is a vast amount of information relevant to the study of political polarization, including the expressed priorities, positions, perspectives, and attitudes of Democrats and Republicans across nearly all American history. This dissertation shows how using this large store of data reveals new and important insights about polarization.

Polarization & Political Communication Theory

The congressional record is a public resource, and so its contents are easily accessible. More difficult is the process of extracting from the record's raw text the relevant information necessary to study political polarization. Theories and research from political communication provide clear frameworks for understanding political rhetoric and how the words spoken by politicians can be used to measure congressional polarization. In this dissertation, three separate theories from political communication – agenda setting, framing, and belief network analysis – are used to understand congressional polarization.

When politicians discuss issues and promote specific causes, they are knowingly or unknowingly conveying to their audience what issues they care about, and perhaps what issues the public ought to care about too. This communication behavior is broadly known as *agenda setting* (McCombs & Shaw, 1972) and is generally understood as a significant driver of public opinion (Weaver, 1991; Camaj, 2019). The ever-changing agendas developed and communicated by politicians provide direct representation of their priorities and goals (Layman & Carsey, 2002), and within congress, the amount of attention given to issues is a valid index party agendas (Jones & Baumgartner, 2004). Polarization in this context is understood as a divergence in the agendas of Democrats and Republicans. *Agenda polarization* represents a growing divide in what congressional partisans care about and discuss.

Beyond variation in what politicians discuss, there are also significant differences in how they talk about issues. Partisan *framing* is yet another communication phenomenon known to affect the thinking and behavior of the public (Feldman & Hart, 2018; Scheufele, 2009; Hemphill,

Culotta, & Heston, 2013), and describes the process by which communicators affect the understanding or interpretation of issues by making salient specific aspects of issues (Entman, 1993). In Congress, issue framing occurs often (e.g. abortion is debated as pro-life vs. pro-choice), and is yet another way in which congressional polarization can be understood through political communication theory. Growing differences in how the parties discuss (frame) issues indicates a widening gulf in their shared understanding of issues. *Frame polarization* represents an increasing tendency for politicians to talk about the same issues but talk past each other.

Many cognitive models depict human knowledge as a web of interconnected, or associated concepts (Collins & Loftus, 1975; Anderson, 1983). In political science and communication, there is a long history of understanding political ideology in the same way (Converse, 1964; Gerring, 1997; Jost, 2006; Kalmoe, 2020). Several lines of research have enriched the understanding of elite and public knowledge by operationalizing belief systems as networks and applying methods from network science (DellaPosta, 2020; Boutyline & Vaisey 2017; Brandt, Sibley, & Osborne 2019) to understand them. In some research, these networks are constructed via the co-occurrence of issues in political messages (Vu, Guo, & McCombs 2014; Chen, Guo, & Su 2020), which ultimately represent the spreading activation of one concept to another (e.g. discussing guns and terrorism together represents an association).
Within congress, the patterns of issue co-occurrence in political speeches and debates can also represent cognitive associations. When these patterns of association are taken together across all of a party's speech, a party belief system, or ideology is approximated. Furthermore, *belief system polarization* can be measured as a growing dissimilarity in the structure of Democrat and Republican belief system networks.

To my knowledge, this is the first program of research to take advantage of the vast amount of information available in the congressional record to study elite political polarization and link agenda setting, framing, and belief network analysis to congressional polarization under a unified methodological framework i.e. text-as-data.

Text-as-data, models of political speech & polarization

In recent years, there has been a move in the social sciences toward computational approaches (Grimmer, Roberts, & Stewart, 2021), and a growing interest in applying machine learning and advanced statistical methods to the study of political language (Grimmer & Stewart, 2013). Recognizing the vast amount of text available for study and the ability of advanced methods to process and code many more documents than previously possible with hand-coding, political communication has seen a rapid rise in the application of computational methods for studying various sources of political language, including speeches (Diaf & Fritsche, 2021), news media transcripts and articles (Hagar, Wachs, & Horvát, 2021), open-ended survey responses (Roberts et al. 2014), online public discourse (Xiong, Cho, & Boatwright, 2019), and legislative debate (Peterson & Spirling, 2018).

One method that has seen significant development and use in political communication over the past decade is the application of topic models (see Grimmer & Stewart 2013 for an extensive review), which, put simply, are unsupervised machine learning models that "discover" latent topics within document collections. Various algorithms and implementations of topic models exist (e.g. Latent Dirichlet Allocation (Lafferty & Blei, 2009), Structural Topic Models (Roberts et al., 2013), Correlation Explanation (Gallagher et al., 2017), Non-negative Matrix Factorization (O'Callaghan et al., 2015)), but all essentially identify topics through patterns of term co-occurrence and clustering techniques. Topic models, unlike alternative content-coding methods, do not require any a priori knowledge about the topics to be discovered or how topics are described within corpora, and because topics are discovered inductively, also do not rely on any human coding input. As such, topic models are incredibly powerful for the analysis of very large bodies of text with few assumptions (Grimmer, Roberts, & Stewart 2021).

One area in which topic models have seen great success is in the study of elite rhetoric, and specifically the study of congressional leaders and their speech patterns. An incredible amount of text data is produced every day by U.S. Representatives and Senators, and scholars have applied topic models to measure and track issue attention both in and out of the halls of Congress. For example, the expressed agendas of partisans were tracked by Grimmer (2010)

using a topic model fit to one year of Senate press releases, and ten years of Senate debate was modeled by Quinn and colleagues (Quinn et al., 2010) to measure and track issue attention. In both studies, the models developed were found to be highly reliable and externally valid, and returned useful insights related to how politicians prioritize issues and communicate to their constituents.

Despite their proven success in extracting meaning from elite rhetoric, topic models are very infrequently used for studying elite polarization. Few studies exist in this domain, but several have proven that topic models can be successfully leveraged for comparative analysis (Hagar, Wachs, & Horvát, 2021). One such study was conducted by Hopkins, Schickler & Azizi (2021), who trained a topic model on over a thousand state party platforms between 1918 and 2017 to directly assess whether the positions of state parties have diverged and become more consistent across the nation. These authors identified 55 relevant political topics (e.g. labor, civil rights, abortion), and, more importantly, found that over time Democrat and Republican state party platforms diverged in their probabilities of mentioning certain issues. Further, the authors found that the polarization identified through this method matched well with additional accounts of issue divergence, lending external validity to the article's conclusions.

Expanding on previous studies of congressional language (Gentzkow, Shapiro, & Taddy 2016), the current dissertation considers over three decades of legislative speech to describe and understand patterns of elite political polarization. Unlike any previous unsupervised analyses of legislative speech in the US, the model developed here accounts for the ever-changing landscape of political issues by fitting a *dynamic topic model* (Greene & Cross, 2017; Blei & Lafferty, 2006). This model allows for the discovery of both broad topics that persist across time (e.g. economy, taxes) and niche topics that are relevant in only a subset of years (e.g. homeland security, disasters). Additionally, the dynamic topic model allows the terms which describe topics to shift across time, reflecting the changing nature of how issues are discussed, and allowing for more accurate document classification.

Both features of this model (dynamic topics and dynamic term descriptions) are necessary for describing the evolution of congressional polarization, as the agendas of the

Democrat and Republican parties and the frames used by the two parties to describe political issues are not and have not remained static. In this dissertation, I contend that with this single model it is possible to study elite polarization through the lens of political communication theory, and both the theories of framing and agenda setting. These two closely related lines of research provide clear frameworks for understanding the divergent evolution of what issues the Democratic and Republican parties have prioritized over time (agenda polarization) and how different aspects of more persistent issues and newer issues alike have been emphasized by the parties (frame polarization).

Further, I demonstrate that this model can also be used to address the question of ideological polarization within congress by interrogating the covariation of issues in the agendas of the political parties over time and formalizing these relationships as networks of issues. This network approach provides a very close approximation to the belief systems conceived by Philip Converse in his definition of ideology (Converse, 1964). Instead of strictly describing ideology along a single dimension on which legislators either fall to the left or right (Poole & Rosenthal, 1985), the belief system perspective describes ideology as a network of interdependent beliefs and policy positions. Ideological polarization, from this perspective, is simply defined as a divergence between the structure of party belief systems.

Overview Of Dissertation

In this dissertation, I study and describe elite polarization by analyzing approximately half a million speeches from the floor of the US House of Representatives from 1983 to 2016. This period was selected for two reasons. First, because previous studies of congressional polarization indicate a meaningful increase in polarization beginning in the 1980s (Voteview, 2022), and second, because this "daily edition" of the record is held apart from the remaining record due to changes to the format of the transcriptions after the 96th congress. The dynamic topic model developed in this research provides rich and meaningful measurements of attention to political issues and language use surrounding issues. Dynamic in this context refers to the

7

changing composition of said topics and the changing language within them over time. I infer the agendas of the Democrat and Republican parties from the relative attention given to specific issues and framing of issues through the interrogation of the language used within these issues. Finally, I construct association networks, or belief networks, by measuring the covariation of topics across legislators within parties; an operationalization that maps directly onto the ideological belief systems theorized by Converse.

In the first chapter of this dissertation, I introduce the need for the dynamic topic model approach and fit a dynamic non-negative matrix factorization model (DNMF) to 34 years of congressional debate. Validation checks are also provided in this chapter to provide confidence in the model's decomposition of the latent topic space and later use in the analytic chapters. Chapter 2 operationalized polarization as a divergence in the political agendas of Democrats and Republicans (agenda polarization) and indeed finds that over the past three decades what Democrats and Republicans talk about have become increasingly distinct. Recognizing that Democrats and Republicans also tend to discuss the same issues but vary in how they talk about them, chapter 3 operationalizes polarization as a growing divergence in how issues are discussed by the parties (frame polarization). This analysis also finds a growing divergence in the perspectives offered by Democrats and Republicans, and an increasing number of issues in which competitive frames are offered. Finally, chapter 4 interrogates ideological polarization by measuring the association of issues via belief network analysis; polarization in this chapter was operationalized as increased dissimilarity in the structure of Democrat and Republican belief systems (belief system polarization). In contrast to alternative accounts of elite ideological polarization and the findings from chapters 2 and 3, I find with this operationalization no differences between the belief systems of Democrats and Republicans and discuss potential explanations for this null result. Finally, I summarize these findings, with an emphasis on the methodological advances made in this dissertation, and present potential future directions for this research. This dissertation provides a framework for leveraging dynamic topic models and their output for comparative analysis in the area of elite polarization and aims to ultimately inspire future advances in this domain.

CHAPTER 1: DYNAMIC TOPIC MODELS OF CONGRESSIONAL SPEECH 1983-2016

Introduction

Legislative speech is a critical resource for the study of elite politics and political

polarization in the United States, as it is on the floor of the House and Senate that all manner of

national issues is raised, debated, framed, and ultimately decided on. Legislative language can

reveal many things about the parties, including what issues are of concern to Democrats and

Republicans, how such issues are understood and contextualized by lawmakers and their

coalitions, and potential areas of division within and between parties (Nguyen et al., 2015). Elite

rhetoric greatly influences public opinion (and vice versa) (Lenz, 2012), finding its way from the

floor of congress, to primetime news, and into the thinking and beliefs of the public (Harris, 2005).

As such it is also important to measure congressional speech and partisan language because

these quantities contribute to a greater understanding of communication flows (Zaller, 1991) and

the function and development of representative democracies.

Methodologies for the study of legislative speech content

Legislative speech is studied using a multitude of research methods, ranging from human

content-coding to inferential models of word-choice (Gentzkow, Shapiro, & Taddy; 2016; Jensen

et al., 2012), and advanced unsupervised computational models. For each approach, there are

both advantages and disadvantages that must be considered and balanced before a method is

used. Here, I briefly review several methods for studying legislative speech with a focus on

content-coding and computational approaches. Further, I indicate a need for a novel method that

can provide rich, externally valid, and meaningful insights relating to political communication and

polarization within the US Congress.

As mentioned by Quinn et al (2010), the most elemental method for deriving meaning

from text is to simply read it. From reading and subjective categorization of texts into types (e.g. "I

think this speech is about the economy and this other speech about national security"), derives a

more rigorous methodology for studying political content, human coded content analysis, which

refers to the process by which individual *coders* or teams of coders manually read and assign

texts to content categories which are assumed to be known to the researchers conducting the

analysis prior to evaluation (Krippendorff and Bock 2009; Ansolabehere, Snowberg, and Snyder

2003). The most well-known effort for analyzing legislative speech using manual coding is the

Policy Agendas Project (Baumgartner, Breunig & Grossman, 2019) which maintains a constantly

updated record of speeches and the policy categories they fall under. The Policy Agendas Project

is a reputable and frequently referenced database in political science, primarily because of the

accuracy and reliability with which it describes texts (John, 2006). The Policy Agendas Project,

much like any other research program based on human coding has notable costs though.

Specifically, manual coding efforts are time-consuming and costly, which often limits the scope of

research, and in the context of legislative speech may limit analysis to only small samples of text

instead of entire corpora.

Computer-assisted methods, like dictionary-based coding and supervised learning are

offered as potential solutions to the shortcomings of manual coding, and examples of their

application to legislative speech indicate that they are useful alternatives for identifying content

and meaning (Lupia, Soroka, & Beatty, 2020; Proksch et al., 2019; Martin et al., 2014; Jensen et

al., 2012). With both approaches, the number of documents able to be classified is increased by

orders of magnitude. This is accomplished by dictionary-based approaches using lists of

representative category-specific terms (dictionaries) developed by researchers, which are then

passed to computers to determine document categories; ordinarily, computer programs rely on

coding rules that compare the relative proportion of terms within a document that come from each

category. Supervised learning algorithms, on the other hand, *learn* the pattern of words within a

"training set" of documents, which contains a fraction of documents from the original corpus that

have been pre-labeled according to the coding scheme established by the experimenter. Once

trained, an algorithm can be applied to unseen documents to infer their categories.

Despite the increased utility of computer-assisted content-coding methods, these

approaches maintain that researchers already have in mind a set of categories to code and

exactly how to code them. As a result, these methods are dependent upon a researcher's

complete and unbiased knowledge of the contents of a corpus and the features within each category that best characterize them. For larger and more diverse collections of text, these assumptions may be increasingly difficult to satisfy.

To address these issues, scholars have turned to unsupervised hierarchical clustering algorithms like Latent Dirichlet Allocation (LDA) (Lafferty & Blei, 2009), non-negative matrix factorization, and other Bayesian methods (Grimmer, 2010). Topic models, as these methods are more generally called, identify the statistical co-occurrence of terms across a collection of text to identify latent "topics", which are in turn interpreted based on the distribution of terms that comprise them. Relying on these unsupervised models, the identification and labeling of legislative speech can be performed simultaneously. For example, Quinn and colleagues (Quinn et al., 2010) applied a topic model to senate speeches between the 105th and 108th congresses (1995-2004) and described multiple politically relevant topics, including one labeled *banking/finance* (terms: *bankruptci, bank, credit, case, ir, compani, file, card, financi, lawyer*) and another labelled *abortion* (terms: *procedur ,abort, babi, thi, life, doctor ,human, ban, decis).* Further, these authors indexed attention to these issues across this period with high fidelity and reliability.

Dynamic Topic Models

The advantage of topic models is their ability to process and label large bodies of text quickly and reliably. This feature allows for meaningful large-scale analysis, and the measurement of attention to issues over longer periods of time without the need to analyze only small samples, specify explicit coding schemes or annotate training sets. Previous applications of topic models to study congressional speech over large durations of time are, however, somewhat limited in the scope of inferences they can provide. More specifically, prior research has relied on models which assume a static set of political issues and a static set of terms to describe them, even as the reality of the political ecosystem is constantly changing. The previously described study by Quinn (2010), for instance, analyzed 9 years of legislative debate, but maintained a static set of topics, when it is possible that a variety of temporally specific and

11

important issues were debated but not captured by the model. Additionally, Sakamoto and

Takikawa (2017) fit a single topic model (k=70) to a much larger period -- three decades of

Senate debate -- making a potentially erroneous assumption that the issue environment across

these years remained stable. At issue here is the disconnect between the assumption of static

issues and the reality of a constantly shifting landscape of political issues and how they are

discussed. This assumption has real consequences for the classification of documents across a

corpus. Topic models derive topics by term co-occurrence, and as such speeches may be

classified as topics that are entirely inappropriate in relation to the context in which they were

given. For example, what is a researcher to make of a topic describing a speech in 1960 which is

described by the terms *internet, hacker,* and *security*? Is this a general topic regarding *national*

security that could be applicable for all subsequent years, or did this model assign the *internet*

security topic to a speech made in Congress long before the internet was popularized?

To address the reality of the dynamic topic space, several methods are available,

including the dynamic topic model (Blei & Lafferty, 2006), structural topic model (Roberts et al.,

2013), and dynamic non-negative matrix factorization (DNMF) (Greene & Cross, 2017). These

approaches, while each leveraging different methods, all allow the distribution of topics and the

distribution of terms within topics to vary across time, thus minimizing the types of errors

previously described. Greene and Cross, for example, studied 15 years of speeches made in the

European Parliament between 1999 and 2014. With their method (DNMF), the authors

discovered a variety of issues whose attention and content map closely onto the reality of politics

in the region. Topics were identified that persisted throughout the 15-year corpus (e.g. economy)

but whose terms changed as world events shaped them (e.g. 2008 financial crisis). More

transient topics were also identified, which accurately reflect the short-lived attention given to

certain issues like disasters and infrequent EU treaty reforms.

While dynamic topic models have been applied to the study of other legislative bodies, no

study has yet to study legislative debate in the U.S. Congress with this method. Additionally, while

previous studies have applied topic models to study polarization in Congress (Sakamoto &

Takikawa, 2017), no study has used dynamic topic models to study elite political polarization in

this externally valid way. In this chapter speeches from the U.S. House of Representatives between 1983 and 2016 were modeled using the dynamic non-negative matrix factorization method developed by Greene and Cross (2017). The current model includes 63 unique dynamic topics, whose content and presence in each year were allowed to vary, thus providing a more externally valid picture of legislative speech during each period. To establish the validity of the current model, several analyses were also performed here, including an analysis that shows that topics in this model track reliably with legislative world events.

The dynamic topic model developed in this chapter lays the groundwork for all subsequent analyses in this dissertation. As such, the validity established in this chapter provides a level of credibility and construct validity to all downstream results. The use of topic models to study political polarization is a relatively new approach; the analysis performed in this chapter is the first foundational step in the broader goal of this dissertation to advance this area of methodological development.

Materials and Methods

This dissertation seeks to understand political polarization in the U.S. House of Representatives between 1983 and 2016. Three separate operationalizations of polarization are proposed and studied in the subsequent chapter of this dissertation, including *agenda polarization*, the divergence in what Democrats and Republicans talk about; frame *polarization*, the growing divergence in *how* Democrats and Republicans talk about issues, and *belief network* polarization, which describes growing differences in how Democrats and Republicans associate political ideas. In this first chapter, I describe a dynamic topic model that was designed to extract meaning from the approximately 500,000 unique speeches made in the House over the 34 years under investigation. The output of this model is the foundation of all subsequent studies of elite political polarization in this dissertation. In the following section, I describe the corpus of political speeches used for this analysis (i.e. the congressional record), the steps that were taken to pre-process, filter, and prepare data for analysis, and the modeling techniques used.

Congressional Record

The congressional record is an extension of a constitutional mandate that "each House shall keep a journal of its proceedings" (Article 1, Section 5, Clause 3), and is a recording of all congressional debate, proceedings, and activities. The record, as an official government record, was first published for the public in 1873 and was preceded by several non-governmental entities' recordings of congressional speeches (Amer et al., 2001). The record consists of four sections: House, Senate, Extension of Remarks, and the Daily Digest. The House and Senate sections are meant to be "substantially verbatim" (Amer et al., 2001) records of all proceedings as they occurred, whereas the Extensions of Remarks section includes both speeches from each chamber and additional remarks inserted into the record that were not actually made on the floor. Such extensions include tributes, memorials, one-minute speeches, and additional extraneous materials that are beyond the scope of congressional business or exceed the one-minute speech rules of the House. The Daily Digest section of the record is situated at the end of the daily record and serves as a summary and index of the House and Senate business conducted that day. In the current research, the House section of the record is studied.

The congressional record is publicly available via Congress.org in both PDF and HTML formats. Gentzkow, Shapiro & Taddy, in their pioneering work on congressional speech, parsed and organized the congressional record in its entirety, and it is this pre-compiled dataset that is used in this dissertation. Briefly, I summarize the methods used by these researchers to prepare the congressional record for their studies and the future work of other scholars. Gentzkow et al. 2019, first began their processing of the congressional record by obtaining the digital text of the record from HeinOnline, which performed optical character recognition (OCR) on scanned versions of the print volumes from 1873 to 2017. The format of the record has changed slightly over time and so has how the record is published. The 43rd to 111th congressional records were published as bound editions, whereas the 97th to 114th were published as daily editions. Gentzkow and colleagues used the bound edition from the 44th to 111th congresses and the daily edition for the remaining congresses. The raw text of each record was parsed to

separate each speech and the authors identified the speaker of each speech; additionally, they

associated each speaker's name with their state and gender (for further details of the method,

see Gentzkow et al., 2019) [1].

In this dissertation, the daily edition (97th to 114th) was studied. This section of the record

was selected because it encompasses a period often characterized by rising polarization.

Additionally, the daily edition provides a consistent format (as opposed to the bound edition),

which significantly simplifies the process of preparing the corpus for analysis. Table S1.1 (see

appendix) describes the raw count of speeches for every congress between 1983 and 2016 (17

total) broken down by section (House, Senate) and party. This table includes only statements

made by Democrats or Republicans and does not include speeches from non-voting delegates.

Importantly, this table includes speeches prior to filtering, meaning that the counts shown include

brief procedural utterances by delegates as well as substantive statements. In total, across all 17

congresses there were roughly 1.2 million entries for the House and 1 million entries for the

Senate (300,000 entries in the Extensions section are not included).

Data Selection & Filtering

Because the congressional record includes all proceedings and speeches in both

chambers, a significant amount of the recorded statements are procedural in nature, or do not

capture language relevant for the analysis of legislative speech. In this section I describe the

criteria used for omitting speeches from the congressional record.

Omitting procedural speeches: First, only speeches made in the House of Representatives were

selected for analysis. Speeches were then omitted if they were determined to be procedural in

nature, as these statements both appear frequently and use of such language is unlikely

uninformative for the study of polarization (Gentzkow, Shapiro, & Taddy 2019). Within the House

---

[1] The party ID for each speaker is not included in the processed record provided by Gentzkow et al. To link each representative to their registered party, I matched all representatives' personal information (name, state, gender) to a database of US elected officials hosted by ProPublica.

and Senate there are a variety of procedural rules and traditions, which are executed through verbal requests or statements. These procedural utterances include, for example, requests for unanimous consent, scheduling, the yielding of time, motions to table amendments, and statements from the presiding officer. A list of procedural phrases was derived from the glossary of legislative terms (See appendix) found on congress.gov (https://www.congress.gov/help/legislative-glossary). Speeches were considered procedural and, following the methods of Gentzkow et al. (2019), removed from the record if 10% or more of the speech (after processing) was included in the list of procedural tokens.

Omitting speeches by length: Procedural speeches in the record could still survive the procedural language filtering if they happened to contain a high proportion of non-substantive language or were simply interjections made by speakers. Short procedural statements tended to occur when representatives gave minor statements before yielding their time. To remedy this issue, a lower bound of 50 words was used to further filter procedural language and interjections. This bound was established through manual tuning, in which values from 0-200 in increments of 25 were used. A lower bound of 50 removed the most procedural speeches without removing too many substantive speeches. A lower bound on the number of words required for a speech also benefits the later topic modeling procedures carried out in this dissertation, as such algorithms often perform poorly on short text (Yin & Wang, 2014). Table S.1.2 (see appendix) provides an updated final count of speeches by party in the House of Representatives. Approximately 600,000 unique speeches are included in the final corpus (591,390).

Text pre-processing

To prepare the congressional record for text analysis and topic modeling, I employed several common pre-processing methods for natural language processing. Pre-processing began with the normalization of all text within a single congress by lowercasing all text, removing all punctuation and numbers, and stripping the remaining white space. Additionally, all state names

were removed from the corpus to better capture more general issues that were not specific to any individual state. Following this procedure, each speech was tokenized into single word tokens and only certain parts of speech were maintained, including nouns, proper nouns, verbs, and adjectives. Maintaining only certain parts of speech serves two purposes: first, this procedure removes common "stop-words" like "the", "and", and other prepositions, pronouns, and articles. Second, this procedure helps to significantly reduce the size of the feature space for topic modeling while maintaining tokens that are rich in meaning.

All tokens were then lemmatized to reduce tokens to their base form. Lemmatization is an effective alternative to stemming procedures, which effectively converts terms to their dictionary form and removes inflectional endings to words. For example, lemmatization will convert the words "saw" and "seeing" to the word "see", whereas a common stemming algorithm like Porter's algorithm will simply stem all of these words to "s". Lemmatization has the added benefit of making the interpretation of processed text easier than stemming, while still reducing the feature space like stemming.

Following lemmatization, bigram tokens were generated based on their co-occurrence across the specific congressional corpus (e.g. all House speeches for the 99th congress). Including bigrams in topic models dramatically improves the interpretability of model results by capturing meaningful phrases (e.g. death tax, clean air) that may not be captured if their paired words were kept as distinct tokens. Longer n-grams were not selected for this analysis, as previous studies appear to show that bigrams are sufficient to capture meaningful variation in partisan speech (Gentzkow & Shapiro, 2010; Hopkins, Schickler, & Azizi, 2021). Bigrams were included in the corpus vocabulary if they were found in a given congressional corpus 50 or more times, and if their phrase score (Mikolov et al., 2013) exceeded a value of 10 (see appendix for formula). These thresholds were set based on pilot analyses using a single congress (110th) in which a range of values were manually assessed for quality. All occurrences of phrases in the corpus were replaced by a joined bigram (e.g. death tax became death_tax).

Speeches were finally split by year, resulting in 34 unique year corpora (e.g. speeches from the House in 1983, speeches from the House in 1984, etc.). Each of these subsets of

speeches were in turn converted into independent document term matrices (DTM). A document term matrix is the most common data structure used in natural language processing and is an effective means for numerically representing large collections of text (Grimmer & Stewart, 2013). Rows in a DTM correspond to documents in a corpus (here a speech), and columns represent tokens (here unigram lemmas or lemma bigrams). An element in a DTM indicates how many times a given token occurred in a speech. A DTM is not human readable, such that the order of words for a given speech is not represented in the matrix. Instead, the DTM is a numerical representation of the corpus to be used for computational analysis. A final step in constructing DTMs for each year and party was to remove highly infrequent and frequent terms; terms that occurred less than 2 times were removed, and terms that occurred in more than 30% of speeches were removed. The former procedure effectively removes highly uncommon language that is likely to not contribute to model fit, while the latter removes procedural language at the start and end of longer speeches (e.g. "madam speaker", "ask unanimous consent", "yield time"). The lower bound for this filtration is a common default for NLP pre-processing and removes tokens that only occur once in the corpus. The upper bound was selected after fitting several pilot topic models with a single congress (110th); a value of 30% effectively removed procedural language not captured by the original dictionary selection approach, and also returned more distinct and separated topics (a more stringent value led to less coherent topics and an overly sparse vocabulary).

After all processing steps were completed, 34 DTMs were returned representing all 575,608 speeches and a total vocabulary of 22,931 unique tokens. Meta-data for each speech was also maintained alongside the final DTMs, which included the name, party, and state of the speaker of each speech and the date on which the speech was given.

Dynamic Topic Models

Dynamic non-negative matrix factorization (DNMF) (Greene & Cross, 2017) was used to identify and cluster latent political topics across 34 years of congressional debate. DNMF

leverages non-negative matrix factorization (NMF) to identify topics within pre-specified windows of time and combines these individual "window topic models" via a second-level decomposition. As noted by Greene and colleagues, while LDA is often the preferred algorithm for topic modeling, NMF is a powerful alternative algorithm that has gained broader acceptance for its ability to simultaneously capture both broad and niche topics within large vocabularies (O'Callaghan et al., 2015), and for its substantially faster computation time.

Unlike LDA, NMF was not originally developed with natural language processing in mind. Instead, NMF is a general multivariate analysis method that has found broad application in a variety of contexts. Most often NMF is applied as a method for dimensionality reduction (Tsuge et al., 2001), although the algorithm can also be used for data clustering (Li & Ding, 2018), and classification analysis (Berry et al., 2007).

Briefly, NMF begins with a matrix V of size $n\ x\ m$ containing only non-negative values. In the context of natural language processing, V is a document term matrix, whose values represent term frequencies within documents. NMF attempts to then produce a reduced rank-k approximation in the form of the product of two non-negative factors, $V \approx WH;$ this is achieved by finding the matrix decomposition with the smallest error between the observed matrix V and its approximation $\hat{V}$ (Greene & Cross, 2017). The resulting W matrix, which is of shape $n\ x\ k$, represents the relative membership weight of all n speeches to each topic k, and the rows of this matrix can be sorted in descending order to identify the most important topics within a given document, Matrix H, which is of shape $k\ x\ m$, represents the term weights for every word in the corpus for each topic, and the rows in this matrix can also be sorted in descending order to understand the most relevant terms within topics and infer their meaning.

There are many options for minimizing this error term, and in this research, error was calculated using the Frobenius norm and minimized using coordinate descent (Cichocki & Phan, 2009). It is important to note that using coordinate descent means that the NMF algorithm is prone to convergence at multiple local minima which may lead to unstable model estimation. To improve model convergence and reduce the noise in model results, I borrow again from the

methods of Greene et al. and initialized all models with Non-negative Double-Singular Value

Decomposition (NNDSVD) (Boutsidis & Gallopoulos, 2008). Additionally, following trends in

applications of NMF for text modeling, the values in V were transformed before modeling via a

weighting function such as term frequency-inverse document frequency (TF-IDF), which has been

shown to improve the quality of model results (O'Callaghan et al., 2015).

As with all unsupervised topic model algorithms, a user-specified number of topics (or

components), $k,$ must be specified with NMF. A significant body of research has explored

automated methods for identifying the "optimal" number of topics for language models, including

extrinsic evaluation metrics like perplexity or held-out-likelihood, and intrinsic measures like topic

coherence. Extrinsic measures, however, often fail to correlate well with human judgment (Chang

et al., 2009), and intrinsic measures like topic coherence are also not always reliable and tend to

vary little when corpora are very large (O'Callaghan et al., 2015). Ultimately, the quality of topic

models are evaluated by humans and for humans (Grimmer & Stewart, 2013), and as such, it is

perhaps best that human judgment is relied on to determine the quality of models and how many

topics to include. Such evaluations must be made through careful reading of term distributions

within topics and careful reading of documents associated with topics (Maier et al., 2018; Quinn

et al., 2010).

DNMF applies NMF to identify topics within pre-specified windows of time and clusters

the topics from these window models to reveal broader topics that persist over time. In essence,

this process automates the manual process of matching similar topics across many topic models

assumed to have relatively similar content. Figure 1.1 provides a schematic of the DNMF

procedure and shows that this method begins first with identifying the windows of time to be

modeled. For each of these windows T_i a single NMF topic model, M_i, is run on the selected

documents resulting in a topic term matrix H_i with k_i window topics; these models and their output

are considered *layer 1* of the dynamic model.

For *layer 2* of the model a new matrix, V', is generated by combining all window topic

term matrices, such that each row of V' is the t top-ranked terms for every row of $H_1, H_2 \dots H_i$.

This new matrix has the same organization as a document term matrix, but instead of documents

as rows, contains topic term distributions. Just as in a DTM, V' is a sparse matrix where

"documents" or topics are expected to share common themes based on their term distributions.

Greene et al provide a clear breakdown of the procedure for generating V':

(1) Start with an empty matrix V'
(2) For each window topic model M_i:
  (a) For each window topic within M_i, select the 10 top-ranked terms from the corresponding row vector of the associated NMF factor H, set all weights for all other terms in that vector to 0. Add the vector as a new row in V'.
(3) Once vectors from all topic models have been stacked in this way, remove any columns with only zero values (i.e. terms from the original corpus which did not ever appear in the top-ranked terms for any window topics).

The resulting level 2 topic term matrix, V', can now be used to run yet another NMF topic model.

Here the resulting W' and H' matrices can be interpreted as such: the values in W' indicate the

strength of the relationship between each dynamic topic and window topic, and the values in H'

represent the most important terms that describe each dynamic topic.

To assign dynamic topics to window topics, the dynamic topic whose document

membership weight is highest is selected for each row of W'. Because the second level NMF is

blind to the time windows from which window topics come from, multiple topics within a single

window can be assigned to the same dynamic topic or dynamic topics can not occur in certain

windows at all. This is seen as a significant advance over other dynamic topic model algorithms

like the dynamic variant of LDA presented by Blei and Lafferty (Blei & Lafferty, 2006), as it

accounts for not only the gradual evolution of topics but also the birth, death, merging and

splitting of topics over time (Gropp et al., 2016).

Application of DNMF to the congressional record

In this research, DNMF was performed with 34 level 1 topic models, one for each year of

debate from the floor of the House of Representatives (1983-2016). NMF topic models were run

independently using the python library scikit-learn NMF decomposition function (Pedregosa et al.,

2011), with a NNDSVD initialization and 5,000 iterations. For every level 1 NMF topic model, a

TF-IDF transformed DTM was first created using the procedures previously described. This matrix was then fed to a model with $k = 45$ topics. To identify this number of topics a range of values for k were tested between 20 and 100 in increments of 5 for every window topic model. 25% of the topic models (9 topic models) were randomly selected for manual assessment, which involved a careful reading of the term distributions for topics within models and random samples of speeches.

Across the sampled models, $k = 45$ tended to return both broad and niche topics. Values of $k < 45$ often missed critical issues that were found in the larger models (e.g. guns, foreign aid) or combined issues that were found to be distinct topics in larger models (e.g. gender discrimination and abortion). Based on the results of these manual readings, $k = 45$ was selected as a default value for all first-level models. This value is selected for *all* level 1 models with an understanding that no value of k is likely to identify the "true" distribution of latent topics within a given corpus. Instead, $k = 45$ is selected for its likelihood to act as a meaningful approximation of the topics discussed in House speeches across the years in question. Later validation tests in this chapter provide further confidence that k = 45 for all years produced semantically coherent topics.

Once all level 1 models were completed, the top 10 terms were selected from each H_i and entered into the level 2 NMF model to generate dynamic topics. The appropriate number of dynamic topics was selected using the same methodology used with the level 1 models; after testing values of k' between 50-150 in increments of 10, a second-level model with 80 topics was found to capture both general topics that were persistent across time and niche topics that were specific to certain periods.

To interpret the results of this final model and make it usable for later analysis, the dynamic topics were labeled based on the 10 most strongly weighted words for each dynamic topic and from a closer reading of the 10 most strongly weighted words for each window topic
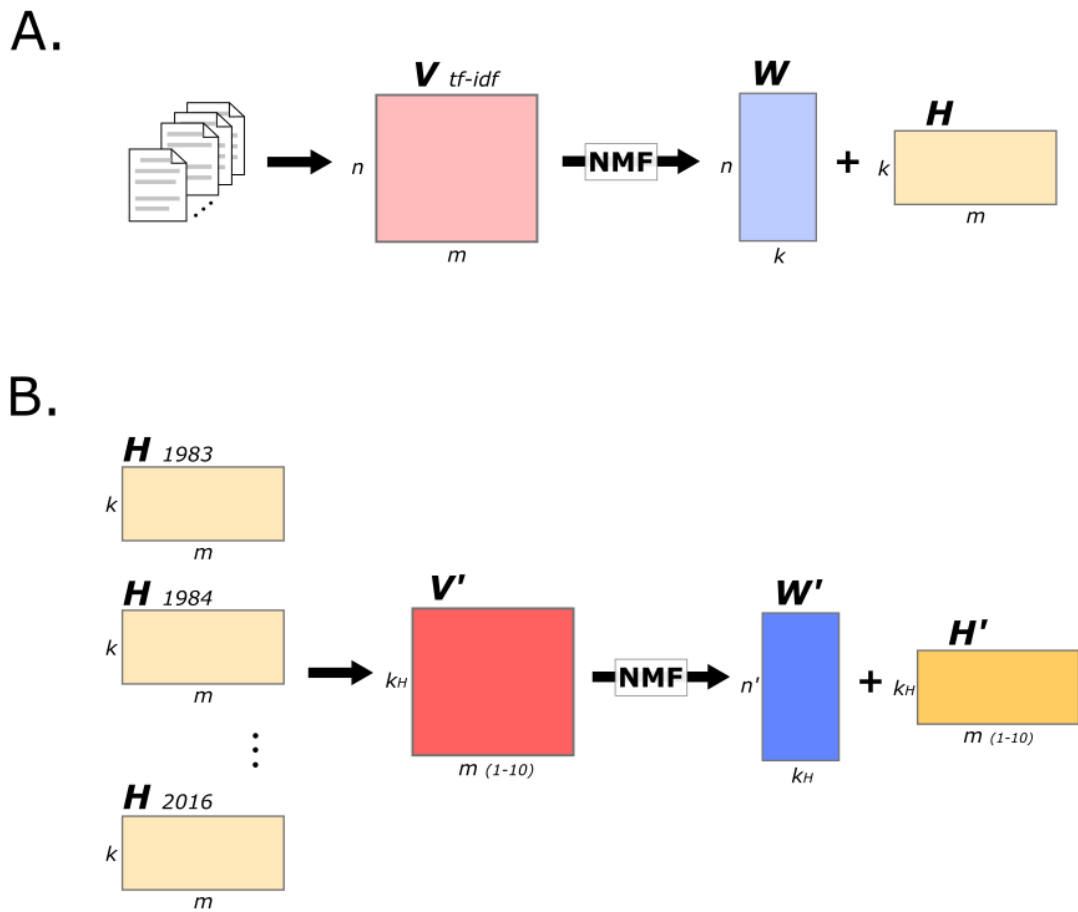
Figure 1.1: Dynamic Topic Model. A) NMF for window topic model, where tf-idf document term matrix *V* of size n x m is decomposed to matrices W and H. B) Dynamic NMF method, where top 10 terms for topics in each window matrix H are combined in topic-term matrix *V'*. *V'* is then decomposed using NMF into *W' and H'*.

assigned to the dynamic topic. Labels were roughly based on the major-issue and sub-issue codes provided in the Comparative Agendas Project (CAP) master codebook (Bevan, 2019). This codebook contains 23 major-issue codes including macroeconomics, defense, civil rights, environment, and more. Each major-issue also contains multiple sub-issue codes, which were used to guide the labeling of topics in this work. A small number of topics were repeated in the level 2 model (e.g. healthcare), however, a closer reading of documents in these topics either revealed greater nuance or represented the fact that these dynamic topics captured language

about the same topic at different periods. For example, three dynamic topics were labeled *taxes*, with one topic representing language about tax cuts and tax brackets (*tax_cut, wealthy, rich, tax_break*), another representing tax relief (*tax_relief, tax_credit, package, relief, tax_break, marriage_penalty*), and a third representing spending of tax dollars (*taxis, revenue, raise_taxis, tax, income, spending*). In later analyses, these topics are collapsed such that, for example, all topics related to taxes were grouped as a single *taxes* topic. Fourteen dynamic topics were not given meaningful codes, as their top terms were either very ambiguous, procedural in nature, or tributes to individuals or sports teams. A full list of dynamic topics is provided in the appendix.

The final step in preparing this analysis was to associate dynamic topics with window topics and further propagate these dynamic topic labels to individual speeches. Dynamic topics were assigned to window topics by finding the dynamic topic with the greatest weight for every row of W'. Next, window topics were assigned to individual speeches by finding the window topic with the greatest weight for every row of W_i. Maintaining a record of these assignments allows for a direct mapping of the dynamic topic to speech (e.g. speech 124 was labeled with window topic 5, which was assigned dynamic topic 34 - *national budget*). Once all dynamic labels were assigned to window labels and subsequently assigned to speeches, it was found that 3.3% were assigned to the ambiguous dynamic topics, 2.8% to a topic representing tribute speeches, and 9.8% were assigned to the procedural topic. Thus, approximately 16% of all speeches in this corpus were not given a meaningful political issue topic. After removing these 93,000 speeches, the corpus remained very large – at 482,608 unique speeches.

Model Output and Validation

The dynamic topic model developed in this research acts as the primary source of information for the remainder of the analysis and research conducted in this dissertation. As such, it is critical to establish that this model provides meaningful and sound measurements of political speech and attention to political topics. Quinn et al. (2010) give clear guidance on how to evaluate measurements derived from language models, and advise that "the evaluation of any

24

measurement is generally based on its reliability (can it be repeated?) and validity (is it right?)."

(Quinn et al., 2010). In this section, I describe several validation steps that were taken to ensure

that the measurements provided by the dynamic topic model are suitable for further study in this

dissertation. Specifically, I assessed the semantic validity (the extent to which topics have

coherent meaning) and external validity of this model (the extent to which measurements derived

from the model track external events) to build confidence in its downstream measurements.

Unlike hand-coded labeling schemes and human-assisted document labeling, which often suffer

from intercoder reliability issues, the labels assigned to documents by topic models are "100%

reliable, completely replicable" (Quinn et al., 2010). Thus, concerns about reliability can be set

aside and energy can be focused on establishing the validity of the measurements derived from

the dynamic topic model.

Semantic Validity

Semantic validity is best described as a measure of how well a categorization or content-

coding of text corresponds with meaningful interpretation, or "the extent to which the results of

coding correspond to how ordinary readers, or better still the [speaker] themselves would

categorize what they say" (Glazier, Boydstun, & Feezell, 2021). In the context of topic modeling,

semantic validity is assessed through the careful reading of topic descriptions (i.e. most

representative terms for a topic) and more importantly, through close readings of documents to

confirm their correspondence with their assigned topic. With dynamic topic models an additional

step is required, where the window descriptions assigned to dynamic topics must also be

assessed for their correspondence to the overarching dynamic topic to which they are assigned

(Greene & Cross, 2017).

Here, three topics (abortion, environment, taxes) were focused on to show that the

dynamic topic model developed in this research is semantically valid across both issues whose

content has remained static across time and issues that have taken on different language and

meaning over time. Dynamic and window topic descriptions were derived by taking the top 10

most strongly weighted terms for each topic. Similarly, documents were selected for assessment

based on their posterior probability of belonging to a given window topic; the top 10 most

probable documents within each time window respectively were selected for closer reading. A

more thorough and complete list of topic descriptions and documents can be found in the

appendix, and to further validate the entire model, interested readers can visit

https://jacobpearl.shinyapps.io/dtm_tracker/ to examine topic descriptions across the entire 34

years corpus.

Abortion: Women's access to abortion and reproductive healthcare is and has been a topic of

significant debate in the United States, and is frequently used as a wedge issue in political

movements, as it not only touches on questions of personal liberty, but also moral theology,

sexual morality, and gender roles (Jelen & Wilcox, 2003). The dynamic topic model in this

research identified an abortion topic that was represented by one or more topics in nearly every

year of the corpus starting in 1988. The dynamic topic description included the terms *abortion,*

family_planning, baby, pregnancy, procedure, mother, prolife, unborn_child, clinic, ban. Window

topic descriptions fit with this dynamic topic description, such that all topics included the word

abortion, and one or more additional contextual terms like *pregnant, baby,* and *family_planning*.

Table 1.1 shows the full list of window topics clustered into the abortion topic by the dynamic topic

model.

Some years included multiple unique topics that were clustered by the level 2 NMF model

under the abortion topic. For example, in 2011 two topics were included, which appear to

represent the two frames of the issue argued by Democrats and Republicans respectively. In

other years (e.g. 2015) multiple aspects of the issue were discussed; the dynamic topic model

also picks up this diversity of language and in these contexts also assigns multiple window topics

to the abortion topic within a given year. In reviewing the window topic distributions, it is clear that

at this level of analysis this topic passes any test of semantic validity.

Evaluation of the most probable speeches under the abortion topic revealed a similar

level of coherence and clear relation to the abortion topic. All speeches were directly related to

the issue of abortion access and funding, and their changing content over time reflected specific

debates occurring on the House floor. For example, in the late 80s speeches on the topic of abortion revolved around the passage of legislation that allowed the District of Columbia to use public funds for abortions, something that was directly opposed by then-President George H.W. Bush. In the early 90, the debate revolved around the Mexico City gag rule order, which barred NGOs from providing abortion counseling services. In 2015 various aspects of the abortion debate were discussed, primarily motivated by H.R. 7 or the "No Taxpayer Funding for Abortion and Abortion Insurance Full Disclosure Act of 2015". At the level of specific documents, the abortion topic again shows that the model produced here has strong semantic validity.

YEAR	WINDOW TOPIC DESCRIPTION
1988	abortion, district_columbia, rape_incest, mother, rape, pregnancy, tile, poor, human, residency_requirement
1989	abortion, family_planning, veto, district_columbia, unfpa, rape_incest, mother, population, dornan, pregnancy
1990	abortion, family_planning, romania, organization, district_columbia, unfpa, romanian, population, contraceptive, plan_parenthood
1991	abortion, family_planning, gag_rule, clinic, title, doctor, medical, patient, mexico_city, regulation
1992	abortion, family_planning, gag_rule, clinic, title, medical, option, madam_chairman, doctor, pregnancy
1993	abortion, clinic, family_planning, violence, title, hyde_amendment, pregnancy, baby, mother, reproductive_health
1995	abortion, family_planning, procedure, title, smith, organization, baby, federal_employee, population, prolife.
1996	abortion, procedure, mother, baby, doctor, medical, partial_birth, partialbirth_abortion, pregnancy, ban
1997	abortion, family_planning, smith, mexico_city, population, promote_abortion, release, prolife, perform_abortion, method procedure, abortion, mother, partialbirth_abortion, ban, doctor, medical, baby, fetus, exception
1998	abortion, contraception, federal_employee, contraceptive, procedure, family_planning, mother, pregnancy, cover, baby
1999	abortion, family_planning, unfpa, young, minor, smith, pregnancy, choose, baby, population
2000	abortion, procedure, baby, ban, mother, family_planning, pregnancy, partialbirth_abortion, doctor, medical
2001	abortion, family_planning, mexico, global_gag, gag_rule, hyde, pregnancy, promote, foreign, choice
2002	abortion, parent, hospital, minor, pregnancy, conscience, young_woman, parental_consent, legal, daughter procedure, partialbirth_abortion, ban, mother, baby, supreme_court, medical, exception, abortion, infant
2003	procedure, partialbirth_abortion, ban, mother, baby, perform, pregnancy, supreme_court, practice, lateterm_abortion
2003	abortion, partial_birth, unfpa, madam, supreme_court, china, mother, ban, family_planning, pregnancy
2005	parent, abortion, minor, young_woman, daughter, notification, parental_notification, young_girl, girl, interstate_abortion
2007	abortion, family_planning, contraceptive, mexico, organization, aid, plan_parenthood, lowey, hiv, international

2008	innocent, abortion_demand, mother, kill, unborn_child, die, human, unborn, sunset_memorial, wade
2009	abortion, plan_parenthood, prolife, baby, funding_abortion, pregnancy, unborn_child, family_planning, stupak, women
2011	abortion, funding_abortion, hyde_amendment, coverage, conscience, prolife, prohibit, baby, subsidize, taxpayer_dollar
	plan_parenthood, title, family_planning, women_health, clinic, abortion, organization, center, abortion_provider, pence
2012	abortion, baby, district_columbia, unborn_child, sex, pregnancy, wade, roe_v, fetus, pain
2013	abortion, baby, gosnell, unborn_child, clinic, unborn, kermit_gosnell, mother, murder, pain
2014	abortion, funding_abortion, hyde_amendment, coverage, private, prolife, unborn, taxpayer_dollar, women_health, mother
2015	baby, mother, video, unborn_child, paincapable_unborn, human, unborn, gosnell, womb, pain
	plan_parenthood, video, women_health, parenthood, clinic, health_center, investigation, defund_plan, healthcare, organ
	abortion, taxpayer, hyde_amendment, roe_v, medical, wade, pregnancy, doctor, prohibit, prolife
2016	abortion, conscience, prolife, participate, employer, religious, healthcare_provider, discrimination, boss, choose

Table 1.1: Window topic descriptions for topic "abortion"

Environment: Another salient political issue that has received much attention in the House of Representatives is environmental regulation and environmental protection. Over the past 30 years, political discourse and legislative action have focused on regulation and deregulation of polluters and the impact of energy production on ecosystems and public health. In more recent years discourse surrounding the environment has taken on a slightly different tone, focusing on the global impact of pollution. The dynamic topic model tracks the *environment* issue across the entire 34-year period under investigation in this dissertation, and the content of this topic reveals both stable and changing language.

The dynamic topic distribution includes the terms, EPA*, environmental, site, superfund, waste, cleanup, clean, environment, clean_air, and coal*. Across all topics, one of several critical terms is included – EPA*, environment, waste* – which indicate a clear relationship between the window topic and the issue of the environment. Additional time-specific terms were also revealed in the window topic distributions, such as *superfund* and *nuclear_power*, which relate to ithe nitial funding of nuclear cleanup projects in the late 80s and 90s. Later window topic distributions begin

to include terms like *climate_change*, which represent a growing interest in more global impacts of pollution.

Shifting attention to specific speeches tagged with the environment topic reveals additional evidence that the dynamic topic model is semantically coherent. Mirroring the information found in the window topic distributions, speeches in the 80s and 90s tended to focus on specific house resolutions to support the EPA in waste management and superfund sites. Speeches in the early '90s tagged by the model as relating to the environment also began to reflect a growing focus on air pollution (e.g. Clean Air Act of 1990) and featured notable speeches on the floor which emphasized important policies such as cap and trade. Finally, speeches from the latter portion of the corpus were found to be consistent with growing interest in the global impact of pollution and specifically mentioned emissions standards, global warming, and climate change.

Overall, the environment topic provides another piece of evidence that the dynamic topic model developed here is semantically valid. The window topic distributions associated with the environment topic reveal coherent and meaningful terms associated with both general and time-specific environmental issues. A closer reading of the speeches from each period lends even greater confidence that the content associated with this topic strongly corresponds to environmental issues.

Taxes: Finally, the ever-present political topic of taxation was selected for evaluation. There are many dimensions of tax policy that are discussed on capitol hill, and this is reflected in the fact that multiple dynamic topics were given the *taxes* topic. Specifically, three topics were given this label, one topic representing language pertaining to tax brackets (*tax_cut, wealthy, rich, tax_break*), another representing tax relief (*tax_relief, tax_credit, package, relief, tax_break, marriage_penalty*), and a third representing spending of tax dollars (*taxis, revenue, raise_taxis, tax, income, spending*).

Upon reading the window topic distributions for all three dynamic topics (see Table 1.3), it is clear that the dynamic topic model clustered window topics together coherently and

29

consistently. Almost all window topic distributions included the term *tax* or some form of the term

(e.g. lemmatized *taxis* and bigrams like *tax_cut*). The inclusion of terms like *rich, middle_class,*

and income indicate that one of the three topics consistently captures legislative debate

surrounding tax brackets and redistributive policy. The presence of terms like *surplus,*

middle_class, economic_growth, and *spending* also showed that this topic captures debate

surrounding how tax dollars are to be spent. Closer inspection of the most representative

speeches for this topic showed that much like the abortion topic and the environment topic,

speeches tagged by the model as *taxes* were closely, if not entirely, related to debates on the

House floor surrounding taxes.

YEAR	WINDOW TOPIC DESCRIPTION
1983	epa, hazardous_waste, waste, regulation, rcra, disposal, land_disposal, environmental, generator, small_generator
1984	superfund, site, epa, cleanup, hazardous_waste, clean, environmental, standard, toxic_waste, environmental_protection
1985	superfund, site, cleanup, epa, hazardous_waste, chemical, waste, compromise, clean, environmental
1986	superfund, cleanup, site, epa, conferee, hazardous_waste, clean, environmental, toxic_waste, chemical
1987	plant, nuclear, safety, nrc, accident, utility, liability, nuclear_power, markey_amendment, nuclear_powerplant
	research, ground_water, epa, acid_rain, water, contamination, environmental, technology, pollution, ssc
1988	medical_waste, waste, ocean, disposal, epa, ocean_dump, dump, beach, fee, penalty
1989	environmental, cleanup, clean, dod, compliance, site, waste, doe, priority, fine
1990	clean_air, acid_rain, emission, compromise, clean, air_pollution, air, energy_commerce, air_quality, health
	environmental, epa, environmental_protection, environment, cabinet, earth_day, pollution, wetland, status, elevate
1991	basis, base_closure, facility, foreign, recommendation, closure, domestic, overseas, list, department_defense
1992	facility, waste, wipp, site, test, environmental, doe, epa, disposal, safety
1994	regulation, epa, environmental, risk_assessment, regulatory, risk, compliance, environmental_protection, unfunded_mandate, analysis
	waste, flow_control, facility, solid_waste, local_government, disposal, landfill, municipal_solid, county, waste_management
1995	epa, water, clean_water, environmental, environment, clean_air, rider, pollution, clean, risk_assessment
1996	environmental, environment, environmental_protection, epa, republican_leadership, regulation, rider, forest, standard, clean_water
	superfund, clean, site, polluter, cleanup, epa, continue_resolution, markey_amendment, superfund_site, shutdown

1997	nuclear_waste, site, waste, compact, transportation, transport, nuclear, facility, department_energy, yucca_mountain
2002	nuclear_waste, yucca_mountain, site, waste, nuclear, repository, safe, shipment, nuclear_fuel, store
2009	capandtrade, electricity, clean_energy, coal, emission, price, carbon, oil, capandtax, environment
2010	estuary, water, bay, ocean, coastal, clean, great_lake, habitat, specie, fish
2011	clean_air, mercury, pollution, emission, air, toxic, study, source, pollutant, cement_kiln
2012	epa, coal, coal_ash, standard, permit, clean_air, power_plant, mining, regulatory, electricity
	fuel, yucca_mountain, waste, nuclear_waste, senator, nuclear, highlevel_nuclear, site, store, position
2013	coal, epa, coal_ash, war_coal, electricity, natural_gas, standard, plant, technology, production
2014	epa, clean_water, standard, epas, plant, coal, board, farmer, science, regulate
	facility, post_office, prison, chemical, cfat, postal_service, locate, guantanamo, site, building
2015	epa, epas, coal_ash, clean_water, science, regulatory, datum, clean_power, regulate, propose_rule
	climate_change, climate, emission, environment, clean_power, greenhouse_gas, paris, carbon_pollution, pollution, global
	coal, electricity, war, coal_ash, power_plant, natural_gas, coal_industry, mine, electric, utility
2016	epa, standard, clean_air, brick, air_quality, epas, chemical, pesticide, clean_water, regulatory
2016	coal, stream, mining, plant, waste, mine, mountaintop, coal_refuse, coal_mining, coal_industry

Table 1.2: Window topics descriptions for "environment"

The semantic validity of any content-coding tool is paramount; if the coding of texts is not consistent and the language included in coding schemes is not coherent, it is difficult to trust any downstream measurements derived from the categorizations. Here I have shown with three politically important examples that the dynamic topic model developed in this dissertation has strong semantic validity. All 80 of the dynamic topics included in the final model were manually evaluated for their coherent meaning, and all passed the same level of rigorous testing shown in the above discussion. Indeed, the close reading of window topic distributions and the speeches coded by the model was a key part of the labeling process itself. As such, the confidence shown in this discussion can also be ensured for the entire model.

External validity

Another approach taken to validate topic models is to match the measurements of attention provided by the model to external events (Grimmer, 2010; Quinn et al., 2010). Attention in the context of a topic model is operationalized as the relative frequency of documents classified

as a specific topic, and the validity of a model can be checked by matching spikes in attention to real-world external events. The utility of this approach is particularly high in the current context, as external events are often addressed or amplified by politicians in one-minute speeches and legislation can be, and often is, introduced in the house at the will of the majority party in response to external events.

The external validity of the dynamic topic model is made clear by the two graphs presented in Figure 1.2. The top graph depicts attention given to the *guns* topic from 1983 to 2016. This topic was chosen because the issue garners significant attention when devastating events occur like the 1999 Columbine High School shooting. Additionally, over the 34 years covered by this corpus, Democratic House members introduced multiple efforts to regulate the sale of firearms and were met with strong opposition by Republicans and the National Rifle Association, and Republicans pushed for the repeal of these bills. A major spike in panel A. of Figure 1.2 indicates significant debate in 1999 following the Columbine shooting, and significant debate can also be seen in 1991 with the introduction of the Brady Bill, which mandated federal background checks on firearm purchases. A major spike in 2016 also tracks with a widely publicized sit-in on the House floor conducted by Democrats in order to pressure Congress to act on combating gun violence, and in 2012 many representatives gave one-minute speeches to mourn the tragic Sandy Hook Elementary School mass shooting. Additional legislative events related to guns and firearms are identified and labeled in the first panel of Figure 1.2.

The bottom graph in Figure 1.2 tracks speeches related to *banking/finance*, and is targeted in this analysis for the expectation that significant attention should be given to it between 2008 and 2012 – a period marked by market crashes and significant debate surrounding

YEAR	WINDOW TOPIC DESCRIPTION
1983	spending, taxis, tax_cut, budget_resolution, revenue, interest_rate, cap, recovery, democratic, reduce_deficit
	withholding, repeal, income, interest_dividend, taxpayer, taxis, compliance, repeal_withholding, withhold, irs
1985	tax_reform, way_mean, taxis, revenue, tax_code, taxpayer, republican, deduction, income, state_local

Year	
1986	tax_reform, taxpayer, taxis, income, deduction, tax_code, investment, revenue, corporation, economy
	pledge, raise_taxis, sign, democrat, taxis, sign_pledge, walter_mondale, mondale, election, willing
1987	tax, revenue, taxis, deficit_reduction, grammrudman, reduction, raise_taxis, taxpayer, reduce_deficit, spending_cut
1988	tax, diesel_fuel, revenue, taxis, taxpayer, repeal, technical_correction, property, way_mean, tax_reform
1989	capital_gain, tax, ira, investment, taxis, reduction, income, tax_cut, rich, revenue
1990	rich, democrats, capital_gain, wealthy, income, democrat, tax_cut, raise_taxis, middle_class, poor
1991	package, economic_growth, democrats, capital_gain, create_job, democratic, democrat, growth_package, democratic_leadership, side_aisle
	taxis, income, rich, capital_gain, average, tax_cut, revenue, middle_class, poor, wealthy
1992	package, tax_credit, recession, unemployment, economic_growth, taxis, unemployed, capital_gain, create_job, incentive
1993	spending_cut, cut_spending, tax_increase, promise, raise_taxis, fee, national_debt, ratio, occur, reduction
1994	taxis, revenue, tax_increase, clintongephardt, raise_taxis, cut_spending, entitlement, promise, low, income
1995	tax_break, wealthy, rich, tax_credit, income, capital_gain, poor, gingrich, middle_class, corporation
	contract, promise, contract_america, keep_promise, item, package, senior_citizen, staff, unfunded_mandate, equity
1996	investment, pension, pension_fund, etis, invest, pension_plan, retirement, return, private, manager
	taxis, raise_taxis, clinton, average, income, earn, gas_tax, liberal, pay_taxis, tax_relief
1997	tax_cut, tax_break, wealthy, rich, gingrich, student_loan, medicare_medicaid, tax_credit, speaker_gingrich, dole
	tax_relief, relief, package, balance_budget, earn, liberal, death, middleclass, permanent, farm
	tax_credit, college, capital_gain, credit, tax_break, rich, wealthy, earn_income, tuition, rate
1998	tax_cut, rich, middle_class, wealthy, democrat, cut_taxis, deficit, democrats, earn, liberal
	tax_code, penalty, couple, marriage_tax, eliminate_marriage, married, marriage, income, married_couple, married_working
	taxis, raise_taxis, income, tobacco, sale, kid, low, revenue, balanced_budget, rate
	surplus, tax_cut, trust_fund, save_social, deficit, budget_surplus, tax_relief, balanced_budget, senior, debt
1999	tax_relief, retirement, relief, package, democrat, earn, liberal, lock, democrats, agenda
	taxis, income, raise_taxis, tax_code, tax_burden, earn, average, revenue, freedom, married
	tax_cut, wealthy, national_debt, irresponsible, rich, save_social, tax_break, deficit, interest_rate, pay_debt
2000	marriage_tax, tax_penalty, penalty, tax_code, eliminate_marriage, married, married_couple, wipe_marriage, marriage, fairness
	marriage_penalty, relief, tax_relief, married_couple, veto, income, estate_tax, marriage, penalty, democratic_alternative
	tax_cut, wealthy, budget_resolution, irresponsible, national_debt, rich, blue_dog, republican_leadership, target, massive
2001	tax_relief, taxpayer, taxis, relief, marriage_penalty, package, estate_tax, tax_credit, income, tax_code
2002	marriage_tax, taxis, tax_penalty, permanent, penalty, couple, eliminate_marriage, married_couple, marriage_penalty, married
	tax_cut, surplus, permanent, deficit, wealthy, bush, trust_fund, priority, recession, promise

2003	tax_credit, pay_taxis, millionaire, check, earn_income, earn, tax_break, republican_leadership, wealthy, income
	tax_relief, package, create_job, economic_growth, taxis, dividend, investment, job_growth, invest, stimulate_economy
2004	tax_relief, amt, tax_credit, relief, marriage_penalty, alternative_minimum, credit, middle_class, permanent, income
	taxis, tax_code, pay_taxis, income_tax, sale_tax, irs, corporation, page, raise_taxis, income
2005	spending, growth, taxis, saving, reduction, raise_taxis, tax_relief, defense, mandatory_spending, fiscal
	tax_cut, wealthy, rich, student_loan, millionaire, reconciliation, food_stamp, surplus, cut_taxis, middle_class
2006	tax_cut, wealthy, taxis, income, rich, average, deficit, revenue, middle_class, capital_gain
2007	taxis, tax_cut, revenue, raise_taxis, tax_relief, amt, rate, alternative_minimum, tax_credit, balance
	tax_increase, average, taxis, raise_taxis, single_large, farm_bill, budget_resolution, next_year, impose, save
2008	renewable_energy, tax_credit, conservation, energy_independence, incentive, comprehensive, invest, energy_efficiency, green, tax_incentive
	taxis, tax_increase, spending, debt, deficit, raise_taxis, tax_cut, balance, revenue, democrats
2009	package, tax_cut, infrastructure, economic_recovery, tax_relief, tax_credit, stimulate_economy, stimulus_package, rebuild, recovery_reinvestment
2010	hire, invest, job_creation, growth, infrastructure, recovery_act, tax_credit, recession, credit, private_sector
	tax_cut, middle_class, wealthy, rich, bush, income, tax_break, tax_relief, tax_credit, estate_tax
2011	tax_break, subsidy, end_medicare, medicaid, choice, wealthy, oil_company, big_oil, reduce_deficit, corporation
	middle_class, income, struggle, wealth, agenda, tax_relief, compromise, strengthen, back, payroll_tax
	tax_cut, revenue, bush, surplus, wealthy, extend, extend_payroll, rich, clinton, millionaire_billionaire
2012	tax_cut, bush, wealthy, income, sign, millionaire, revenue, tax_break, average, tax_credit
2013	deficit, taxis, create_job, revenue, middle_class, proposal, reduce_deficit, growth, tax_code, balanced
2014	permanent, tax_credit, credit, deficit, tax_code, tax_reform, offset, expire, charitable, bonus_depreciation
	internet, taxis, moratorium, permanent, taxation, free, fee, communication, grandfather, revenue
2016	tax, taxis, carbon_tax, tax_code, climate_change, infrastructure, oil, corporation, fairtax, revenue

Table 1.3: Window topics descriptions for "taxes"

regulation of financial institutions and the national debt – and between 1986 and 1993, a period encompassing the savings and loan crisis and subsequent legislative efforts to address corporate and market regulations. Indeed, such a pattern is seen in the data. Beginning in 2008 substantial debate can be seen on matters of banking and finance, and between 1986 and 1991 significant debate was found that reflected discussion of consumer finance and the savings and loan crisis. The model also picked up on spikes in debate in October and November 1991, in which the

House debated two substantial pieces of legislation, the Federal Deposit Insurance Corporation improvement act, and the financial safety and consumer choice act, which directly addressed the preceding market crisis. Various other House debates considering issues of banking and finance were also captured by the model and are labeled in the second panel of Figure 1.2.

Together tests provide even more confidence that the document clustering provided by the current model can be used for further analysis and a fruitful base of information on which to establish further investigations of partisan rhetoric and political attention.

Discussion

In this chapter, I introduced the dynamic topic model method and developed a semantically and externally valid model of congressional speech from 1983 to 2016. This chapter lays the groundwork for the rest of the current dissertation, providing the necessary quantitative and qualitative information for answering novel questions about congressional polarization. The topic model for this dissertation summarizes political speech in such a way that the development of polarization can be interrogated with further statistical models, natural language processing, and network analysis.

Unlike previous topic models of congressional speech (Quinn et al., 2010), the dynamic topic model allows topics to evolve, emerge and disappear over time by fitting multiple window topic models to pre-specified periods, and clustering their output (Greene & Cross, 2017). This approach was shown to be appropriate for the analysis of congressional speech over three decades, as political issues not only come and go over time but also take on different and time-specific frames or sub-issues in separate political eras. For example, the current model captured and grouped speeches about abortion access, and critically, revealed that the most relevant terms for the issue changed and reflected current debates as the topic persisted across the breadth of the entire time period included in this study.

Apart from the face validity of the dynamic topic model approach taken here (e.g. the topics identified by the model capture meaningful issues), the current model was also shown in this chapter to be both semantically and externally valid. This chapter showed the results of

several careful validation steps taken to ensure that the model output was coherent and reflected

the reality of the political landscape between 1983 and 2016. Specific topics – abortion,

environment, taxes, guns, and banking/finance– were focused on in this write-up, but all 80 topics

were manually assessed with the same level of scrutiny as those reported here.
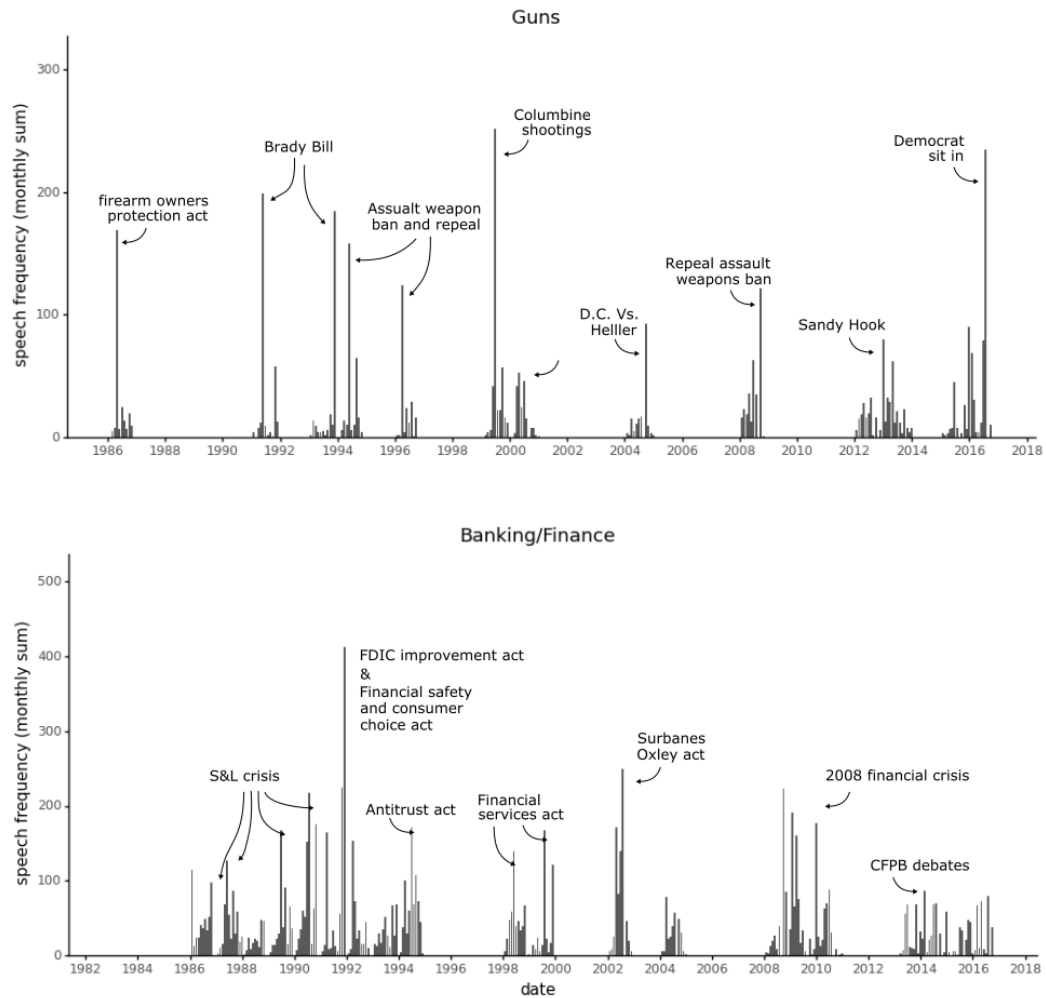
Figure 1.2: Topic attention to two dynamic topics between 1983 and 2016. The top panel shows attention to the Guns topic, and the middle shows attention to the banking and finance topic.

The current model, like any topic model, is limited in several ways that may impact further

analysis. First, the output of any topic model is dependent upon the number of topics specified

prior to their development. It is never the goal of a topic model to identify the "true" distribution of

all topics within a corpus, as the number of topics is truly dependent upon the level at which

analysis is conducted. At a coarse level, a political corpus may contain only two topics (social issues, economic issues), but as the specificity of issues is taken into greater consideration it is easy to see how the number of possible issues can grow substantially. The utility of a topic model to capture the information relevant to a specific question is of paramount importance (Grimmer, Roberts, & Stewart, 2021), however, a question can be answered from many different perspectives. As such, the current model is limited in that it includes only one of many decompositions of the current corpus; it is certainly the case that instead of 80 dynamic topics, a model with 20 or 200 may also perhaps return meaningful output. A fine balance must be found in the selection of the number of topics. Here, a model with 80 topics was found to capture a diverse set of externally valid and coherent topics. From the results of the validation tests performed in this chapter, the subsequent inferences in this dissertation can be trusted to be based on valid measurements of political attention and language.

CHAPTER 2: CONGRESSIONAL AGENDA POLARIZATION

Introduction

Critical to the work and success of politicians is the communication of issue priorities. Indeed, the effectiveness of a political campaign is partially driven by a politician's ability to emphasize and thus make salient particular issues that resonate with the public (Vavreck, 2009). In communication research, "agenda setting" describes the more general process by which communicating agents (often elites or the media) draw attention to specific issues and thus make them more "important" in the public's mind (McCombs & Shaw, 1972). Before the actual "setting" of agendas in the minds of the public, political leaders decide and establish which issues to emphasize (Scheufele, 2009).

The agendas of politicians and their parties are representative of their ideologies and policy goals (Layman & Carsey, 2002). The expressed priorities of legislators, and the differences between them, also provide unique information about potential divisions in all political arenas, including in the halls of Congress. Indeed, recent scholarship contends that individual legislators' expressed priorities, as indexed by "symbolic activities" like bill introductions and co-sponsorship, provide a truer representation of party polarization than the more often relied on measure of roll-call voting behavior -- as the former is often reactive, binary, and constrained by top-down pressures from party leadership, while the latter is proactive, varied in strength, and a bottom-up behavior driven by individuals' ideals (Sulkin & Schmitt, 2014). Furthermore, the expression of issue priorities is not limited by congressional procedure in the same way as roll-call voting, and thus captures more information regarding differences in party preferences. Within the House of Representatives, majority control of the legislative calendar limits what bills reach the floor for consideration, and many if not most bills never even receive a final vote. The "negative agenda control" exerted by the majority in the House, is often used to outright block the agendas of the minority party (Gailmard & Jenkins, 2007), thus obfuscating a great deal of information about partisan differences in policy priorities and agendas.

For the study of agendas and agenda polarization, speeches made on the House floor are also not subject to the same limitations as formal bill introduction or voting behavior. Through one-minute speeches and remarks given during morning business, legislators may voice concern for issues that may not be related to any to-be-debated legislation. In this way, representatives can circumvent the firm legislative control of the majority party to promote their party's agenda. In addition to these speeches, debate can also reveal more about the agenda priorities of the parties beyond their votes, as speeches on the floor, like other symbolic behaviors, are not compulsory or reactive like voting, and instead imply greater than average interest in the matter at hand (Sulkin & Schmitt, 2014). In this way, the amount of debate dedicated to an issue on the part of both supportive and dissenting members indicates what issues individual politicians, and their party find important, or in other words, include in their agenda. For example, the introduction of (and subsequent expressed support of) legislation to regulate the sale of high-capacity rifle magazines is a clear signal that a Democratic representative includes gun issues in her agenda. Similarly, fervent opposition to such regulations would also represent a Republican's inclusion of gun issues in their agenda and to what extent the issue is cared about as well.

Comparative analysis of legislative word choice further supports the proposition that language reveals important dimensions of partisan polarization that roll-call votes do not capture. After fitting a descriptive model of Senate speech between 1995 and 2014, Lauderdale and Herzog (2016) found that differences in word choice between Democrats and Republicans have increased faster than roll-call voting. Gentzkow, Shapiro & Taddy (Gentzkow et al., 2019), taking a much broader look at the entire congressional record, showed that although partisan differences in language track with the polarization shown by DW-NOMINATE, they are only moderately correlated. Differences in political speech at the level of word choice do not necessarily imply differences in party agendas though, and conversely, common verbiage does not imply agreement on issue priorities – the quantity of interest here. To return to the example of firearm legislation, both parties may prioritize this issue, but use very different language to discuss it (see chapter 3); one party evoking the second amendment and the other public safety.

Conversely, both parties may also discuss the national budget using the same terms, but one may disproportionately focus on the issue as a critical agenda item relative to the other.

In this study, polarization was operationalized as a divergence in what Democrats and Republicans discuss in their congressional agendas, a concept I refer to here as *agenda polarization*. Here, supervised machine learning was leveraged to develop a valid and robust model of partisan agenda polarization. Following the logic laid out by Peterson and Spirling (2018) in their analysis of 78 years of debate from the British House of Commons, agenda polarization is operationalized as the ability of a supervised classification model to distinguish between Democrat and Republican speeches based on the issues discussed within them. The better able a model can "learn" to differentiate between Democrat and Republican speeches, the more we can say that the contents of said speeches are polarized. The contents of speeches in Peterson and Spirling's work were described by word frequency distributions; this study uses the probabilistic distribution of topics within speeches. Thus, while the previous study measures polarization in terms of word choice, the current study measures it in terms of a very different quantity, issue attention. An additional advantage of focusing on issue attention as the unit of analysis here is that the coefficients from classification models can provide a picture of what issues were most divisive within a given year, an analysis that would require additional levels of inference if word-choice was used (i.e. what political issue does the word "attack" belong to?).

Besides diverging party agendas, another feature of polarization focused on is party consolidation. Indeed, one of the defining insights of DW-NOMINATE is that the Democrat and Republican parties have become notably more homogenous in their voting behavior. While previous research has integrated legislative speech and topic models with DW-NOMINATE scores (Kim, Londregan, & Ratkovic, 2015; Nguyen, 2015), no study to date has used topic models to specifically understand party agenda consolidation. Here the topic distributions of every individual legislator are compared pairwise to derive a measure of intra-party agenda consolidation.

A machine learning classifier was trained on the topic distributions derived from the topic model developed in chapter 1 to distinguish between Democrat and Republican speeches within

each year of congressional debate. Plotting the classification performance over time, I show in this chapter that indeed, between 1983 and 2016 the Democratic and Republican parties' agendas have meaningfully diverged. Additionally, I unpack the predictive coefficients of the models in this analysis to understand what topics were most important for differentiating between the two parties. The results in this analysis reflect historical legislative events, further validating them. I also show that during the period under investigation, the differences between individual legislators' agendas within both the Democrat and Republican parties did not change substantially, an indication of little party consolidation.

Materials and methods

To study the evolution of polarization, as operationalized as (dis)similarity in political agendas, I compared the topic probability distributions for speeches between Democrats and Republicans for every year respectively. More specifically, for every year, the corresponding W matrix (rows = speeches, columns = topics), was used as input to a simple classification model trained to distinguish between the pattern of topic probabilities for Democrats and Republicans. After each model was trained, the prediction accuracy was tested and interpreted as a measure of polarization. Put plainly, prediction accuracy here provides a measure of how distinguishable the overall distribution of topics are between Democrats and Republicans in any given year (Peterson & Spirling, 2018). The model is asked to predict the "most likely" label for a given speech based upon the learned pattern of topic probabilities across all speeches. Once the accuracy of each year was calculated, the pattern of agenda polarization across the 34 years under investigation was evaluated.

Next, I describe the methods employed in this chapter in more detail, and also describe additional analyses conducted to further elucidate the driving factors which have influenced partisan agenda polarization.

Data: Approximately half a million speeches from the US House of Representatives (1983-2016) were used to describe patterns of political polarization through the lens of agenda setting theory.

41

The dynamic topic model developed in chapter 1 was fit to identify 80 dynamic topics (63 of which were unique), whose substantive labels were linked to 45 window topics in each year of congressional speech respectively. The 45 window topics and their probability distributions for each speech are represented by the dynamic topic model with a W_i matrix, of size n_i x 45, where n_i represents the number of speeches in year *i*. For each speech, a probability distribution of length 45 describes the probability of said speech belonging to each topic.

As discussed in Chapter 1, a representation of political speech with a fixed number of 45 topics across years may not capture all political content (e.g. the number of topics may change from year to year such that multiple topics within a year may receive the same, broader dynamic topic label), however careful piloting and testing of models of varying size revealed that 45 topics described congressional speech well across the corpus. Additionally, the decision to fix the number of topics across years is advantageous for the classification approach used in this chapter, as varying numbers of features across models with large and sparse input can introduce bias and lead to spurious findings in a classification context (Gentzkow et al., 2019; Peterson & Spirling, 2018). It is also important to reiterate here that the dynamic topic model approach allows for the meaning of the 45 topics in each year to vary, such that topic 4 in 1995 will not have the same meaning as topic 4 in 1996. Rather, the dynamic topic model associates topics across (and within) years through a second-level decomposition (see chapter 1). Thus, the political issues in each year are not static and instead are allowed to vary across time, much like political issues do in the real world. Later inferences regarding the contribution of specific issues to model performance is performed through the projection of dynamic topic labels to window topic model coefficients (a processed described in detail in chapter 1).

Every row of W_i represents a single speech given on the House floor. For each speech, the corresponding author and their party affiliation was extracted from the speech meta-data, and paired with W_i. Party affiliation was then converted from its string form (e.g. D, R) to binary form, such that Democrat speeches were represented by 1 and Republican speeches 0. The proportion of Democrat and Republican speeches varied across years and was never perfectly balanced (see appendix). To account for class imbalance, which can negatively affect model training and

bias model fitting, a weighting procedure was employed during model training, where the class

(party) weight for each speech was given by $\frac{n}{2n_p}$, where *n* indicates the number of speeches for

that year and n_p indicates the number of speeches for class (party) *p.*

Machine learning classifier, party divergence: For each year, a least absolute shrinkage and

selection operator (LASSO) logistic regression was tuned and trained to distinguish between the

agendas of Democrats and Republicans. For each year respectively, the vector of binary class

labels was predicted using the feature matrix W_i. Models were developed and implemented using

the Python machine learning toolkit scikit-learn (Pedregosa et al., 2011). Stratified 10-fold cross-

validation was used to tune, train and test models for each year. The L1 regularization

hyperparameter was tuned using a range of values between 0.1 and 1, in steps of 0.1. The best

performing (based on average out-of-sample classification accuracy) model was then selected as

the output of this procedure and refit to the entire set of speeches to obtain feature weights and

draw inferences about what topics contributed most to the polarization seen in any given year.

The average out-of-sample classification accuracy for the best model was used to index

polarization in a given year.

Model Stability: Several steps were also taken to ensure that the classification accuracy results

for each year and the model coefficients obtained in this analysis were statistically significant, and

that results were robust to alternative modeling choices. First, a permutation procedure was

performed for every year to show that classification results deviated from random noise. This

procedure involved taking the best-performing model for a given year and refitting it 200 times on

a set of shuffled class labels. Next, to ensure that the results obtained here were not due to the

choice of the classification model, a linear kernel support vector classifier (SVC) was also fit for

each year using the same methods as before. Lastly, four alternative topic models with 20, 60,

80, and 100 topics were again fit using the LASSO methods described above. This step was

taken to show that the results obtained in this analysis are robust to alternative decompositions of the topic space in any given year.

Party agenda consolidation: To develop a measure of party agenda consolidation, the average distribution of topics for each representative was compared with every other representative within their party for a given year. This process involved first subsetting all speeches for a given representative and averaging the topic probability distributions across speeches. Next, the average topic distributions for every representative within a party were compared to every other representative within their party using the Jensen-Shannon (JS) distance metric, which measures the dissimilarity of two probability distributions of the same size and ranges between 0 and 1. Similarity was simply calculated as 1 minus the JS distance metric, and thus party consolidation was calculated as the average of all pairwise similarities within a party.

Results

Agenda Polarization: Figure 2.1 shows the classification accuracy for 34 regularized logistic regressions (LASSO), trained on the probability distribution of topics within a given year. The top panel of Figure 2.1 shows a steady rise in agenda polarization between 1983 and 2016, operationalized as the ability of a classification algorithm to reliably differentiate between Democrat and Republican speeches.

For every year, the model was able to reliably differentiate between the party agendas of Democrats and Republicans. More specifically, for every year the classification models performed above chance (balanced accuracy = 0.5). Model permutation tests revealed that the accuracy scores reported for each year were indeed nonspurious (see appendix for a version of Figure S2.1 with null models plotted as well Figure S2.2A), and after fitting several alternative models it was found that the inferences made here regarding agenda polarization were not the result of the choice of classification algorithm or decomposition of the latent topic space (Appendix Figures S2.3A and S2.4A). All trends observed in the main results were closely replicated with all alternative modeling decisions.

The first notable feature of Figure 2.1 is the steady rise in agenda polarization seen between 1983 and 2016. The most notable increase in agenda polarization is seen beginning in 2005 and continuing through 2011. According to these results, there was over a 10% improvement in model performance from 2005 to 2011, after which agenda polarization slightly decreased. However, differences in party agendas remained higher than in any previous period before 2005. These findings, and specifically the dramatic rise in polarization between 2005 and 2010 and subsequent decline, indicate a substantial divergence in the content of Democrat and Republican agendas, and this operationalization is consistent with not only previous research investigating trends in legislative speech polarization (Gentzkow et al., 2019) but also polarization indexed by DW-NOMINATE and roll-call voting. Additionally, the rise in agenda polarization observed between 1993 and 1995 is consistent with the roll-out of an aggressive conservative legislative agenda led by House Speaker Newt Gingrich known as the *contract with America* (Gingrich et al., 1995).

Intra-party agenda consolidation: Elite polarization is also characterized by party consolidation. Party consolidation was defined as the average similarity between topic distributions for every
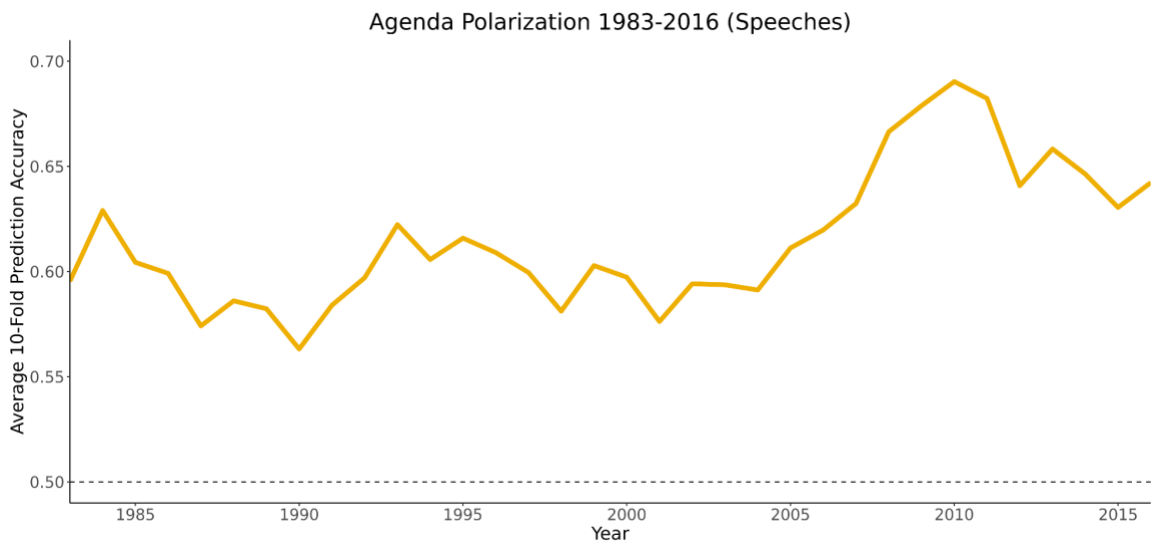

Figure 2.1: Agenda polarization operationalized as average 10-fold out-of-sample classification accuracy for partisan speeches.

legislator within a given party Figure 2.2 shows the average similarity between all legislators within both Democrat (blue line) and Republican (red line) parties respectively. Overall, this analysis revealed moderate agenda consolidation for both parties (average JS-divergence of 0.57), with a slight decrease in consolidation over time. A notable spike in party consolidation for both parties was found in 1995, again consistent with historical accounts of this periods.

Contribution of individual agenda topics: One of the advantages of the classification approach over previous distance-based methods (Sakamoto & Takikawa, 2017), is the ability to directly investigate the contribution of features to the fit of a given model. Indeed, one of the practical reasons for choosing a LASSO model (and linear kernel SVC) for this analysis was to maintain linear coefficients for ease of interpretation and to also perform feature selection. Here the most strongly weighted logistic regression coefficients (which survived regularization), were plotted in order to understand what topics contributed most to differences in Democrat and Republican agendas across time, and the direction of these effects. The feature weights from the 200 null models were also used in this analysis to check the statistical significance of all feature weights.

Focusing on periods of higher polarization, e.g. 1994, 1995, 2010, and 2011, Figure 2.3 presents a more detailed picture of what features were favored by Democrats and Republicans and the magnitude of these coefficients. The important contribution of these graphs is in their ability to show dramatic differences in how Democrats or Republicans favored certain political topics. For example, consistent with the content of the *Contract with America*, partisanship, taxes, and national debt were favored by Republicans in the mid '90s and these topics had a strong influence in differentiating between the parties. Further, the reaction of Democrats to negotiate tax policy in 1995 is reflected in the coefficients for that year. In 2010 and 2011 The Republican agenda also had an outsized effect on the estimates of polarization made here. Partisanship, unemployment, and health care were favored substantially by Republicans in these years. Again, these features are consistent with external events from the period, like the ardent opposition of Republicans to the Affordable Care Act and concern over unemployment in

the wake of the 2008 financial crisis. The results shown in these plots provide valid inferences about what issues cut most decisively between Democrats and Republicans. Confidence in the reliability of these inferences is derived from the fact that the coefficients discussed here all fall well below the threshold set for statistical significance ($\alpha < 0.01$).

Table 2.1 shows the top three largest positive (Democrats) and negative (Republican) topic coefficients (all statistically significant at $p < 0.01$) for every year between 1983 and 2016. It is important to note here that a strong positive or negative coefficient indicates a relative imbalance in topic prevalence for one party or the other and not an absolute prevalence. This fact means that a topic of importance to the agendas of both parties will not have a strong weight, as there is no meaningful variation in the topic prevalence between parties which contributes to differentiating between them. For convenience, I use the language of "favored" to describe topics with relatively greater attention by one party relative to another. Specific topics like civil rights, labor, and health insurance are commonly favored by Democrats throughout the corpus, while Republicans tend to favor topics like small business, taxes, and partisanship. The consistency of these regularly favored topics with common knowledge of Democrat and Republican policy preferences lends face validity to these results.

Discussion

The goal of this chapter was to use the output of topic models to investigate partisan polarization through the lens of institutional agenda building. Here the topic distributions for every speech given by House Democrats and Republicans from 1983 to 2016 were used as input to a classification algorithm trained to differentiate between the parties for each year respectively. Agenda polarization was operationalized as the ability of a classifier to correctly "learn" the difference between Democrat and Republican topic distributions and accurately label unseen speeches based on their topic distributions. In addition to this primary analysis, party agenda consolidation was also investigated in this period by calculating the similarity of topic distributions for every representative within a party.
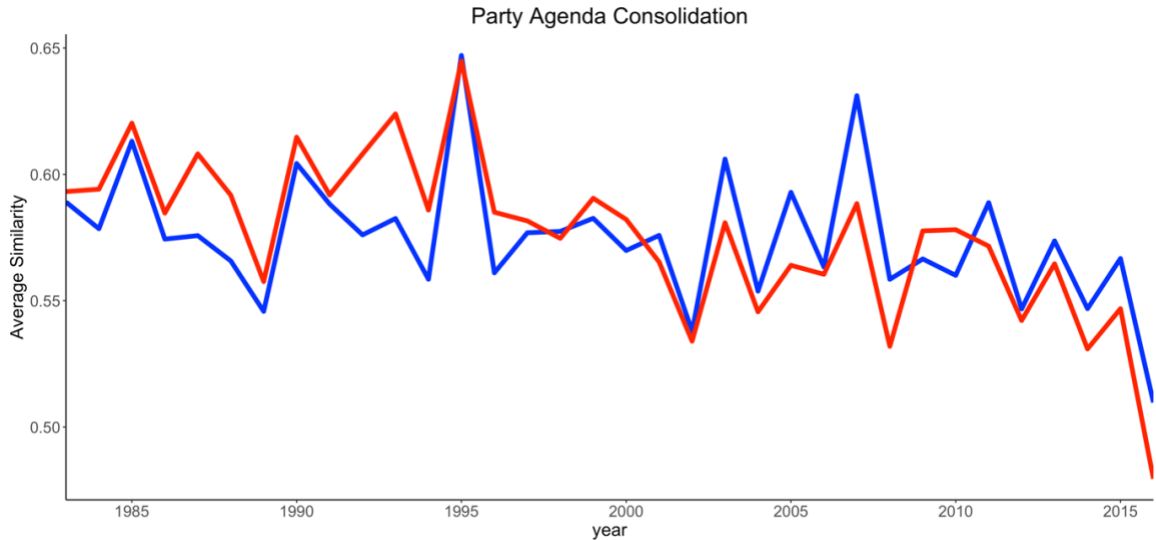
47

Figure 2.2: Party agenda consolidation for Democrats (blue line) and Republicans (red line) between 1983 and 2016. Consolidation is indexed by the average pairwise similarity between each legislator's topic distribution.

Using these analytic approaches, a pattern of inter-party agenda polarization and moderate intra-party agenda consolidation was indeed revealed. Between 1983 and 2016, the formal agendas of Democrats and Republicans became increasingly distinct, and thus more readily discernible. This was shown to be particularly true beginning in 2005 and peaking in 2010. Historical accounts of party agendas were also reflected in these results, as evidenced by a rise in predictive accuracy during Newt Gingrich's introduction of the *Contract with America* in 1994. Party consolidation also became stronger between 1994 and 1995, as evidenced by the dramatic rise in the average similarity between legislators within both the Democratic and Republican parties. Consolidation, however, was generally stable across time, with party members moderately consistent in their expressed agendas. These results have important implications for the understanding of not only elite polarization but also polarization in the mass public. Indeed, public polarization is in part affected by the expressed priorities of political leaders (Layman & Carsey, 2002), and so greater distinction between parties and more consistent messaging surrounding these differences has the potential to further drive a gulf between every-day Democrats and Republicans understanding of what matters

48
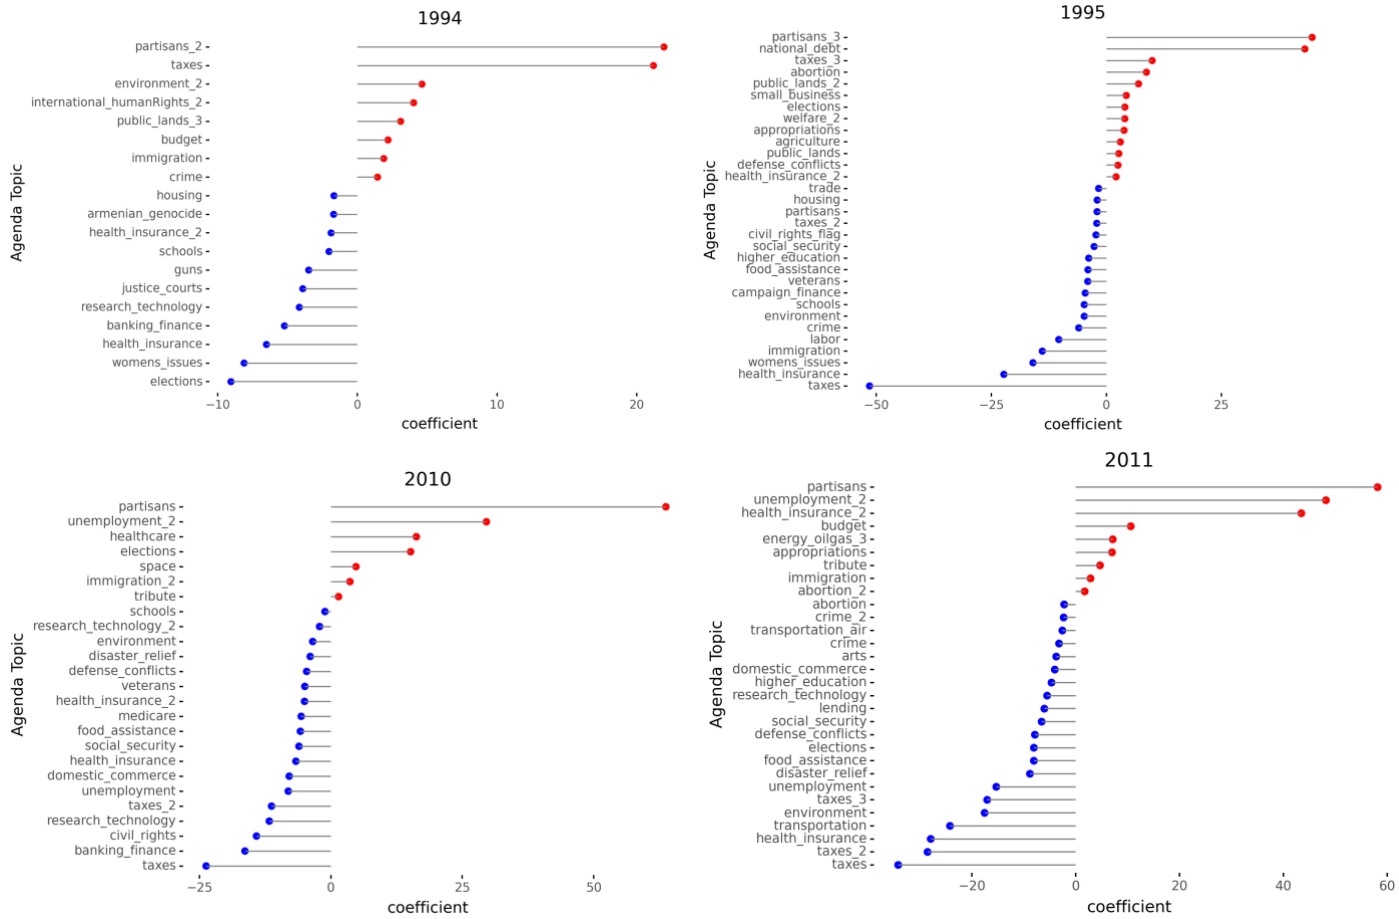
Figure 2.3: Model Coefficients for classifiers 1994, 1995, 2010, 2011. Larger coefficients in red indicate topics favored by Republicans which had an important influence on the performance of the model. Conversely, coefficients in blue indicate those topics favored by Democrats which had an important influence on the performance of the model (Direction of coefficients are flipped for consistency with the canonical left-right comparison between Democrats and Republican

In addition to the main results presented in this chapter, several supplemental analyses were performed to ensure that the results were valid and stable. A distribution of null models was developed for every year to establish that the main results were not spurious. These null models also allowed for the rigorous interrogation of model coefficients, and further inferences regarding the topics which drove model performance in each year. During model development, there are many choices an experimenter makes which can potentially bias results; to account for experimenter variables, a series of alternative models (SVC) and model inputs (different size topic models), were also fit to the data and evaluated for their similarity to the main results. The results of these alternative models were very consistent with those obtained in the primary analysis in this chapter (see figures in appendix). Regardless of modeling decisions, the same patterns of polarization were obtained. The validation procedures described above lend confidence to the results and inferences included in this chapter.

Despite the validation steps taken in this chapter, there remain limitations to the approach taken here that may impact future efforts. First, the inferences made about what topics contributed most to model performance are presented through my interpretation and naming of each topic. As with any topic model, topic names are given through close reading of their descriptive terms and subsequently labeled through human judgment. As such, the interpretation of model coefficients in this chapter is limited by the labels given to them. In addition to this, there are of course additional modeling choices that could have been made during classification, which may have still performed better and returned alternative results. For example, Peterson and Spirling include in their supplemental analyses an ensemble tree model, which performed better than their initial results (Peterson & Spirling, 2018). However, these authors showed that their boosted algorithm returned the same pattern of results as their main findings. The current study only implemented two competing models, but future work could continue in validating the current results with additional model choices. Finally, the current approach required substantial computational resources. For each year (34) 10-fold cross-validation was performed over tens of thousands of documents, and this procedure was again

Year	Democrat	Republican
1983	employment, defense weapons, unemployment	taxes, research tech, defense weapons
1984	employment, defense weapons, research science	partisanship, schools, constitution
1985	public land, social security, defense weapons	partisanship, elections, central america
1986	central america, civil rights, defense weapons	taxes, central america, budget
1987	employment, defense_weapons, unemployment	taxes, research_tech, defense_weapons
1988	employment, public land, labor	human rights, partisanship, human rights 2
1989	womens issues, environment, public land	health insurance, human rights, elections
1990	civil rights, higher ed, transportation	crime, drugs, campaign finance
1991	health insurance, unemployment, higher ed	budget, taxes, small business
1992	partisanship, human rights, welfare	partisanship 1, partisanship 2, banking/finance
1993	wages, budget, health insurance	partisanship 1, social security, partisanship 2
1994	elections, womens issues, health insurance	partisanship, taxes, environment
1995	taxes, health insurance, womens issues	partisanship 1, national debt, axes
1996	taxes, wages, guns	partisanship, taxes, small business
1997	health insurance, trade, campaign finance	social security, taxes, schools
1998	campaign finance, healthcare, schools	taxes 1, appropriations, taxes 2
1999	taxes, medicare, guns	taxes 1, partisanship, taxes 2
2000	taxes, public health, health insurance	taxes, appropriations, national debt
2001	medicare, air travel, public health	appropriations, taxes, small business
2002	taxes, national debt, health insurance	taxes, schools, disasters
2003	budget, unemployment, military conflict	taxes, appropriations, abortion
2004	human rights, unemployment, partisanship	appropriations, partisanship 2, taxes
2005	taxes, partisanship, national debt	taxes, appropriations, small business
2006	oil gas, partisanship, wages	immigration, appropriations, small business
2007	military conflict, budget, national debt	taxes 1, military conflict 2, taxes 2
2008	partisanship, unemployment, taxes	oil gas, taxes 2, abortion
2009	health insurance, lending, higher ed	partisanship, human rights, banking/finance
2010	taxes, banking/finance, civil rights	partisanship, unuemployment, healthcare
2011	taxes, taxes 2, health insurance	partisanship, unuemployment, health insurance
2012	health insurance, welfare, taxes	health insurance 2, elections, unemployment
2013	civil rights, research tech, partisanship	environment, healthcare, human rights
2014	unemployment, welfare, immigration	health insurance, environment, immigration 2
2015	civil rights, immigration, environment	abortion, courts, tributes
2016	guns, disasters, civil rights	small business, tributes, constitution

Table 2.1: Favored topics for Democrat and Republican agendas 1983-2016. Three topics represented for Democrats and Republicans indicate the three most favored topics in that given year for each party respectively.

performed 200 times for every year respectively. This analysis and validation procedure was performed on the Amazon Web Services cloud platform with parallel computing on 17 CPUs; the entire analysis with these resources took 10 minutes. Without these generous resources, the analysis would have taken much longer, and future research may take this into consideration before undertaking a study of this scale.

The utility of topic models to summarize and decompose legislative agendas is an already proven method (Quinn et al., 2010), and the use of model outputs to study polarization is further supported by the literature (Sakamoto & Takikawa, 2017). Further, the application of machine learning and classification analysis to measure polarization using legislative text is a valid approach (Peterson & Spirling, 2018). In this chapter, I show that these two methods can be linked to perform a comparative analysis in the realm of elite political polarization. With this method, I found that, indeed, Democrats and Republicans have polarized substantially in what issues they prioritize in congress. The topic models developed in this dissertation provide a reliable foundation on which to compare the legislative agendas of Democrats and Republicans. Future work may seek to expand this analysis to the entire congressional record, and both legislative bodies. Additionally, future work could leverage these methods to model and track party division and partisan realignment over time (e.g. southern Democrats in the 1960s). In this chapter, I lay the groundwork for such studies and provide evidence for the utility of this analytic approach by showing patterns of polarization between 1983 and 2016 that are highly consistent with alternative approaches to the study of elite polarization.

CHAPTER 3: CONGRESSIONAL FRAME POLARIZATION

Introduction

In the previous chapter, a model of congressional speech was used to measure and track partisan agenda polarization in the US House of Representatives across three decades. That investigation highlighted growing differences between the issue priorities of Democrats and Republicans, and further, identified issues that received disproportionate partisan attention at different points in time.

Partisan "issue ownership" is not ubiquitous across political topics, however, and more often polarization is manifest instead through competing perspectives on the same issues. To measure congressional polarization through political speech, it is thus necessary to not only consider *what* issues are attended to differently by the parties but also *how* the parties understand issues and present divergent perspectives.

Strategic messaging is a mainstay of partisan politics. Party leadership, media organizations, and other elites have, over the past several decades, honed the craft of directing public understanding of issues by carefully crafting issue-related messages, slogans, and stories (Druckman, 2001). Recognizing the impact of these communication behaviors, political communication scholars have characterized and measured these strategic decisions and their effects on public opinion under the general conceptual framework of *framing theory* and *framing effects* (Chong & Druckman, 2011). Formally, framing behaviors involve "select[ing] some aspects of a perceived reality and mak[ing] them more salient in a communication text, in such a way to promote a particular problem definition, causal interpretation, moral evaluation, and/or treatment recommendation" (Entman, 1993). In communication research, a looser definition of framing behaviors is often used, which describes framing as the process by which communicators emphasize a specific "subset of potentially relevant considerations" related to an issue (Druckman, 2001 p. 1042). In context, such emphasis frames are used in politics to direct public attention to specific elements of national issues and events and sway support for or against

53

legislation, politicians, or government actions. For example, immediately following the terrorist

attacks on September 11th, 2001, then-President George W. Bush and his administration

emphasized the immorality of the terrorist attacks to garner public support for a military campaign

in the Middle East: "This will be a monumental struggle of *good versus evil*, but good will prevail"

(Entman, 2003).

In the competitive arena of politics, framing efforts are, however, rarely unchallenged,

and often opposing partisans provide competing frames to promote alternative perspectives

(Sniderman & Theriault 2018; Druckman, 2001). Indeed, even the "war on terror" was confronted

with counter-frames as American involvement in the Middle East expanded. Skeptical of the Bush

administration's motives for invading Iraq, anti-war activists explained the move as motivated by

western economic expansion with a powerful slogan: "No blood for oil" (Colgan, 2013). Such

competing frames are indicative of differences in how issues are understood by individuals and

groups (Jiang et al., 2021). As partisan differences become more distinct, and more issues come

to be seen and discussed through different lenses, agreement and cooperation on how to

address issues are likely to become more difficult (Feldman & Hart, 2018). In this way, more

fervent frame competition across an increasing number of issues can also be understood as a

form of political polarization. Here, these patterns – greater differences in partisan emphasis

frames and a greater number of issues with competing frames – are together referred to as *frame

polarization*.

On the House and Senate floors, issue frames are frequently deployed during debates

and non-legislative speeches to contextualize bills, present frameworks for understanding party

positions, and in the case of minority parties, subvert majority agendas (Sellers, 2000).

Competition in Congress over *how* issues should be understood has historically spanned many

topics, ranging from abortion (pro-life versus pro-choice) (Esacove, 2004; McCaffrey & Keys,

2000), to immigration (Kang & Yang, 2021) (immigrants versus illegal aliens), tax policies (estate

tax versus death tax) (Lakoff, 2004) and healthcare (ACA versus Obamacare and death panels)

(Hopkins, 2018), For each issue respectively, party frames are deliberately crafted for maximum

effect and integration into party platforms. Indeed, since the introduction of public opinion polling

to congressional strategy, Republicans and Democrats alike have established organizations whose singular responsibility is to craft frames for debate and messaging (e.g. Democratic Message Group and Republican Theme Team) (Harris, 2005b 2005a).

Whether the Democrat and Republican messaging groups successfully move their members towards consistent frames on specific issues is an open question that is also relevant to frame polarization. As mentioned in chapter 2, previous estimations of congressional polarization emphasize the increasing consistency in legislative behaviors (Poole & Rosenthal, 1985). The internal consolidation of party behaviors further contributes to a widening gap and shrinking overlap between the parties. In chapter 2 party agenda consolidation was studied as an element of party polarization. Here, frame consolidation, operationalized as greater similarity in word choice between partisans in the same party, is also studied using the same logic. In addition to measuring polarization as a growing difference *between* party frames, it is also studied here as increased consistency in party frames *within* parties.

The many accounts of framing in the later part of the 20th century and early 21st century indicate that competing frames have become increasingly common across political issues. Additionally, partisan positions and frames appear to be both expressed with more distinct and polarizing language, and more coordinated within the parties (Brock, 2017; Gruszczynski & Michaels, 2012). No comprehensive analysis, however, has yet to rigorously test these disparate observations. Specifically, no study has yet to analyze a broad array of political topics simultaneously to answer whether competing frames have become more divisive and widespread across issues and whether legislators within parties have become more consistent in how they discuss issues.

Text-as-data studies of the congressional record have shown that the language used by Democrats and Republicans have become more distinct over the history of congress (Gentzkow, Shapiro, & Taddy, 2016; Jensen et al., 2012). Gentzkow, Shapiro, and Taddy modeled word choice for the entire congressional record as a function of party and showed that terms and phrases in congress do indeed take on partisan meanings, and when taken together show a pattern of polarization. Jensen et al (2012) revealed a similar pattern of word choice, but

additionally measured the relative intensity with which Democrats and Republicans used distinct language across time. As such, these authors revealed patterns of strategic messaging that track with important political events in Congress and changes in majority control. While foundational studies of the congressional record, both research programs modeled word choice without consideration of how language takes on different meanings in specific contexts. Such polysemy is important to account for in studies of political speech, as terms may be more or less polarizing based on what they are being used to discuss. For example "kill" may be more polarizing when Democrats discuss mass murders in their push for gun regulation, or when Republicans use it to express moral disgust when seeking to limit access to abortion services. In the previous studies, "kill" would receive only one meaning across all issues, while in this study issues are studied independently and thus the same term may take on multiple meanings.

Beyond methodological blind spots, these large-scale studies of congressional language generally neglect to link their findings to the much broader body of framing research. In this chapter, I conduct a similarly large-scale analysis of partisan word choice in Congress, but do so through the lens of framing. Using the topic assignments provided by the dynamic topic model developed in chapter 1, partisan speech in the House of Representatives was categorized into 63 unique issues, and differences in word choice was measured for each issue independently. Over 34 years of congressional debate, I find that partisan differences in issue-specific frames have on average grown larger and that this is in part a function of more issues taking on competing partisan frames. I also show that in the last three decades, the political parties have remained relatively stable in terms of encouraging their members to "get with the party message" (Lakoff, 2004) and that the frame consolidation has occurred more for Democrats relative to Republicans. Lastly, a supplemental analysis provided information about when parties deploy strong frames, showing that Democrat and Republican frames are more vigorously deployed when parties do not hold a majority in the House, a pattern consistent with previous accounts of partisan messaging in Congress and legislative competition (Lee, 2016; Sellers, 2000).

In the following sections, I present a detailed analysis of partisan framing in congress from 1983 to 2016, showing that the inductive analytic approach taken in this chapter allows for

both the analysis of broader trends in frame polarization and more specific interrogation of the words and terms used by Democrats and Republicans when framing issues.

Materials and Methods

In this chapter, polarization is operationalized as a divergence in what aspects of issues Democrats and Republicans focus on when engaging in political debate. Specifically, *frame polarization* is measured as a growing difference (cosine distance) in the term-frequency distributions of Democrats and Republicans when they discuss the same issue. Unlike previous models of partisan word-choice within congress, this study measures polarization within political issues as opposed to across all language use in each congress. This operationalization has several advantages over previous models, including that it accounts for polysemy (e.g. *death* may be used in discussing abortion, taxes, war, and retirement), and allows for the interrogation of frame polarization for each and every issue identified by the topic model developed in chapter 1 across time. Overall frame polarization for a given year is simply measured as the average frame polarization, operationalized as the cosine distance between party term-frequency distributions in a topic, across political issues in a given year. With an average frame polarization for every year in this dataset (1983-2016), a pattern of increasing frame polarization within the House of Representatives is shown.

Data: Speeches from the U.S. House of Representatives between 1983 and 2016 were used in this analysis. As the first step in this analysis, every speech was classified into one of 63 unique political issue topics derived from the dynamic topic model developed in chapter 1. Speeches were classified by simply identifying the topic for each speech that had the highest probability according to the topic model. As discussed in chapter 1, every topic in a year is associated with a dynamic topic that persists across years, and this dynamic topic label was used to label each speech in this dataset. As such, speeches across all years were categorized into a static set of topics across all 34 years under investigation (see chapter 1 for more details).

57

The base unit of measurement for this study was term frequency; all statistics were derived from these units. As in other chapters, not all terms spoken by politicians were included for this analysis, however, and were filtered in two phases. First, a list of procedural terms (e.g. *mr_speaker, amendment, table, minute, schedule, yield_time, etc.)* were removed from all speeches (see appendix for complete list), and a document frequency filter was employed which limited terms to those which occurred in at least 5% of speeches and which occurred in no more than 90% of speeches. This filter is meant to remove terms that are rarely used and too often used to likely be useful in diagnosing differences in topic-specific word use.

Following frequency filtering and procedural language removal a great deal of terms still remained in each collection of speeches, many of which were non-informative (e.g. *people, america, debate, etc.).* Following Jensen et al. (Jensen et al., 2012) and Gentzkow and Shapiro (Gentzkow & Shapiro, 2010), terms were ordered using Pearson's χ^2 statistic, which takes the form:

$$\chi^2{}_{tp} = \frac{(f_{tr} f_{\sim td} - f_{td} f_{\sim tr})^2}{(f_{tr} + f_{td})(f_{tr} + f_{\sim tr})(f_{td} + f_{\sim td})(f_{\sim tr} + f_{\sim td})} \, ,$$

where f_{tp} is the term frequency of term *t* for party $p \in (d, r)$ used in a given year-topic corpus, and $f_{\sim tp}$ is the frequency of all terms excluding term *t* used in a given year-topic corpus by party *p*. Pearson's χ^2 is used here because this statistic identifies terms that are very partisan while controlling for their frequency of use. Jensen et al provide a very powerful example of this method's utility:

> "[Suppose] Congressman Paul Ryan (R-Wisc.) mentions his daughter's full name once in the Congressional Record , it will be scored as a very partisan phrase because it will have been used only by Republicans - namely, Paul Ryan. However, it will have a low probability of being included in our restricted sample because it was said only once and thus does not have a very high probability of being Republican." (Jensen et al., 2012, pg. 10).

58

Terms with a $\chi^2 > 0$ were maintained in the corpus, and thus the final list of terms included in this analysis were those most likely to contribute to partisan framing for a topic in year *i*.

Once the final list of terms was selected, the number of speeches including term *t* for each politician was counted. The final product of preparing this data was a matrix of size *m x t* (and a list of party labels for each row)*,* where *m* represents the number of politicians who discussed the given topic in year *i*, and a second matrix of size *2 x t*, which represents the total sum of speech frequency for every word *t* for Democrats and Republicans respectively.

Frame Polarization: Frame polarization for every topic within each year *i,* was calculated using a simple measure of distance: 1 - cosine similarity, or *cosine distance.* The frequency distributions of terms for the Democratic party and the Republican party were used for this measurement, such that greater cosine distance indicated more distinct patterns of word choice. Cosine similarity was chosen as a measure of similarity between these vectors as the statistic is not affected by the magnitude of vectors, and it was the case in this study that the number of words maintained after filtering was not static within a topic across years. For every year-topic combination, cosine distance was calculated between the two party vectors, and within a year all topic values were averaged to produce a summary value of the extent to which partisans presented competing frames across their agendas.

To obtain p-values for every year-topic comparison, a permutation procedure was performed. For each test, party labels were shuffled for matrix *m x t* 200 times, the two-party sum vectors were calculated, and were tested against one another using cosine distance. P-values in this procedure represent the probability of obtaining a cosine similarity value equal to or greater than the true statistics by chance. Again, these null values were averaged within a year and used as null distribution for the average frame polarization measure for each year (see appendix Figure S3.1). Additionally, a linear mixed-effects model was also fit to this data; this model included a fixed effect term for time, and random intercepts and slopes for time estimated for each topic. A mixed effects model also made it possible to study the second component of frame polarization:

59

whether the number of competitive frames has increased. Coefficients in this model provide a linear estimate of the dissimilarity of language use between Democrats and Republicans between 1983 and 2016.

Party Frame Consolidation: Historical accounts of framing in congress indicate that both the Democrat and Republican parties have engaged in greater framing partly because of individual legislators aligning their framing of issues with the messaging developed by their party. Furthermore, polarization is understood by previous investigations of congressional polarization in terms of intra-party consistency. To investigate whether such a trend has developed, consolidation of party messages was studied by computing the average pairwise similarity of word choice within party for each topic. For every topic every year, the normalized word-frequency distributions for every legislator were calculated to account for differences in baseline speech frequency. The cosine similarity of every legislator's distribution with every other legislator in their party was then calculated, and these values were subsequently averaged to derive a measure of party message consolidation.

Term importance: Previous models of partisan word choice have offered alternative operationalizations of polarization using term frequencies. Jensen and colleagues (Jensen et al., 2012) performed their analysis by calculating an average term-party association, or *partisanship* value across terms in their corpus, which allowed for both a general overview of polarized word choice across time and a detailed look at what words were most strongly related to Democrats and Republicans respectively. The methods of Jensen et al. (2012) were replicated here to obtain directional estimates of term weights, or in other words, to identify what words for each year-topic combination were most likely to be used by Democrats and most likely to be used by Republicans.

The speaker frequency matrix *m x t* was first normalized for each word, such that each term frequency distribution across speakers had a mean of 0 and variance of 1. Next, Republican and Democrat labels were given similarly normalized numeric contrast codes of 1 for Republicans

and -1 for Democrats. Next, for every word a simple correlation was calculated. These correlation values were interpreted as such: a word with a correlation value of 0 indicates equal use among Democrats and Republicans, a word with a positive value indicates a higher frequency of use for Republicans, and a negative value indicates a higher frequency of use for Democrats. In this way, each word in a year-topic corpus was given a *partisanship* score. Further, taking the weighted sum of term coefficients within a topic provides a relative measure of how aggressively Democrats and Republicans employed partisan language (positive values indicate more Republican, negative values more Democrat), thus making it possible to identify *when* framing is most likely used by each party respectively.

Results

Frame Polarization: Operationalizing frame polarization as the dissimilarity of word use within topics, a clear increase in polarization from 1983 to 2016 was identified. Figure 3.1 shows the average polarization – as measured by cosine distance – between Democrats and Republicans for this period. Based on this figure, Democrats and Republicans began focusing on increasingly distinct aspects of political issues beginning in 1995, with a steady rise in frame polarization across all subsequent years. The steady increase in divergence of word choice post-1980 is also identified in prior research on partisan word choice in congress (Jensen et al., 2012), lending convergent validity to the current results. This paired with the reliability of the results through permutation testing procedures (see appendix Figure S3.1), lends substantial confidence to the trend identified here. A supplemental analysis (Figure S3.2) also reveals a similar pattern to Figure 3.1 when the median of topic distance scores is used to account for the sensitivity of averaging to outliers and provide additional evidence that political issues more generally have developed competing partisan frames over time. Additionally, the fixed effect term for time in the mixed-effects model indicates a statistically significant positive linear relationship between time and frame polarization ($\beta = 0.03$, *p* < 0.001).

Averaging frame polarization across all topics provides a general view of how competitive partisan frames have become more readily vocalized by Democrats and Republicans in

61

Congress, but for a deeper understanding of the nature of this type of polarization, analysis of the

individual topic polarization trends is needed. Indeed, it is an empirical question whether the

frame polarization observed in figure 3.1 is driven only by outlier topics, or whether competing

partisan frames have become more common across a larger number of issues over time. To

interrogate this question, the random effects of time on frame polarization (i.e. topic-specific
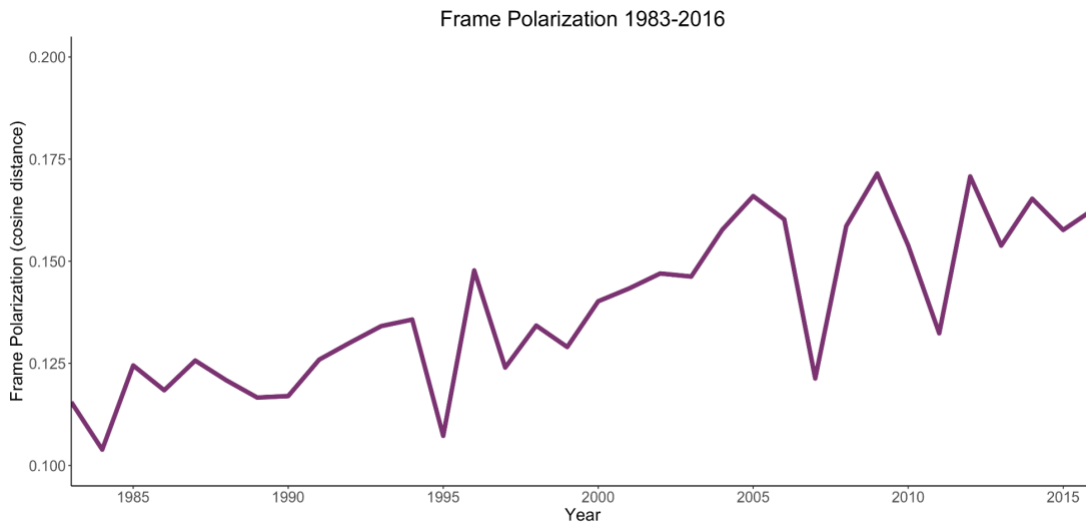
cosine distance)

Figure 3.1: Average frame polarization as measured by cosine distance.

were studied. Figure 3.2 shows these effects, indicating that indeed, the word choice for many of

the issues captured by the topic model has become more polarized over time. Many of the

estimated effects, however, show only a mild effect of time on topic-specific polarization, and a

set of notable and face valid issues show more substantial, but still mild frame polarization. The

current model estimated that Abortion showed a 95% increase in dissimilarity from 1983 to 2016,

guns a 70% increase, the national debt an 83% increase, and unemployment an estimated 93%

increase; all of these issues were estimated to have become significantly more polarized in terms

of how Democrats and Republicans discuss them. The remainder of the model coefficients also

have strong face validity, reflecting many contentious political issues. It is of course important to

note that frame polarization for most topics did not follow a linear trend over time. Instead, topic-

specific polarization ebbs and flows for different issues at different times (see appendix for time-course plots of all topics). The random-effects plotted in Figure 3.2, however, still indicate that for most issues there was an overall positive linear relationship between frame polarization and time over the 34 years studied here.
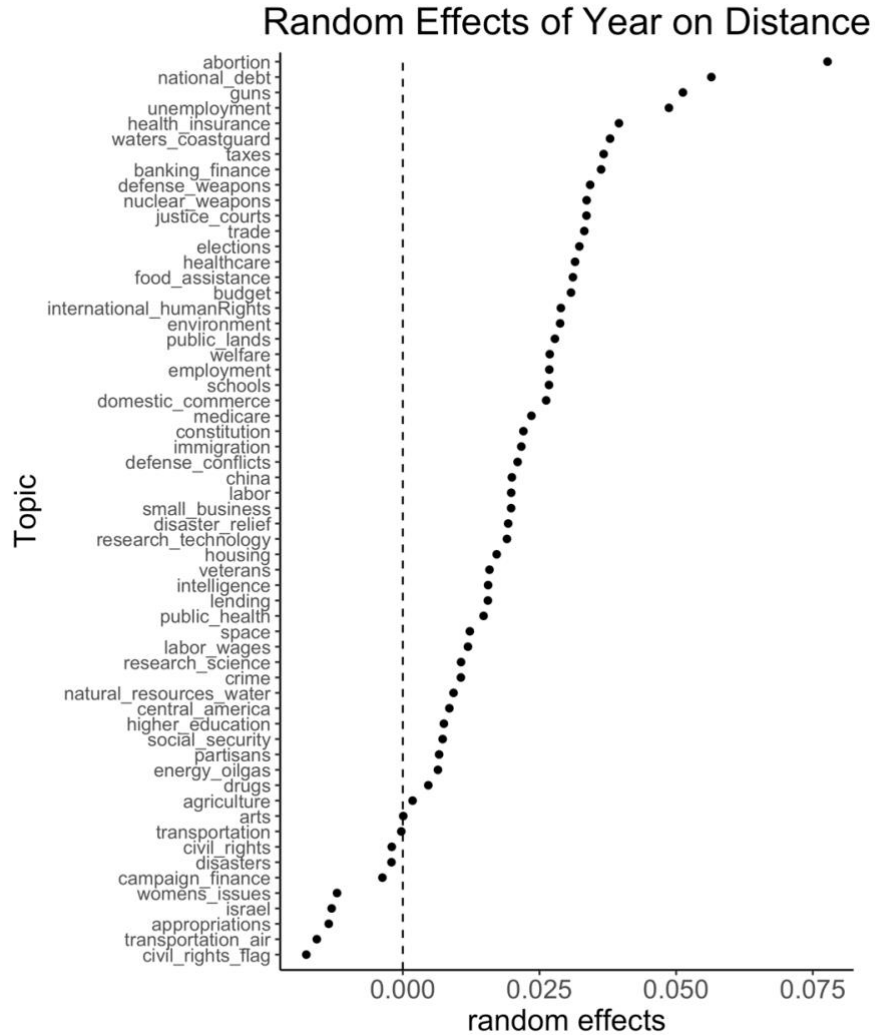

Figure 3.2. The random effect of time (year) on topic-specific distance.

Party Frame Consolidation. Next, I examine whether individual legislators within the Democrat and Republican parties have become more consistent in their framing of political issues over the past three decades. Overall, the answer here is mixed. Figure 3.3 shows the average frame

63

consolidation across all topics, where greater values indicate greater frame consolidation within parties. Figure 3.3 shows that for Democratic legislators issue framing has become more consolidated; For Republicans, across all issues, a notable decrease in consolidation was found after 1990, with a plateau beginning in 1995. A simple linear model confirms the results found in Figure 3.3, indicating that the interaction of party and time was statistically significant (β=-0.006, *p* < 0.001). Additional mixed-effects models were fit for Democrats and Republicans separately to estimate the topic-specific effect of time on party consolidation for each party. Consistent with Figure 3.3, Figure 3.4 shows that Democrats became more consistent in their language use for a wide variety of topics, while Republicans became increasingly inconsistent on most topics. Again, important to note in these figures is the scale of the estimated effects. While the pattern of results is consistent with Figure 3.3, the magnitude of these topic-specific effects is very small, ranging from an absolute value of 0 to 0.01 (cosine similarity has a range of 0 to 1). Such small linear effects indicate little change in consolidation for individual topics, however, it is again important to note the non-linear trends observed in frame consolidation for the two parties, and thus the limitations of linear models in fully explaining the development of frame consolidation over time (see appendix for issue-specific time courses).

Term Importance and Partisanship: Using the methods of Jensen et al (2012), term partisanship was identified within each topic within a given year. This approach allowed for the current investigation to go one step deeper in understanding the growing frame polarization shown in Figure 3.1. Specifically, this approach made it possible to identify not only what topics attracted the most frame competition in a given year (ranked cosine distance), but also what were the specific terms used by Democrats and Republicans in relation to the issues. Table S3.1 in the appendix shows the top three most polarized issues for each year, and the most left- and right-biased terms used for each topic respectively. Of note is the continued appearance of the *abortion* topic in the top three most polarized issues. Terms such as *access, decision, family_planning, constitutional_right,* and *woman* appear in the list of terms most associated with

64

Democrats, while *womb, unborn, baby, kill,* and pro*life* are used by Republicans. Additionally,

consistent with popular accounts of US partisan

framing (Lakoff, 2004), *taxes* appeared multiple times as a topic that attracts competing partisan

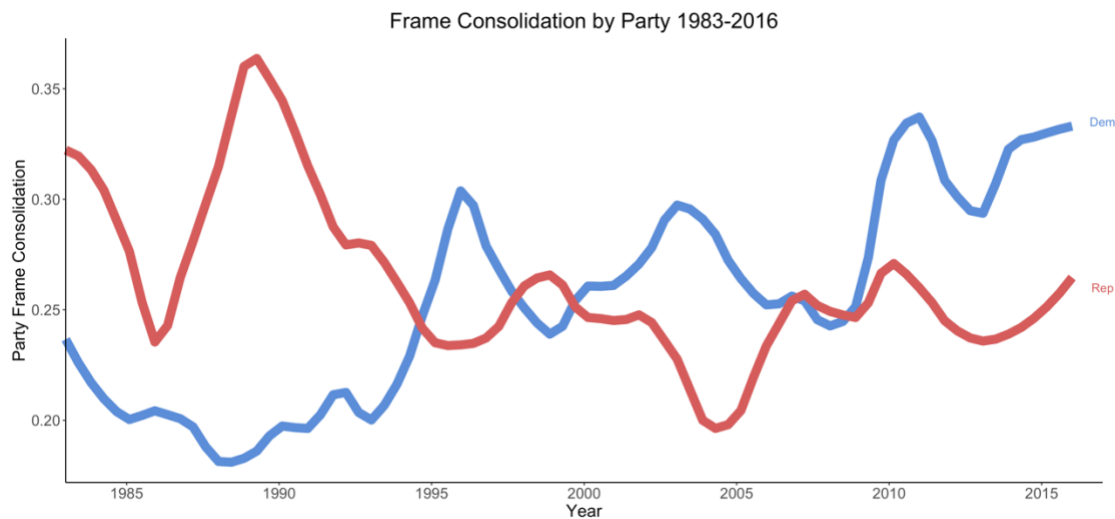frames, where both Democrats and Republicans discuss tax rates for the wealthy but in different

Figure 3.3. Party frame consolidation, measured as the average cosine similarity between all legislators within a party. Lines represent the average of all topic similarity scores. The red line represents average Republican consolidation, and blue line represents average Democrat consolidation.

terms (e.g. *estate tax* vs. *death tax*) and focus attention on different impacts of taxation (e.g.

social_security, *medicare* vs. *tax_burden, job_growth, relief*).

Abstracting once again, a supplemental analysis indicated when partisan frames are

most likely to be deployed. The average of all term (signed) correlation values was calculated for

each year to gain an understanding of which party was engaging in the most effort to frame

issues overall in a given year. Figure 3.5 shows an interesting pattern of partisan language use,

where parties engage in more concerted efforts to frame issues when they do not hold the

majority in the House. When Republicans take control of the House, partisan terms across all

issues are on average more likely to be used by Democrats; conversely when Democrats hold

the majority in the House, Republicans are more likely to use partisan language.

65

This pattern dovetails with Francis Lee's "insecure majorities" thesis, which partly contends that

minority parties are motivated to construct strong partisan messaging in opposition to the majority

party – who are more likely to engage in legislative actions and not messaging – to gain back
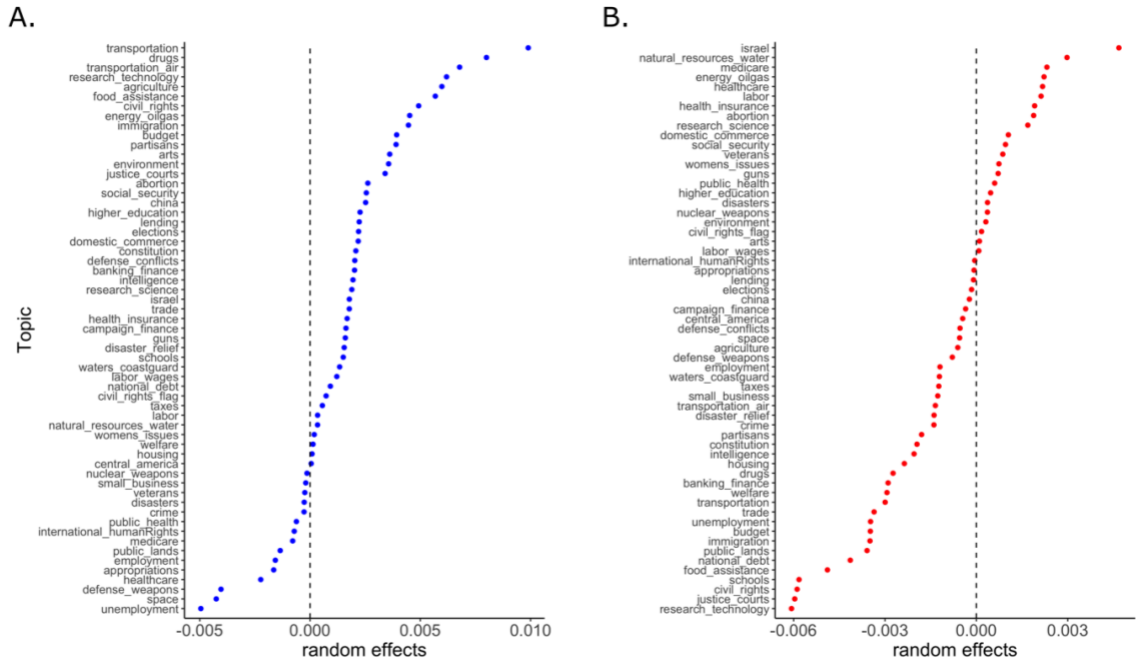
support and rebuild their majority (Lee, 2016).

Figure 3.4 Topic-specific frame consolidation for Democrat and Republican parties.
match closely with the known pro-choice and pro-life frames employed by liberals and

Discussion

In this chapter polarization in the US House of Representatives was investigated through

the lens of partisan framing of political issues. Using automated content analysis, the language

used by Democrats and Republicans was compared, and a measure of dissimilarity for each of

63 unique political topics across 34 years of the congressional debate was derived. For political

topics on average, Democrats and Republicans have increasingly engaged in competitive

framing. Further, the analysis here revealed that increasing frame polarization was not only driven

by single outlier topics but instead – across a larger number of topics – Democrat and Republican

frames have become increasingly divergent and readily employed.

Previous studies of elite framing often focus their attention on one or a handful of political

issues and impose pre-specified content coding schemes or keyword searches to extract frames

and measure competition (Chong & Druckman, 2011; Iyengar & Simon, 1993; Jang & Barnett,

1994). These methods, while useful for direct tests of specific issues and frames, do not provide a

general picture of how Democrats and Republicans have become increasingly reliant on specific

messaging strategies to distinguish their party positions across issues more generally.

Conversely, studies of polarized word choice in congress tend to study all speech in a given year

without consideration of tissue-specific differences in word meaning and with little

acknowledgment of partisan framing (Gentzkow, Shapiro, & Taddy 2016; Jensen et al., 2012). In

this chapter, a middle ground approach was taken to study partisan word-choice

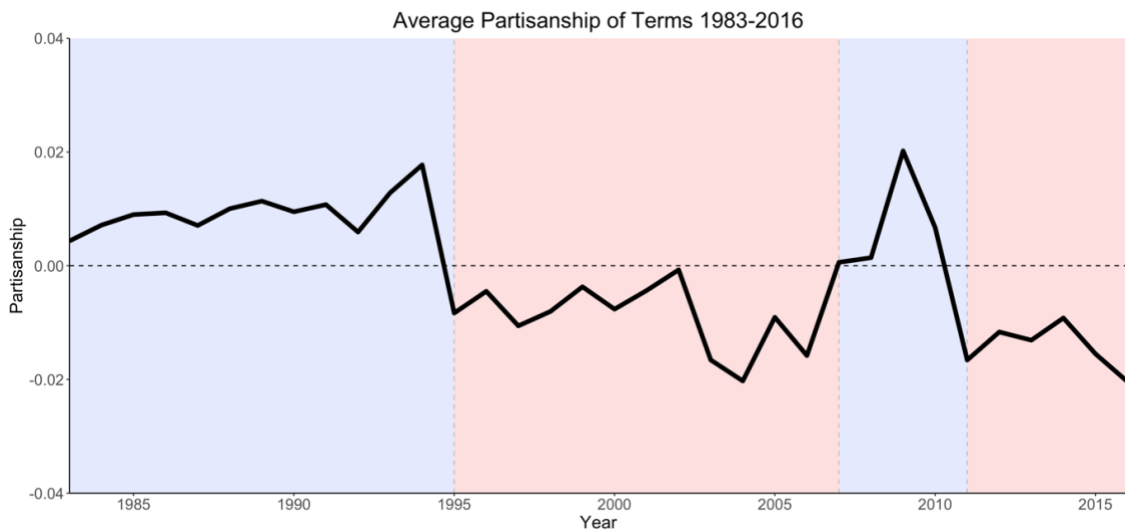Figure 3.5. Total frame partisanship. A weighted sum of all term raw correlation values within a year is computed. Years with weighted sums > 0 indicate that Republicans used more partisan language when discussing issues (framing), and weighted sums < 0 indicate the same for Democrats. Blue portions of the graph represent periods of Democrat control in the House, and red portions represent Republican control in the House.

as framing within many political issues. Here, an inductive approach was taken to identify political issues and frames using the dynamic topic model developed in chapter 1. From this model, speeches were categorized and analyzed independently for differences in word choice between Democrats and Republicans. This approach advances methods for studying partisan framing and polarization by allowing for both the analysis of general trends in framing over time and targeted investigations of what frames were used for specific issues across time. Indeed, the current methodological approach allows one to both ask *what issues garnered competing frames* and *what were the frames employed* simultaneously without a priori assumptions of what those issues and frames are. Table 3.1 reveals several issues known to engender strong competing partisan frames, like taxes and abortion, and details the identification of well-document terms employed by both Democrats and Republicans surrounding these issues (e.g. *estate_tax* vs. *death_tax*, *pro-life* vs. *pro-choice*).

Both the Democrat and Republican parties have dedicated bodies for developing issue-specific messages and disseminating these messages to their party. It is the goal of these organizations to define party positions by keeping their party members "on message". Here it was asked whether the two parties over the past three decades have become increasingly capable of consolidating legislative frames about issues; a behavior that is also indicative of party polarization. Party consolidation varied by party, with Democrats showing a slight increase in consolidation across issues, and Republicans showing a decrease after 1990 and a subsequent plateau. Overall, frame consolidation, however, for both parties was quite low – cosine similarity never reached beyond 0.5), a result also consistent with the careful research of Douglas Harris, who also found that less than one-third of partisan speeches in the House tend to be on message when compared against party leadership (Harris, 2005a).

Finally, it was also found here that partisan framing appears to be a function of congressional power dynamics. By aggregating term associations across topics (which were either negative for Democrats or positive for Republicans), it was shown that when a party loses its majority in the House, they engage in greater efforts to use stronger and more distinct frames for issues. This finding is not independent of previous theoretical accounts and fits well with

research finding that minority party members are apt to engage in deliberate partisan messaging campaigns in order to both oppose the legislative agenda of the majority and regain support from their base and subsequent control of their chamber (Lee, 2016).

The results of this analysis are not without limitations. Of particular note is a persistent issue in studies that employ topic models for content analysis: document misclassification. Careful analysis of Table S3.1 shows that while many terms hold face validity in association with their topic, there are some that appear to not fit the issue. Such inconsistencies are potentially the result of document misclassification, where mismatched documents are classified to the wrong topic for one party more than the other. This imbalance may erroneously identify terms in these documents as partisan, when in fact they are the result of misclassification. Such inconsistencies are a sign that future studies in this vein may benefit from a more restricted analysis of political speech within a smaller time range or with a larger topic model.

Identification of the most polarizing topics was identified by ranking the distance between word-choice distributions, but this method has potential drawbacks. Specifically, ranking topics based on this metric only considers symmetric divergence between Democrats and Republicans, when it is known that while most issues are framed by both parties to some degree, some also engender greater competitive framing at different times and from one party more than another (Klar, 2022). Future studies may seek to identify the most frame-polarizing topics through alternative methods which might account for one-sided framing.

Lastly, the current investigation of party frame consolidation was limited in the depth with which it interrogated patterns of "on message" framing. For example, Harris (2005) investigated a variety of mediating factors after finding that only a small number of legislators tend to be on message. For example, Harris interrogated "who" is likely to be on message, and found that while most legislators are not, legislators who contribute to party messaging organizations are intuitively most likely to employ party frames consistent with party leadership. The current investigation could be expanded to include such an analysis and would further contribute to the study of party dynamics, power, and messaging effects.

Political parties engage in competitive framing to draw attention to specific aspects of issues and sway public opinion with targeted messaging. Competitive framing captures differences in *how* the parties discuss issues and can provide a direct window into the thinking of party members. Here it was shown that the words used by Democrats and Republicans over the past 34 years have become increasingly distinct – an indicator of party polarization and that growing divergence in partisan frames is not simply a function of one or two political issues becoming more polarized. Instead, as time has passed, Democrats and Republicans appear to have employed greater competing language on an increasing number of political issues. Party unity on issue messaging has also remained relatively low, but while Republicans have remained moored in their ability to consolidate party frames, Democrats have become more consolidated in their messaging. Further, efforts to frame issues and use more partisan language appear to be affected by congressional dynamics, where electoral politics motivate minority parties to engage in stronger partisan messaging in the hopes of regaining a majority. These results together indicate that as congressional partisan competition continues to grow, so too may the number of issues framed and the fervency with which they are deployed.

CHAPTER 4: CONGRESSIONAL BELIEF NETWORK POLARIZATION

Introduction

Among the many definitions of ideology (Gerring, 1997), several research programs have converged on a conceptualization that views ideology as a system of interrelated attitudes and beliefs bound together in a causal or non-random fashion (Converse 1964; Gerring, 1997; Jost, 2006; Kalmoe, 2020). Much like other models of knowledge and meaning (DiMaggio, 2011; Anderson, 1983; Mohr, 1998), *belief systems* provide a concrete representation of how concepts are psychologically organized or related, essentially giving a picture of what concepts are associated in the minds of individuals or groups. With such a structure a host of theoretically relevant hypotheses may be tested, including assumptions about the overall structure of ideology (Boutyline & Vaisey, 2017), and what concepts are central to ideologies (Brandt & Sleegers, 2020).

Belief systems are increasingly studied as networks and using the tools of network science; this broader analytic approach is referred to as Belief Network Analysis (BNA) (Boutyline and Vaisey, 2017). Under this operationalization, political beliefs and attitudes are "nodes" and the psychological links between beliefs are "edges". Figure 4.1 provides a visual example of a hypothetical belief network. Typically studies in this domain define belief systems for groups using responses to survey questions, where individual questions (e.g. attitudes towards abortion, family values, international cooperation, liberal-conservative self-placement, federal programs, etc.) act as nodes, and the correlations between questions as edges (Baldassarri & Goldberg, 2014; Baldassarri & Bearman, 2007; Boutyline & Vaisey, 2017).

Using the tools of network science, studies of mass public belief have tested several theories of ideology and public polarization. For instance, researchers have tested the underlying assumption that a coherent belief system is one with high constraint — or strong interdependence between concepts – and that greater constraint can indicate greater polarization. When interrogating individual issues within a network, this form of polarization is studied as a function of how strongly pairs of issues align, such that one could reliably know where a person stands on

one issue (e.g. abortion) from knowing their position on another (e.g. guns) (Baldassarri & Gelman, 2008; Boutyline & Vaisey, 2017). When issue alignment is studied across all issue pairs, constraint is understood as a global, structural phenomenon, i.e. the entire belief system is more interdependent and less cross-cutting. Using community detection and a global measure of network density, DellaPosta showed that across a broad range of issues, public opinion has become increasingly constrained and that the breadth of public opinion consolidation has grown to encompass a greater number of issues (DellaPosta, 2020).

What issues lie at the center of belief systems, or ideology, is another question that is readily answered using methods from network science. Within networks, the number and strength of ties between nodes vary, and as such nodes can be ranked by their relative connectivity or "centrality". Belief network analysis of public opinion surveys has identified that several concepts hold central positions within networks depending on the node-set. Using questions from the 2000 ANES (American National Election Surveys) Boutyline and Vaisey (2017) showed that liberal-conservative self-placement held a central role in public belief systems; Brandt et al. (2019) identified that emotional attachment to political groups was most central to belief systems in a representative panel from New Zealand; and DellPosta (2020) identified a time-dependent number of issue clusters that were central to US belief systems between 1972 and 2016, which reliability contained social issues like attitudes towards race, gender, and civil liberties. Generally speaking, these studies indicate that symbolic identities (e.g. social groups) relative to policy preferences are more central to public belief systems (Fishman & Davis, 2019).

So far, belief network analysis has been applied to the media and the mass public, but not, as far as I know, to political elites. However, the belief systems of political elites are equally as important to map and compare, as knowing how elected officials organize their ideological positions can help to better understand information flows and the relationship between public and elite ideological thinking (Guo & McCombs, 2011). Indeed, Converse makes explicit the utility of understanding ideology as belief systems for rigorously studying how elites transmit "what goes with what" (Converse, 1964, p. 9) to the public or vice versa; if elite and public systems of association mirror one another, this is another indicator of a healthy representative democracy.

The notable lack of elite belief network analysis research is perhaps due to the low availability of large-scale, representative opinion surveys of elected officials, and more specifically of congressional leaders. There is, however, a vast amount of relational information latent in the speech patterns of congressional leaders that can be used to construct belief networks, where the covariation of issue attention across speeches may indicate a psychological association between them. This method for constructing belief networks is rooted in theories that conceive of knowledge as an associative network (Anderson,1983), in which concepts are tied in a causal manner such that the activation of one can have a "spreading activation" to others (Collins & Loftus, 1975). Applied to legislative speech, an associative network of this type would imply that when attention to issues like national debt and taxes, or education and labor covary in the speech patterns of legislators, these issues are conceptually related.

Figure 4.1: Example belief network. Nodes in this network represent political issues, and edges represent relationships between issues.

In this chapter, ideological polarization is operationalized as a divergence in the global topological structure of belief systems (e.g. the larger pattern of associations across an entire network). While previous BNA studies of polarization have operationalized polarization as greater interdependence among issues (e.g. constraint), (DellaPosta, 2020; Kozlowski & Murphy, 2019;

Baldassarri & Gelman, 2008), or greater centrality of partisan identities (Brandt, Sibley, & Osborne, 2019), the current operationalization measures polarization directly as a difference between belief systems.

The methods used in this chapter to construct belief networks and compare them are borrowed from a related area of research in political communication, network agenda setting (Guo & McCombs, 2011), which interrogates the process by which systems of association are communicated by media or elite institutions and received and internalized by the public (Guo & McCombs, 2011; Brosius & Weimann, 1996; Vargo & Guo, 2016). Examples of co-occurrence networks in this literature include those constructed from the covariation of news issues (e.g. health, economy, wars, national security, etc.) in weekly summaries of news coverage (Vu, Guo, and McCombs, 2014), co-occurrence of issues on Twitter in the same day (Vargo and Guo, 2016), and issue co-occurrence in the same article (Chen, Guo, and Su, 2020). The process of network agenda setting is measured in this literature by calculating the similarity of association networks found in the media and those derived from public opinion surveys.

Here, the same analytic strategy used in network agenda setting research is used to also understand elite ideological polarization. More specifically, the methods employed in network agenda setting research are applied to the study of elite ideological polarization by flipping the quantity of interest (network similarity) and measuring the dissimilarity of Democrat and Republican belief systems. Using this approach, no statistically significant differences in global belief system topology were identified between 1983 and 2016.

Materials and Methods

Ideological polarization is operationalized in this chapter as a divergence between the global structure of Democrat and Republican belief systems. The pairwise covariation of political issues across legislators' agendas was measured here to construct issue networks that could be compared using methods borrowed from network agenda setting research. In the following sections, I detail the procedures for creating these networks and the comparative analysis method used here for measuring ideological polarization.

Belief Networks: Like the previous chapters in this dissertation, the methods here begin with the nearly 500,000 speeches made by legislators in the House of Representatives from 1983 to 2016 and the issue topic distributions assigned to these speeches by the dynamic topic model developed in chapter 1. Within each year, speeches were assigned one of 63 unique dynamic topics, and for each House representative within that year, the total number of speeches categorized to each topic was tallied. As such, for each legislator within a year, a topic frequency distribution was generated which described the extent to which that legislator discussed each issue.

The total number of dynamic topics varied by year, such that in a given year some dynamic topics did not appear or some dynamic topics were assigned to multiple window topics (see chapter 1 for more detail on this feature). In the case that a dynamic topic had multiple window topics assigned to it in a year, the frequency of speeches for these individual window topics was combined. Procedural and tribute dynamic topics were repeated in multiple years, and this fact combined with other repeated dynamic topics within a given year, is reflected in the varying number of topics that were ultimately included in this analysis per year (a table of the number of topics for each year is presented in the appendix; Table S4.1)).

Democrat and Republican topic frequency distributions were split into separate datasets. Within each party matrix, cosine similarity was calculated pairwise between every topic based on the topics' distributions across legislators. In this way, a measure of the relationship between topics was calculated which is interpreted as the propensity for two issues to co-occur across a single party's set of speeches. Put differently, a strong relationship between two issues indicates here that when a Democrat gives attention to issue A, they are also likely to attend to issue B. This procedure resulted in two fully connected, weighted adjacency matrices, or network graphs, which were used to represent each party's belief network or ideology.

The last step in constructing partisan belief networks was to then threshold each partisan network for each year so that only statistically significant relationships between issues were maintained. This step was performed through a dyadic thresholding procedure, wherein the

associated p-value for every edge was indexed. This procedure is common in network science research (Ronen et al., 2014; Mukerjee et al., 2022), and is used to focus analysis on relationships in networks that deviate from random noise

P-values for individual edges were derived by permuting one of the two topic vectors 1,000 times and recalculating the cosine similarity. Networks in this analysis contained edge sets ranging in size, from approximately 900 edges to over 1,500. As such, repeated hypothesis tests (i.e. H0: topic 1 and topic 2 are not related), in these networks are likely to result in many spurious rejections of the null. To account for multiple comparisons, the family-wise error rate (FWER) was controlled for during hypothesis testing at an alpha level of 0.05. This step is critical in the construction of large networks and is a common technique implemented in other fields (Zalesky, Fornito, & Bullmore, 2010).

Belief System Polarization and Edge Comparison. Polarization in this chapter is operationalized as the dissimilarity between Democrat and Republican belief systems. As in previous network agenda setting research, networks were compared by measuring the relationship between the global pattern of edge weights between networks. Figure 4.2 provides a simple graphical depiction of this process. For each year, the set of unique edges for Democrat and Republican networks was flattened into one-dimensional arrays and the cosine distance between these vectors was calculated. Cosine distance ranges from 0 to 1, thus values in this analysis closer to 1 indicate that party belief networks are less similar and thus more polarized.

Borrowing from a widely used method in network science for network comparison (the network comparison test) (Van Borkulo et al., 2017), the statistical significance of the global polarization measure (see appendix for supplemental analysis interrogating differences in individual edge strength as well) was evaluated through a permutation procedure. For every year, the party labels for representatives were randomly shuffled before splitting legislator topic distributions into Democrat and Republican matrices, thus breaking the dependency between party and issue relationships. The network creation procedure and comparison methods from above were then again performed to create thresholded null belief networks for Democrats and

Republicans, and null global (and edgewise, see appendix) polarization measures. This process

was repeated 1,000 times every year, and the distribution of null results was compared to the true

test statistics. The resulting p-values were corrected for multiple comparisons again with an
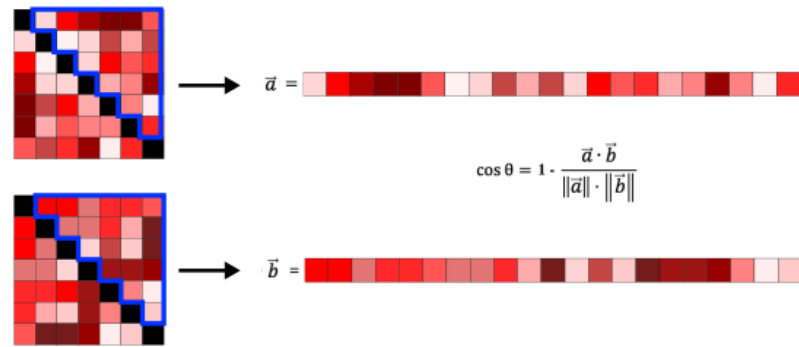
FWER of 0.05.

$$\cos \theta = 1 - \frac{\vec{a} \cdot \vec{b}}{\|\vec{a}\| \cdot \|\vec{b}\|}$$

Figure 4.2: procedure for measuring belief system polarization. The upper triangle of networks is taken and flattened into 1-D vectors, $\vec{a}$, and $\vec{b}$. Cosine distance is then calculated using these vectors, representing the dissimilarity between the global pattern of edge weights.

Results

Belief System Polarization: For every year, the entire pattern of edge weights for Democrats and

Republicans was compared for their dissimilarity. This measure was meant to identify whether the

overall structure of Democrat and Republican belief systems have become more distinct over

time. After deriving p-values through permutation tests, it was found that Democrat and

Republican networks were not reliably different in any year. Table 4.1 provides the cosine

distance measure and associated uncorrected p-value for every year, and this table reveals that

while some years did approach significance prior to correction for multiple comparisons (i.e. 34

tests were performed and would thus need to be corrected for spurious rejection of the null - that

the pattern of edge weights are the same), no p-value was small enough to survive corrections.

Putting aside statistical significance momentarily, the results of this analysis still only found an

overall low to moderate difference between the global structure of Democrat and Republican

belief systems, with an average cosine distance (0 indicating completely similar, 1 completely different) of 0.35, and a maximum distance value of 0.54 in 1985. Again, however, no results in this analysis did achieve statistical significance, and so inferences derived from these results about partisan ideological polarization are unreliable.

Discussion

Various models in psychology (Collins & Loftus, 1975; Eysenck, 2001), message processing (Lang, 2000), and political knowledge (Jost & Sterling, 2020), contend that knowledge is organized through a web of causal associations. In political science, ideology has also been conceptualized as a system of interconnected beliefs (Converse, 1964), and operationalizing these belief systems as networks has provided a variety of novel insights related to public understanding of the political environment and of mass ideological polarization (Boutyline & Vaisey, 2017; DellaPosta, 2020; Brandt, Sibley, & Osborne, 2019). Little research, however, has also applied belief network analysis to the study of elite ideological polarization. The goal of this chapter was to fill this gap and interrogate whether the belief systems of Democrats and Republicans have become less similar, or polarized, over 34 years of congressional debate, from 1983 to 2016.

Here, belief networks were constructed through the covariation of political issues in the agendas of Democrats and Republicans. The frequency distribution of topics identified by the dynamic topic model developed in chapter 1 were interrogated for their co-occurrence across partisan speeches to better understand their conceptual associations. In this way, it was possible to measure the likelihood that Democrats would discuss one issue – say the national budget – by knowing how frequently Democrats discussed another: healthcare. Put differently, the covariation of issues within party agendas made it possible to understand "what goes with what" (Convers, 1964) in the collective thinking of parties.

Ideological polarization was simply measured as a difference in the overall pattern of associations between Democrat and Republican networks, a method frequently employed in agenda setting research (Guo and McCombs, 2011; Guo et al., 2016). In contrast to DW-

NOMINATE, the most relied on metric of ideological polarization in congress (McCarty, Poole, & Rosenthal, 2016), and the previous chapters in this dissertation, this investigation found that when belief systems are constructed using the co-occurrence of topics within congressional speech, no significant differences in the structure of Democrat and Republican belief networks were identified. These divergent results are potentially due to differences in operationalization, wherein the current study centered its understanding of ideology on the symbolic expression of political beliefs (i.e. speeches) while DW-NOMINATE relies solely on the vote count of elected officials. However, it is more likely that the lack of evidence of ideological polarization here is due to limitations in the current data and methodology.

Specifically, the data in this chapter were particularly sparse, with many legislators never mentioning specific issues. This fact paired with the relatively low number of observations (i.e. approximately 435 representatives divided among the two parties) and large number of comparisons (network edge sets ranged between approximately 900 and 1,700), made it very unlikely for network edges to reach statistical significance after correcting for multiple comparisons. Indeed, low power is a common issue in network comparison research, especially in medium to large networks such as semantic networks or brain networks (Van Borkulo et al., 2017; Zalesky et al., 2010), One potential remedy for this issue in the current context is to implement alternative models like the correlated topic model, which implement alternative distribution priors (e.g. multinomial gaussian as opposed to Dirichlet or NMF) and allow for topics to be associated at the level of speeches instead of at the level of legislators. This alternative modeling choice would significantly expand the number of observations within a year and potentially improve statistical power enough to identify more densely connected networks.

Belief System Dissimilarity

1983-2016

Year	Dissimilarity	p-value
1983	0.3954580	0.232
1984	0.3303867	0.950
1985	0.5483945	0.340
1986	0.4186281	0.672
1987	0.3492639	0.262
1988	0.2901767	0.602
1989	0.4145610	0.386
1990	0.3533483	0.538
1991	0.3781226	0.414
1992	0.2766096	0.918
1993	0.3695114	0.570
1994	0.3220692	0.852
1995	0.4875587	0.912
1996	0.3254766	0.822
1997	0.4060790	0.356
1998	0.3586906	0.072
1999	0.3145485	0.682
2000	0.2882502	0.638
2001	0.2723422	0.862
2002	0.3777179	0.078
2003	0.3294716	0.292
2004	0.3049498	0.732
2005	0.3612598	0.936
2006	0.2437233	0.960
2007	0.4344651	0.778
2008	0.3913979	0.252
2009	0.2850451	0.982
2010	0.4416626	0.324
2011	0.2863078	0.914
2012	0.3935807	0.030
2013	0.3099702	0.776
2014	0.3101886	0.084
2015	0.3102624	0.172
2016	0.2895911	0.096

Table 4.1: Belief System Dissimilarity 1983-2016 as measured by cosine distance (range of 0-1). No years showed a difference between the overall structure of Democrat and Republican belief systems that survived corrections for multiple comparisons.

Furthermore, while the ability of the dynamic topic model to allow topics to emerge and disappear across time is seen as an overall improvement in representing political reality, it is perhaps a disadvantage in the study of dynamic networks and the polarization of such networks. Because the topics varied from year to year, it was not feasible to analyze the development of belief systems within parties over time and compare a consistent set of quantities between partisan networks across time. Cosine distance is not affected by the size of compared vectors, and so the current statistical results were comparable across time, however, the changing content of networks does pose a challenge for meaningfully interpreting differences between networks across time. Setting a static set of concept nodes would provide novel insights into the changing structure of party belief systems and could make it possible to interrogate several theoretically relevant phenomena like issue consolidation, constraint (DellaPosta, 2020), and issue centrality over time (Brandt, Sibley, & Osborne 2019). A future direction for research in this area could be to develop dynamic topic models that are specifically designed to maintain the same issues across time – either through guided models e.g. (Gallagher et al., 2017)) or through hierarchical clustering (Quinn et al., 2010) – and then again analyze belief networks and their evolution.

Despite these current practical limitations, the approach employed here is still a significant advance for the study of elite polarization. As mentioned in previous chapters, the application of dynamic topic models for the study of elite polarization is a novel approach. Furthermore, capturing the output of topic models to construct networks of topic associations is an exciting direction for research in political polarization. Alternative models, like the correlated topic model (Blei & Lafferty, 2007), and the structural topic model (Roberts et al., 2014) have set the groundwork for constructing topic networks from document collections, and previous research in political communication has derived frame networks from topic models in more circumscribed datasets (Walter & Ophir, 2019). No study, however, in political communication has attempted to construct time-evolving networks from topic models, and further, no study has linked association networks in this domain to the rich literature on belief systems and ideology. Future

improvements to the current methods, like those mentioned here, will surely provide more robust

and interesting insights into the study of elite ideology and ideological polarization.

CONCLUSION

Current understanding of congressional polarization is based primarily on the voting behavior of US legislators. Such research indicates that since the 1960's the Democrat and Republican parties have rapidly distanced themselves from one another. Voting behavior, however, only provides a limited perspective on elite polarization, and it was the primary thrust of this dissertation to go beyond the formal voting behaviors of elected officials and plumb the thinking and goals of politicians through careful analysis of their speech. Indeed, for some, it is believed that language is a window into the mind, and in this dissertation, the language of politicians was interrogated to better understand whether Democrats and Republicans are not just polarized in what they say "yea" or "nay" to, but in what they express on the floor of Congress between their votes.

Language cannot provide insights into political polarization alone. Specific frameworks and theories are necessary to structure the conceptualization and operationalization of polarization. In this work, I studied congressional polarization through the lens of political communication, and specifically through the theories of agenda setting, framing, and belief network analysis. Using a flexible unsupervised machine learning method, the dynamic non-negative matrix factorization topic model, I showed that it was possible to reliably extract from the Congressional Record these quantities of interest, and in comparing them found – in two out of the three cases – clear evidence of polarization.

In this dissertation, I studied the speech of congressional leaders in the House of Representatives from 1983 to 2016. Representatives in this period made nearly half a million substantive speeches on the floor of the House across these 34 years. In chapter 1, I introduced and rigorously tested the dynamic topic model for its semantic and external validity. Previous applications of topic models to congressional language have shown that topic models are a useful tool in classifying large amounts of political speech and tracking political issues (Grimmer, 2010; Quinn et al., 2010). However, to study language use across much longer periods, it is necessary to account for the fact that, within any domain, issues emerge and disappear from public

83

attention, and more perennial issues change in how they are discussed and understood (Greene & Cross, 2017; Blei & Lafferty, 2006). The dynamics of political language in congress are not an exception to this fact, and so the dynamic topic model is needed.

Sixty-three unique topics were identified across the entire corpus. To develop confidence in the model, the semantic validity and external validity of the model were tested. Topics from the model were internally coherent across time, and although specific topic term distributions shifted across years, topics were consistently interpretable and meaningful. Dynamic topics also were found to track with both legislative events (e.g. introduction of widely debated bills like the Bradey bill) and world events (e.g. the Columbine high school shootings), indicating that the model was providing confident measurements of topic prevalence within and across the corpus.

The development of the dynamic topic model method for the study of congressional speech, while an advance in itself, was ultimately a means to an end in this dissertation. The primary aim of this research was to apply this model to the study of elite polarization. To accomplish this goal, the dynamic topic model was used as a unifying method for the measurement of the three aforementioned concepts from political communication (agenda setting, framing, belief systems) that have independently received significant attention (Cobb, Ross, & Ross, 1976; Glazier & Boydstun, 2012; Shah, McLeod, & Gotlieb, 2009; Chong & Druckman, 2007; Esacove, 2004; Brandt and Sleegers, 2020; Boutyline & Vaisey, 2017), but much less attention as related phenomena (Guo & McCombs, 2011; McCombs & Valenzuela, 2017).

In chapter 2, the distribution of issues, or agendas of Democrats and Republicans, were calculated from the output of the dynamic topic model developed in chapter 1. Party polarization in this chapter was operationalized as a divergence in these agendas: *agenda polarization.* From this analysis it was found that the Democrat and Republican agendas have become increasingly distinct; a classification model trained on the topic distribution matrices for Democrats and Republicans was increasingly able to accurately distinguish between the speeches given by the parties respectively (Peterson & Spirling, 2018). This analysis indicates that Democrats and Republicans have become more distinct in *what* they discuss, or what issues they give priority to.

Consistent with research finding increasing issue ownership in political campaigns (Egan, 2013; Petrocik et al., 2003) and shrinking bipartisan co-sponsorship in congress (Sulkin & Schmitt, 2014), these results add to a growing body of work pointing to ever-increasing disagreement between elite Democrats and Republicans about what issues are most important for the country. This form of polarization does not remain contained within congress, and research also indicates that the polarized agendas of elected officials also influence what issues the mass public finds most important, further polarizing public attitudes as well (Jones & Baumgartner, 2004).

In addition to theoretical contributions, chapter 2 also provides a novel analytic approach to studying elite polarization. Specifically, while using the topic distributions derived from unsupervised topic models for critical analysis of polarization is not in itself a novel approach (Sakamoto & Takikawa, 2017), and the application of supervised machine learning classification for distinguishing between partisan language across time is not new (Peterson & Spirling, 2018), this is the first study to combine these two methods in political communication. Furthermore, this analysis is the first to apply these methods across such a large body of text in the context of American politics. Finally, a critical component of this analysis was the closer inspection of model coefficients to gain a better understanding of issue ownership within specific years. The introduction of permutation procedures to establish the reliability of model coefficients for inference is a novel approach in this domain and provided a rigorous method for understanding what issues drive polarization.

While chapter 2 showed that Democrats and Republicans have become increasingly polarized in terms of what issues they focus on, it is still the case that many issues are frequently addressed by both parties but framed differently. Therefore, in chapter 3 partisan polarization was operationalized as a growing divergence not in what issues Democrats and Republicans talk about, but how partisans discuss issues and understand them. More specifically, chapter 3 sought to understand whether differences in partisan language have become more distinct within issues, and whether differences in frames have become more common across issues. These two dimensions of partisan frame competition together were referred to as *frame polarization*.

Results from chapter 3 also indicated a growing pattern of elite polarization, such that on average partisan frames have become more distinct across issues, and that this form of polarization has grown within most issues since 1983. Ultimately, these results find that it is increasingly common for Democrats and Republicans to differ in their understanding of issues. Debating a single issue, partisans from different sides of the aisle are today far more likely to talk past one another than productively address one another's concerns. A particularly salient example of this failure to communicate and understand alternative perspectives is the debate on abortion access. The results found in this dissertation identified the increasing preference for pro-life emphasis frames by Republicans and pro-choice frames by Democrats. An unwillingness to discuss alternative perspectives and further promote only one side of an issue not only impacts the productivity of Congress but also influences the thinking of the public and their behaviors (Feldman & Hart, 2018). The combined effect of elected officials promoting increasingly polarized frames and the fact that the public is apt to engage in selective exposure to political information (Iyengar et al., 2008; Bolin & Hamilton, 2018; Arceneaux et al., 2012) has the potential to further polarize not just what the public cares about, but more generally how they understand the world and political issues.

Chapter 3, like chapter 2, also advances methodological approaches for the study of elite polarization. Unlike previous approaches to studying differences in partisan language across issues (Gentzkow, Shapiro, & Taddy, 2019; Chong & Druckman, 2011; Jensen et al., 2012), the dynamic topic model used in this dissertation makes no assumptions about what issues ought to be studied or what words will be included in any comparative analysis of partisan speech. Instead, topic models are an inductive method, identifying issues through latent patterns in speech across documents. In this way, the application of topic models for the study of frames and partisan language polarization is a novel approach, and the application of the dynamic topic model is an even further advance, as it allows for more precise identification of both time-dependent and more time-invariant topics within specific periods.

In the final analysis of this dissertation, political ideology was operationalized as belief networks, and polarization was studied as a simple divergence in the organization of Democrat

and Republican belief systems. Chapter 4 borrowed theory and methods from a variety of research domains, including political psychology (Converse, 1964), network science (Boutyline & Vaisey, 2017), and agendas setting research (Guo & McCombs, 2011), to conceptualize and operationalize ideological polarization. The frequency distribution of topics discovered by the dynamic topic model were interrogated for their co-occurrence across partisan speeches to better understand their conceptual associations. When taken together these associations were understood as belief systems and were compared between parties to measure polarization. Unlike previous accounts of elite ideology and previous findings in this dissertation, no pattern of polarization was identified by the belief network analysis in chapter 4. Despite the null findings, this chapter continues a broader trend in political communication research to leverage topic models to develop association networks (Walter and Ophir 2019). Furthermore, this chapter provides a novel approach to understanding ideological polarization in Congress and indicates that alternative methods, like correlated topic models (Blei & Lafferty, 2007) or semantic network analysis (van Atteveldt, 2008), could be used to model the congressional record and extract meaning and compare relevant quantities to political polarization with greater statistical power.

The work in this dissertation is a first step in applying methods from natural language processing to study elite political polarization under theories in political communication. Future directions will certainly improve on the approaches taken here and will tackle other questions and corpora relevant to the study of modern elite polarization. For example, future research may seek to include more recent legislative speeches to study party polarization, especially given the powerful influence of President Donald Trump's rhetoric during his presidential term and beyond. Furthermore, future work may seek to understand the growth of divisions within parties during this period using the frameworks applied here. Indeed, factions have emerged in both the Democratic party (e.g. progressives like Alexandria Ocasio-Cortez et al.) and the Republican party (e.g. QAnon candidates like Marjorie Taylor Greene et al.), which rely on very different rhetorical strategies and take ownership of a divergent set of issues from the majority of their respective parties. Such research may provide an even deeper understanding of who is driving divisions between the two parties that house both moderate and more extreme political actors.

Future work will also certainly find it advantageous to explore alternative modeling decisions from the ones used here. More precisely, in this dissertation, I chose to follow the methods of Greene and Cross (2017) and fit a dynamic non-negative matrix factorization topic model to congressional speech. Differences in Democrat and Republican speech patterns were analyzed after model fitting and were studied using the output of this model . One model of particular interest for future work is the Structural Topic Model developed by Margaret Roberts and colleagues. The structural topic model makes it possible to include covariates in the actual fitting of topic models and opens the possibility of modeling elite rhetoric while simultaneously taking into consideration variation in topic prevalence and content across time and party. This model was not used in the current research because it was not computationally feasible to run the structural topic model on the nearly half-million speeches included in the corpus. Future work may seek to replicate the current analysis by sampling from the congressional record and fitting a structural topic model. The structural topic model would make it possible to quantify changes in both topic prevalence and content over time in a single powerful procedure. Furthermore, topic correlation matrices are easily derived from the structural topic model and are computed from the covariation of topics across speeches (as opposed to legislators like in the analysis in chapter 4 of this dissertation). This fact would also likely provide the statistical power necessary to develop robust and more densely connected belief networks, and thus a more interesting and reliable foundation on which to analyze ideological polarization using belief network analysis.

To summarize this work: this dissertation provides numerous theoretical and methodological advances for the study of congressional elite polarization. Specifically, in this dissertation, the language of US legislators provided a direct window into the nature of political polarization. This research found that politicians have not only become more polarized in what they talk about but also in how they talk about issues. Much of what politicians say is communicated to the public, and it is well established by the broader field of political communication that these messages and the ideas they represent are coopted by the public. This dissertation indicates that our elected officials have, over the past 30-plus years, moved further

and further away from the deliberative democracy we wish them to represent and are setting an example that directly impacts polarization and democratic norms in the public as well.

It was also shown here that the dynamic topic model approach introduced in this research can provide useful and externally valid measurements of political attention in Congress, and more importantly, showed that topic models can provide a unifying methodology for the study of various operationalizations of political polarization. This research is considered in the broader movement of social science towards advanced computational approaches and natural language processing (Grimmer & Stewart, 2013; Grimmer, Roberts, & Stewart, 2021). In this context, I hope that this dissertation inspires future applications of the dynamic topic model in the analysis of large political corpora, and motivates researchers to continue to apply these models in creative ways beyond the simple measurement of issue attention. In 2022, it has become clear that the words politicians use do truly matter and have the potential to inspire both great and terrible things. With every day text data becomes the most abundant resource we as communication scholars have at our disposal. As such, the possibilities for future research interrogating the power of language are seemingly endless.

Congress	Extensions		House		Senate	
	D	R	D	R	D	R
98th	11,973	7,008	47,467	38,671	34,426	47,316
99th	11,811	7,855	40,963	39,417	43,468	48,184
100th	10,736	7,536	35,920	35,944	56,537	35,365
101st	11,307	8,022	34,796	36,612	47,534	27,882
102nd	11,660	7,222	37,759	33,548	45,732	25,133
103rd	8,531	5,052	34,230	37,243	46,231	28,455
104th	6,702	5,746	40,234	51,956	38,611	39,706
105th	7,665	6,095	31,041	41,100	24,512	38,295
106th	7,823	6,457	30,930	39,713	29,354	37,199
107th	8,813	6,895	27,479	31,928	36,228	26,804
108th	9,995	8,520	36,527	35,139	30,153	33,433
109th	10,779	7,983	36,152	36,018	22,707	28,743
110th	11,859	9,353	50,840	42,589	30,172	17,846
111th	10,092	7,671	35,126	28,962	21,148	12,625
112th	10,399	9,484	31,566	34,287	19,880	11,737
113th	9,000	8,411	25,968	29,175	16,817	9,838
114th	6,651	9,275	21,015	26,855	8,455	12,562

Table S.1: Count of unfiltered speeches, utterances and insertions for all three sections of the congressional record (Extension of Remarks, House, and Senate) for each party from the 98th to 114th congress.

| | House | |
Congress	D	R
98th	22,597	16,996
99th	20,618	18,110
100th	18,551	16,779
101st	17,760	15,327
102nd	18,825	15,886
103rd	16,935	18,028
104th	22,320	25,820
105th	16,975	19,197
106th	19,756	23,189
107th	17,325	18,873
108th	17,163	14,665
109th	17,506	15,314
110th	22,013	17,342
111th	15,078	14,396
112th	15,011	14,500
113th	12,657	12,934
114th	10,657	12,287

Table 2: Final count of speeches included in analysis after filtering for House and Senate across all 17 congresses.

Procedural Stop words

talk, thing, colleague, hear, floor, think, thank, insert, section, act_chair, amendment, clerk, clerk_designate, pursuant, minute, desk, amendment_text, amendment_desk, rule, debate, process, offer_amendment, majority, order, pass, extension, urge, urge_colleague, defeat_previous, yield_balance, member, committee, chairman, mr, subcommittee, rank_member, mr_chairman, oversight, yield_minute, yield_time, gentlewoman, gentleman, gentlelady, h_r, time_consume, legislation, measure, rollcall, rollcall_vote, vote_aye, vote_nay, nay, debate, point_order, chair, clause, clause_rule, germane, sustain, remark, conference, pass, oppose, offer, opposition, ask, speaker, bill, follow_prayer, approve_date, pledge_journal, morning_hour, today_adjourn, proceeding, deem_expire, reserve, complete, permit_speak, authorize_meet, session_senate, office_building, entitle, conduct_hearing, m_room, consent, ask_unanimous, dirksen_senate, senate_proceed, intervene_action, consider, notify_senate, senate, legislative_session, legislation, legislature, further_motion, motion, lay_table, motion_reconsider, reconsider, hearing, leader, p_m, a_m, period_morning, period_afternoon, executive_session, follow, senate_proceed, morning_business, authorize, motion_concur, concur, session, hour, control, follow_morning, senate_resume, follow, monday, tuesday, wednesday, thursday, friday, ask_unanimous, motion_reconsider, amendment, consent, motion_proceed, cloture, proceed, motion_invoke, cloture_motion, invoke, no_, modify, program, percent, increase, fund, funding, suspension, count, yesterday, tomorrow, act, previous_question, present, record, resolution, house_concurrent, house_joint, previous_question, yield_such, introduce, call, re, recognize, commend, cosponsor, express, print, action, pursuant_house, h_re, continue, sponsor, yield, thank_gentleman, second, friend, comment, appreciate, gentleman_california, statement, distinguished, gentleman_texas, thank_gentlewoman, gentleman_ohio, gentleman_illinois, gentleman_pennsylvania, gentleman_florida, gentleman_michigan, want_commend, bring, special_order, house_representative, leadership, bring, consideration, matter, other_body, adjourn, legislative, version, move, meet, resolve, motion_instruct, appropriation_bill, madam_speaker, yield_such, reserve_balance, bipartisan, support_h, previous_question, introduce, important, good_friend, rise_today, pleased, sponsor, rise, like_thank, representative, second, want_thank, leadership, join, allow, consideration, discharge_further, ask_immediate, j_re, joint_resolution, immediate_consideration, week, senate_joint, designate, designate_week, re, h_j, america, american, work, law, want, issue, get, try, take, let, question, answer, report, say, know, come, tell, people, country, language, conference, conference_report, need, see, commission, let, tell, day, united_state, deal, point, address, look, congress, congressional, go, come, put, agree, yield, alaska, alabama, arkansas, american_samoa, arizona, california, colorado, connecticut, district_of_columbia, delaware, florida, georgia, guam, hawaii, iowa, idaho, illinois, indiana, kansas, kentucky, louisiana, massachusetts, maryland, maine, michigan, minnesota, missouri, mississippi, montana, north_carolina, north_dakota, nebraska, new_hampshire, new_jersey, new_mexico, nevada, new_york, ohio, oklahoma, oregon, pennsylvania, puerto_rico, rhode_island, south_carolina, south_dakota, tennessee, texas, utah, virginia, virgin_islands, vermont, washington, wisconsin, west_virginia, wyoming

Bigram Phrase Score Formula

$$\frac{(bigram\ count - \min count) * length\ vocabulary}{word\ a\ count * word\ b\ count},$$

Where *bigram count* is the number of occurance of worda_wordb, and *min counti* is a set threshold for cooccurrence.

Dynamic Topic	Terms
tribute	miss, love, congressman, john, wife, post_office, father, world_war, sacrifice, honor
veterans	veteran, va, veteran_affair, compensation, world_war, hospital, medical_care, medical, vietnam_veteran, veteran_administration
taxes	tax_relief, tax_credit, package, relief, tax_break, marriage_penalty, credit, tax_code, wealthy, income
natural_resources_water	water, river, water_supply, dam, water_resource, lake, great_lake, drought, clean_water, construction
small_business	small_business, sba, owner, contract, loan, capital, regulation, entrepreneur, paperwork, firm
research_science	research, science, disease, nih, university, center, cancer, national_institute, scientist, scientific
abortion	abortion, family_planning, baby, pregnancy, procedure, mother, prolife, unborn_child, clinic, ban
food_assistance	food, hunger, nutrition, hungry, meal, food_stamp, agriculture, snap, poor, feed
agriculture	farmer, farm, agriculture, agricultural, crop, farm_bill, producer, price, payment, sugar
housing	housing, public_housing, hud, unit, homeless, affordable_housing, voucher, rent, resident, lowincome
taxes	taxis, revenue, raise_taxis, tax, income, spending, tax_code, average, tax_reform, tax_increase
veterans	va, veteran_affair, facility, hospital, claim, wait, backlog, vas, medical, contract
employment	employee, employer, federal_employee, pension, retirement, civil_service, post_office, mandate, employment, hire
crime	crime, law_enforcement, victim, violence, officer, criminal, prison, police, murder, domestic_violence
israel	israel, peace, israeli, palestinian, middle_east, jewish, hama, ally, region, gaza
healthcare	patient, doctor, physician, hospital, medical, provider, hmo, treatment, health_care, quality
justice_courts	court, judge, legal, justice, supreme_court, judicial, lawsuit, lawyer, death_penalty, attorney
social_security	social_security, trust_fund, senior_citizen, retirement, surplus, disability, retiree, income, cola, recipient
domestic_commerce	consumer, price, internet, credit, credit_card, competition, product, fcc, market, sugar
national_debt	debt, national_debt, borrow, bankruptcy, credit_card, interest_rate, deficit, debt_limit, surplus, owe
public_lands	land, national_park, park, county, property, acre, river, management, wilderness, forest_service
procedural	minority, staff, understanding, discussion, schedule, side_aisle, discuss, tonight, next_week, correct

procedural	substitute, motion_recommit, alternative, defeat, original, nature_substitute, compromise, underlying_bill, instruction, waiver
transportation	transportation, amtrak, highway, infrastructure, bridge, safety, rail, road, transit, construction
drugs	drug, fda, treatment, heroin, drug_abuse, marijuana, law_enforcement, colombia, cocaine, select_committee
space	space, nasa, space_station, science, mission, satellite, space_shuttle, launch, earth, astronaut
energy_oilgas	oil, energy, natural_gas, price, coal, production, gas, fuel, supply, drill
procedural	title, amend, code, move_suspend, waiver, relate, technical, whole, clarify, purpose
trade	trade, agreement, free_trade, trade_agreement, export, colombia, mexico, fast_track, labor, nafta
nuclear_weapons	iran, nuclear, sanction, iranian, nuclear_weapon, regime, ally, threat, weapon, middle_east
procedural	object, reserve_right, objection, withdraw_reservation, chief_sponsor, minority, unanimous_consent, inform, reservation, unanimousconsent_request
international_humanRights	cuba, cuban, castro, democracy, human_right, russia, regime, embargo, haiti, russian
medicare	medicare, senior, prescription_drug, senior_citizen, elderly, health_care, premium, hospital, beneficiary, medicaid
immigration	border, immigration, mexico, border_patrol, illegal_alien, agent, immigrant, illegal_immigration, illegal, homeland_security
NA	authorization, continue_resolution, figure, appropriation_committee, budget_resolution, authorization_bill, waiver, item, conferee, foreign_aid
intelligence	intelligence, intelligence_community, terrorist, national_security, threat, permanent_select, cia, terrorism, fisa, intelligence_agency
health_insurance	coverage, health_insurance, insurance, premium, cover, affordable_care, insurance_company, uninsured, affordable, employer
tribute	team, win, coach, university, game, championship, congratulate, player, season, high_school
waters_coastguard	coast_guard, vessel, port, ship, maritime, boat, mission, merchant_marine, safety, fishery
unemployment	unemployment, unemployed, unemployment_benefit, recession, unemployment_rate, extend_unemployment, unemployment_insurance, lose_job, package, unemployment_compensation
taxes	tax_cut, wealthy, rich, tax_break, surplus, middle_class, revenue, bush, income, capital_gain
environment	epa, environmental, site, superfund, waste, cleanup, clean, environment, clean_air, coal
arts	art, nea, national_endowment, museum, artist, humanity, cultural, art_humanity, music, culture
higher_education	student, college, school, university, high_education, student_loan, educational, high_school, graduate, pell_grant

procedural	gentleman_yield, correct, understanding, schedule, intention, finish, oclock, next_week, intent, sure
constitution	constitution, supreme_court, constitutional, marriage, district_columbia, article, religious, religion, unconstitutional, found_father
armenian_genocide	armenian, cyprus, turkey, turkish, genocide, armenian_genocide, armenia, island, greek, peace
budget	balanced_budget, deficit, spending, budget_resolution, revenue, fiscal, surplus, balance_budget, deficit_reduction, balance
schools	teacher, classroom, parent, kid, head_start, school_district, teach, quality, test, public_school
womens_issues	woman, breast_cancer, women, victim, era, equal, age, cancer, young, violence
guns	gun, ban, firearm, gun_violence, weapon, assault_weapon, criminal, background_check, handgun, brady_bill
defense_conflicts	troop, afghanistan, iraqi, soldier, iraq, mission, defense, afghan, taliban, war
disasters	fire, forest, firefighter, forest_service, national_forest, wildfire, timber, tree, burn, acre
appropriations	fiscal_year, appropriation, account, budget_resolution, request, department_defense, whole, union, allocation, amend
partisans	democrat, democrats, white_house, veto, party, republican, investigation, democratic, side_aisle, continue_resolution
transportation_air	airport, safety, aviation, faa, airline, air, flight, passenger, air_traffic, transportation
defense_weapons	missile, test, weapon, nuclear, sdi, arm_control, mx, treaty, deploy, capability
procedural	reclaim_time, correct, next_week, discussion, madam_chairman, gentleman_wisconsin, happy, sure, withdraw, understanding
procedural	end_september, table, amendment_thereto, district_columbia, disagree, union, related_agency, amend, revise_extend, revenue
central_america	nicaragua, contra, sandinista, central_america, el_salvador, nicaraguan, peace, communist, region, democracy
public_health	aid, disease, africa, hiv, cancer, treatment, breast_cancer, african, prevention, global
china	china, chinese, human_right, trade, taiwan, mfn, beijing, chinese_government, export, tiananman_square
international_humanRights	human_right, freedom, soviet_union, soviet, religious, vietnam, soviet_jew, jew, lithuania, emigration
trade	japan, export, japanese, steel, import, product, company, foreign, trade, buy
native_americans	indian, tribe, india, native_american, tribal, indian_tribe, reservation, pakistan, census, settlement
labor	worker, union, employer, labor, contract, notice, strike, construction, osha, plant

|	
civil_rights_flag	flag, symbol, constitution, freedom, constitutional_amendment, supreme_court, desecration, burn, desecrate, liberty
disaster_relief	emergency, disaster, fema, flood, flood_insurance, damage, offset, hurricane, relief, victim
welfare	welfare, welfare_reform, parent, poverty, poor, welfare_recipient, child_care, mother, kid, medicaid
procedural	page_line, strike, follow_page, insert_follow, dollar_amount, amount_insert, control_minute, chair_recognize, member_oppose, oppose_control
taxes	irs, investigation, tax, internal_revenue, taxpayer, information, tax_code, abuse, target, audit
energy_oilgas	pipeline, keystone_pipeline, oil, permit, canada, safety, approve, natural_gas, construction, keystone
civil_rights	civil_right, discrimination, black, king, voting_right, minority, quota, african_american, caucus, equality
labor_wages	minimum_wage, wage, increase_minimum, raise_minimum, earn, middle_class, poverty, fair, poor, income
campaign_finance	campaign_finance, campaign, candidate, soft_money, shaysmeehan, contribution, pac, special_interest, influence, ad
research_technology	technology, innovation, manufacturing, science, research_development, cybersecurity, private_sector, export, invest, patent
elections	election, voter, democracy, campaign, candidate, voting, ballot, poll, elect, democratic
banking_finance	bank, financial, banking, institution, financial_institution, loan, wall_street, capital, saving_loan, check
procedural	earmark, county, request, project, appropriation_committee, transparency, specific, contain, list, campaign_contribution
lending	loan, mortgage, lender, interest_rate, student_loan, borrower, foreclosure, homeowner, credit, lending

Table S2.1: Dynamic topic labels and top-terms.
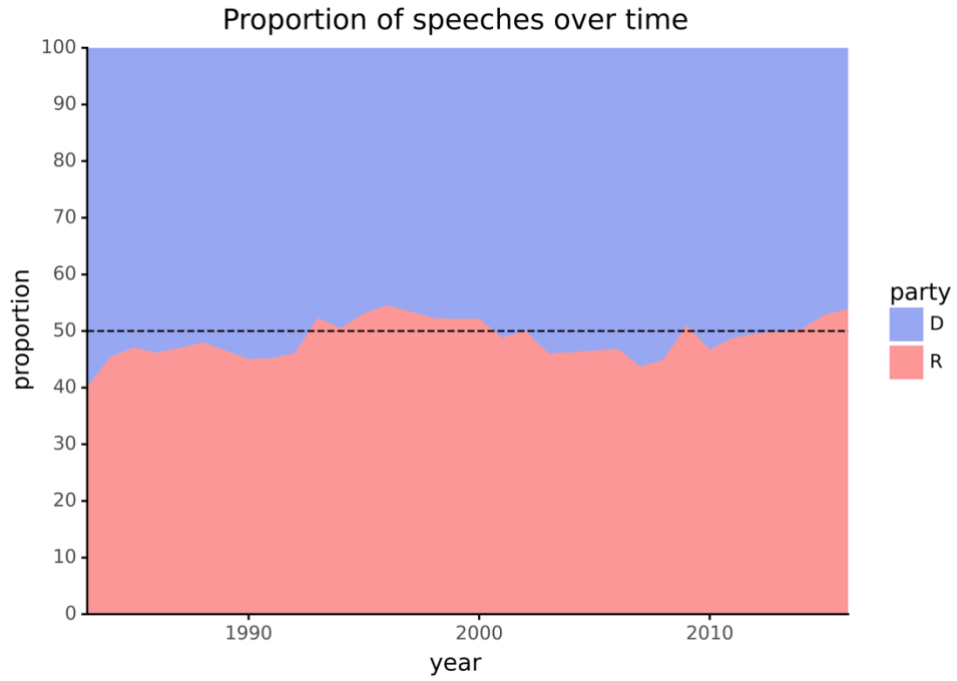
Figure S2.1: Proportion of speeches for Democrats and Republicans across time. In the classification analysis, the relatively low imbalance observed in this plot was corrected for through a weighting procedure.
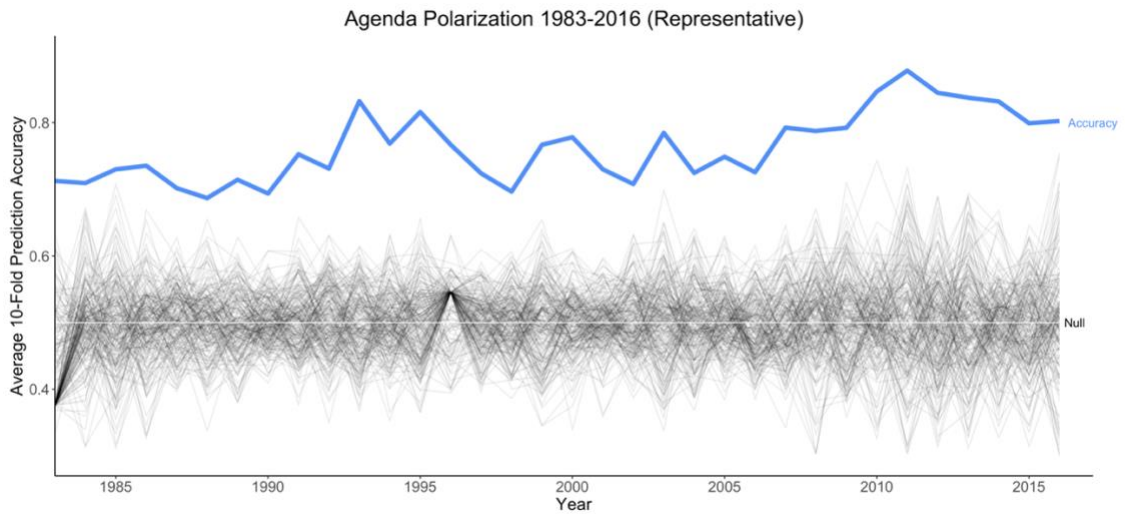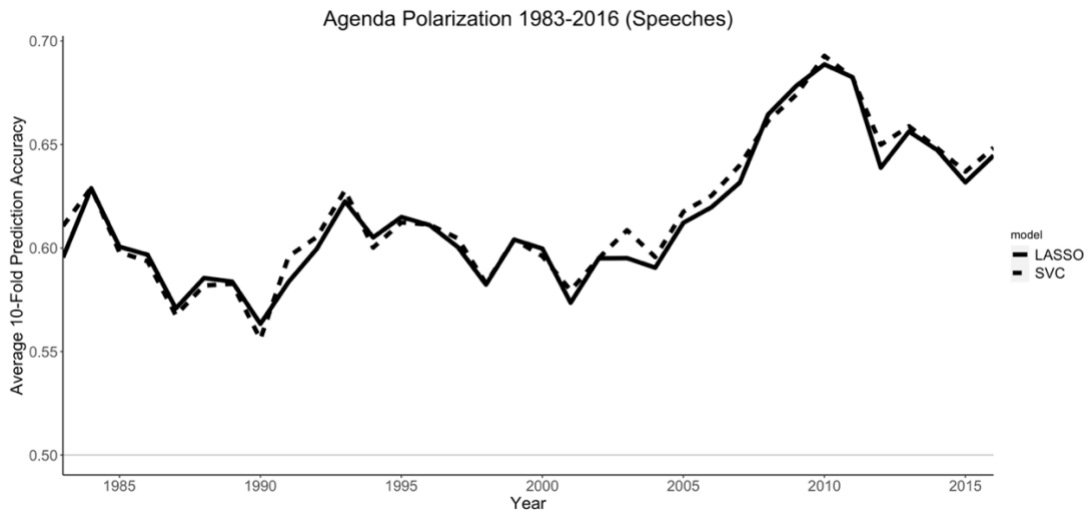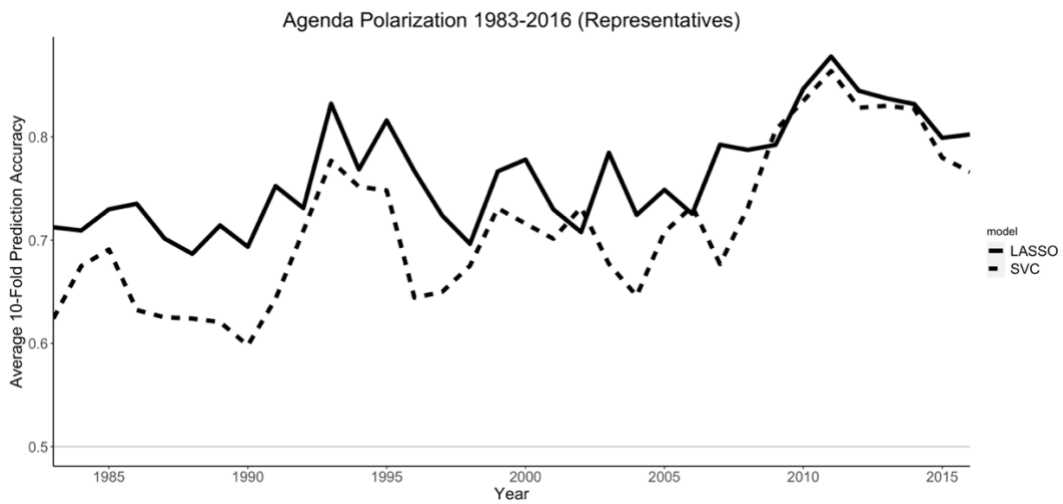
A.
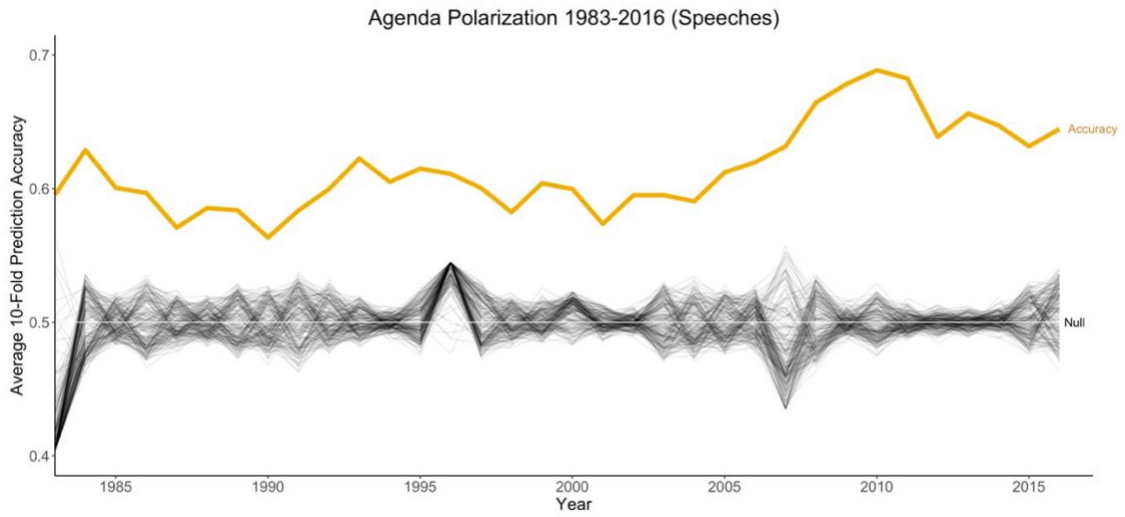
B.

Figure S2.2: Agenda polarization with null distributions, operationalized as average 10-fold out-of-sample classification accuracy for A) speeches and B) a supplemental analysis in which speech probabilities were aggregated at the level of legislator and evaluated for predictive accuracy. Gray lines represent 200 null models fit using shuffled party class labels.
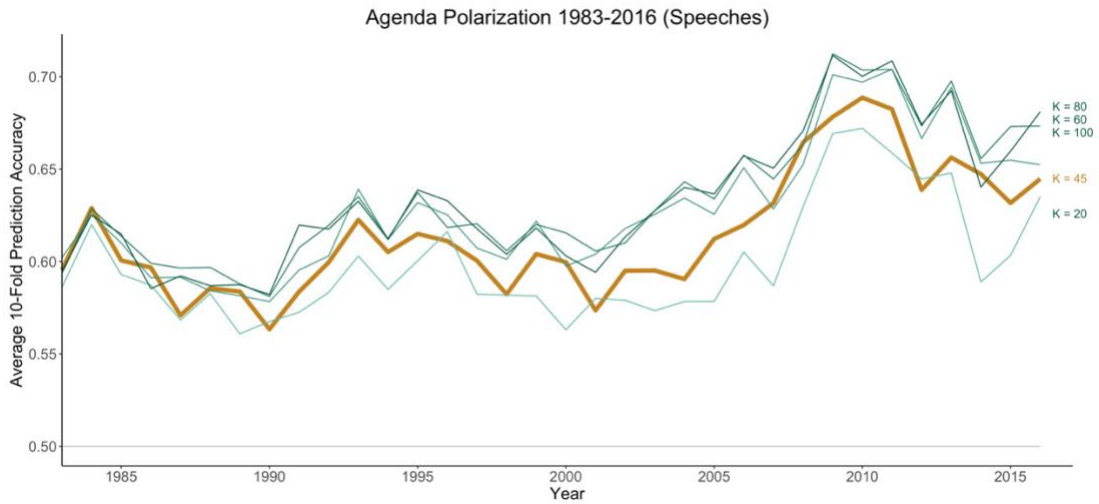
A.

B.

Figure S2.3: Agenda polarization with alternative classification algorithm, operationalized as average 10-fold out-of-sample classification accuracy for A) speeches and B) a supplemental analysis in which speech probabilities were aggregated at the level of legislators and evaluated for predictive accuracy. Black line represents main results with LASSO, dashed line represents results obtained with a linear kernel support vector classifier.
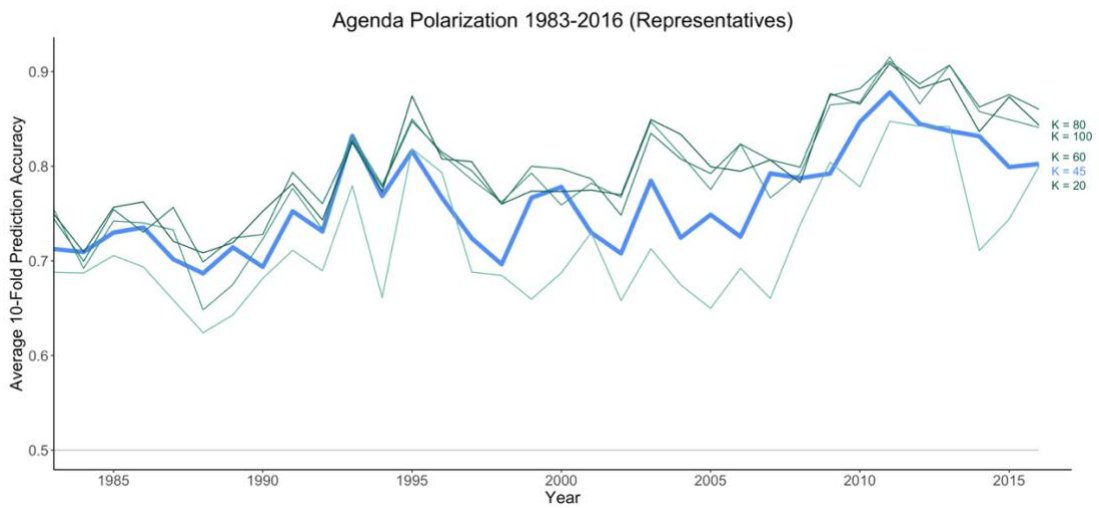
A.

B.

Figure S2.4: Agenda polarization with different K topic models. operationalized as average 10-fold out-of-sample classification accuracy for A) speeches and B) a supplemental analysis in which speech probabilities were aggregated at the level of legislators and evaluated for predictive accuracy. Green lines indicate alternative topic models with values of k 20, 60, 80, and 100.
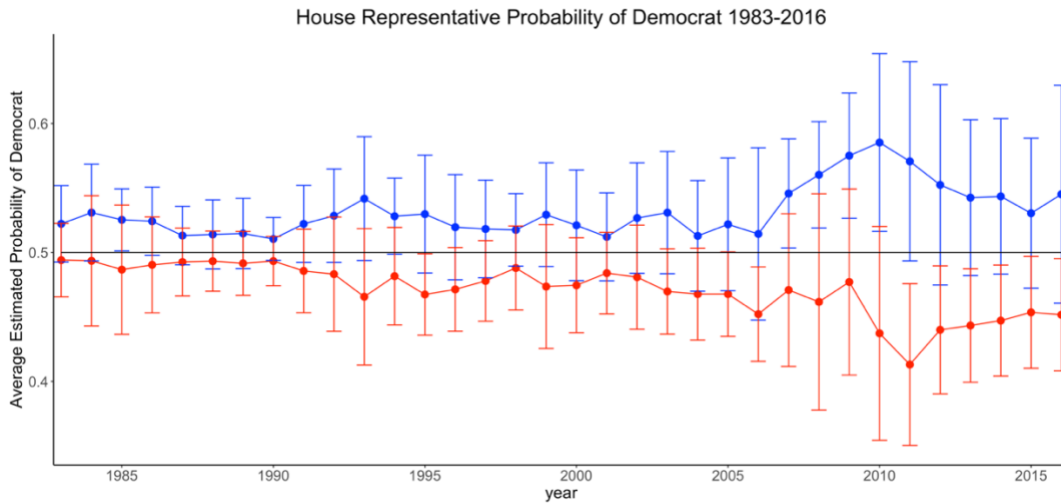
Figure S2.5: Average and SD of predicted partisanship of legislators based on aggregated speech probabilities. Figure shows a general symmetric polarization trend across time.
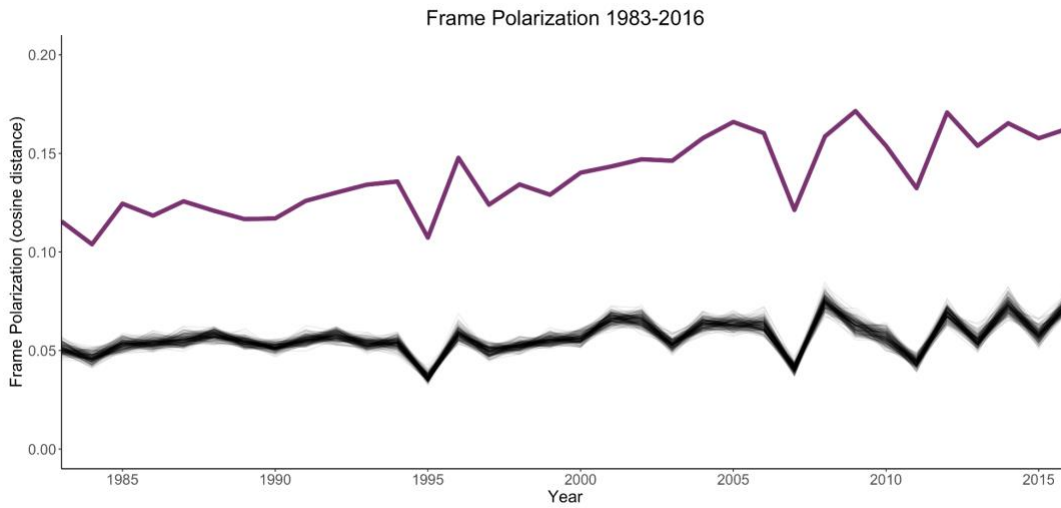

Figure S3.1: Average frame polarization as measured by cosine distance with 200 null models derived from shuffling legislators' party labels.
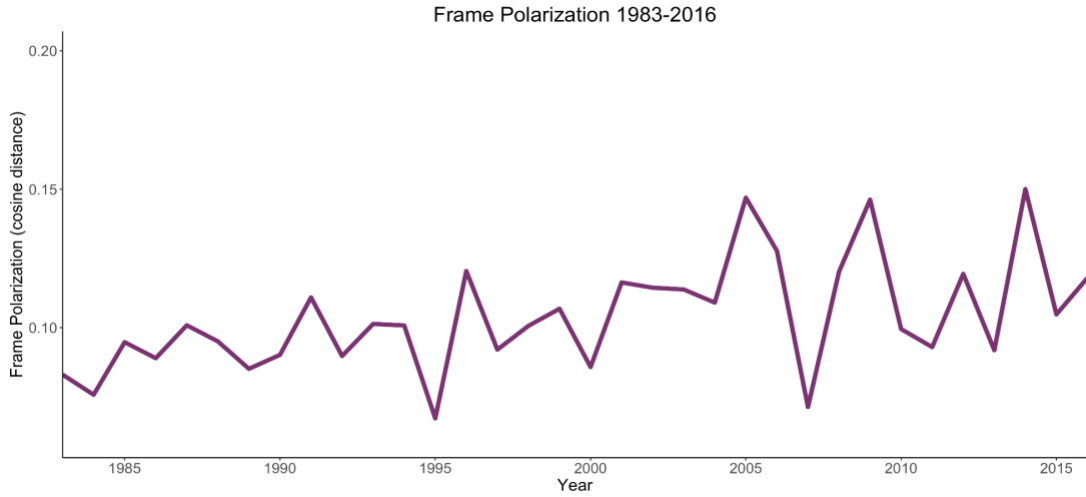
101

\Figure S3.2: Median frame polarization as measured by cosine distance. Pattern of results mirror closely those derived from averaging cosine distance.
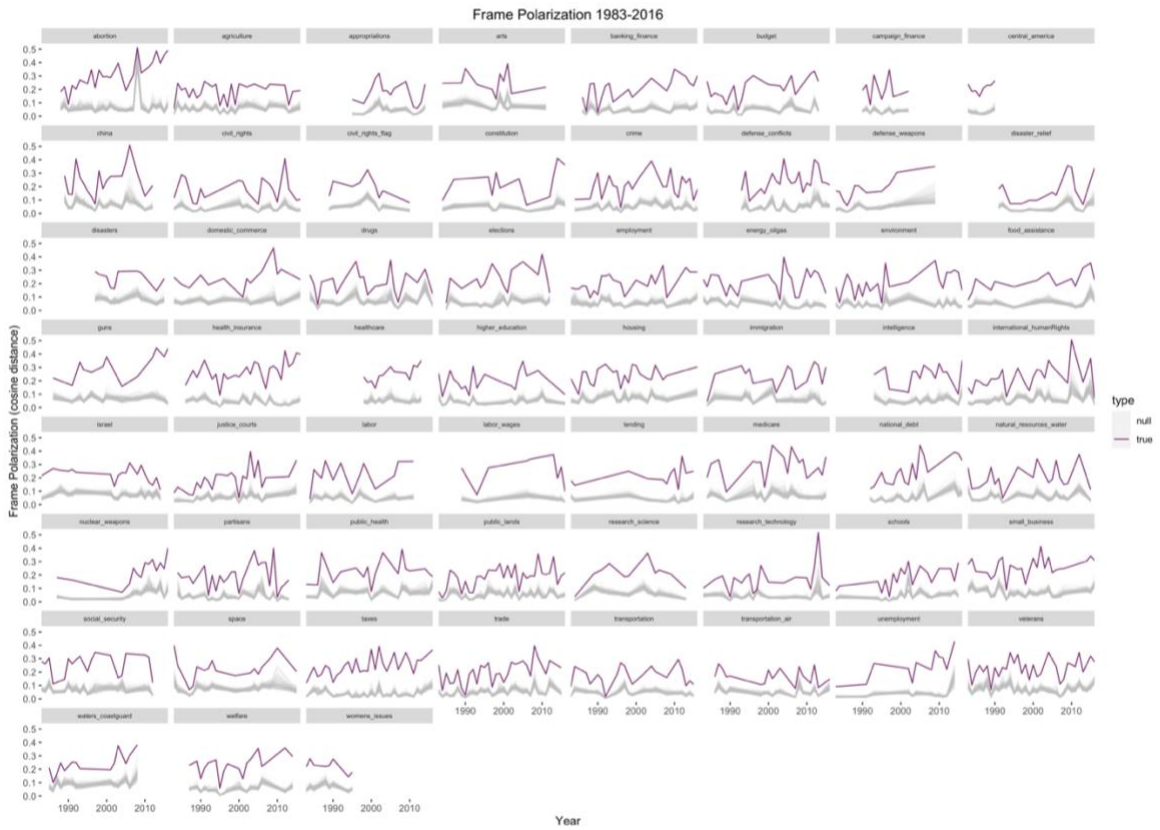

Figure S3.3: Topic specific cosine distance over time. Purple lines indicate true distance values, grey band represents distribution of 200 null models.
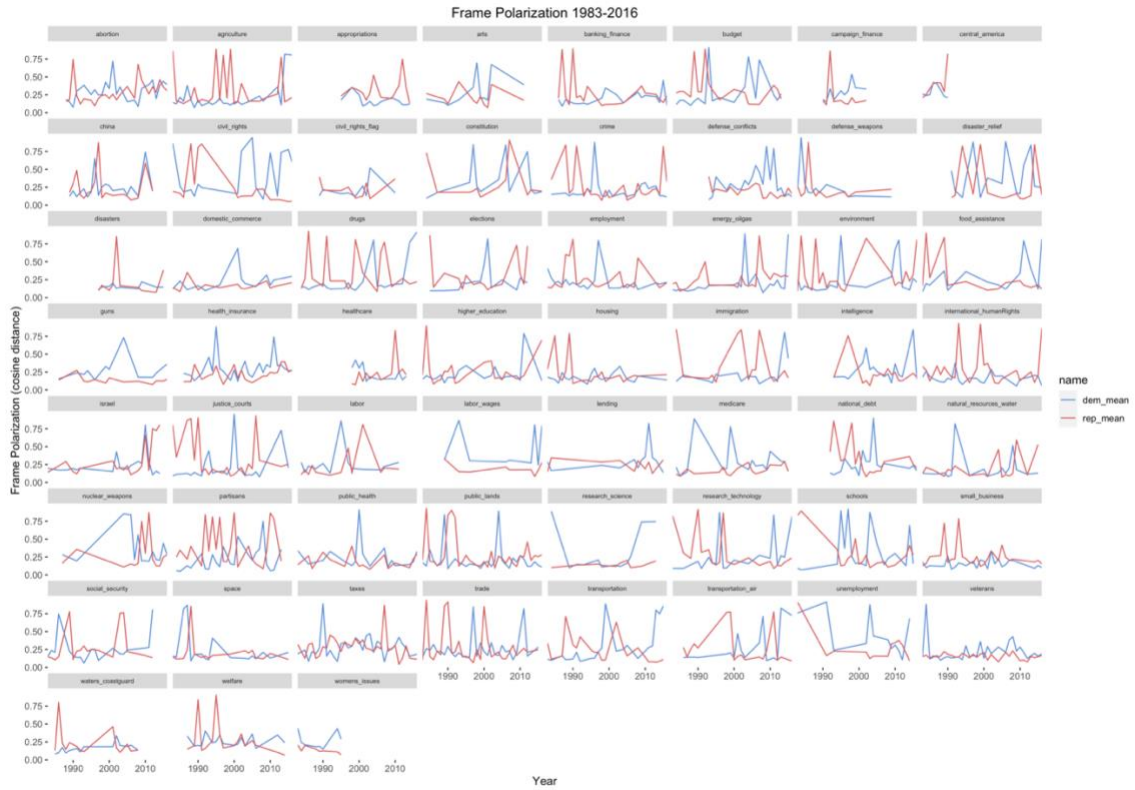
102

Figure S3.4: Topic specific intra-party cosine similarity (frame consolidation). Red lines indicate Republican frame consolidation, blue lines indicate Democrat frame consolidation.

Predictors	distance		
	Estimates	CI	p
(Intercept)	0.22	0.21 – 0.23	**<0.001**
year	0.03	0.02 – 0.03	**<0.001**
Random Effects			
σ^2	0.01		
τ_{00} topic	0.00		
τ_{11} topic.year	0.00		
ρ_{01} topic	0.63		
ICC	0.13		
N topic	59		
Observations	1089		
Marginal R^2 / Conditional R^2	0.093 / 0.213		

Table S.1: Mixed effect model for frame polarization. Model indicates a significant positive effect of time on frame polarization.

Predictors	value		
	Estimates	CI	p
(Intercept)	-7.57	-10.00 – -5.14	**<0.001**
name [rep mean]	12.55	9.11 – 15.99	**<0.001**
year	0.00	0.00 – 0.01	**<0.001**
name [rep mean] * year	-0.01	-0.01 – -0.00	**<0.001**
Observations	2178		
R^2 / R^2 adjusted	0.025 / 0.023		

Table S.2: Simple linear model for frame consolidation, indicating a significant interaction effect of time and party.
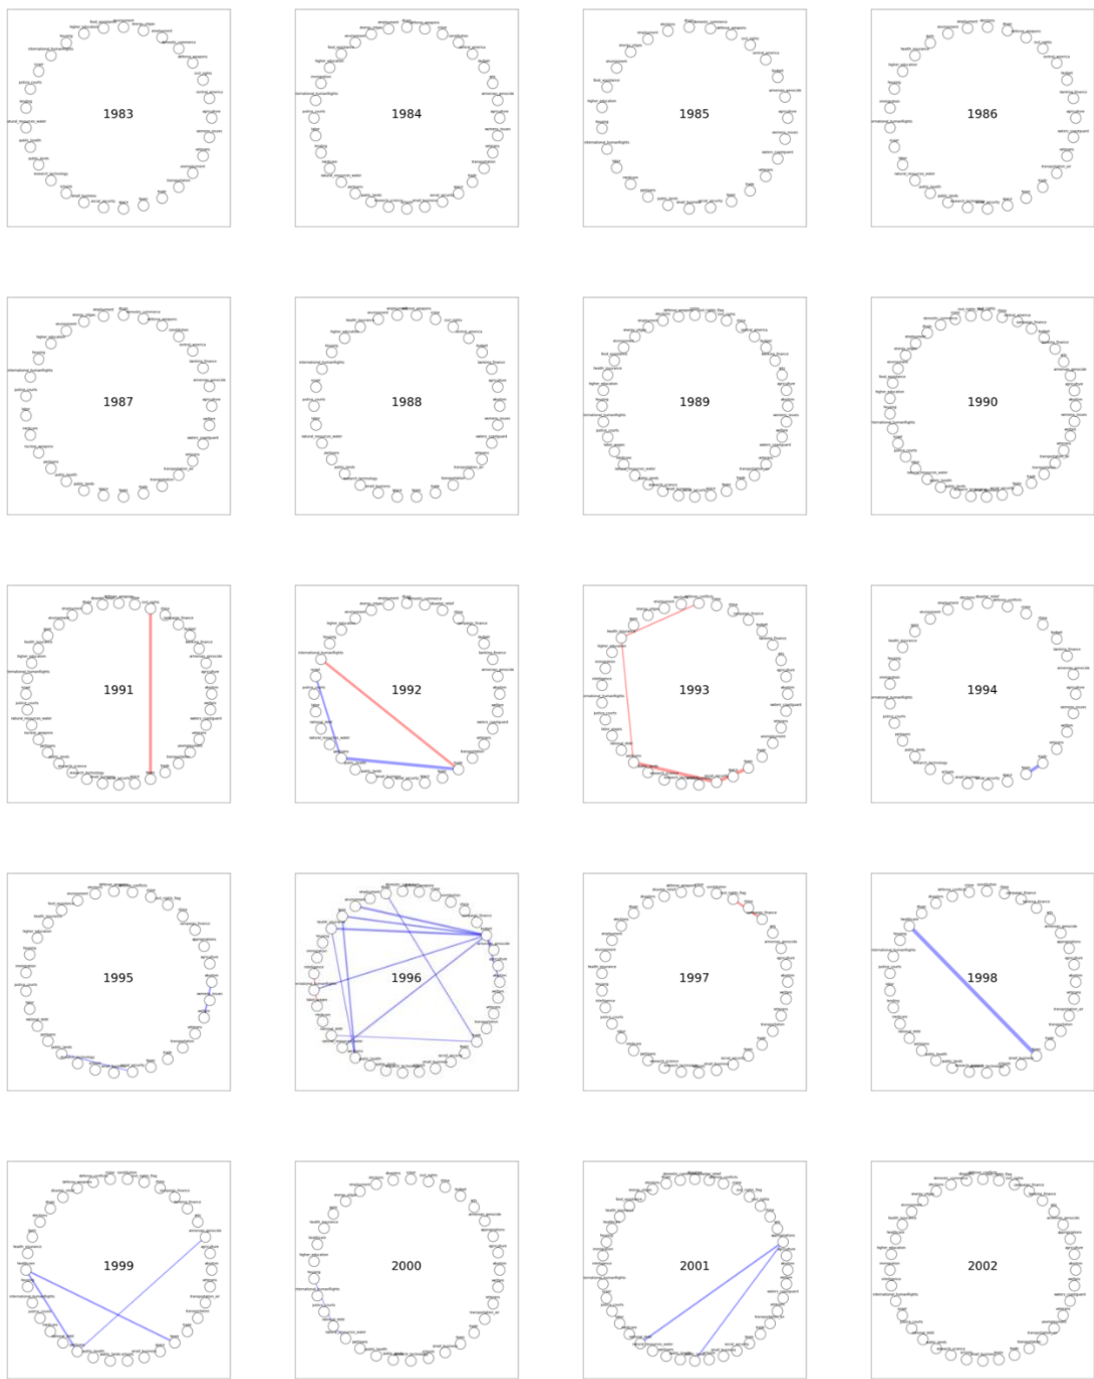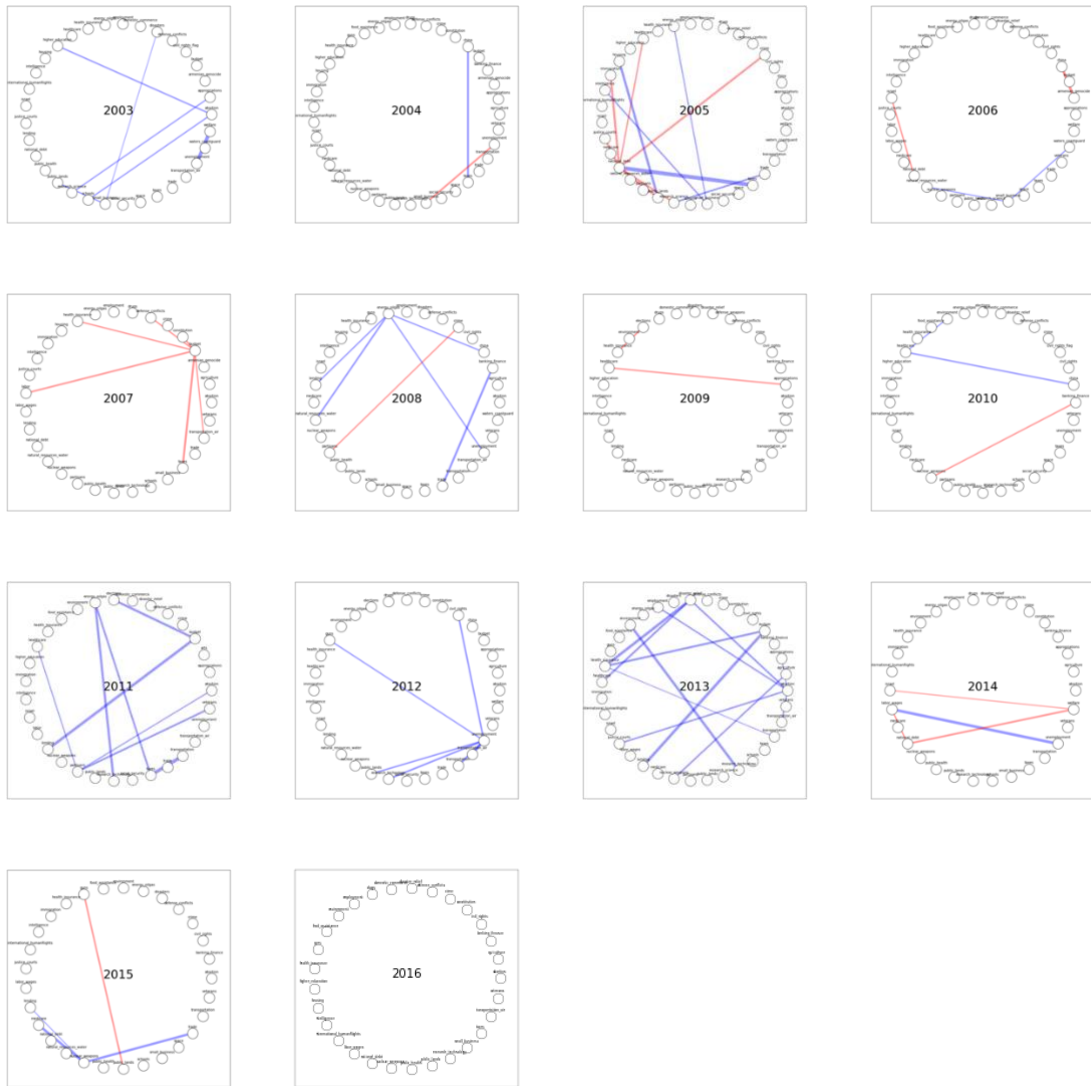
Figure S4.1: Individual belief network edge differences (part 1) 1983 – 2002. In addition to analyzing differences in the overall structure of belief systems, individual edge strength was compared for every edge within networks across time. These results were also dramatically affected by high-dimensionality and relatively low numbers of observations. Despite this, edges for some years did survive correction, however no discernable pattern was observed. Blue edges weights indicate greater strength In Democrat Belief networks, red edges indicate greater weight in Republican belief networks; Edge thickness indicates magnitude of difference scaled between 0 and 1 for each network respectively. Edges were tested for significance against permutation tests (n=1000) and corrected for multiple comparison with the FWER set to alpha = 0.05.

WORKS CITED

Abramowitz, A. I., & Saunders, K. L. (1998). Ideological realignment in the US electorate. *The Journal of Politics*, *60*(3), 634–652.

Amer, M. L., Government, & Division, F. (2001). *A User's Guide to the Congressional Record*. https://fas.org/sgp/crs/misc/98-265.pdf

Anderson, J. R., & Others. (1983). A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior*, *22*(3), 261–295.

Ansolabehere, S., Snowberg, E. C., & Snyder, J. M. (2003). Statistical bias in newspaper reporting on campaign finance. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.463780

Arceneaux, K. (2012). Cognitive biases and the strength of political arguments. *American Journal of Political Science*, *56*(2), 271–285.

Bakker, B. N., Lelkes, Y., & Malka, A. (2020). Understanding Partisan Cue Receptivity: Tests of Predictions from the Bounded Rationality and Expressive Utility Perspectives. *The Journal of Politics*, *82*(3), 1061–1077.

Baldassarri, D., & Bearman, P. (2007). Dynamics of Political Polarization. *American Sociological Review*, *72*(5), 784–811.

Baldassarri, D., & Gelman, A. (2008). Partisans without Constraint: Political Polarization and Trends in American Public Opinion. *AJS; American Journal of Sociology*, *114*(2), 408–446.

Baldassarri, D., & Goldberg, A. (2014). Neither ideologues nor agnostics: alternative voters' belief system in an age of partisan politics. *AJS; American Journal of Sociology*, *120*(1), 45–95.

Bateman, D. A., & Lapinski, J. (2016). Ideal Points and American Political Development: Beyond DW-NOMINATE. *Studies in American Political Development*, *30*(2), 147–171.

Baumgartner, F. R., Breunig, C., & Grossman, E. (2019). *Comparative policy agendas: Theory, tools, data*. Oxford University Press.

Berry, M. W., Browne, M., Langville, A. N., Pauca, V. P., & Plemmons, R. J. (2007). Algorithms and applications for approximate nonnegative matrix factorization. *Computational Statistics & Data Analysis*, *52*(1), 155–173.

Bevan, S. (2019). The Creation of the Comparative Agendas Project Master Codebook. *Comparative Policy Agendas: Theory, Tools, Data*, 17.

Blei, D. M., & Lafferty, J. D. (2006). Dynamic topic models. *Proceedings of the 23rd International Conference on Machine Learning - ICML '06*. the 23rd international conference, Pittsburgh, Pennsylvania. https://doi.org/10.1145/1143844.1143859

Blei, D. M., & Lafferty, J. D. (2007). A correlated topic model of Science. *The Annals of Applied Statistics*, *1*(1), 17–35.

Bolin, J. L., & Hamilton, L. C. (2018). The News You Choose: news media preferences amplify views on climate change. *Environmental Politics*, *27*(3), 455–476.

Bond, R., & Messing, S. (2015). Quantifying Social Media's Political Space: Estimating Ideology from Publicly Revealed Preferences on Facebook. *The American Political Science Review*, *109*(1), 62–78.

Boutsidis, C., & Gallopoulos, E. (2008). SVD based initialization: A head start for nonnegative matrix factorization. *Pattern Recognition*, *41*(4), 1350–1362.

Boutyline, A., & Vaisey, S. (2017). Belief Network Analysis: A Relational Approach to Understanding the Structure of Attitudes. *The American Journal of Sociology*, *122*(5), 1371–1447.

Brandt, M. J., Sibley, C. G., & Osborne, D. (2019). What Is Central to Political Belief System Networks? *Personality & Social Psychology Bulletin*, *45*(9), 1352–1364.

Brandt, M. J., & Sleegers, W. (2020). *Evaluating Belief System Networks as a Theory of Political Belief System Dynamics*. https://doi.org/10.31234/osf.io/b8hfa

Brock, C. (2017). Framing child nutrition programs: The impact of party and district characteristics on elite framing. *Social Science Quarterly*, *98*(2), 628–643.

Brosius, H.-B., & Weimann, G. (1996). Who Sets the Agenda: Agenda-Setting as a Two-Step Flow. *Communication Research*, *23*(5), 561–580.

Camaj, L. (2019). Motivational Theories of Agenda-Setting Effects: An Information Selection and

      Processing Model of Attribute Agenda-Setting. *International Journal of Public Opinion*

      Research, *31*(3), 441–462.

Chang, J., Boyd-Graber, J., Wang, C., Gerrish, S., & Blei, D. M. (2009). Reading tea leaves: How

      humans interpret topic models. *Neural Information Processing Systems*, *22*, 288–296.

Chen, H.-T., Guo, L., & Su, C. C. (2020). Network Agenda Setting, Partisan Selective Exposure,

      and Opinion Repertoire: The Effects of Pro- and Counter-Attitudinal Media in Hong Kong.

      The Journal of Communication, *70*(1), 35–59.

Chong, D., & Druckman, J. N. (2011). Identifying frames in political news. *Sourcebook for Political*

      Communication Research: Methods, Measures, and Analytical Techniques, 238–267.

Cichocki, A., & Phan, A.-H. (2009). Fast local algorithms for large scale nonnegative matrix and

      tensor factorizations. *IEICE Transactions on Fundamentals of Electronics*

      Communications and Computer Sciences, *E92-A*(3), 708–721.

Cobb, R., Ross, J.-K., & Ross, M. H. (1976). Agenda building as a comparative political process.

      The American Political Science Review, *70*(1), 126–138.

Coleman, J. J. (1997). The Decline and Resurgence of Congressional Party Conflict. *The Journal*

      of Politics, *59*(1), 165–184.

Colgan, J. D. (2013). Fueling the Fire: Pathways from Oil to War. *International Security*, *38*(2),

      147–180.

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing.

      Psychological Review, *82*(6), 407–428.

Converse, P. E. (2006). The nature of belief systems in mass publics (1964). *Critical Review* ,

      18(1–3), 1–74.

DellaPosta, D. (2020). Pluralistic Collapse: The "Oil Spill" Model of Mass Opinion Polarization.

      American Sociological Review, *85*(3), 507–536.

Diaf, S., & Fritsche, U. (2021). Topic Scaling: A Joint Document Scaling -- Topic Model Approach

      To Learn Time-Specific Topics. In *arXiv [cs.IR]*. arXiv. http://arxiv.org/abs/2104.01117

DiMaggio, P. (2011). Cultural networks. *The Sage Handbook of Social Network Analysis*, 286–310.

Druckman, J. N. (2001). On the Limits of Framing Effects: Who Can Frame? *The Journal of Politics*, *63*(4), 1041–1066.

Egan, P. J. (2013). *Partisan Priorities: How Issue Ownership Drives and Distorts American Politics.* Cambridge University Press.

Entman, R. M. (1993). Framing: Towards clarification of a fractured paradigm. *McQuail's Reader in Mass Communication Theory*, 390–397.

Entman, R. M. (2003). Cascading Activation: Contesting the White House's Frame After 9/11. *Political Communication*, *20*(4), 415–432.

Esacove, A. W. (2004). Dialogic framing: The framing/counterframing of "partial-birth" abortion. *Sociological Inquiry*, *74*(1), 70–101.

Eysenck, M. W. (2001). *Principles of Cognitive Psychology*. Psychology Press.

Feldman, L., & Hart, P. S. (2018). Climate change as a polarizing cue: Framing effects on public support for low-carbon energy policies. *Global Environmental Change: Human and Policy Dimensions*, *51*, 54–66.

Fiorina, M. P. (2016). The political parties have sorted. *Hoover Institution Essay on Contemporary American Politics. Series*, *3*.
https://www.hoover.org/sites/default/files/research/docs/fiorina_3_finalfile.pdf

Fishman, N., & Davis, N. T. (2019). *Change we can believe in: Structural and content dynamics within belief networks.*
https://files.osf.io/v1/resources/6r3ym/providers/osfstorage/5d7030c45a289c00194a8a3b?action=download&version=2&direct&format=pdf

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.

Gallagher, R. J., Reing, K., Kale, D., & Ver Steeg, G. (2017). Anchored Correlation Explanation: Topic Modeling with Minimal Domain Knowledge. *Transactions of the Association for Computational Linguistics*, *5*, 529–542.

Gentzkow, M., & Shapiro, J. M. (2010). What drives media slant? Evidence from U.s. daily

    newspapers. *Econometrica: Journal of the Econometric Society*, *78*(1), 35–71.

Gentzkow, Matthew, Shapiro, J. M., & Taddy, M. (2019). Measuring Group Differences in High-

    Dimensional Choices: Method and Application to Congressional Speech. *Econometrica*,

    87(4), 1307–1340.

Gentzkow, Matthew, Shapiro, J., & Taddy, M. (2016). *Measuring Polarization in High-Dimensional*

    Data: Method and Application to Congressional Speech (id:11114). eSocialSciences.

    https://ideas.repec.org/p/ess/wpaper/id11114.html

Gerring, J. (1997). Ideology: A Definitional Analysis. *Political Research Quarterly*, *50*(4), 957–994.

Gingrich, N., Armey, D., & Republican National Committee. (1995). *Contract with America*.

    Random House Audio Publishing Group.

Glazier, R. A., & Boydstun, A. E. (2012). The President, the Press, and the War: A Tale of Two

    Framing Agendas. *Political Communication*, *29*(4), 428–446.

Glazier, R. A., Boydstun, A. E., & Feezell, J. T. (2021). Self-coding: A method to assess semantic

    validity and bias when coding open-ended responses. *Research & Politics*, *8*(3),

    20531680211031750.

Greene, D., & Cross, J. P. (2017). Exploring the Political Agenda of the European Parliament

    Using a Dynamic Topic Modeling Approach. *Political Analysis: An Annual Publication of*

    the Methodology Section of the American Political Science Association, *25*(1), 77–94.

Grimmer, J. (2010). A Bayesian Hierarchical Topic Model for Political Texts: Measuring

    Expressed Agendas in Senate Press Releases. *Political Analysis: An Annual Publication*

    of the Methodology Section of the American Political Science Association, *18*(1), 1–35.

Grimmer, J., Roberts, M. E., & Stewart, B. M. (2021). Machine Learning for Social Science: An

    Agnostic Approach. *Annual Review of Political Science*, *24*(1), 395–419.

Grimmer, J., & Stewart, B. M. (2013). Text as Data: The Promise and Pitfalls of Automatic

    Content Analysis Methods for Political Texts. *Political Analysis: An Annual Publication of*

    the Methodology Section of the American Political Science Association, *21*(3), 267–297.

Gropp, C., Herzog, A., Safro, I., Wilson, P. W., & Apon, A. W. (2016). Scalable Dynamic Topic
Modeling with Clustered Latent Dirichlet Allocation (CLDA). In *arXiv [cs.IR]*. arXiv.
http://arxiv.org/abs/1610.07703

Gruszczynski, M. W., & Michaels, S. (2012). The evolution of elite framing following enactment of
legislation. *Policy Sciences*, *45*(4), 359–384.

Guo, L., & McCombs, M. (2011). Network agenda setting: A third level of media effects. *Annual
Conference of the International Communication Association, Boston, MA*.
http://www.leiguo.net/publications/guo_nas_2011_ica.pdf

Guo, L., Vargo, C. J., Pan, Z., Ding, W., & Ishwar, P. (2016). Big Social Data Analytics in
Journalism and Mass Communication: Comparing Dictionary-Based Text Analysis and
Unsupervised Topic Modeling. *Journalism & Mass Communication Quarterly*, *93*(2), 332–
359.

Hagar, N., Wachs, J., & Horvát, E.-Á. (2021). Publishing patterns reflect political polarization in
news media. In *arXiv [cs.SI]*. arXiv. http://arxiv.org/abs/2101.05044

Harris, D. B. (2005a). Orchestrating party talk: A party-based view of one-minute speeches in the
house of representatives. *Legislative Studies Quarterly*, *30*(1), 127–141.

Harris, D. B. (2005b). House Majority Party Leaders' Uses of Public Opinion Information.
Congress & the Presidency, *32*(2), 133–155.

Hemphill, L., Culotta, A., & Heston, M. (2013). Framing in Social Media: How the US Congress
Uses Twitter Hashtags to Frame Political Issues. In *Available at SSRN 2317335*.
https://doi.org/10.2139/ssrn.2317335

Hopkins, D. J. (2018). The Exaggerated Life of Death Panels? The Limited but Real Influence of
Elite Rhetoric in the 2009–2010 Health Care Debate. *Political Behavior*, *40*(3), 681–709.

Hopkins, D. J., Schickler, E., & Azizi, D. (2021). *From Many Divides, One? The Polarization and
Nationalization of American State Party Platforms, 1918-2017*.
https://papers.ssrn.com/abstract=3772946

Iyengar, S., Hahn, K. S., Krosnick, J. A., & Walker, J. (2008). Selective Exposure to Campaign

      Communication: The Role of Anticipated Agreement and Issue Public Membership. *The*

      Journal of Politics, *70*(1), 186–200.

Iyengar, S., & Simon, A. (1993). News Coverage of the Gulf Crisis and Public Opinion: A Study of

      Agenda-Setting, Priming, and Framing. *Communication Research*, *20*(3), 365–383.

Jang, H.-Y., & Barnett, G. A. (1994). Cultural Differences in Organizational Communication: a

      Semantic Network Analysis 1. *Bulletin of Sociological Methodology/Bulletin de*

      Méthodologie Sociologique, *44*(1), 31–59.

Jelen, T. G., & Wilcox, C. (2003). Causes and Consequences of Public Attitudes Toward

      Abortion: A Review and Research Agenda. *Political Research Quarterly*, *56*(4), 489–500.

Jensen, J., Naidu, S., Kaplan, E., Wilse-Samson, L., Gergen, D., Zuckerman, M., & Spirling, A.

      (2012). Political Polarization and the Dynamics of Political Language: Evidence from 130

      Years of Partisan Speech [with Comments and Discussion]. *Brookings Papers on*

      Economic Activity, 1–81.

Jensen, K. B., Rothenbuhler, E. W., Pooley, J. D., & Craig, R. T. (Eds.). (2016). *The International*

      Encyclopedia of Communication Theory and Philosophy (1st ed., Vol. 8, pp. 1–13). Wiley.

Jiang, X., Su, M.-H., Hwang, J., Lian, R., Brauer, M., Kim, S., & Shah, D. (2021). Polarization

      Over Vaccination: Ideological Differences in Twitter Expression About COVID-19 Vaccine

      Favorability and Specific Hesitancy Concerns. *Social Media + Society*, *7*(3),

      20563051211048412.

John, P. (2006). The policy agendas project: a review. *Journal of European Public Policy*, *13*(7),

      975–986.

Jones, B. D., & Baumgartner, F. R. (2004). Representation and Agenda Setting. *Policy Studies*

      Journal: The Journal of the Policy Studies Organization, *32*(1), 1–24.

Jones, D. R. (2015). Declining Trust in Congress: Effects of Polarization and Consequences for

      Democracy. *Forum*, *13*(3), 375–394.

Jost, J. T. (2006). The end of the end of ideology. *The American Psychologist*, *61*(7), 651.

Jost, J. T., & Sterling, J. (2020). The language of politics: ideological differences in congressional communication on social media and the floor of Congress. *Social Influence*, *15*(2–4), 80–103.

Kalmoe, N. P. (2020). Uses and abuses of ideology in political psychology. *Political Psychology*, *41*(4), 771–793.

Kang, Y., & Yang, K. C. C. (2021). Communicating Racism and Xenophobia in the Era of Donald Trump: A Computational Framing Analysis of the US-Mexico Cross-Border Wall Discourses. *Howard Journal of Communications*, 1–20.

Kim, I. S., Londregan, J., & Ratkovic, M. (n.d.). *Voting, Speechmaking, and the Dimensions of Conflict in the US Senate*. https://www.princeton.edu/~ratkovic/public/sfa_all.pdf

Klar, S. (2022). *Political Dynamics of Framing*. http://dx.doi.org/

Kozlowski, A. C., & Murphy, J. P. (2019). *Issue Alignment and Partisanship in the American Public: Revisiting the "Partisans without Constraint" Thesis*. https://doi.org/10.31235/osf.io/jex9k

Krippendorff, K., & Bock, M. A. (2009). *The content analysis reader*. SAGE Publications.

Lafferty, J., & Blei, D. (2009). Topic Models. In *Text Mining* (pp. 71–93). https://doi.org/10.1201/9781420059458.ch4

Lakoff, G. (2004). Don't Think of an Elephant! Know Your Values and Frame the Debate (White River Junction. *VT: Chelsea Green Publishing*.

Lang, A. (2000). The limited capacity model of mediated message processing. *The Journal of Communication*. https://academic.oup.com/joc/article-abstract/50/1/46/4110103

Lauderdale, B. E., & Herzog, A. (2016). Measuring Political Positions from Legislative Speech. *Political Analysis: An Annual Publication of the Methodology Section of the American Political Science Association*, *24*(3), 374–394.

Layman, G. C., & Carsey, T. M. (2002). Party Polarization and "Conflict Extension" in the American Electorate. *American Journal of Political Science*, *46*(4), 786–802.

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

Lenz Gabriel, S. (2012). *Follow the Leader? How Voters Respond to Politicians' Policies and Performance*. Chicago, IL: University of Chicago Press.

Levendusky, M. S. (2010). Clearer Cues, More Consistent Voters: A Benefit of Elite Polarization. *Political Behavior*, *32*(1), 111–131.

Li, T., & Ding, C.-C. (2018). Nonnegative Matrix Factorizations for Clustering: A Survey. In *Data Clustering* (pp. 149–176). https://doi.org/10.1201/9781315373515-7

Lupia, A., Soroka, S., & Beatty, A. (2020). What does Congress want from the National Science Foundation? A content analysis of remarks from 1995 to 2018. *Science Advances*, *6*(33), eaaz6300.

Maier, D., Waldherr, A., Miltner, P., Wiedemann, G., Niekler, A., Keinert, A., Pfetsch, B., Heyer, G., Reber, U., Häussler, T., Schmid-Petri, H., & Adam, S. (2018). Applying LDA Topic Modeling in Communication Research: Toward a Valid and Reliable Methodology. *Communication Methods and Measures*, *12*(2–3), 93–118.

Martin, S., Saalfeld, T., Strøm, K. W., Slapin, J. B., & Proksch, S.-O. (2014). Words as Data. In *The Oxford Handbook of Legislative Studies*. Oxford University Press.

McCaffrey, D., & Keys, J. (2000). Competitive framing processes in the abortion debate: polarization-vilification, frame saving, and frame debunking. *The Sociological Quarterly*, *41*(1), 41–61.

McCarty, N. (2016). In Defense of DW-NOMINATE. *Studies in American Political Development*, *30*(2), 172–184.

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

McCombs, M. E., & Shaw, D. L. (1972). The agenda-setting function of mass media. *Public Opinion Quarterly*, *36*(2), 176.

McCombs, M., & Valenzuela, S. (2017). Agenda-Setting Theory. In Kate Kenski And Kathleen (Ed.), *The Oxford Handbook of Political Communication*. Oxford University Press.

Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., & Dean, J. (2013). Distributed Representations of Words and Phrases and their Compositionality. In C. J. C. Burges, L.

Bottou, M. Welling, Z. Ghahramani, & K. Q. Weinberger (Eds.), *Advances in Neural Information Processing Systems* (Vol. 26). Curran Associates, Inc. https://proceedings.neurips.cc/paper/2013/file/9aa42b31882ec039965f3c4923ce901b-Paper.pdf

Mohr, J. W. (1998). Measuring meaning structures. *Annual Review of Sociology*. https://www.annualreviews.org/doi/abs/10.1146/annurev.soc.24.1.345?casa_token=2RrK uWlYAcMAAAAA:3tBhqCg4no1fz3X_urWpkjH9i-lGdg0w38KMRkZJEqaUuKgJSGz4Z2iY-5VzRklaQl8AqNg2CQhK_Fc

Mukerjee, S., Yang, T., Stadler, G., & González-Bailón, S. (2022). What counts as a weak tie? A comparison of filtering techniques to analyze co-exposure networks. *Social Networks*, *68*, 386–393.

Nguyen, V.-A., Boyd-Graber, J., Resnik, P., & Miler, K. (2015). Tea party in the house: A hierarchical ideal point topic model and its application to republican legislators in the 112th congress. *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 1438–1448.

O'Callaghan, D., Greene, D., Carthy, J., & Cunningham, P. (2015). An analysis of the coherence of descriptors in topic modeling. *Expert Systems with Applications*, *42*(13), 5645–5657.

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., & Others. (2011). Scikit-learn: Machine learning in Python. *The Journal of Machine Learning Research*, *12*, 2825–2830.

Peterson, A., & Spirling, A. (2018). Classification Accuracy as a Substantive Quantity of Interest: Measuring Polarization in Westminster Systems. *Political Analysis: An Annual Publication of the Methodology Section of the American Political Science Association*, *26*(1), 120–128.

Peterson, E., Goel, S., & Iyengar, S. (2018). Echo chambers and partisan polarization: Evidence from the 2016 presidential campaign. *Unpublished Manuscript. Https://5harad.*

Com/Papers/Selecfive-Exposure. Pdf.

https://pdfs.semanticscholar.org/3e71/8f98d619e6ef6841990b14856b1348c4b176.pdf

Petrocik, J. R., Benoit, W. L., & Hansen, G. J. (2003). Issue Ownership and Presidential
Campaigning, 1952-2000. *Political Science Quarterly*, *118*(4), 599–626.

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

Proksch, S.-O., Lowe, W., Wäckerle, J., & Soroka, S. (2019). Multilingual sentiment analysis: A
new approach to measuring conflict in legislative speeches. *Legislative Studies Quarterly*,
44(1), 97–131.

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.

Roberts, M. E., Stewart, B. M., Tingley, D., Airoldi, E. M., & Others. (2013). The structural topic
model and applied social science. *Advances in Neural Information Processing Systems
Workshop on Topic Models: Computation, Application, and Evaluation*, *4*, 1–20.

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

Rodriguez, C. G., Moskowitz, J. P., Salem, R. M., & Ditto, P. H. (2017). Partisan selective
exposure: The role of party, ideology and ideological extremity over time. *Translational
Issues in Psychological Science*, *3*(3), 254.

Ronen, S., Gonçalves, B., Hu, K. Z., Vespignani, A., Pinker, S., & Hidalgo, C. A. (2014). Links
that speak: the global language network and its association with global fame.
Proceedings of the National Academy of Sciences of the United States of America,
111(52), E5616-22.

Sakamoto, T., & Takikawa, H. (2017). Cross-national measurement of polarization in political
discourse: Analyzing floor debate in the U.S. the Japanese legislatures. *2017 IEEE
International Conference on Big Data (Big Data)*, 3104–3110.

Scheufele, D. A. (2009). Agenda-Setting, Priming, and Framing Revisited: Another Look at

Cognitive Effects of Political Communication. *Mass Communication & Society*, *3*(2), 297–

316.

Sellers, P. J. (2000). Manipulating the message in the U.s. congress. *Harvard International

Journal of Press/Politics*, *5*(1), 22–31.

Shah, D. V., McLeod, D. M., & Gotlieb, M. R. (2009). Framing and agenda setting. *The Sage

Handbook Of*.

https://books.google.ca/books?hl=en&lr=&id=CMO1aEWrzacC&oi=fnd&pg=PT91&ots=s

BIFMI0arF&sig=N0vMeuqqhm5yoOjHrEaFdqMH36Q

Sniderman, P. M., & Theriault, S. M. (2018). CHAPTER 5: The Structure of Political Argument

and the Logic of Issue Framing. In *Studies in Public Opinion* (pp. 133–165). Princeton

University Press.

Stonecash, J. M., Brewer, M. D., & Mariani, M. D. (2018). *Diverging parties: Social change,

realignment, and party polarization*. Routledge.

Sulkin, T., & Schmitt, C. (2014). Partisan Polarization and Legislators' Agendas. *Polity*, *46*(3),

430–448.

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

Tsuge, S., Shishibori, M., Kuroiwa, S., & Kita, K. (2001). Dimensionality reduction using non-

negative matrix factorization for information retrieval. *2001 IEEE International Conference

on Systems, Man and Cybernetics. e-Systems and e-Man for Cybernetics in Cyberspace

(Cat.No.01CH37236)*, *2*, 960–965 vol.2.

Van Borkulo, C. D., Boschloo, L., Kossakowski, J., Tio, P., Schoevers, R. A., Borsboom, D., &

Waldorp, L. J. (2017). Comparing network structures on three aspects: A permutation

test. *Manuscript Submitted for Publication*, *10*.

https://www.researchgate.net/profile/Claudia-Van-Borkulo-

2/publication/357057237_Comparing_Network_Structures_on_Three_Aspects_A_Permu

tation_Test/links/61bdd53a1d88475981fac9b1/Comparing-Network-Structures-on-Three-

Aspects-A-Permutation-Test.pdf

Vargo, C. J., & Guo, L. (2016). Networks, Big Data, and Intermedia Agenda Setting: An Analysis of Traditional, Partisan, and Emerging Online U.S. News. *Journalism & Mass Communication Quarterly*, *94*(4), 107769901667997.

Vavreck, L. (2009). *The Message Matters*. Princeton University Press.

Vu, H. T., Guo, L., & McCombs, M. E. (2014). Exploring "the World Outside and the Pictures in Our Heads": A Network Agenda-Setting Study. *Journalism & Mass Communication Quarterly*, *91*(4), 669–686.

Walter, D., & Ophir, Y. (2019). News Frame Analysis: An Inductive Mixed-method Computational Approach. *Communication Methods and Measures*, *13*(4), 248–266.

Weaver, D. (1991). ISSUE SALIENCE AND PUBLIC OPINION: ARE THERE CONSEQUENCES OF AGENDA-SETTING? *International Journal of Public Opinion Research*, *3*(1), 53–68.

Xiong, Y., Cho, M., & Boatwright, B. (2019). Hashtag activism and message frames among social movement organizations: Semantic network analysis and thematic analysis of Twitter during the #MeToo movement. *Public Relations Review*, *45*(1), 10–23.

Yin, J., & Wang, J. (2014). A dirichlet multinomial mixture model-based approach for short text clustering. *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 233–242.

Zalesky, A., Fornito, A., & Bullmore, E. T. (2010). Network-based statistic: identifying differences in brain networks. *NeuroImage*, *53*(4), 1197–1207.

Zaller, J. (1991). Information, Values, and Opinion. *The American Political Science Review*, *85*(4), 1215–1237.