PSJ Policy Studies Journal

The Policy Studies Journal, Vol. 00, No. 00, 2015

Modeling Issue Definitions Using Quantitative Text Analysis

Matthew C. Nowlin

Issue definitions, the way policy issues are understood, are an important component for understanding the policymaking process. Research on issue definitions has been divided between a macro level that examines collective issue definitions and a micro level focusing on the ways in which policy actors frame policy issues. This article develops a model of issue definitions that assumes issues are multidimensional, competition exists among policy actors in defining issues, and that collective issue definitions can be understood as the aggregation of individual issue definitions. This model is then estimated using quantitative text analysis. While various approaches to text analysis and categorization have been used by scholars, latent Dirichlet allocation (LDA), a specific type of topic modeling, is used to estimate issue definitions. Using LDA, witness testimony taken from Congressional hearings that occurred from 1975 to 2012 about the issue of used nuclear fuel (UNF) is examined and seven distinct dimensions of the UNF debate are estimated. The construct validity of these dimensions is checked by testing them against two major policy changes that occurred in the UNF domain. I conclude with a discussion of the strengths and weakness of topic modeling, and how this approach could be used to test hypotheses drawn from several of the major policymaking theories.

KEY WORDS: issue definitions, policy process, text analysis

Introduction

Many of the policy issues facing modern societies are complex and contain multiple attributes. This complexity often means that policy issues can potentially be defined or understood in a variety of different ways. Policy scholars have long noted that how issues are defined can have profound impacts on the politics of the issue and the ultimate policy choices (Baumgartner & Jones, 1993; Kingdon, 1984; Rochefort & Cobb, 1993; Stone, 1989; Wood & Doan, 2003; Wood & Vedlitz, 2007). Policy process scholars use various concepts and focus on different aspects of the policy process across multiple theories; however, several of these theories are interested in the way that issues come to be understood, and thereby defined by elite policy actors and the mass public. Issue definitions (and redefinitions) can play a role in getting an issue on the policymaking agenda, facilitating punctuated policy changes, policy

Published by Wiley Periodicals, Inc., 350 Main Street, Malden, MA 02148, USA, and 9600 Garsington Road, Oxford, OX4 2DQ.

design choices, and coalition formation. Given the importance of issue definitions for understanding the policymaking process, finding a way to theoretically and empirically model issue definitions could produce important insights based on empirical tests of hypotheses drawn from the multiple policy process theories and frameworks.

Issue definitions are largely communicated, and contested, through debate and argument that occurs within political institutions and public forums. Often these arguments are transcribed as part of a public record, and examining these records empirically requires the use of sophisticated methods of quantitative text analysis. Several approaches to quantifying texts, across multiple disciplines, have been developed. These approaches include human (or hand) coding, computer-assisted (supervised learning), and fully automated (unsupervised learning), with each approach having advantages and disadvantages. The choice of which approach to use should always be based on the research question and the theory, or theories, being tested. In this article, a theoretical model of issue definitions is developed and empirically estimated using an established, and fully automated, approach of text analysis. Specifically, latent Dirichlet allocation (LDA), a commonly used topic model (Mohr & Bogdanov, 2013), is applied to determine how used nuclear fuel (UNF) management has been defined over time. The data analyzed is drawn from the opening statements of witnesses at 140 Congressional hearings about UNF that took place between 1975 and 2012.

The management of UNF—the radioactive material used in the production of nuclear energy—has been a scientifically, technically, and politically vexing policy issue for several decades. Currently in the United States, UNF is stored on-site at the 65 nuclear power plants that produce nuclear energy. This material was intended to be shipped and stored at Yucca Mountain, located about 100 miles from Las Vegas, Nevada. However, in February 2010 the Obama Administration asked the Department of Energy (DOE) to withdraw the license application for Yucca Mountain that had previously been submitted to the Nuclear Regulatory Commission, leaving the question of how this material should be managed unanswered. The examination of the issue of UNF management allows important insights into the policymaking process, particularly within policy areas that are steeped in scientific and technical information, controversial, and possess a long time horizon.[1] These qualities make considerations of how policy issues are understood by policy actors even more vital.

This article develops a theoretical model of issue definitions that could be used to test hypotheses from multiple policy process frameworks. The model is based on scholarship from the policy process literature and the framing literature within political science. In addition, LDA and its potential use as a method for estimating issue definitions are discussed. The next section describes the theoretical approach to understanding issue definitions. Following that, various ways that texts can be analyzed quantitatively, including the advantages and disadvantages of each, is discussed. The subsequent section discusses LDA, the approach used to estimate issue definitions. Then the issue of UNF is examined using LDA and seven dimensions, deemed to have semantic validity, are derived based on witness testimony in Congressional hearings. Next, the construct validity of those dimensions is checked by predicting shifts in several of the dimensions based on two major policy changes. I find that these policy changes impact the degree to which the various dimensions are discussed in a systematic way, providing evidence for the validity of LDA as a way to model issue definitions. Finally, I conclude with a discussion of how this approach may be used to test hypotheses drawn from the various policy process theories.

Issue Definitions and the Policy Process

Policy scholars have long been concerned with how policy issues are understood and the implications of those understandings for the policymaking process and for policy outcomes. Terms such as "problem definitions," "policy images," and "frames" have each been used to describe the way that policy issues are defined. Overall, research on issue definitions has operated at two levels, (1) a *macro* level that examines aggregate, or collective, issue definitions and (2) a *micro* level examining attempts by policy actors to redefine or reframe policy issues (Baumgartner & Mahoney, 2008; Wood & Vedlitz, 2007). The approach described below offers a way to combine these two aspects of issue definitions by aggregating micro-issue definitions to estimate the collective definition of an issue.

In policy research, collective issue definitions have been shown to influence the ways in which "conditions" that exist in the world become salient "problems" (or issues) on the policymaking agenda (Baumgartner & Jones, 1993; Kingdon, 1984; Rochefort & Cobb, 1993; Stone, 1989; Wood & Doan, 2003). Multiple issues compete for agenda access, and collective issue definitions can play a large role in determining which issues reach the agenda. For example, Rochefort and Cobb (1993) argue that problems are defined along several dimensions including problem causation, nature of the problem, characteristics of the problem population, and the nature of the solutions. How problems are understood along these dimensions can impact whether they reach the policymaking agenda. In addition, Punctuated Equilibrium Theory (PET) notes the important role that policy images—a type of collective issue definition—play in getting issues on the agenda and subsequently facilitating policy change. Within PET, changing policy images interact with policy venues to produce large-scale punctuated policy changes (Baumgartner & Jones 1993). For example, the nuclear energy subsystem experienced rapid shifts as the policy image changed from positive to negative and attracted new policy actors across new policy venues. In general, policy issues reach the agenda when the salience (i.e., importance) of that issue increases. Often this increased salience is a result of the issue being redefined. While collective issue definitions help to explain how problems reach the policymaking agenda, attempts to define or frame a policy issue do not end once they are on the policymaking agenda.

The many issues competing for agenda access have, in past research, been characterized as dimensions in a multidimensional policy space (e.g., Downs, 1957; Hinich & Munger, 1994). However, given the complexity and ambiguity associated with many policy issues, the issues themselves are multidimensional (Jochim & Jones, 2013; Zahariadis, 2007). The dimensions of particular policy issues are the attributes or characteristics of the issue (Jones, Talbert, & Potoski, 2003).[2] Micro-level, or individual-level issue definitions are the ways in which the multiple dimensions of a specific issue are arranged by policy actors. Much like issues competing for agenda access, actors attempt to highlight one or more of the dimensions of the issue in an attempt to reframe or redefine the issue. Research on media framing provides a useful way to distinguish between the collective issue definitions of agenda setting, known as "first-order" agenda setting, and the micro-issue definitions based on the dimensions of a specific issue, which is termed "second-order" agenda setting (Boydstun, Glazier, & Phillips, 2013; Ghanem, 1997; McCombs, 2005; McCombs, Llamas, Lopez-Escobar, & Rey, 1997; Weaver, 2007). In the media framing literature, policy issues on the first-order agenda constitute *objects* and each of those objects consists of multiple *attributes* that make up the second-order agenda.

Due to bounded rationality and the resulting attention scarcity, only certain attributes (i.e., dimensions) of an issue may be salient at any point in time (Jones & Baumgartner, 2005). The combination of complexity (i.e., multidimensionality) and bounded rationality is what allows policy actors to frame (reframe) an issue. Framing occurs at the micro level of issue definitions and can be defined as the manipulation of dimensions in a multidimensional issue space. These manipulations can include adding dimensions to the issue space through emphasis framing, or shifting the salience of a single dimension through equivalency framing (Chong & Druckman, 2007).[3] These micro-issue definitions, when aggregated, become the macro-issue definition. The process of how micro definitions become collective definitions is the basis of the theoretical model of issue definitions.

Modeling Issue Definitions

The model of issue definitions proposed here seeks to lay a theoretical and empirical foundation for estimating how issues are defined. As noted, past literature divides issue definitions between a micro level where individuals attempt to reframe an issue to their advantage and a macro level representing the overall collective way that the issue is defined (Baumgartner & Mahoney, 2008; Wood & Vedlitz, 2007). The issue definition model presented here argues that the collective issue definition is the aggregation of definitions by policy actors. The basic assumption of the model is that issues are complex and contain multiple attributes and that these attributes can be thought of as existing in a multidimensional issue space. Issues are defined collectively based on the weights assigned to the various dimensions by policy actors. However, it is likely that policy actors, particularly in conflictual policy areas, will have differing issue definitions and be actively engaged in attempting to reframe the issue. These attempts to redefine a policy issue involve policy actors manipulating dimensions through highlighting one (or a few) dimensions in an attempt to draw attention to those particular dimensions. Indeed, these varying definitions and reframing attempts are what structure conflict in contentious policy areas and are what determine the overall collective definition of the issue.[4] Therefore, modeling the various issue definitions can provide insight into the nature of policy conflict and resolution.

As noted, the model of issue definitions assumes a multidimensional issue space, policy actors that vary in their definition of the issue, and that a policy issue is defined collectively by aggregating the issue definitions of individual policy actors. An issue definition for policy issue a, denoted θ_{a}, is an aggregation of the various issue dimensions weighted by the salience of each dimension.[5] The model is expressed simply and more formally as:

$$\theta_{\mathrm{a}} = \sum_{i=1}^{n} s_i$$

where s_i is the salience of dimension i for issue a.

Note that the model above is static, or the issue definition at time t. However, the salience of dimensions can vary over time making issue definitions dynamic. The dynamic element of issue definitions is a function of the changing salience s_i that results from competition between policy actors over time.[6] Witness testimony in Congressional hearings will be used to estimate the issue definition model, which requires the quantification of "text-as-data."

Quantitative Text Analysis

The analysis of text-as-data is a growing area of interest in political science and public policy research. More scholars are moving toward quantitative methods to examine the growing amount of text data that is becoming widely available through social media and the increasing digitization of historical documents (Hopkins & King, 2010). Quantitative text analysis seeks to categorize and draw inferences from texts in a systematic fashion. There are multiple ways and methods to quantify text and each approach contains strengths, weaknesses, and trade-offs (Quinn, Monroe, Colaresi, Crespin, & Radev, 2010). Broadly speaking, scholars apply quantitative text analysis techniques for either the *classification* or *scaling* of texts (Grimmer & Stewart, 2013). Classification, or clustering, places texts into one of several possible categories, and scaling locates actors in a policy space (e.g., along a liberal-conservative dimension). My interest is in clustering—specifically clustering texts by issue dimension—therefore, I will focus on those techniques.[7] The various classification techniques include human coding, as well as computer-assisted approaches such as supervised learning, where categories are known prior to analysis, and unsupervised approaches where categories are unknown (Grimmer & Stewart, 2013).

Classification Techniques

Multiple techniques have been used by scholars to analyze and classify text documents, with each approach having its strengths and weaknesses. The most widely used approach is human coding. Human coding involves the hand-coding of documents, usually by multiple coders, based on categories determined by the researchers and assumed to be unchanging. Coders then assign the text documents to one of the categories, often following a codebook developed prior to the analysis.[8] For example, some of the early work on the Advocacy Coalition Framework (ACF) used hand-coded Congressional testimony to determine the policy beliefs and coalition affiliations of policy actors (Jenkins-Smith & St. Clair, 1993; Jenkins-Smith, St. Clair, & Woods, 1991). In another example, the Policy Agendas Project codes Congressional hearings, executive orders, State of the Union speeches, Supreme Court cases, and other policy documents as being about 1 of 19 major policy topics (e.g., health, education, environment) and 225 subtopics nested within the major policy topics.[9] More recently, scholars have developed ways to discern the major elements of policy designs by coding policies according to a rubric based on the Grammar of Institutions developed by Crawford and Ostrom (1995; these studies include Basurto, Kingsley, McQueen, Smith, & Weible, 2010; Siddiki, Basurto, & Weible, 2012; Siddiki, Weible, Basurto, & Calanni, 2011). Human coding, particularly when using multiple coders, can be a valid and reliable approach to code texts because multiple coders permit various reliability tests. In addition, human coders following a codebook could conceivably code anything of interest to researchers, such as beliefs, issue attributes, and attempts to reframe an issue, whereas some of the semiautomated and automated approaches discussed below may be limited in scope. However, human coding is often cost prohibitive due to the resources needed to develop a codebook and to train and pay coders. In addition, the categories, or the other item(s) to be coded, need to be clearly defined and well understood by each coder. As a result of the challenges inherent in human coding, several computer-assisted approaches have been developed. It should be noted that these approaches do not supplant careful reading and understanding of the texts being studied; rather they augment and enhance the ability of humans to classify a large volume of texts across multiple categories.

Computer-assisted clustering approaches are divided into two categories based on the automation of the approach. The first type, *supervised learning*, is semiautomated and the second type, *unsupervised learning*, is fully automated. Supervised learning combines human coding and automated approaches. A subset of the texts are hand coded, typically the texts are placed into categories by one or more human coders, and then the remaining texts are categorized using a machine-learning algorithm (Collingwood & Wilkerson, 2012). The hand-coded texts constitutes the "training" data, and the excluded texts are the "test" data. Often the idea behind this approach is to get a sense of the overall proportion of documents that belong within a particular category (Grimmer & King, 2011; Hopkins & King, 2010).[10] This approach combines some of the strengths of human coding with the additional speed of computer-aided coding, making possible the coding of a large number of texts. However, this approach is limited in that, typically, each text can be placed in only one of several possible categories and agreement is needed between coders as to in which category the document should be placed.

While supervised learning methods require the use of some precoded data, unsupervised methods, such as topic modeling, are fully automated approaches requiring no prior analysis or setting of categories. Probabilistic topic models involve the automated processing of a large corpus (i.e., body) of texts to determine underlying latent themes or topics (Blei, 2012; Mohr & Bogdanov, 2013; Steyvers & Griffiths, 2007). In short, "Topic modeling algorithms are statistical methods that analyze the words of the original texts to discover the themes that run through them, how those themes are connected to each other, and how they change over time" (Blei, 2012, pp. 77–78). Topic modeling has been used to examine first-ordering agenda setting through Senate press releases (Grimmer, 2010, 2013), floor speeches in Congress (Quinn et al., 2010), Supreme Court opinions (Rice, 2012), and newspaper articles about federal funding for the arts (DiMaggio, Nag, & Blei, 2013), among other applications. Given this success with first-order agenda setting, I use a topic modeling approach to estimate second-order agenda setting or issue definitions. Specifically, the approach I adopt to model issue definitions is LDA modeling, which is a type of topic model.

Latent Dirichlet Allocation. The issue definition model posits that issues are defined by the dimensions of the policy issue weighted by the salience of those dimensions. I use LDA to empirically determine the issue dimensions and the salience of each dimension. LDA is an automated data clustering technique, and the core assumption of LDA is that a single document can contain multiple latent clusters of topics (Blei, Ng, & Jordan, 2003), where a topic is determined probabilistically by the co-occurrence of words. LDA can be evaluated by the semantic validity of the topics (the ways the terms associated with each topic cohere) and by construct validity (how the topics correspond to exogenous events).

As noted, LDA models have been used previously to estimate first-order agenda setting, which is the salience of particular issues within the multidimensional policy space. For example, Quinn et al. (2010) estimated the topics of legislative speeches given in the U.S. Senate from 1997 to 2004 and found 42 distinct topics (i.e., policy issues) drawn from the Senate speeches. In addition, they demonstrated that topic frequency varied in expected ways. For example, abortion was discussed in floor speeches in conjunction with votes related to abortion and the number of speeches about defense rose in response to September 11th and the Iraq War. Examples of LDA for second-order agenda setting are limited (but see Borang et al., 2014; Hopkins, 2013), however, it offers several strengths for modeling issue definitions.

One weakness of human coding and supervised learning methods is that texts tend to be classified or clustered into only a single grouping or category. However, many of the types of documents used to determine issue dimensions—newspaper articles, interest groups statements, and Congressional testimony—are likely to touch on more than one of the possible dimensions. For example, a news article about climate change could discuss the scientific consensus, possible extreme weather, and economic costs of regulation; however, human coders, and supervised approaches would require that article to be classified as "mostly" about one of those topics.[11]

Therefore, LDA models are a useful approach to model issue definitions given that it is important to be able to assume that policy issues exist in a multidimensional space and that policy actors will likely touch on more than a single dimension. As noted by Blei (2012), "the distinguishing characteristic of latent Dirichlet allocation [is that] all the documents in the collection share the same set of topics, but each document exhibits those topics in different proportion" (p. 79). For this reason, LDA offers a distinct advantage over human coding and semiautomated approaches in estimating issue definitions because the model makes probabilistic assessments about the proportion of each topic within each document.

In brief, LDA assumes a hierarchical structure that contains a corpus—a body or collection of documents—documents that are a collection of topics, and topics that are a collection of words (Steyvers & Griffiths, 2007). As such, LDA models are structured hierarchically as:

$$\text{Corpus} \rightarrow \text{Documents} \rightarrow \text{Topics} \rightarrow \text{Words}$$

Similarly, the issue definition model, where the documents are the opening statements from a witness at a Congressional hearing, and the topics are the issue dimensions is structured as:

$$\text{Opening Statements} \rightarrow \text{Issue Dimensions} \rightarrow \text{Words}$$

LDA can estimate θ_a, the collective issue definition, *as the proportion of each witness statement that is about each issue dimension*. When estimated in this way, it is assumed that the proportion of the witnesses' testimony about K issue dimension is a proxy for the salience the policy actors attach to the K dimension. In addition, the salience of the various dimensions represents attempts by policy actors to frame—individually define—the policy issue.

An additional assumption of LDA is that the documents and the words are observed, but a latent, or hidden, structure of topics, topic distributions per documents, and word distributions per topic exists. Therefore, LDA can be thought of as a dimension reduction approach where multiple words are estimated to be associated with a few latent topics. The model estimates words within topics and topics within documents simultaneously, with the goal of inferring the latent structure of topic proportions within documents. In more formal terms, LDA is a generative model based on *observed* variables (words) and *hidden* variables (topics) that defines a joint probability distribution. The joint probability distribution is then used to calculate, by Bayes' rule, a conditional or posterior distribution of the hidden variables given the observed variables (Blei, 2012, pp. 79–80).

Using formal notation adopted from Blei et al. (2003), a *word* is denoted w, and a *document* is a collection of N words, $\mathbf{w} = (w_1, w_2, \ldots, w_n)$, where w_n is the nth word in the document. Note the word ordering is not important and that LDA assumes a "bag of words" approach, where the co-occurrence of terms and not the order of terms is used to determine underlying topics. Finally, a *corpus* is a collection of M documents denoted $\mathbf{D} = (\mathbf{w}_1, \mathbf{w}_2, \ldots, \mathbf{w}_M)$. Following this notation, the generative process—the assumed process that generated the documents, topics, and words—can be described as follows (Blei et al., 2003, p. 996):

- Choose $N \sim \text{Poisson}(\xi)$
- Choose $\theta \sim \text{Dir}(\alpha)$

- For each of the N words w_n
  - Choose a topic $z_n \sim \text{Multinomial}(\theta)$
  - Choose a word w_n from $p(w_n \mid z_n, \beta)$, a multinomial probability conditioned on the topic z_n.

Topic proportions, denoted θ, are the sum of the probabilities for each topic in the dth document. In the theoretical issue definition model, θ was the summation of the issue dimensions weighted by salience, and that concept is operationalized as the topic proportions estimated by LDA.[12] These proportions are assumed to be drawn from a Dirichlet prior, $\theta \sim \text{Dir}(\alpha)$ where α is the Dirichlet distribution's shape parameters. The number of topics K, and by extension the dimensionality of $\text{Dir}(\alpha)$ and the topic variable z, is assumed *a priori* and is also assumed as fixed. Note that a proportion is estimated for each of the K topics within each document; however, these estimated proportions can be quite small (e.g., 0.001) which would indicate that topic was not prevalent within that document. The standard Dirichlet prior of α is $50/K$ (Griffiths & Steyvers, 2004). Finally, the distribution of words is parameterized by β.

As noted, the observed and latent variables form a joint distribution that given the parameters α and β, topic mixture θ, a set of N topics $\mathbf{z}$, and a set of N words $\mathbf{w}$ can be expressed as:

$$p(\theta, \mathbf{z}, \mathbf{w} \mid \alpha, \beta) = p(\theta \mid \alpha) \prod_{n=1}^{N} p(z_n \mid \theta) p(w_n \mid z_n, \beta)$$

This joint distribution is used to calculate a posterior distribution of topic probabilities for each document expressed as:

$$p(\theta, \mathbf{z} \mid \mathbf{w}, \alpha, \beta) = \frac{p(\theta, \mathbf{z}, \mathbf{w} \mid \alpha, \beta)}{p(\mathbf{w} \mid \alpha, \beta)}$$

The numerator is the joint distribution of all random variables and the denominator is the probability of obtaining the observed corpus under any topic model. These probabilities could be summed over all possible topic structures; however, given the large number of possible topic structures this becomes "intractable to compute" (Blei, 2012, p. 81). Therefore, this total needs to be approximated using either sampling based or variational approximations (Blei, 2012). Gibbs sampling is used to estimate the posterior distribution (Griffiths & Steyvers, 2004).

In less formal terms, LDA is a data-reduction technique or cluster analysis that assumes that a text document has a mixture of latent topics. Words that co-occur repeatedly across documents are assigned a high probability of being associated with a topic. Then, based on the array of words in a document, LDA estimates the proportion of each document that refers to each topic. One important way of evaluating the results of LDA is examining whether the words that were given a high probability of being associated with a topic cohere in a reasonable fashion. In other words, when grouped together do the terms seem to relate to the same topic. The next section describes the data-collection process, and the topics present during Congressional hearings about UNF.

Defining Used Nuclear Fuel

Data for the analysis came from 140 Congressional hearings,[13] drawn from the *ProQuest Congressional* database about UNF between the years 1975 and 2012. The hearings were found using the search terms "nuclear waste," "spent nuclear fuel," "used nuclear fuel," and "Yucca Mountain."

Hearing transcripts, in PDF format, were available for each hearing. Overall, these hearings included a total of 1,322 witnesses, of which 1,271 gave opening statements, and the remainder were either silent or only answered questions. The verbal opening statements of each of these witnesses were extracted from the hearing transcript, each statement became a separate document, and these documents were combined into a corpus.[14] The next section describes the dimensions of the UNF issue and examines their semantic validity. Following that I examine the construct validity of the dimensions by predicting the mean proportion of each dimension following two major policy changes: the Nuclear Waste Policy Act (NWPA) of 1982 and the Nuclear Waste Policy Amendments Act (NWPAA) of 1987.

Document Preprocessing

While LDA does not require any preanalysis such as categorizing texts, some preprocessing of the documents prior to modeling is necessary. The first steps involved removing numbers and punctuation and placing terms in lowercase. The next step was to remove all standard stop-words (e.g., and, for, in, is, it, the) from the corpus. An additional number of stop-words unique to the policy issue (e.g., nuclear, waste, spent, fuel) and unique to the context of Congressional hearings (e.g., chairman, committee, thank, testimony) were also removed. These stop-words are removed because they are common to nearly every witness and, therefore, provide no valuable information.[15] While it is important to remove words that are extremely common, it is just as important to remove words that are rarely used, or used by only a few witnesses. Removing these words is important because rarely used words are not likely to be associated with any topic, thereby introducing unwanted noise into the model. Thus, all terms that were not in at least 1 percent of the documents were removed. Next, the Porter algorithm (Porter, 1980) was used to stem the remaining terms. Stemming refers to the process of removing the suffix of terms to equate similar terms while reducing the overall number of terms. For example, the stem word "decis" stands for decision or decisions. Finally, I created a document-term matrix where rows represent the documents, the columns represent the terms, and the cells contain the number of times each term appears in each document. This matrix contained the 1,271 documents as the rows and the 2,941 terms.[16] After preprocessing the documents the word length of the documents ranged from a minimum of 17 to a maximum of 4,822, with a mean word length of 553.31. Once this preprocessing is complete, the next step is determining the number of topics.

*The Selection of **K** Topics*

Since the *K* number of topics are assumed *prior* to modeling, which is a common feature of cluster algorithms (Grimmer & King, 2011), care must be taken when assigning the number of topics. Determining the *K* number of topics is one of the major challenges when using topic models (Blei & Lafferty, 2009). Compared to supervised approaches where the onus is on the researcher to develop categories *prior* to analysis, unsupervised approaches, such as topic models, place the burden on the researcher to ensure that the categories derived by the model are justifiable according to clear guidelines that are spelled out by the researcher.

There are several possible guidelines that scholars can use when selecting *K*, most of which are substantive and conceptual (Quinn et al., 2010, p. 216). First and foremost, however, the topics must have semantic validity (Grimmer & Stewart, 2013; Quinn et al., 2010). Semantic validity means that each topic has a clear and coherent meaning that can be discerned by the association of terms to that topic (Krippendorff, 2003). Substantive knowledge of the policy issue or area can also be used to determine *K*. Given that automated techniques augment human understanding without replacing it, having a basic understanding of the issue and how it has evolved is important. In particular, substantive knowledge of the policy area is extremely helpful in determining semantic validity.

In addition, *K* should be selected based on the research question, the hypotheses being tested, and be guided by theory. For example, in this article LDA is used to estimate the issue definition model by determining the issue dimensions that have been used to define the UNF issue over several decades. Given that the model of issue definitions is based on the collective issue definition, *K* is smaller than if the interest was in particular framing strategies of interest groups, since the collective issue definition is biased toward framing strategies that were successful.

Finally, sensitivity analysis—varying *K*—should always be used when performing LDA. Researchers should use an iterative process of modeling with varying *K* topics and each iteration should be examined using the guidelines suggested above.

UNF Dimensions

To determine the distribution of dimensions at Congressional hearings regarding used nuclear fuel, I performed LDA on the opening statements of witnesses appearing at those hearings. As noted, LDA assumes a proportion per document of latent topics based on the distribution of terms, where terms that frequently co-occur are understood as belonging to the same topic. Terms are assigned to topics probabilistically, and a term can have a high probability of being associated with more than one topic. The observed data (terms) are used to estimate the unobserved topics. In other words, the model finds the parameters α and β that maximize the log-likelihood of the data in the 1,271 $\times$ 2,941 matrix (Blei et al., 2003). For my purposes, topics are assumed to be synonymous with the dimensions of an issue.

The criteria used to evaluate the results of the models using different numbers of topics was based on semantic validity, knowledge of the issue, and the proposed theoretical model of issue definitions. Given that the interest was in measuring the various dimensions of the UNF issue across several decades, K should represent dimensions that are likely to exist throughout the time period being examined. Sensitivity analysis was used, and models estimated with varying K topics, including 2, 4, 5, 6, 7, 8, 10, 20, and 50 were examined, and the model that was estimated with $K = 7$ dimensions was ultimately selected. With $K \leq 6$ several terms overlapped, making topics somewhat less clearly differentiated and with $K \geq 8$ there tended to be clusters that either did not have semantic validity—were not clear as to what possible topic they represented—or were too narrow and could best be combined under one dimension. Given these results seven topics were chosen as the best fitting number of dimensions. Finding seven dimensions within a single issue seems to fit with other work on issue definitions. For example, recent work by Rose and Baumgartner (2013) found five dimensions (or frames) of poverty policy, Baumgartner et al. (2008) found seven dimensions for the death penalty,[17] and Hopkins (2013) found twelve dimensions for health care policy. The results of the LDA process are shown in Table 1.

Table 1 lists each of the seven dimensions and the 30 terms that had the highest probability of being associated with each dimension. For example, *site* is listed first under programmatic and *program* is listed second, indicating that the term *site* had the highest probability of being associated with that topic and *program* had the second highest probability of being associated.[18] As is evident in Table 1, terms tended to cluster into seven discernible dimensions that I then termed as the following: programmatic, safety/regulation, Yucca Mountain, site selection, scientific/technical, storage, and transportation. As noted, the same term can be associated with more than one topic, and this is common with LDA (Griffiths & Steyvers, 2004). For example, the term *site* has a high probability of being associated with both the programmatic, site selection, and scientific/technical dimensions. However, when determining the semantic validity of a topic, it is the probability of several terms being associated with the topic that illustrates its meaning. For example, apart from *site*, terms such as *program*, *DOE (The Department of Energy)*, *repositori*, *act*, and *process* all have a high probability of being associated with the programmatic dimension, whereas the site selection dimension includes *site*, but also includes *nation*, *senat*, *energi*, *time*, *process*, and *concern*. Below each of the seven dimensions is briefly described.

Programmatic: The programmatic dimension deals largely with the development and implementation of programs related to the selection and characterization of a suitable site for waste disposal. This is evidenced by terms such as "site," "program," "review," "plan," and "process."

Safety/Regulation: The safety/regulation dimension deals with the development and implementation of proper regulatory standards for dealing with

Table 1. Dimensions of UNF and the Terms Most Associated with Each Dimension

Dimensions	Terms
Programmatic	site, program, DOE, repositori, act, process, review, plan, licens, NRC, polici, issu, public, character, decis, propos, provid, recommend, technic, requir, commiss, depart, concern, comment, develop, environment, particip, final, manag, feder
Safety/Regulation	radioact, dispos, highlevel, EPA, environment, standard, lowlevel, develop, materi, program, manag, radiat, protect, level, oper, dump, research, public, erda, contain, ocean, product, uranium, futur, intern, agenc, activ, unit, regul, requir
Yucca Mountain	DOE, program, mountain, yucca, fund, nevada, project, repositori, depart, issu, energi, cost, board, continu, billion, scientif, act, report, court, site, current, million, util, begin, secretari, applic, nation, polici, manag, meet
Site Selection	site, nation, senat, energi, time, process, concern, polit, feel, depart, repres, hear, bill, tri, countri, legisl, governor, issu, nevada, citizen, reason, land, simpli, decis, happen, suggest, hope, govern, serious, talk
Scientific/Technical	repositiori, site, geolog, studi, test, time, data, system, technic, develop, water, salt, dispos, evalu, design, potenti, rock, form, isol, requir, program, perform, activ, investig, result, process, environ, concept, research, suitabl
Storage	storag, facil, reactor, manag, dispos, feder, power, plant, oper, reprocess, util, govern, energi, cost, interim, provid, industri, commerci, perman, capac, polici, develop, licens, nation, technolog, electr, time, generat, store, administr
Transportation	transport, materi, local, shipment, respons, safeti, regul, feder, radioact, citi, cask, counti, hazard, accidpublic, health, rout, govern, risk, depart, nrc, requir, communiti, york, ship, propos, concern, impact, emerg, involv

used nuclear fuel. This can be seen by the importance of such terms as "radioact," "dispos," "standard," and "manag."

Yucca Mountain: This dimension deals with Yucca Mountain in Nevada, the site ultimately chosen for UNF disposal, as evidenced by terms such as "mountain," "yucca," and "nevada."

Site Selection: Site selection is largely about the *politics* surrounding the process of choosing possible site locations to store used nuclear fuel. Important terms include "site," "nation," "concern," and "polit." The term "polit" is the stem for terms like *politics* and *political*.

Scientific/Technical: The scientific/technical dimension deals with the science of used fuel storage (e.g., "studi," "test," "data") and the type of natural medium that could be used "salt," "geolog," "rock" for storage.

Storage: This dimension deals with questions about storage (e.g., "storag") and waste management, but largely centers on concerns about on-site storage at nuclear power plants vs. off-site storage of waste. This is evidenced by terms like "facil," "reactor," "capac."

Transportation: The transportation dimensions deals with transportation of spent nuclear fuel from the production site to the storage site (e.g., "transport," "shipment," "rout").

Table 2. Descriptive Statistics of UNF Dimensions

Dimension	Mean	sd	Min	Median	Max
Programmatic	0.199	0.154	0.007	0.148	0.797
Safety/Regulation	0.114	0.112	0.012	0.073	0.747
Yucca Mountain	0.137	0.116	0.008	0.095	0.743
Site Selection	0.166	0.120	0.010	0.137	0.588
Science/Technical	0.127	0.109	0.012	0.091	0.818
Storage	0.144	0.124	0.010	0.101	0.705
Transportation	0.113	0.124	0.013	0.070	0.737

As discussed above $K = 7$ was chosen following multiple iterations of the model with varying values for K. With $K \leq 6$ the programmatic, safety/regulation, and scientific/technical dimensions tended to overlap, yet those dimensions should be seen as distinctive. With $K \geq 8$ the storage dimension tended to divide into long-term storage and interim storage. Questions related to interim storage of UNF were discussed largely in the mid to late 1990s; therefore, interim storage as a separate dimension from larger questions of UNF storage was deemed to be specific to that time period. Another example was the scientific/technical dimension where questions of the proper geological medium for storage tended to be its own dimensions when $K \geq 8$. However, deep geologic disposal was settled as the preferred medium in the late 1970s to early 1980s, making questions of storage medium too specific to that time period.

Table 2 provides descriptive statistics for each dimension. As noted, each opening statement has an estimated topic proportion for each of the seven dimensions, yet this estimated proportion can be quite small. Table 2 displays the overall mean, standard deviation, median, and minimum and maximum values for each dimension.

The programmatic dimension has a mean of 0.199, which indicates that the average proportion of the programmatic dimension across all the statements is 0.199. The programmatic dimension also has the highest mean, which indicates that it was the most salient dimension over the full span of hearings. Site selection had the next highest mean at 0.166, followed by storage and Yucca Mountain at 0.144 and 0.137, respectively. The science/technical dimension had a mean of 0.127, and finally safety/regulation and transportation had mean proportions of 0.114 and 0.113. However, the high standard deviations of each dimension indicates a sizable amount of variation around the mean. This is also seen in the min and max values that range from 0.007 to 0.818. A min value of 0.007, which is near zero, indicates that the dimension was not discussed much in at least one statement and a higher max value, such as 0.818, means that an estimated 82 percent of at least one statement was about that particular dimension. The dimension with the lowest max value was site selection at 0.588, whereas the dimension with the highest was scientific and technical at 0.818. Site selection had the lowest "peak" proportion (0.588) but the second highest mean proportion (0.166), indicating that this issue dimension was rarely the sole topic addressed by any particular witness, but was likely discussed in tandem with other dimensions.

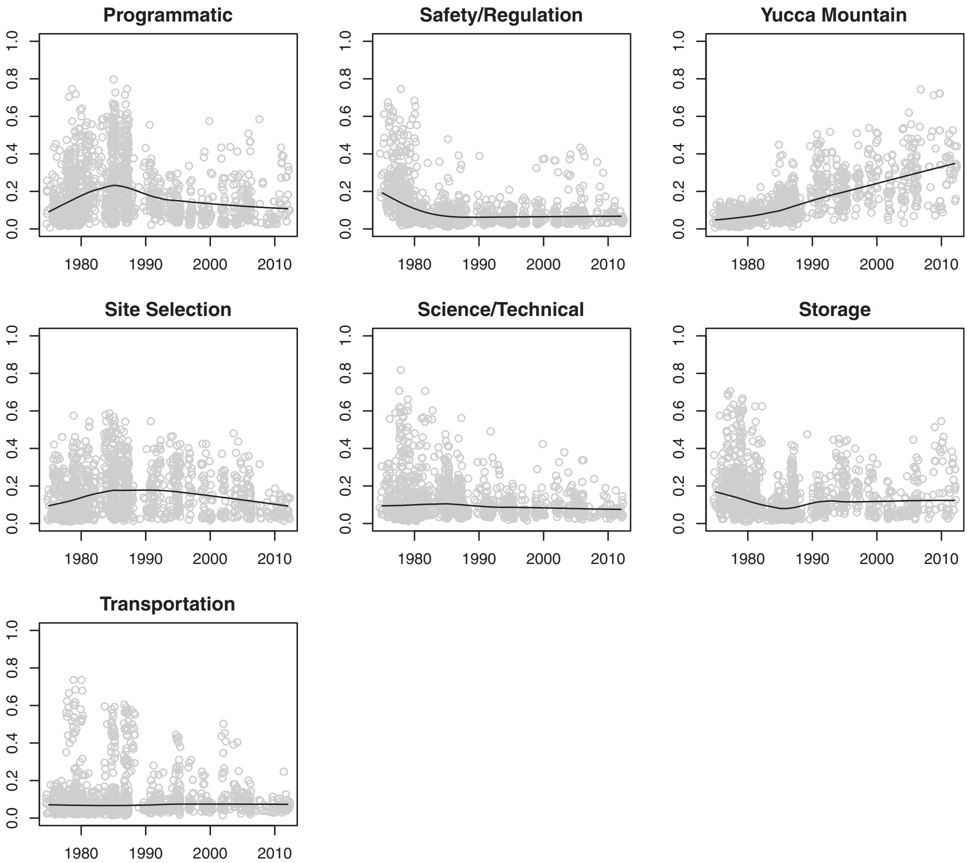

Figure 1. Used Nuclear Fuel Dimension Proportions by Year (1975–2012).

Figure 1 shows how the proportion of attention paid to the different dimensions tended to vary over the period from 1975–2012. The panels within Figure 1 show a scatterplot of the proportion of each dimension for each statement by year. Each scatterplot contains a lowess line—a nonparametric measure of fit—for each dimension showing the proportion of discussion devoted to that dimension over time.

As can be seen, the dimensions vary over time in their salience. The programmatic dimension seems to have risen for each year and peaked around 1985–86. This is as expected given that is it the period after the NWPA of 1982 and just prior to the NWPAA of 1987, which designated Yucca Mountain as the only site to be considered. The safety/regulation dimension was more salient in earlier years as opposed to later years, because this was when Congress was first examining how to deal with used nuclear fuel. Yucca Mountain became increasingly salient following the NWPAA, over which period the Yucca site was the focus of UNF disposal policy. Site selection seems to have been most prominent following 1982, the year in which the NWPA was passed, and then starts to slowly taper off. The science and technical dimension seems to have remained relatively flat; however, it seems to have much more variation in the earlier years than in later ones. The storage dimension seems to have declined from 1975 until rising again after 1987 and then staying at that level. Finally, transportation seems to have remained relatively flat over the entire period, although with periods of more variability.

Overall, the salience of each dimension seems to rise and fall as would be expected given the policy developments occurring within the various years. This pattern of variation within dimensions provides some evidence for the validity of these dimensions, suggesting that they capture the nature of the debate concerning UNF policy over the 1975–2012 period. The next section more closely examines the validity of the dimensions by using OLS to model the dimensions following two major exogenous events: the passage of the NWPA and the passage of the NWPAA.

Validity Check

One way to examine the validity of the issue definition model is to examine whether the measure(s) correspond to exogenous events as would be expected.[19] To do this, I examine the impact of two major pieces of legislation on the salience of each of the dimensions.

The NWPA of 1982 was the first major piece of legislation that dealt directly with used nuclear fuel. Among its provisions, it required DOE to study five potential sites and recommend three to the president by 1985; allowed a veto for the potential host state, however, this veto could be overridden by Congress; and established away-from-reactor storage once storage capacity at a plant site had been met. The second major piece of legislation, the NWPAA of 1987, designated Yucca Mountain in Nevada as the only site to be considered for a UNF repository (Vandenbosch & Vandenbosch, 2007).

Knowing a little bit about these major policy changes and the dimensions, some expectations about the behavior of the dimensions following each major change can be derived. For example, I expect that the salience of the programmatic dimension would increase following the NWPA, since the programmatic dimension is largely focused on the development and implementation of the programs that are called for in the NWPA. I also expect that site selection would become more salient after the passage of NWPA, since that legislation initiated the process of studying various potential sites. Finally, I would expect that the Yucca Mountain dimension would become more salient after the passage of the NWPAA.

To examine these salience changes, I used OLS to model the salience of each dimension following each of the two major policy changes.[20] The dependent variable is the mean proportion of each dimension aggregated by year. Hearings about UNF did not take place in 1996, 1998, and 2001 so $N = 35$, for the 38 years in the dataset (1975–2012) minus the three years with no hearings. The two independent variables are dummy variables that represent the two policy changes. The NWPA variable has a value of 0 for 1975 to 1981 and 1987 to 2012, and a value of 1 for 1982 to 1986. The NWPA was superseded by the NWPAA, which is why it has zero values from 1987

Table 3. Salience of UNF Dimensions Following Policy Change

	Program	Safety	Yucca	Site	Sci/Tech	Storage	Trans
Intercept	0.18****	0.19***	0.06**	0.13***	0.15***	0.19***	0.10***
	(0.02)	(0.02)	(0.03)	(0.02)	(0.01)	(0.03)	(0.02)
NWPA 1982	0.11***	−0.12***	0.04	0.08**	−0.00	−0.12**	0.00
	(0.04)	(0.03)	(0.05)	(0.03)	(0.02)	(0.05)	(0.04)
NWPA Ammendment 1987	−0.03	−0.11***	0.21***	−0.00	−0.05***	−0.03	0.01
	(0.03)	(0.02)	(0.03)	(0.02)	(0.01)	(0.04)	(0.03)
Adj. R^2	0.34	0.45	0.55	0.18	0.23	0.10	−0.06
Num. obs.	35	35	35	35	35	35	35

***$p < 0.01$, **$p < 0.05$, *$p < 0.1$

to 2012. The NWPAA variable has a value of 0 for 1975 to 1986 and a value of 1 for 1987 to 2012. Table 3 presents the results.

As shown in Table 3, the programmatic dimension, as expected, became more salient following the passage of the NWPA, but showed no significant change following the NWPAA. The safety/regulation dimension became less salient following each major change. This may be because questions of the proper safeguards of this material were addressed in both the NWPA and the NWPAA, thereby making them less salient following each major policy change. Also as expected, the Yucca Mountain dimension became significantly more salient following the NWPAA. The site selection dimension, as expected, became more salient following the NWPA. The scientific and technical dimension became less salient following the NWPAA, perhaps because the focus shifted to Yucca Mountain and away from broader technical questions. The storage dimension declined following the NWPA, likely because the NWPA settled the question of away-from-reactor storage, although away-from-reactor storage remains a highly controversial aspect of the UNF debate. Finally, neither policy change had any significant impact on the transportation dimension. Overall, the dimension measures seemed to perform as expected, with *a priori* expectations about the programmatic, site selection, and Yucca Mountain dimensions being met and the other dimensions performing in ways that are quite plausible given the history of the UNF debate.

Conclusion

Policy scholars have long noted the importance of issue definitions for the development of public policy. However, in the literature, issue definitions have been examined either at the macro level, focusing on questions of agenda setting, or the micro level, examining ways in which policy actors try to frame a policy issue. This article developed a model of issue definitions based on the underlying attributes (i.e., dimensions) of the issue and how they are combined at the micro level to produce collective issue definitions. The suggested approach to issue definitions requires the use of sophisticated methods to analyze text. LDA, as a fully automated approach, takes advantage of the increase in the availability of digitized policy-related documents and presents policy scholars with unique opportunities to analyze how issues are defined.

To estimate the dimensions of the UNF policy debate, I performed LDA analysis on the opening statements of the 1,271 witnesses that appeared in 140 Congressional hearings about UNF from 1975 to 2012. As described above, this approach was adopted after consideration of several other possible approaches. One advantage of the LDA approach is that it assumes multiple topics (i.e., dimensions) within a single document (i.e., testimony). This allowed for the estimation of the relative weights that were assigned to each of the seven dimensions for each witness. The results of the LDA model presented issue dimensions that possessed both semantic validity—a coherent meaning derived from the co-occurrence of terms—and external validity-responsiveness to exogenous measures. This indicates that the issue definition model, as proposed, and LDA more broadly, should have utility for policy scholars interested in inferring topics (i.e., issue dimensions) from text documents. A drawback of this approach, however, is that dimensions must be defined rather broadly to avoid clusters that are too context dependent (i.e., narrow or time-specific). In others words, to be able to compare dimension salience over time and across different policy actors, the dimensions must be consistent and appear across multiple time periods. While this approach is extremely useful to estimate the broad and collective dimensions of an issue, it could miss some of the nuance associated with particular attempts to reframe a policy issue. For example, national security elements were likely discussed when a) reprocessing of UNF was being considered in the mid-1970s and b) when Yucca Mountain was approved shortly after the 9–11 terrorist attacks. A topic modeling approach focused, like the one presented here, on the collective dimensions may miss some of these nuances. However, LDA could be used to identify these more nuanced frames, or attempts to reframe. Specifically, an increase in the K number of topics to be estimated could show attempts to reframe that were unsuccessful. Best practice for examining attempts at framing would likely include human coding as a validity check on LDA analysis with the larger K. Finally, determining the number of K dimensions should be done through an iterative process, and the choice of the best number of K dimensions can be difficult. In contrast to other approaches to coding documents, this difficulty is faced following the analysis, and is a trade-off with the difficulty of developing (and updating) a codebook prior to the analysis.

The approach to issue definitions and LDA presented here is intended to complement and build on, not replace, previous work. Indeed, the issue definition model could be incorporated into several of the leading theories of the policy process. The primary focus of the model is aggregating micro-issue definitions into collective issue definitions using LDA. The collective issue definitions could then be used to examine questions of agenda setting. For example, Rochefort and Cobb (1993) argued that problems are defined across several categories such as problem causation, nature (e.g., severity, novelty, crisis), and solution and are more (less) likely to appear on the policymaking agenda based on how problems fit into those categorizations. It is likely that one or more of the dimensions of an issue would fall into one or more of those categories. In addition, the work by Kingdon (1984) could be used, and dimensions of the issue could be classified as being in one of the three streams. Finally, the issue definition model could be used as a way to operationalize policy images to track instances of changing images—which would be represented by different dimensions becoming more salient—that may lead to policy punctuations.

The model of issue definitions proposed and estimated using LDA is largely intended for questions that go beyond agenda setting. For example, the model could be used to estimate information signals that are processed by policymaking institutions, within the theory of information processing (Jones & Baumgartner, 2005; Workman, Jones, & Jochim, 2009). In that approach, information is deemed as signals and these signals could be operationalized as the highlighting of different dimensions by policy actors within a policy debate. Finally, the issue definition model and LDA could be used to operationalize policy core beliefs in the ACF. The ACF argues that policy actors exhibit a three-tiered structure of beliefs with core beliefs, such as political ideology, being the most abstract and applicable across multiple policy domains; followed by policy core beliefs, which are more narrow in scope and contained within a single policy domain; and secondary beliefs, which are the most narrow and refer to things like particular policy instruments (Sabatier & Jenkins-Smith, 1993; Sabatier & Weible, 2007). A second argument of the ACF is that policy actors in subsystems are best understood as acting within coalitions that are bound together by policy core beliefs and coordinate their actions. Applied to the ACF, the issue definition model could establish shared policy core beliefs by determining which dimensions are highlighted by the same policy actors, *assuming that the highlighting of a particular dimension or set of dimensions by a policy actor is evidence of policy core beliefs*. This would be similar to the approach used by Jenkins-Smith et al. (1991) and Jenkins-Smith and St. Clair (1993) that used testimony at Congressional hearings to determine the make-up of advocacy coalitions. Therefore, the issue definition model and LDA would be best suited to determine coalition memberships over time based on expressed policy core beliefs occurring in forums that are carefully documented.

Matthew C. Nowlin is an assistant professor in the Department of Political Science at the College of Charleston.

Notes

1. Indeed, the Environmental Protection Agency is required to develop exposure regulations that stretch out to 1 million years.
2. Note that these attributes have also been termed “frames” (e.g., Baumgartner, De Boef, & Boydstun, 2008; Rose & Baumgartner, 2013). However, I am using the terms frames, attributes, and dimensions, interchangeability.
3. Riker (1986) termed these types of strategic maneuvers heresthetics, and one particular type of heresthetic is the manipulation of dimensions—the restructuring of the issue space through the addition and/or subtraction of dimensions—related to a particular issue (Riker, 1990).
4. Baumgartner and Mahoney (2008) note that collective definitions are also determined by exogenous forces, however in the model proposed here the collective issue definitions are the summation of attempts to frame an issue by individual policy actors. Further work should seek to explore ways to incorporate exogenous forces into the model.

5. This is similar to the model of frames discussed in Chong and Druckman (2007), which focused on emphasis framing (i.e., highlight one or a few dimensions of an issue) by individual policy actors.
6. Also, this model could be used to examine first-order agenda setting, in which case θ_a would be the policymaking agenda and S_i would be the salience of each issue.
7. For more about approaches for scaling texts, see Laver, Benoit, and Garry (2003); Lowe (2008); Lowe, Benoit, Mikhaylov, and Laver (2011).
8. Although the codebook can be adjusted in an iterative process during the analysis phase.
9. See http://www.policyagendas.org/for more information.
10. There are a growing number of software applications, based in the R statistical programming language that use this type of approach. These include *ReadMe* and *RTextTools*.
11. For some research questions and hypotheses this would be desirable. In addition, the unit of analysis for human coders could be adjusted down to the sentence level to capture multiple topics within a document. However, at that level of analysis agreement between coders may be more difficult to obtain.
12. Note that these estimated proportions sum to 1.
13. I excluded appropriation hearings from the subsequent analysis.
14. Some previous work has coded the written statements of witnesses (e.g., Esterling, 2011), however I chose to use the verbal statements because witnesses, due to time constraints, are likely to express the most salient points in the verbal statement. In addition, only the opening statements were used and not the question and answer portions of the hearings, so that only those dimensions thought salient by the witness, and not the questioner, would be determined.
15. This is standard practice when doing this type of text analysis (see Grimmer, 2010; Quinn et al., 2010).
16. For more detail on the steps discussed above see Feinerer, Hornik, and Meyer (2008).
17. However, those seven dimensions of the death penalty issue were aggregated based on the discovery of 65 "frames" used in *New York Times* article abstracts. As noted, the issue definition model is focused on collective issue definitions and, therefore, is focused on a smaller number of *K* topics. However, LDA could be used to analyze strategic framing by selecting a higher number of *K* topics.
18. 30 terms was chosen for illustrative purposes, but dimensions/topics should likely be discernible from fewer terms.
19. As noted earlier, Quinn et al. (2010) found that their categorization of Congressional floor speeches matched important events in expected ways.
20. The focus here is only to see if some evidence of validity of the dimension measures can be gained by examining their relationship with two major pieces of legislation. Future research should use more sophisticated time series techniques and examine the relationship between the dimensions and policy change in much more depth.

References

Basurto, Xavier, Gordon Kingsley, Kelly McQueen, Mshadoni Smith, and Christopher M. Weible. 2010. "A Systematic Approach to Institutional Analysis: Applying Crawford and Ostrom's Grammar." *Political Research Quarterly* 63 (3): 523–37.

Baumgartner, Frank R., and Bryan D. Jones. 1993. *Agendas and Instability in American Politics*. Chicago, IL: University of Chicago Press.

Baumgartner, Frank R., and Christine Mahoney. 2008. "The Two Faces of Framing: Individual-Level Framing and Collective Issue Definition in the European Union." *European Union Politics* 9 (3): 435–49.

Baumgartner, Frank R., Suzanna L. De Boef, and Amber E. Boydstun. 2008. *The Decline of the Death Penalty and the Discovery of Innocence*. New York, NY: Cambridge University Press.

Blei, David, and John Lafferty. 2009. "Topic Models." In *Text Mining: Classification, Clustering, and Applications*, eds. Ashok Srivastava, and Mehran Sahami. Boca Raton, FL: Chapman & Hall/CRC, 71–94.

Blei, David M. 2012. "Probabilistic Topic Models." *Communications of the ACM* 55 (4): 77–84.

Blei, David M., Andrew Y. Ng, and Michael I. Jordan. 2003. "Latent Dirichlet Allocation." *The Journal of Machine Learning Research* 3: 993–1022.

Borang, Frida, Rainer Eising, Heike Kluver, Christine Mahoney, Daniel Naurin, Daniel Rasch, and Patrycja Rozbicka. 2014. "Identifying Frames: A Comparison of Research Methods." *Interest Groups and Advocacy* 3 (2): 188–201.

Boydstun, Amber E., Rebecca A. Glazier, and Claire Phillips. 2013. "Agenda Control in the 2008 Presidential Debates." *American Politics Research* 41 (5): 863–99.

Chong, Dennis, and James N. Druckman. 2007. "Framing Theory." *Annual Review of Political Science* 10 (1): 103–26.

Collingwood, Loren, and John Wilkerson. 2012. "Tradeoffs in Accuracy and Efficiency in Supervised Learning Methods." *Journal of Information Technology & Politics* 9 (3): 298–318.

Crawford, Sue E. S., and Elinor Ostrom. 1995. "A Grammar of Institutions." *The American Political Science Review* 89 (3): 582–600.

DiMaggio, Paul, Manish Nag, and David Blei. 2013. "Exploiting Affinities Between Topic Modeling and the Sociological Perspective on Culture: Application to Newspaper Coverage of U.S. Government Arts Funding." *Poetics* 41 (6): 570–606.

Downs, Anthony. 1957. "An Economic Theory of Political Action in a Democracy." *The Journal of Political Economy* 65 (2): 135–50.

Esterling, Kevin M. 2011. "'Deliberative Disagreement' in U.S. Health Policy Committee Hearings." *Legislative Studies Quarterly* 36 (2): 169–98.

Feinerer, Ingo, Kurt Hornik, and David Meyer. 2008. "Text Mining Infrastructure in R." *Journal of Statistical Software* 25 (5), 1–54. [Online]. http://www.jstatsoft.org/v25/i05. Accessed August 13, 2012.

Ghanem, Salma I. 1997. "Filling in the Tapestry: The Second Level of Agenda Setting." In *Communication and Democracy: Exploring the Intellectual Frontiers in Agenda-Setting Theory*, eds. Maxwell McCombs, Donald Shaw, and David H. Weaver. Mahwah, NJ: Lawrence Erlbaum Associates, 3–14.

Griffiths, Thomas L., and Mark Steyvers. 2004. "Finding Scientific Topics." *Proceedings of the National Academy of Sciences of the United States of America* 101 (1 Supplement): 5228–35.

Grimmer, Justin. 2010. "A Bayesian Hierarchical Topic Model for Political Texts: Measuring Expressed Agendas in Senate Press Releases." *Political Analysis* 18 (1): 1–35.

———. 2013. "Appropriators Not Position Takers: The Distorting Effects of Electoral Incentives on Congressional Representation." *American Journal of Political Science* 57 (3): 624–42.

Grimmer, Justin, and Gary King. 2011. "General Purpose Computer-Assisted Clustering and Conceptualization." *Proceedings of the National Academy of Sciences of the United States of America* 108 (7): 2643–50.

Grimmer, Justin, and Brandon M. Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* 21 (3): 267–97.

Hinich, Melvin J., and Michael C. Munger. 1994. *Ideology and the Theory of Political Choice*. Ann Arbor, MI: University of Michigan Press.

Hopkins, Daniel J. 2013. The Exaggerated Life of Death Panels: The Limits of Framing Effects in the 2009-2012 Health Care Debate [Online]. http://dx.doi.org/10.2139/ssrn.2163769. Accessed July 9, 2014.

Hopkins, Daniel J., and Gary King. 2010. "A Method of Automated Nonparametric Content Analysis for Social Science." *American Journal of Political Science* 54 (1): 229–47.

Jenkins-Smith, Hank C., and Gilbert K. St. Clair. 1993. "The Politics of Offshore Energy: Empirically Testing the Advocacy Coalition Framework." In *Policy Change and Learning: An Advocacy Coalition Approach*, eds. P. A. Sabatier, and H. Jenkins-Smith. Boulder, CO: Westview, 149–75.

Jenkins-Smith, Hank C., Gilbert K. St. Clair, and Brian Woods. 1991. "Explaining Change in Policy Subsystems: Analysis of Coalition Stability and Defection over Time." *American Journal of Political Science* 35 (4): 851–80.

Jochim, Ashley E., and Bryan D. Jones. 2013. "Issue Politics in a Polarized Congress." *Political Research Quarterly* 66 (2): 352–69.

Jones, Bryan D., and Frank R. Baumgartner. 2005. *The Politics of Attention: How Government Prioritizes Problems*. Chicago, IL: University Of Chicago Press.

Jones, Bryan D., Jeffery Talbert, and Matthew Potoski. 2003. "Uncertainty and Political Debate: How the Dimensionality of Political Issues Gets Reduced in the Legislative Process." In *Uncertainty in American Politics*, ed. Barry C. Burden. Cambridge: Cambridge University Press, 118–38.

Kingdon, John W. 1984. *Agendas, Alternatives and Public Policies*. New York, NY: Longman.

Krippendorff, Klaus. 2003. *Content Analysis: An Introduction to Its Methodology*. 2nd ed. Thousand Oaks, CA: Sage Publications, Inc.

Laver, Michael, Kenneth Benoit, and John Garry. 2003. "Extracting Policy Positions from Political Texts Using Words as Data." *The American Political Science Review* 97 (2): 311–31.

Lowe, Will. 2008. "Understanding Wordscores." *Political Analysis* 16 (4): 356–71.

Lowe, Will, Kenneth Benoit, Slava Mikhaylov, and Michael Laver. 2011. "Scaling Policy Preferences from Coded Political Texts." *Legislative Studies Quarterly* 36 (1): 123–55.

McCombs, Maxwell. 2005. "A Look at Agenda-Setting: Past, Present and Future." *Journalism Studies* 6 (4): 543–57.

McCombs, Maxwell, Juan Pablo Llamas, Esteban Lopez-Escobar, and Federico Rey. 1997. "Candidate Images in Spanish Elections: Second-Level Agenda-Setting Effects." *Journalism & Mass Communication Quarterly* 74 (4): 703–17.

Mohr, John W., and Petko Bogdanov. 2013. "Introduction-Topic Models: What They Are and Why They Matter." *Poetics* 41 (6): 545–69.

Porter, M. F. 1980. "An Algorithm for Suffix Stripping." *Program: Electronic Library and Information Systems* 14 (3): 130–37.

Quinn, Kevin M., Burt L. Monroe, Michael Colaresi, Michael H. Crespin, and Dragomir R. Radev. 2010. "How to Analyze Political Attention with Minimal Assumptions and Costs." *American Journal of Political Science* 54 (1): 209–28.

Rice, Douglas. 2012. "Measuring the Issue Content of Supreme Court Opinions Through Probabilistic Topic Models." In *Annual Meeting of the Midwest Political Science Association* [Online]. http://qssi.psu.edu/files/NF5Rice.pdf. Accessed January 17, 2014.

Riker, William H. 1986. *The Art of Political Manipulation*. New Haven, CT: Yale University Press.

———. 1990. "Heresthetic and Rhetoric in the Spatial Model." In *Advances in the Spatial Theory of Voting*, eds. James M. Enelow, and Melvin J. Hinich. New York, NY: Cambridge University Press, 46–65.

Rochefort, David A., and Roger W. Cobb. 1993. "Problem Definition, Agenda Access, and Policy Choice." *Policy Studies Journal* 21 (1): 56–71.

Rose, Max, and Frank R. Baumgartner. 2013. "Framing the Poor: Media Coverage and U.S. Poverty Policy, 1960–2008." *Policy Studies Journal* 41 (1): 22–53.

Sabatier, Paul A., and Hank C. Jenkins-Smith. 1993. *Policy Change and Learning: An Advocacy Coalition Approach*. Boulder, CO: Westview Press.

Sabatier, Paul A., and Christopher M. Weible. 2007. "The Advocacy Coalition Framework: Innovations and Clarifications." In *Theories of the Policy Process*, ed. Paul A. Sabatier. Boulder, CO: Westview Press, 189–222.

Siddiki, Saba, Xavier Basurto, and Christopher M. Weible. 2012. "Using the Institutional Grammar Tool to Understand Regulatory Compliance: The Case of Colorado Aquaculture." *Regulation & Governance* 6 (2): 167–88.

Siddiki, Saba, Christopher M. Weible, Xavier Basurto, and John Calanni. 2011. "Dissecting Policy Designs: An Application of the Institutional Grammar Tool." *Policy Studies Journal* 39 (1): 79–103.

Steyvers, Mark, and Tom Griffiths. 2007. "Probabilistic Topic Models." In *Handbook of Latent Semantic Analysis*, eds. Thomas K. Landauer, Danielle S. McNamara, Simon Dennis, and Walter Kintsch. New York: Taylor & Francis.

Stone, Deborah A. 1989. "Causal Stories and the Formation of Policy Agendas." *Political Science Quarterly* 104 (2): 281–300.

Vandenbosch, Robert, and Susanne E. Vandenbosch. 2007. *Nuclear Waste Stalemate: Political and Scientific Controversies*. Salt Lake City, UT: University of Utah Press.

Weaver, David H. 2007. "Thoughts on Agenda Setting, Framing, and Priming." *Journal of Communication* 57 (1): 142–47.

Wood, B. Dan, and Alesha Doan. 2003. "The Politics of Problem Definition: Applying and Testing Threshold Models." *American Journal of Political Science* 47 (4): 640–53.

Wood, B. Dan, and Arnold Vedlitz. 2007. "Issue Definition, Information Processing, and the Politics of Global Warming." *American Journal of Political Science* 51 (3): 552–68.

Workman, Samuel, Bryan D. Jones, and Ashley E. Jochim. 2009. "Information Processing and Policy Dynamics." *Policy Studies Journal* 37 (1): 75–92.

Zahariadis, Nikolaos. 2007. "The Multiple Streams Framework: Structure, Limitations, Prospects." In *Theories of the Policy Process*, ed. Paul A. Sabatier. Boulder, CO: Westview Press, 65–92.