

Topic Modeling the President: Conventional and Computational Methods

J.B. Ruhl, John Nay** & Jonathan Gilligan****

ABSTRACT

Law is generally embodied in text, and lawyers have for centuries classified large bodies of legal text into distinct topics—that is, they “topic model” the law. But large bodies of legal documents present challenges for conventional topic modeling methods. The task of gathering, reviewing, coding, sorting, and assessing a body of tens of thousands of legal documents is a daunting proposition. Yet recent advances in computational text analytics, a subset of the field of “artificial intelligence,” are already gaining traction in legal practice settings such as e-discovery by leveraging the speed and capacity of computers to process enormous bodies of documents, and there is good reason to believe legal researchers can take advantage of these new methods as well. Differences between conventional and computational methods, however, suggest that computational text modeling has its own limitations. The two methods used in unison, therefore, could be a powerful research tool for legal scholars.

To explore and critically evaluate that potential, we assembled a large corpus of presidential documents to assess how computational topic modeling compares to conventional methods and evaluate how legal scholars can best make use of the computational methods. We focused on presidential “direct actions,” such as executive orders, presidential memoranda, proclamations, and other exercises of authority the President can take alone, without congressional concurrence or agency involvement. Presidents have been issuing direct actions throughout the history of the republic, and although the actions have often been the target of criticism and controversy in the past, lately they have become a tinderbox of debate. Hence, although long ignored by political scientists and legal scholars, there has been a surge of interest in the scope, content, and impact of presidential direct actions.

* David Allen Daniels Distinguished Chair of Law, Vanderbilt University Law School; Co-founder of Skopos Labs, Inc., a software company developing applications of artificial intelligence in the legal and financial industries.

** Post-Doctoral Research Fellow, Information Law Institute, New York University; Adjunct Professor of Law, New York University School of Law; Affiliate, Berkman Klein Center, Harvard University; CEO & Co-founder of Skopos Labs, Inc.

*** Associate Professor, Department of Earth and Environmental Sciences, Vanderbilt University.

The authors are thankful to participants in the 2017 Vermont Law School Environmental Law Colloquium for their helpful comments on an early draft of this work, and to the research assistance of Vanderbilt Law School students Ryan Dewey, Laura Komarek, and Caroline Schube.

Legal and policy scholars modeling direct actions into substantive topic classifications thus far have not employed computational methods. To compare the results of their conventional modeling methods with the computational method, we generated computational topic models of all direct actions over time periods other scholars have studied using conventional methods, and did the same for a case study of environmental-policy direct actions. Our computational model of all direct actions closely matched one of the two comprehensive empirical models developed using conventional methods. By contrast, our environmental-case-study model differed markedly from the only empirical topic model of environmental-policy direct actions using conventional methods, revealing that the conventional methods model included trivial categories and omitted important alternative topics.

Provided a sufficiently large corpus of documents is used, our findings support the assessment that computational topic modeling can reveal important insights for legal scholars in designing and validating their topic models of legal text. To be sure, computational topic modeling used alone has its limitations, some of which are evident in our models, but when used along with conventional methods, it opens doors towards reaching more confident conclusions about how to conceptualize topics in law. Drawing from these results, we offer several use cases for computational topic modeling in legal research. At the front end, researchers can use the method to generate better and more complete topic-model hypotheses. At the back end, the method can effectively be used, as we did, to validate existing topic models. And at a meta-scale, the method opens windows to test and challenge conventional legal theory. Legal scholars can do all of these without “the machines,” but there is good reason to believe we can do it better with them in the toolkit.

TABLE OF CONTENTS

INTRODUCTION	1245
I. DEFINING AND CLASSIFYING PRESIDENTIAL DIRECT ACTIONS	1253
A. <i>The Family of Direct Actions</i>	1253
B. <i>Conventional Direct-Action Topic Models</i>	1258
1. <i>Meta-Topic Models</i>	1259
2. <i>Environmental Topic Models</i>	1264
II. COMPUTATIONAL TOPIC MODELS	1272
A. <i>Computational Topic Modeling Basics</i>	1274
B. <i>Study Methods and Design</i>	1281
III. FINDINGS AND ASSESSMENTS	1284
A. <i>The Meta-Topic Model</i>	1284
1. <i>Findings</i>	1285
2. <i>Assessment</i>	1290
B. <i>Environmental Topic Models</i>	1292
1. <i>All Direct Actions Model</i>	1294
a. <i>Findings</i>	1294

b. <i>Assessment</i>	1300
2. Executive Orders Model	1306
a. <i>Findings</i>	1308
b. <i>Assessment</i>	1309
IV. RESEARCH USE CASES	1310
A. <i>Front End: Generating Model Hypotheses</i>	1311
B. <i>Back End: Validating Models</i>	1312
C. <i>Meta-Scale: Building and Challenging Theory</i>	1312
CONCLUSION	1313

INTRODUCTION

[O]ne of the things that I have learned in the last two years is that the President can do an awful lot of things by executive action

—President William J. Clinton¹

Presidential direct actions—the flow of executive orders, presidential memoranda, proclamations, declarations, executive agreements, national security directives, signing statements, and similar official missives emanating from the White House—are a President’s means of flexing legal and policy muscle without congressional concurrence or agency initiative.² The political seesaw that has defined control of the White House over the past three decades has fueled

¹ *The Clinton Record; Interview with Clinton: Political Landscape*, N.Y. TIMES (July 28, 1996), <http://www.nytimes.com/1996/07/28/us/the-clinton-record-interview-with-clinton-political-landscape.html> [https://perma.cc/FG93-9SYE].

² See generally PHILLIP J. COOPER, BY ORDER OF THE PRESIDENT: THE USE AND ABUSE OF EXECUTIVE DIRECT ACTION (2d ed. 2014) (providing a comprehensive overview of each of these instruments, referring to them collectively as “executive direct actions”). Other scholars refer to these instruments as “unilateral presidential directives.” See, e.g., GRAHAM G. DODDS, TAKE UP YOUR PEN 4 (2013). However, for many (but not all) direct actions, the President is acting *directly* through the action—i.e., without need of congressional consent or agency involvement—but the action is taken pursuant to a statutorily delegated authority and thus is not a purely *unilateral* exercise of power. Presidential proclamations establishing national monuments under authority of the Antiquities Act are a classic example of this, as discussed in more detail *infra* in Part I. Other commentators refer to all direct actions as “executive orders,” capturing their essence as being an action taken by the President without legislative action. See *id.* at 16–17 (commenting on this practice). Yet, direct actions are formally divided into several types, including those expressly promulgated as executive orders, and thus using the term “executive order” to refer to all direct actions can lead to confusion and obfuscate differences among the types. We adopt Cooper’s “direct action” nomenclature for these reasons—dropping the “executive” as it can be presumed for our purposes—to avoid the problems that come with using the term “unilateral” or “executive order” while capturing the idea that these are mechanisms for the President to act directly, without congressional consent or agency involvement. Regardless of what one calls them, it is almost always the case that many actors from within the White House, and often from agencies as well, are involved in the negotiation and drafting of direct actions for the

aggressive use of direct actions, putting them front and center in the public eye and in the debate over presidential concentration of power.³ Presidential scholars thus have begun to study direct actions as an important component of a President's tenure and legacy, providing a window into how a particular President, acting presidentially, shaped a policy legacy. Presidential scholars also use the actions as a medium for tracing patterns and trends in the Office of the President over time. By contrast, direct actions as a class of presidential action have not received much attention from legal scholars from the perspective of the actions' legality, process, and reception in Congress and the courts,⁴ but President Trump's use of direct actions has stimulated more research in those respects.⁵

Most studies of presidential direct actions are descriptive or theoretical, using selected actions and historical context as representative case studies to develop accounts and theories of the presidency.⁶ A few researchers have used empirical methods to classify Presidents and the Presidency into topics and eras as a foundation for analysis of

President's final say and signature. See Andrew Rudalevige, *The Contemporary Presidency: Executive Orders and Presidential Unilateralism*, 42 PRESIDENTIAL STUD. Q. 138, 142–44 (2012).

³ COOPER, *supra* note 2, at 20 (“[P]residential direct action has been at the root of some of the most intense debates in American history.”); *id.* at 118 (discussing the extensive use and evolution of direct actions by Presidents since President Reagan); DODDS, *supra* note 2, at 1–4 (“Public awareness of unilateral presidential directives has been growing . . .”).

⁴ One notable exception is Professor Kevin Stack's series of articles exploring the exercise and judicial review of presidential direct actions implementing statutorily delegated powers. See generally Kevin M. Stack, *The Reviewability of the President's Statutory Powers*, 62 VAND. L. REV. 1171 (2009); Kevin M. Stack, *The President's Statutory Powers to Administer the Laws*, 106 COLUM. L. REV. 263 (2006); Kevin M. Stack, *The Statutory President*, 90 IOWA L. REV. 539, (2005) [hereinafter Stack, *The Statutory President*]. Several legal academics and practitioners have voiced concerns over presidential abuse of direct actions. See, e.g., Tara L. Branum, *President or King? The Use and Abuse of Executive Orders in Modern-Day America*, 28 J. LEGIS. 1, 2 (2002); John C. Duncan, Jr., *A Critical Consideration of Executive Orders: Glimmerings of Auto-poiesis in the Executive Role*, 35 VT. L. REV. 333, 344–45 (2010); Todd F. Gaziano, *The Use and Abuse of Executive Orders and Other Presidential Directives*, 5 TEX. REV. L. & POL. 267, 297–316 (2001).

⁵ For example, legal scholars have differed sharply over the legality of President Trump's presidential declarations reducing the size of two large national monument areas. Compare Richard H. Seamon, *Dismantling Monuments*, 70 FLA. L. REV. (forthcoming 2018) (manuscript at 51), <https://ssrn.com/abstract=3054682> [<https://perma.cc/F83E-4GNH>] (concluding that President Trump has the authority to “undo the acts of his predecessors” by “reduc[ing] or rescind[ing] monuments they created”), with Mark Squillace et al., *Presidents Lack the Authority to Abolish or Diminish National Monuments*, 103 VA. L. REV. ONLINE 55, 56 (2017) (arguing that “the President lacks the legal authority to abolish or diminish national monuments”).

⁶ See, e.g., DODDS, *supra* note 2; RICARDO JOSÉ PEREIRA RODRIGUES, *THE PREEMINENCE OF POLITICS: EXECUTIVE ORDERS FROM EISENHOWER TO CLINTON* (2007).

presidential exercise of authority.⁷ With tens of thousands of direct actions on the books,⁸ however, conducting a comprehensive empirical study of all direct actions to develop a granular model of the topics they address is a daunting undertaking. Researchers have used several familiar methods to work around this classic problem of how to work with massive bodies of text. One way is to limit the number of documents categorically. Indeed, most empirical studies of direct actions to date have included only executive orders.⁹ A prominent example is Lyn Ragsdale's ten-topic classification of executive orders issued from 1949 through 1997, presented in his indispensable *Vital Statistics on the Presidency*.¹⁰ Alternatively, a researcher might review all the documents but sort them into a coarse classification system to reduce the labor of producing a more granular classification, as Adam Warber did for over 5,000 executive orders issued from 1936 through 2001, classifying the content of each as either symbolic, routine, policy, or hybrid.¹¹ Another approach uses conventional random sampling to reduce the number of documents and thereby allow more granular topic-coding methods, such as Kenneth Mayer's classic study of about 1,000 of the more than 5,800 executive orders issued during the period from 1936 through 1999.¹² Other researchers have reduced the study set to a manageable number by carving out particular themes for evaluation, as Jonathan West and Glen Sussman did for executive orders relating to environmental policy for the period from 1933 through 1995.¹³

All of these conventional research methods come at a cost. Focusing on executive orders to the exclusion of all other direct actions necessarily skews any topic model, suppressing the influence of other direct-action mechanisms on our broader understanding of presidential use of direct-action authorities. executive orders are considered the most prominent of direct actions but are by no means the only

⁷ See, e.g., KENNETH R. MAYER, *WITH THE STROKE OF A PEN: EXECUTIVE ORDERS AND PRESIDENTIAL POWER* (2001); see Jonathan P. West & Glen Sussman, *Implementation of Environmental Policy: The Chief Executive*, in *THE ENVIRONMENTAL PRESIDENCY* 77–111 (Dennis L. Soden ed., 1999).

⁸ Due to poor recordkeeping prior to the mid-1900s, the exact number of direct actions is not known. See DODDS, *supra* note 2, at 15–17. We discuss how we assembled our database of direct actions in Part II.

⁹ See *infra* Part I.

¹⁰ See LYN RAGSDALE, *VITAL STATISTICS ON THE PRESIDENCY* 353–56 (1998).

¹¹ See ADAM L. WARBER, *EXECUTIVE ORDERS AND THE MODERN PRESIDENCY* 39–41 (2006).

¹² See MAYER, *supra* note 7, at 79.

¹³ See West & Sussman, *supra* note 7, at 80.

mechanisms through which Presidents have exercised significant and substantial policy muscle.¹⁴ A coarse topic model like Warber's constrains the breadth and depth of content analysis. For example, Warber sorts almost 40% of the executive orders from his study time period into the "policy" category, but does not provide more detailed substantive classifications of the various policy themes.¹⁵ Random sampling can serve as a starting point, but it is not generally statistically sound for document analysis because documents differ so much from one another that an enormous sample size is required before one can be confident that the sample is representative of the corpus.¹⁶ Both random sampling and limited-theme sampling rely on the human researcher to interpret the thrust of particular documents to develop the topic model, as Mayer did to derive his top eight topics¹⁷ and as West and Sussman did to decide first what qualified as "environmental" executive orders and then to divide them into twelve topics.¹⁸ In short, all the workarounds for the large-text-corpus problem come with methodological baggage.

Developments in computational text analysis methods over the past decade offer a different approach to topic modeling for a large text corpus. Using natural language processing and machine learning algorithms to detect semantic structure patterns, enormous bodies of text units or documents¹⁹ can be classified into semantically similar clusters without human direction, requiring only that the researcher later assign a label to the clusters based on the key words and the documents the analytics identify as the core of a semantic cluster.²⁰ If the text corpus spans a time period, the analytics also can trace the ebb and flow of particular topics in the model as a component of the corpus over time by modeling how time affects the prevalence of a topic. In essence, these methods flip the research process, using the

¹⁴ See DODDS *supra* note 2, at 5–10.

¹⁵ See WARBER, *supra* note 11, at 39, 55–60, 140–45.

¹⁶ Generating representative samples from textual corpora is fraught with subtle challenges that require careful stratified sampling designs rather than simple proportional random sampling. See Douglas Biber, *Representativeness in Corpus Design*, 8 LITERARY & LINGUISTIC COMPUTING 243, 243–48 (1993); Ted Dunning, *Accurate Methods for the Statistics of Surprise and Coincidence*, 19 COMPUTATIONAL LINGUISTICS 61, 61–62, 71 (1993).

¹⁷ See MAYER, *supra* note 7, at 80–81.

¹⁸ See West & Sussman, *supra* note 7, at 85–86.

¹⁹ Although these modeling exercises involved entire discrete documents in the form of whole direct actions, a document for purposes of computational topic modeling could be any unit of text, such as sections of bills, statutes, or regulations divided at a selected scale.

²⁰ We describe computational topic modeling methods in detail *infra* in Part II.

“machine” to classify and trace topics in the text corpus first, and then relying on the human to interpret the results.

In their recent application of these computational tools to the corpus of U.S. Supreme Court opinions, legal scholar Michael Livermore and his co-authors referred to this method as moving from the “top down” approach of relying on the human researcher to use theory and expertise to develop the topic model, to a “ground up” approach relying on the machine to extract the topics directly from the documents with no preconceived theory or model.²¹ An additional distinction is one based on timing. The conventional research method, at least over the past several decades, has employed computational methods but at the *back end* of the research project, when data are crunched through statistical analyses such as linear regression. By contrast, the computational topic modeling method uses computational technology at the *front end* to construct the topic model before the researcher dives in for deeper analyses.²²

The question, of course, is whether the front-end/“ground up” constructed topic model makes sense to a human researcher in the relevant field and leads to a new understanding of the subject matter. To satisfy that test, the computational method need not produce the same topic model that a human researcher would produce using conventional research methods. Indeed, the point of using computational text analytics is to leverage computational power operating vastly beyond a human’s capacity, thereby opening up the possibility of identifying semantic structures in the text corpus that a human researcher would not detect. So, the front-end/“ground up” model might be different but also better in some respects in terms of classifying the text set into topic clusters. Or, the alternative computational model, if not a full substitute for the human researcher’s model, might help the researcher refine a topic model produced through conventional “top down” random sampling and coding methods. Even if the two methods produce the same model, one advantage of computers over humans is undeniable—if the computational method produces a useful topic model, the computer can take on vastly larger text sets and do the job much faster.

²¹ See Michael A. Livermore et al., *The Supreme Court and the Judicial Genre*, 59 ARIZ. L. REV. 837, 856 (2017) (“This approach of defining the genre from the ground up (from the observed documents), rather than from the top down (based on a theory of judicial legitimacy), has some useful advantages.”).

²² This is related to the idea of “grounded theory.” See Eric P. S. Baumer et al., *Comparing Grounded Theory and Topic Modeling: Extreme Divergence or Unlikely Convergence?*, 68 J. ASS’N INFO. SCI. & TECH. 1397, 1399 (2017).

Nevertheless, computers do not cognitively understand anything about the text set, the topics the computer produces, or the broader context within which the text and topics exist, meaning a computational topic model cannot evaluate its fit with the real world. Only a human can do that. The *semantic* clusters the computational method separates into topics also might not produce *substantive* topics of any coherence for the legal domain. There are other differences between the two methods that affect how they can be deployed and their respective results interpreted.²³ Neither can fully replicate what the other is capable of producing.

Hence, rather than thinking of the computer as replacing the human, in an ideal world, the two working together would be better than either working alone. Humans can interpret real world meaning and fit of topics far better than algorithms can, but given the challenges posed by sampling, together with the time it takes for a person to read and interpret a document, computational topic modeling may help a human researcher choose more useful documents to analyze, assure a more representative selection of documents, and enable the researcher to obtain greater value for the time invested in close reading. To explore whether and how computational topic modeling can be leveraged to realize that possibility, we assembled and analyzed a corpus of four predominant types of direct actions—executive orders, presidential memoranda, proclamations, and presidential determinations.²⁴ We compared our results and interpretations to those Mayer and Ragsdale reached in their respective executive order studies and to those West and Sussman reached in their study of environmental-policy executive orders, as well as to the historical and theoretical accounts of direct actions both generally and regarding environmental policy.

Of course, one might reasonably ask, why topic model at all? But even by the second week of law school, a law student could give an answer: to help us make sense of it all! Topic models are a means of organizing large bodies of knowledge into coherent structures that help us navigate the corpus of information. Consider Westlaw's familiar Topic and Key Number system, which Westlaw claims is "an indispensable part of learning how to do effective legal research" and describes as follows:

²³ See *infra* Part II.

²⁴ We explain the differences between these four types of direct actions *infra* in Section I.A.

The Topic and Key Number System is a big outline or index that works like this:

1. The American system of law is broken down into Major Topics—there are more than 400 topics, such as Civil Rights, Pretrial Procedure, and Treaties.
2. Each of those topics is divided, in greater and greater detail, into individual units that represent a specific legal concept—like steps in an outline. There can be up to eight steps in the hierarchy to reach the narrowest concept.
3. Each of the narrowest concepts (and there are approximately 100,000 of them!) has a unique number that allows you to find it on the outline. This number is called a specific Key Number.²⁵

We are not proposing that direct actions be divided into 100,000 topics. The point is clear—topic models help lawyers organize the law and legal institutions into coherent categories. In the context of direct actions and other text documents, such as legislative bills, agency rules, and compliance filings, topic modeling also can provide insight into what institutions work on. For example, given the high profile direct actions are taking on, it may be useful to know what they are about broadly before offering assessments of their impacts. Yet, for any of these information domains, there is no single inevitable topic model. For example, starting over from scratch could lead to many different versions of the 400 Major Topics and 100,000 narrowest concepts included in Westlaw’s Key Number System, some more useful than others.

This point—that there are multiple possible coherent models of any corpus of legal text—suggests that there is more to topic modeling than organizing and categorizing. More deeply, topic models can also help lawyers *conceptualize* law and legal institutions. For example, if two human researchers using conventional methods developed two vastly different topic models of presidential direct actions, they could both be completely accurate categorizations, but one might be far more useful in providing insight into the role of the President for any particular purpose. A topic model built around fields of policy (e.g., war, trade, labor, environment) may help in conceptualizing constitutional distribution of powers, whereas one built around the functions of the actions (e.g., communicating policy preferences to agencies,

²⁵ Description of Westlaw’s Topic and Key Number System, WESTLAW, <https://lawschool.westlaw.com/marketing/display/RE/24> [<https://perma.cc/VSS3-LTYV>].

managing internal agency affairs, sanctioning foreign governments) could help in conceptualizing the President as a political actor.

To the extent computational topic modeling can help construct better topic models for lawyers and legal scholars to use in any of these senses, it should be evaluated for that purpose. Indeed, computational text modeling has begun to take hold widely in practical legal applications such as e-discovery platforms²⁶ and caselaw search engines.²⁷ By contrast, legal scholars have only just begun to experiment with applying computational text modeling techniques in their research,²⁸ with its efficacy compared to conventional legal empirical studies methods yet to be assessed in application. To simulate that evaluation, this Article reports the design, findings, and assessment of a side-by-side comparison of conventional and computational topic model research techniques and results applied across the same large legal-text corpus compiled over time. Our primary objective is to demonstrate and evaluate computational text modeling as a research tool for legal scholars. Incidental to that methodological purpose, we also offer some conclusions regarding what our computational study reveals about presidential use of direct actions generally and in shaping environmental policy. In that regard, however, we do not purport to offer a comprehensive review of presidential direct actions, either generally or for environmental policy, nor are we laying out an instruction manual for computational text modeling. Rather, after providing the necessary background for each, we go to the heart of the matter by comparing text models of direct actions using conventional and computational methods.

The Article proceeds in four parts. In Part I we provide the context of direct actions and summarize the existing historical, theoretical, and empirical studies classifying topics, Presidents, and eras for all direct actions and, as an in-depth case study, for direct actions relating to environmental policy. Part II presents the basics of computational topic modeling and explains our study methods. Part III presents our findings and assesses how our results compare to the prior studies and what can be drawn substantively and methodologically from the comparisons. We close in Part IV with observations about how legal schol-

²⁶ See, e.g., *Ringtail Demo Request*, RINGTAIL, https://www.ringtail.com/demo-request?gclid=EA1aIQobChMI1szG1qDz1wIViVcNCh059QELEAAAYASAAEgImXvD_BwE [<https://perma.cc/A3R3-4GZN>] (e-discovery software).

²⁷ See, e.g., CASETEXT, <https://casetext.com> [<https://perma.cc/D62P-TMCN>].

²⁸ See, e.g., Livermore et al., *supra* note 21, at 841–42, 862 (discussing the few other legal studies employing computational topic modeling).

ars can use computational topic modeling analytics to inform their research, presenting use cases of the computational model for legal scholars.

Our study demonstrates that computational topic modeling can substantially contribute to theoretical and empirical legal studies, but also that the computational method is no panacea. Two key results from our models support this moderated bottom line. First, as anyone using computational methods will confirm, their performance improves dramatically as the amount of relevant data increases. As the time-period and direct-action-type parameters of our models expanded, our topic models performed more robustly. Computational techniques, therefore, will be less useful to legal scholars working with a small text corpus. On the other hand, when computational topic models do perform robustly—as we conclude several of our models did—they can provide important insights into the text corpus, allowing researchers to test an existing topic classification, rethink the topic divisions, or generate a set of classifications as a starting point. But even when they perform well, these topic classifications require subjective human interpretation to give them meaning. Computational methods of topic modeling thus will not substitute for conventional methods, but the reverse also is true. Together, therefore, they can provide a powerful research platform for exploring the meaning and content of large bodies of legal-text documents, as well as for validating or challenging broader conceptions of how law and legal institutions are thematically structured.

I. DEFINING AND CLASSIFYING PRESIDENTIAL DIRECT ACTIONS

A. *The Family of Direct Actions*

Phillip Cooper's *By Order of the President*, published in its second edition in 2014, is a masterful overview of direct actions, working through each major type to describe its features and uses and to assess its place in history over time.²⁹ Like many direct action researchers (including us), for source material he draws heavily from the American Presidency Project (“APP”), which is maintained by John Woolley and Gerhard Peters and hosted online by the University of California at Santa Barbara.³⁰ Although the APP includes empirical and analytical evaluations of direct actions, Cooper provides a far deeper account

²⁹ COOPER, *supra* note 2.

³⁰ John T. Woolley & Gerhard Peters, AM. PRESIDENCY PROJECT, <http://www.presidency.ucsb.edu/index.php> [<https://perma.cc/Y24S-Q9ND>].

of the different direct action types, covering seven major instruments: executive orders, presidential memoranda, presidential proclamations, signing statements, presidential determinations, national security directives, and Executive agreements.³¹ Although, depending on how one classifies them, there are almost thirty different types of presidential direct actions and their boundaries are fuzzy at best,³² Cooper's seven types correspond closely to the APP categorizations, from which we drew our data.

Executive Orders. Considered the most formal and prominent of the direct actions,³³ executive orders are written directives to government officials and agencies of the executive branch delegating authority to the agencies to implement specific statutes or instructing them to take action, stop a specified activity, or change policy or management direction.³⁴ The State Department began numbering executive orders in 1907, and since the Federal Register Act of 1935, executive orders are, in almost all cases, required to be published in the *Federal Register*.³⁵

Presidential Memoranda. Cooper refers to these as “executive orders by another name,” in the sense that “[a]s a practical matter, the memorandum is now being used as the equivalent of an executive order, but without meeting the legal requirements for an executive order,” such as numbering and publishing.³⁶ Modern Presidents have routinely used both executive orders and memoranda interchangeably, and the conventional view is that there is no substantive difference in legal force or effect.³⁷

Presidential Proclamations. These instruments, which must be published in the *Federal Register*, state conditions, trigger implementation of laws, and recognize symbolic events, such as declaring a natural disaster or declaring a day or week of recognition.³⁸ Whereas

³¹ See COOPER, *supra* note 2, at 2, 123–24; see also MAYER, *supra* note 7, at 35 (“The major classes of presidential policy instruments are executive orders, proclamations, memoranda, administrative directives, findings and determinations, and regulations.”).

³² DODDS, *supra* note 2, at 5–10.

³³ See MAYER, *supra* note 7, at 35 (“[T]here is little doubt that presidents and their staffs consider executive orders to be the most important statements of executive policy.”). *But see*, e.g., COOPER, *supra* note 2, at 120–23 (describing presidential memoranda as having become indistinguishable from executive orders); WARBER, *supra* note 11, at 140 (describing executive orders and substantive proclamations as indistinguishable).

³⁴ COOPER, *supra* note 2, at 21.

³⁵ *Id.* at 22, 24.

³⁶ *Id.* at 115, 120–23.

³⁷ *Id.* at 120–21.

³⁸ *Id.* at 172.

executive orders and memoranda generally are directed to federal agencies and officials within the executive branch, proclamations generally are aimed outward, to foreign, state, local, and private institutions.³⁹

Presidential Determinations. Although similar to presidential memoranda, determinations generally are focused on foreign policy and are numbered chronologically by fiscal year.⁴⁰ They are usually made pursuant to statutes that require the President “to make findings concerning the status of a foreign country or some activity in the foreign policy field,” at which point some action or other condition is triggered under the statute.⁴¹

National Security Directives. These are formal notifications to government agencies or officials regarding presidential decisions in the field of national security to coordinate military policy, foreign policy, intelligence policy, or other security policies, usually those managed through the National Security Council.⁴²

Executive Agreements. Cooper describes Executive agreements as “[t]he substance of a treaty without the constitutional process.”⁴³ Indeed, he notes that the State Department defines two kinds of international agreements, treaties and Executive agreements, the latter being “other international agreements” the President enters into pursuant to a treaty, legislation, or “the constitutional authority of the president.”⁴⁴

Signing Statements. These are written comments a President issues at the time of signing legislation. Although most merely comment briefly and favorably on the bill signed, the more controversial statements express concerns and limitations. For example, the statement might claim that the legislation infringes on the constitutional powers of the Presidency, announce interpretations of language used in the legislation, or instruct executive branch officials how to implement the new law, including by ignoring it.⁴⁵

Most new Presidents swiftly make use of these forms of direct action—often on the first day they occupy the White House—and frequently do so to undo a predecessor’s direct actions.⁴⁶ This has often

³⁹ *Id.* at 173.

⁴⁰ *Id.* at 123.

⁴¹ *Id.* at 123–24.

⁴² *Id.* at 208.

⁴³ *Id.* at 282.

⁴⁴ *Id.*

⁴⁵ *Id.* at 325.

⁴⁶ *Id.* at 117–18 (discussing President Clinton); *id.* at 68 (discussing President George W.

attracted accusations that the President is playing “power grab” politics,⁴⁷ but the historical fact is that for well over a century Presidents have been using direct actions extensively, and with meaningful force and effect, throughout their terms.⁴⁸ The inflection point, however—when presidential use of direct actions increased by an order of magnitude—was the Administration of President Theodore Roosevelt, who issued almost as many executive orders (1,081) as did all the Presidents in the 112 years before him combined (1,262).⁴⁹ Use of executive orders has ebbed and flowed since then but has remained relatively stable since the mid-1900s. The average number of executive orders has been under 100 per year since Harry S. Truman, albeit with other direct-action types slowly gaining in frequency over time.⁵⁰

To illustrate the point, Figure 1 provides a histogram of the four most potent and broadly deployed direct actions—executive orders, presidential memoranda, proclamations, and determinations—issued from January 1929 to June 2017. We separated proclamations into substantive and symbolic categories by classifying those with terms suggestive of a nonsubstantive purpose in the title, such as “week” in a proclamation declaring National Boating Week (TrivialProc in Figure 1), as symbolic. We then combined substantive proclamations with determinations, given that these types of direct actions have similar purposes (Proc_or_Det in Figure 1).

Bush); *id.* at 32–33, 68 (discussing President Obama); see WARBER, *supra* note 11, at 47–61 (surveying this practice through time).


⁴⁷ See COOPER, *supra* note 2, at 3 (“Rule by presidential decree has been the subject of serious controversy since the administration of George Washington . . .”).

⁴⁸ See COOPER, *supra* note 2, at 20 (“[There is] certainly nothing new about making quick use of the executive order to enact policy and communicate political messages.”); DODDS, *supra* note 2, at 151 (“[Theodore Roosevelt’s] successors generally followed his precedent of regularly using unilateral presidential directives for a wide variety of purposes.”).

⁴⁹ See DODDS, *supra* note 2, at 121. Consistent with many other assessments, Dodds concludes that “[t]he nature of the use of unilateral presidential directives changed dramatically with Theodore Roosevelt . . .” *Id.* at 27.

⁵⁰ See Woolley & Peters, *supra* note 30.

FIGURE 1. EXECUTIVE ORDERS, PRESIDENTIAL MEMORANDA, SUBSTANTIVE PROCLAMATIONS/DETERMINATIONS, & TRIVIAL PROCLAMATIONS, JANUARY 1929–JUNE 2017⁵¹



As a result of gaps in the APP database of digitized direct actions, explained in more detail below, our dataset is incomplete prior to President Truman. Nevertheless, the histogram clearly shows that the annual number of executive orders declined soon after World War II and has been roughly stable since 1952. By contrast, presidential memoranda and substantive proclamations and determinations have increased in number since then,⁵² though they have stabilized since 1980, and nonsubstantive symbolic actions have risen dramatically since 1980, becoming the dominant type. Direct action scholars have noted these trends and offered various explanations.⁵³ Whatever the explanation, however, these trends suggest that studying only execu-

⁵¹ Data are incomplete prior to President Truman (1945).

⁵² See generally COOPER, *supra* note 2, at 120–23 (describing the rise, since President Reagan, in use of presidential memoranda to accomplish the same ends as executive orders).

⁵³ See DODDS, *supra* note 2, at 215–22.

tive orders does not provide a complete picture of presidential use of direct actions,⁵⁴ a theme we explore further below.

Clearly, direct actions are and always have been a component of presidential exercise of power. But knowing how often they have been used does not indicate which policy domain any particular direct action was targeting. Nor do numbers of direct actions reveal anything about substance. As Cooper observes, the media and political commentators have made news stories out of how many direct actions a President has issued over a span of time (especially at the beginning of a term), but “this is an unhelpful exercise because the issue is more about content than quantity,” and thus “running the numbers tells us relatively little.”⁵⁵ Numbers also tell us very little about impact. Indeed, as Mayer notes, most political scientists, particularly those studying the Presidency, have minimized the significance of direct actions, portraying them as mostly addressing routine administrative matters limited in scope and reach and cherry picking for further study only the ones they believe are important.⁵⁶ To be sure, most direct actions are, by any measure, mundane on the surface—they are used to move public lands among agencies, set civil service pay, and declare national days of recognition.⁵⁷ But many are nontrivial, and in the aggregate, Mayer argues, they reveal how a President uses constitutional, statutory, and other powers to act without congressional or agency involvement.⁵⁸ Topic modeling exercises assessing large bodies of direct action documents thus have begun to take hold in presidential studies. We turn in the next Section to examine how other researchers have used conventional research methods to gain this deeper insight into presidential use of direct actions.

B. Conventional Direct-Action Topic Models

Cooper observes that “until recently, the literature on the presidency has largely ignored the tools of presidential direct action,”⁵⁹ yet

⁵⁴ See COOPER, *supra* note 2, at 16 (discussing the rising use of direct actions other than executive orders).

⁵⁵ *Id.* at 16.

⁵⁶ See MAYER, *supra* note 7, at 10.

⁵⁷ See WARBER, *supra* note 11, at 37–40 (describing most executive orders in his study period as symbolic or routine).

⁵⁸ See MAYER, *supra* note 7, at 79–86 (contesting the view of many political scientists that executive orders are “merely a routine tool, not . . . an instrument for making important policy decisions”).

⁵⁹ COOPER, *supra* note 2, at 2. See MAYER, *supra* note 7, at 11 (“If executive orders are such an important element of presidential power, why have political scientists paid so little attention to them?”).

now they are a tinderbox of controversy. Given the growing recognition of their role in defining a President, an important step in assessing the use and impact of direct actions is to develop a more complete picture of the policy domains on which Presidents have focused this form of presidential authority—that is, to build a topic model. Scholars have approached this at meta-levels, providing a broad view of direct action deployment, and at more granular scales, taking one policy domain from the meta-set, such as environmental policy, and dissecting it into subtopics. In this Section, we review several prominent meta-topic and environmental-topic direct-action studies, which will serve as the comparators for our computational-topic-modeling study described in Parts II and III.

1. *Meta-Topic Models*

In building out his central thesis that “recent presidents have pushed the boundaries of presidential power” through “their mixing and matching of direct action tools,”⁶⁰ Cooper delves deeply into each of the direct-action instruments, explaining their historical uses to extract the essence of each and why Presidents use it in lieu of others for particular objectives. His book is an indispensable guide to the history, strategies, tactics, and politics of direct actions, and in it Cooper demonstrates beyond question how central a role direct actions have played in American law and policy. Yet his approach of assessing how and why Presidents use different direct-action instruments produces only a rough topic model at best, in the form of section headings.

For example, Cooper includes the following as some of his examples of how Presidents use executive orders: to issue binding pronouncements to units of the executive branch; to make policy in fields generally conceded to the President; to initiate or direct regulation; to delegate authority to other agencies or officers; to reorganize agencies, to eliminate existing organizations, or create new ones; or to manage federal personnel.⁶¹ Similarly, for presidential memoranda Cooper’s section headings suggest several topics: to present a presidential veto; to make hortatory declarations; to initiate a policy purpose; and, one of his main points, to accomplish similar purposes as an executive order.⁶² He does the same for the other direct-action types, producing a long list of potential topics for further study.⁶³ A re-

⁶⁰ COOPER, *supra* note 2, at x.

⁶¹ See COOPER, *supra* note 2, at 25–38.

⁶² See *id.* at 123–39.

⁶³ See generally *id.*

searcher could, for example, compile a complete list of Cooper's headings and rework it into a more compact set of administrative and substantive law and policy topics through which to conduct an empirical frequency assessment. Cooper, however, uses the topics primarily to plumb the history of strategic uses of direct actions of various kinds and to highlight their differences, not to develop a robust topic model.

Although other scholars have included some form of topic classification to dissect the direct-action story, most have been limited in scope and do not purport to have generated a complete topic model.⁶⁴ Two direct-action studies have gone further in using empirical methods to develop a more complete and precise meta-topic model of direct actions. Mayer's *With the Stroke of a Pen*,⁶⁵ published in 2001, is most well-known and figures prominently in all subsequent studies of direct actions. Although his detailed historical and theoretical analyses include all direct-action types, his empirical study focuses exclusively on executive orders. He drew a random sample of 1,028 executive orders from the full corpus of approximately 5,800 issued from March 1936 through December 1999.⁶⁶ He then created eight "exhaustive and mutually exclusive categories" and sorted each document into the category he determined "best described the order's primary focus."⁶⁷ In the other empirical topic model, *Vital Statistics on the Presidency*, Lyn Ragsdale included *all* executive orders for the period from 1949 through 1997—just over 3,000 in total—and classified them into ten topics "on the basis of title description and, in some cases, the text of the orders."⁶⁸ His study is more data driven than Mayer's, however, with comparatively little substantive analysis.⁶⁹ Both studies show proportions of orders falling into each topic in total and over time, by decade for Mayer⁷⁰ and by year and administration

⁶⁴ See, e.g., RODRIGUES, *supra* note 6, at 31–273 (studying direct actions through three extensive case studies on use of executive orders to advance equal employment, regulatory review, and environmental policy); WARBER, *supra* note 11, at 76–86 (focusing on phases of the Presidency and direct actions, briefly offering some sense of three executive order policy domain topics: military and war policy (with subtopics); administrative reforms; and distributive, redistributive, and regulatory policy).

⁶⁵ MAYER, *supra* note 7.

⁶⁶ *Id.* at 79.

⁶⁷ *Id.* at 80.

⁶⁸ RAGSDALE, *supra* note 10, at 353–56, 356 n.

⁶⁹ *Id.* at 308.

⁷⁰ MAYER, *supra* note 7, at 81–82.

for Ragsdale.⁷¹ Table 1 shows the topics and total proportions for both topic models.

TABLE 1. COMPARISON OF MAYER VERSUS RAGSDALE, EXECUTIVE ORDERS TOPIC MODELS

MAYER 1936–1999		RAGSDALE 1949–1997	
Topic	%	Topic	%
Executive branch administration	25.5	Personnel/agency requests	25.3
Civil service	19.6	Governance/economic management	15.3
Public lands	15.6	Defense	15
Defense and military policy	11.9	Foreign trade and diplomacy	14.9
Foreign affairs	11.3	Natural resources/environment	12.7
War and emergency powers	7.1	Social welfare/civil rights	9.2
Labor policy	5.4	Ceremonial/cultural	4.4
Domestic policy	3.8	Federalism	4.4
		Agriculture	1.3
		Foreign aid	1.1

Neither author, however, explains how he developed his topic model.⁷² Did they construct them before reviewing the documents, and then sort documents into the topic, or did they create and revise topics as they read the documents? Or perhaps they started with a model and improvised along the way. Either way, the two “exhaustive and mutually exclusive”⁷³ topic models are different in several material respects, making comparisons difficult.⁷⁴ For example, Mayer’s lacks a distinct agriculture topic, possibly including orders dealing with agriculture in his domestic policy topic. Ragsdale’s lacks a distinct labor policy topic, perhaps including it in his social welfare/civil rights topic. Mayer’s public lands topic might correspond to Ragsdale’s natural resources/environment topic—the proportions are roughly the same—although much environmental and natural resources policy has nothing to do with public lands. And although it does appear that the sets of top-two topics in both models arguably roughly correspond between the models, the proportions of the two distinct topics do not.

⁷¹ RAGSDALE, *supra* note 10, at 353–56. Ragsdale reports totals broken down into Democratic and Republican Presidents. We recomputed for all executive orders combined.

⁷² See MAYER, *supra* note 7, at 79–80; RAGSDALE, *supra* note 10, at 304–05, 353–56.

⁷³ MAYER, *supra* note 7, at 80.

⁷⁴ See *supra* Table 1.

Given these differences, the claim that either topic model is “exhaustive and mutually exclusive” is subject to question. Neither model seems to be an exhaustive set of direct-action topics, particularly given that both are limited to executive orders, and the topics seem too malleable to be mutually exclusive. Mayer explains, for example, that some orders “addressed multiple issues or crossed policy boundaries,” but he nonetheless assigned such orders to only one topic based on his assessment of “the category that best described the order’s primary focus.”⁷⁵ And on what basis did Mayer pick eight topics and Ragsdale ten? Why not five, or twelve? For example, Mayer’s labor policy could be incorporated into his domestic policy topic, or a public lands topic could be carved out of Ragsdale’s natural resources/environment topic. Where did Mayer put orders dealing with the environment that were not public lands orders? Possibly in the domestic policy topic, but why not create a distinct topic?

Indeed, another topic model of direct actions is found in the National Archives’ chapter links to its index of presidential proclamations and executive orders issued from April 1945 through January 1989.⁷⁶ It has forty topic chapters, with the foreign relations and national defense chapters divided into five and three subchapters, respectively. Many of the chapters correspond to Mayer’s and Ragsdale’s topics—there are chapters on agriculture, labor, and public lands, for example—but many do not, such as the Archives’ chapters on the Panama Canal, banks and banking, and food and drugs, which find no corollaries in Mayer’s or Ragsdale’s models.

As the expert compiler and indexer of these documents, perhaps one should consider the National Archives’ topic classification as the gold standard. Alas, the National Archives appears to have put little thought into its topic model—the chapter organization merely duplicates the chapter organization of the Code of Federal Regulations (“C.F.R.”). It is also not clear the National Archives is as expert at classifying as one might think. For example, its “Protection of Environment” chapter contains only thirteen documents,⁷⁷ far below the

⁷⁵ MAYER, *supra* note 7, at 80.

⁷⁶ Proclamations and executive orders not in effect as of January 20, 1989, are not included. See *Index: Chapter Links to the Codification of Presidential Proclamations and Executive Orders*, NAT’L ARCHIVES, <https://www.archives.gov/federal-register/codification/chapter.html> [<https://perma.cc/3J6K-W4JE>]; see also *Numeric Codification of Presidential Proclamations and Executive Orders*, NAT’L ARCHIVES, <https://www.archives.gov/federal-register/codification/numeric.html> [<https://perma.cc/LQ3A-GUZR>].

⁷⁷ See Chapter 40—*Protection of Environment: Chapter Links to the Codification of Presi-*

number legal and political science scholars put in that category,⁷⁸ and the proclamation declaring the independence of the Philippines appears in, of all places, the public lands chapter.⁷⁹

Even if one puts trust into using the C.F.R. titles organization to classify direct actions and going with the National Archives' sorting of documents into it, studying direct actions using nearly fifty topics could dissect the corpus into units too small to be of analytical value.⁸⁰ The Panama Canal chapter, for example, contains only seven documents. One could easily collapse the National Archives' topics into a reduced set of eight or ten larger themes and produce a model every bit as coherent as Mayer's or Ragsdale's. But that is the point—different researchers will likely come up with different topic models, any of which could provide a reasonable picture of presidential direct-action themes.⁸¹

So, which of the two more compact topic models is better, Mayer's or Ragsdale's? It is hard to say. On the one hand, Ragsdale reviewed all the executive orders in his time frame,⁸² whereas Mayer randomly sampled and reviewed only 20% from nearly the same time frame.⁸³ On the other hand, Mayer read each order he sampled to determine its content and significance,⁸⁴ whereas Ragsdale classified primarily by the document's title.⁸⁵ Mayer claims that random sampling, because it reduces numbers and thus facilitates deeper reviews, “allows for a more detailed (and tractable) investigation into the question of what fraction of [Executive] orders can be considered signifi-

dential Proclamations and Executive Orders, NAT'L ARCHIVES, <https://www.archives.gov/federal-register/codification/chapter-40.html> [<https://perma.cc/4B3Y-REPF>].

⁷⁸ See *infra* Part III.

⁷⁹ *Chapter 43—Public Lands: Chapter Links to the Codification of Presidential Proclamations and Executive Orders*, NAT'L ARCHIVES, <https://www.archives.gov/federal-register/codification/chapter-43.html> [<https://perma.cc/7QVZ-2HCR>]. As discussed *infra* Part II, this is an example of a researcher being “boxed in” by a fixed predetermined topic model using a one-document, one-topic method. The C.F.R. has no appropriate topic for declaring a territory's independence; thus, the Archives chose the closest fit. See *id.*

⁸⁰ See *Index: Chapter Links to the Codification of Presidential Proclamations and Executive Orders*, NAT'L ARCHIVES, <https://www.archives.gov/federal-register/codification/chapter.html> [<https://perma.cc/J946-YBYL>].

⁸¹ See *id.* The Executive Office of the President also maintains a filing system with categories for executive orders, but because the indices it uses change over time, it is not a good candidate for developing a topic model spanning long periods of the Presidency. See Rudalevige, *supra* note 2, at 146.

⁸² See RAGSDALE, *supra* note 10, at 304–05, 353–56.

⁸³ See MAYER, *supra* note 7, at 79.

⁸⁴ See *id.* at 80.

⁸⁵ See RAGSDALE, *supra* note 10, at 353–56.

cant,”⁸⁶ and it is true that his book engages in extensive empirical analyses and significance classifications not found in Ragsdale’s work.⁸⁷ But does that make Mayer’s *topic model* better? If his topic model is in some way flawed, the flaw can carry through to his other empirical and descriptive analyses. For example, Mayer devotes a full chapter of historical and theoretical analysis to the theme of how executive orders played a key role in promoting civil rights; civil rights appear as a distinct topic in Ragsdale’s model but, with no explanation why, *not* in Mayer’s.⁸⁸

The point of these comparisons and questions—and many more could be made and posed—is that both of their topic models, as well as the National Archives’ model, deeply reflect the human classifier’s perspectives and assumptions. They are the product of the “top down” approach in which the text-corpus classifications are molded more by the researcher’s constructed model than by the text of the documents. This is not to say Mayer’s or Ragsdale’s topic models are not insightful or useful—they certainly are. This Article addresses the question of whether adding “ground up” computational text modeling to the tool kit can make them and similar topic models more insightful and useful.

2. *Environmental Topic Models*

We can drill down further on the features and limits of “top down” topic modeling by focusing on one topic—the environment—which appears as a distinct topic in Ragsdale’s model but not in Mayer’s. Environmental policy has received considerable attention in other studies of direct actions, perhaps because the environment played prominently in Theodore Roosevelt’s Administration, which was the first to use direct actions extensively.⁸⁹ Also, since the 1970s, environmental policy has been the subject of what Richard Lazarus characterizes in *The Making of Environmental Law* as a “pathological cycle” of back-and-forth policy perspectives in successive administrations,⁹⁰ a trend that has continued since his book’s publication in 2004.

Ironically, Lazarus’s book, one of the most comprehensive and insightful histories of environmental law published, barely mentions

⁸⁶ MAYER, *supra* note 7, at 79.

⁸⁷ *See id.* at 83–108.

⁸⁸ *See id.* at 182–217.

⁸⁹ Dodds argues that “[t]he nature and the use of unilateral presidential directives changed dramatically with Theodore Roosevelt” and devotes an entire chapter to that theme. *See* DODDS, *supra* note 2, at 27, 120–51.

⁹⁰ *See* RICHARD J. LAZARUS, *THE MAKING OF ENVIRONMENTAL LAW* 89 (2004).

direct actions. To be sure, the White House plays a prominent role in his history, but primarily in relation to congressional politics and oversight of agencies through means other than direct actions. Lazarus's references to direct actions are few and far between. He briefly mentions presidential use over time of the authority conferred under the Antiquities Act to designate national monuments from existing federal public lands.⁹¹

Beyond that, he discusses only a few direct actions in any detail—President Clinton's Executive Order on environmental justice and the series of executive orders, beginning with President Reagan's, on agency rulemaking review.⁹² President Nixon's environmental message to Congress, while not a direct action, receives some attention as well.⁹³ Although President Reagan's Executive Order on agency rulemaking review receives its own book index entry,⁹⁴ the index contains no entry for direct actions—not even executive orders—as a general category.⁹⁵ Additionally, the chapter notes are virtually devoid of references to direct actions.⁹⁶ Lazarus's final three chapters of the book “reflect on the present state of U.S. environmental law and speculate about its future,”⁹⁷ yet there is no mention of any direct action in those chapters, much less a discussion of the role direct actions could play. In short, Lazarus's history of environmental law essentially leaves direct actions out of the story.⁹⁸

Perhaps Lazarus is right to have mostly excluded direct actions—maybe they have not played a role in shaping environmental policy. Yet, several authors devote substantial attention to environmental policy direct actions, either as a case study of direct actions generally or as a component of a broader assessment of presidential influence on environmental policy. An example of the former is Ricardo Rodrigues, who uses environmental policy as one of his three topic case studies of direct actions in his 2007 book, *The Preeminence of Politics*. Although the book's subtitle is *Executive Orders from Eisenhower to Clinton*, Rodrigues starts the environmental policy case study with

⁹¹ See *id.* at 33.

⁹² *Id.* at 100–01, 139.

⁹³ *Id.* at 76.

⁹⁴ *Id.* at 305. The order also appears under the entry for President Reagan. *Id.* at 313.

⁹⁵ *Id.* at 295–318. President Clinton's order on environmental justice and use of the Antiquities Act to proclaim national monuments are mentioned under this index entry. *Id.* at 299.

⁹⁶ *Id.* at 255–94.

⁹⁷ *Id.* at 168.

⁹⁸ Nevertheless, as we explain below, Lazarus's coverage of the role of the President in shaping environmental law and policy is by far the most comprehensive, leading us to adopt his phases of presidential emphasis and influence for our study.

President Nixon. He works from there, one President at a time, through President Clinton, using a comprehensive and detailed historical assessment to build support for his thesis that “[t]he history of presidential use of executive orders to advance environmental policy is marked by a struggle for policy leadership between the executive and legislative branches of government,” with Presidents using direct actions in a play to take “turf” from Congress.⁹⁹ Rodrigues does not provide any form of a topic model, however; rather, his chronological history is aimed at demonstrating the influence of three factors in this power struggle—“divisiveness in Congress, public support for the issue, and the consistency of a president’s policy with the preferences of one’s supporting political coalition.”¹⁰⁰

An example of work using environmental policy direct actions as one of several mediums through which to study presidential environmental policy is Robert Shanley’s 1992 book, *Presidential Influence and Environmental Policy*.¹⁰¹ In one chapter from the book, Shanley’s “focus is upon a handful of executive orders in which presidents exercised a discretionary role,”¹⁰² but most of that discussion (like Lazarus’s) is devoted to President Reagan’s orders requiring White House review of agency regulations¹⁰³ and to agency assessment of the impact of their rules on property takings (Executive Order 12,630).¹⁰⁴ Beyond that, Shanley’s chapter on direct actions does not purport to provide any form of a topic model. Ironically, neither of those orders is, on its face, about environmental policy; rather, their significant impact on environmental rulemaking has led many scholars to follow Shanley’s lead and treat them as environmental-policy direct actions.¹⁰⁵ As Rodrigues puts it, “Although introduced as a program affecting all regulations, most accounts have related that President Reagan’s regulatory relief package targeted environmental regulations in particular.”¹⁰⁶ This does suggest that assigning a topic to a direct action based solely on its text can miss the reality of its impact in practice, a theme we return to below.

⁹⁹ RODRIGUES, *supra* note 6, at 269.

¹⁰⁰ *Id.* at 270.

¹⁰¹ ROBERT A. SHANLEY, *PRESIDENTIAL INFLUENCE AND ENVIRONMENTAL POLICY* (1992).

¹⁰² *Id.* at 49.

¹⁰³ Exec. Order No. 12,291, 3 C.F.R. § 127 (1982) (revoked 1993).

¹⁰⁴ Exec. Order No. 12,630, 3 C.F.R. § 554 (1988); see SHANLEY, *supra* note 101, at 61–84.

¹⁰⁵ See, e.g., SHANLEY, *supra* note 101, at 100–01; West & Sussman, *supra* note 7, at 87.

¹⁰⁶ RODRIGUES, *supra* note 6, at 225. Nevertheless, Rodrigues—we believe accurately—does not include these executive orders in his list of environmental policy orders. See *id.* at 275 tbl.13.1.

The only example of a broader empirical study of environmental-policy direct actions with an aim toward developing a topic model is a chapter that Jonathan West and Glen Sussman published in Dennis Soden's 1999 book, *The Environmental Presidency*.¹⁰⁷ Although they do not explain their methodology in any detail, they produce a table counting total executive orders and those related to environmental policy issued in each presidential term beginning with Franklin D. Roosevelt ("FDR") in 1933 through William Clinton in 1995.¹⁰⁸ According to their model, after FDR, over 11% of all executive orders issued through 1995—394 of 3,387—were aimed at topics within the scope of environmental policy, with the average per presidential term ranging from 7% to 15%.¹⁰⁹ Putting FDR back into the mix bumps the aggregate average to 22%—1,581 of 7,120.¹¹⁰

West and Sussman also present a table showing the breakdown by presidential term of environmental-policy executive orders classified across twelve policy-content topics.¹¹¹ Table 2 shows the percentage distribution of each of the ten topics in aggregate over their study period.

TABLE 2. WEST & SUSSMAN ENVIRONMENTAL TOPIC MODEL FOR FDR (1933)–CLINTON (1995)¹¹²

Topic	Number	%
Land Use	626	40.5
Animal/Plant	256	16.5
Parks/Forests	253	16.5
General	102	6.5
Water	81	5.0
Energy	70	4.5
Oil	65	4.0
Mineral/Coal	42	3.0
Radioactivity/Nuclear	34	2.0
Preservation	6	0.4
Air	5	0.3
Waste	4	0.3
TOTAL	1,544	100

¹⁰⁷ See West & Sussman, *supra* note 7, at 77–112.

¹⁰⁸ See *id.* at 80.

¹⁰⁹ See *id.* at 80 tbl.4.1. We have aggregated their data.

¹¹⁰ See *id.*

¹¹¹ See *id.* at 85 tbl.4.3.

¹¹² This is our computation of and ranking by percentages based on West and Sussman's Table 4.3. *Id.*

The picture West and Sussman paint of presidential use of executive orders in the environmental policy sphere is difficult to square with Lazarus's account of the making of environmental law. Whereas Lazarus discusses only a handful of environmental direct actions of any kind, West and Sussman claim that over 1,500 executive orders bearing on a broad swath of environmental policies have been issued, beginning with FDR's first Administration through the Clinton Presidency.¹¹³ How could only a few of over 1,500 environmental executive orders have registered in Lazarus's history? One possible explanation is that West and Sussman used a broad definition of "environmental" that sorted too many of FDR's executive orders into the category. They report that FDR issued 1,144 environmental executive orders in his first three terms, accounting for 31% of all his executive orders issued in that period and almost 75% of all the environmental orders they studied.¹¹⁴

This is likely to come as a surprise to modern environmental lawyers. Lazarus does not even mention FDR in his history of environmental law, nor does Rodrigues in his chapter on environmental direct actions, or Shanley in his chapter on environmental executive orders. For Rodrigues and Shanley, the reason why is simple—they both begin their studies with President Nixon, who presided in the White House during the flurry of new environmental statutes enacted in the early 1970s.¹¹⁵ Lazarus devotes a few pages to environmental law before Nixon, but with sparse references to Presidents, much less to any direct actions.¹¹⁶ Was FDR truly the Environmental President everyone else has overlooked?¹¹⁷

The answer is that it all depends on how the researcher designs the scope and time frame of the topic model. On scope, for any classification of topics in environmental law, the first question is what is environmental law? Are the West and Sussman categories of "land use," "parks/forests," and "animal/plant" best characterized as topics of environmental law, natural resources law, or land use law? And what about energy, coal, and oil—why aren't they *energy law*? Schol-

113 See *supra* Table 2.

114 This is our computation based on West & Sussman, *supra* note 7, at 80 tbl.4.1.

115 See RODRIGUES, *supra* note 6, at 181–82; SHANLEY, *supra* note 101, at 49.

116 See LAZARUS, *supra* note 90, at 50–53.

117 See generally Andrea K. Gerlak & Patrick J. McGovern, *The Twentieth Century: Progressivism, Prosperity, and Crisis*, in THE ENVIRONMENTAL PRESIDENCY, *supra* note 7, at 41 (providing extensive coverage of FDR's environmentalism).

ars have hotly debated those very divisions,¹¹⁸ and where one divides the line necessarily influences the topic model and sorting of documents into it. To illustrate the point, Rodrigues, who never defines what he means by “environmental” policy, counts nine environmental policy executive orders issued by President Reagan, listing each in a table,¹¹⁹ whereas West and Sussman count twenty-six and list none.¹²⁰

Of course, taking a broad view of a field is a reasonable approach, but West and Sussman also do not define what they mean by “environmental,” as if it is somehow intuitive or universally understood, and do not explain how they arrived at their twelve topics. The role FDR plays depends largely on this boundary line. Their explanation for FDR’s outsized presence is brief but to the point:

The three substantive areas with the largest number of executive orders are land use, parks and forests, and animal and plant life. These topics were especially popular during the presidency of Franklin D. Roosevelt, when considerable attention was devoted to modification of public land use; establishment of migratory bird, wildlife, and waterfowl refuges; and enlargement of national parks.¹²¹

Indeed, backing those three categories out of FDR’s first three terms, his number of environmental executive orders drops from 1,144 to 172—about fifty-seven per term.¹²² Going further, defining “environmental” even more narrowly by focusing on core topics such as air, water, and waste leads to even more dramatically different results—in FDR’s first three terms, he issued zero executive orders on air, zero on waste, and forty-five on water. Moreover, some of the topics in their model seem trivial over the entire time span. For example, three West and Sussman topics arguably sitting at the core of environmental law—preservation, air, and waste—account for a combined total of fifteen executive orders for the *entire* sixty-three-year study period. This low representation suggests these topics were not important direct-action themes for any President in their pool—more a measure of what West and Sussman thought Presidents could or should address

118 See Jody Freeman, *The Uncomfortable Convergence of Energy and Environmental Law*, 41 HARV. ENVTL. L. REV. 339, 342–43 (2017).

119 RODRIGUES, *supra* note 6, at 226. Reagan is the only President for whom Rodrigues reports such a count, and he makes no effort to sort the nine orders into a more granular classification.

120 West & Sussman, *supra* note 7, at 82.

121 *Id.*

122 This is our computation based on West and Sussman’s data. West & Sussman, *supra* note 7, at 80 tbl.4.1, 85 tbl.4.3.

rather than what Presidents in fact did address—in which case it is not clear they are useful as distinct topics in the model.

Time frame, given evolving national and global contexts, can also influence how to interpret the theme of any direct action. For example, although FDR's numbers are high even when scope is narrowed compared to, say, President Ford's twenty-three executive orders across all twelve categories,¹²³ some of the difference could be attributed to wartime, such as FDR's issuance of thirty-four executive orders dealing with energy and twenty-six dealing with oil. Are these best thought of as wartime policy orders or environmental policy orders?¹²⁴ The same could be said of FDR's conservation orders, many of which were issued during the Great Depression in connection with economic relief programs such as the Tennessee Valley Authority ("TVA") and Civilian Conservation Corps.¹²⁵

More to the point, if West and Sussman had started with President Truman,¹²⁶ their topic proportions would have looked substantially different, as shown in Table 3. Although land use remains the largest category, its percentage falls from 40.5% to 25.9%. The "general" category topic rises to second place, increasing from 6.6% to 17.3%, whereas the animal/plant category plummets from 16.6% to 3.8%. Including or excluding FDR thus produces a very different picture indeed. Hence, as with the meta-topic models of direct actions, West and Sussman's environmental executive orders study also suggests how "top down" topic modeling is prone to researcher idiosyncrasies.

¹²³ *Id.* at 80 tbl.4.1.

¹²⁴ See Gerlack & McGovern, *supra* note 117, at 45.

¹²⁵ See *id.* at 65–66.

¹²⁶ See Dennis L. Soden & Brent S. Steel, *Evaluating the Environmental Presidency*, in *THE ENVIRONMENTAL PRESIDENCY*, *supra* note 7, at 313, 337–39 (starting with Truman in their assessment of environmental Presidents).

TABLE 3. WEST & SUSSMAN ENVIRONMENTAL TOPIC MODEL, TRUMAN–CLINTON (1995)¹²⁷

Topic	Number	%
Land Use	102	25.9
General	68	17.3
Parks/Forests	43	10.9
Oil	39	10.0
Energy	35	8.9
Water	35	8.9
Radioactivity/Nuclear	32	8.1
Animal/Plant	15	3.8
Mineral/Coal	15	3.8
Air	4	1.0
Waste	3	0.8
Preservation	3	0.8
TOTAL	394	100

Yet, taking any of the above tables as the definitive story still leaves open the question of why direct actions play such a small role in Lazarus’s history of environmental law. One possibility is that most direct actions are not major in scope and effect. The fact of the matter is that many executive orders and other direct actions are minor if not trivial, particularly in the three categories West and Sussman identify as dominating the field.¹²⁸ Only a few mundane sentences are needed in an executive order to move a boundary line of a national park or wildlife refuge a few feet. As West and Sussman observe, “these devices typically deal with routine matters that generate little controversy rather than represent major policy thrusts.”¹²⁹ More definitively, Warber’s study of all executive orders from FDR through Clinton found almost 60% falling in his routine category, with another 3% being symbolic.¹³⁰ Of course, that leaves close to 40% in his policy category, and Shanley’s and Rodrigues’s histories of environmental direct actions from Presidents Nixon through Clinton make strong cases that direct actions have been a prominent means for Presidents to flex policy muscle in the field.¹³¹ The answer may simply come down to

¹²⁷ Table 3 shows our computations of and rankings by percentages based on West and Sussman’s work. See West & Sussman, *supra* note 7, at 85 tbl.4.3.

¹²⁸ See West & Sussman, *supra* note 7, at 82.

¹²⁹ *Id.* at 79.

¹³⁰ WARBER, *supra* note 11, at 39 tbl.2.1.

¹³¹ See RODRIGUES, *supra* note 6, at 269; SHANLEY, *supra* note 101, at 64–65.

researcher focus—Lazarus was more interested in the Legislature and agencies, whereas Shanley, Rodrigues, and West and Sussman were more interested in the President.

* * *

As our reviews of both the meta-topic direct action studies and the environmental-topic direct-action studies have revealed, direct actions present a rich body of material through which to evaluate the Presidency through time and themes, yet researchers vary widely in how they use and assess them for that purpose. All of the direct-action topic model studies discussed above share two traits, however—they impose the researcher’s “top down” historical and theoretical model, and they employ conventional research methods for selection, classification, and evaluation of direct actions. In the next part, we introduce the “ground up” method of computational text modeling and explain how we used it to design a meta-topic study and environmental topic study of direct actions.

II. COMPUTATIONAL TOPIC MODELS

Almost all law is expressed in natural language text; therefore, natural language processing (“NLP”) is a key component of automated methods for understanding law at scale.¹³² NLP uses machine learning techniques to convert unstructured text into a formal representation that computers can understand and analyze.¹³³

“Machine learning” refers to a subfield of computer science concerned with computer programs that are able to learn from experience and thus improve their performance over time. . . . [T]he idea that the computers are “learning” is largely a metaphor and does not imply that computer[] systems are artificially replicating the advanced cognitive systems thought to be involved in human learning. Rather, we can consider these algorithms to be learning in a *functional* sense: they are capable of changing their behavior to enhance their performance on some task through experience.¹³⁴

¹³² See John J. Nay, *Gov2Vec: Learning Distributed Representations of Institutions and Their Legal Text*, 2016 PROC. EMNLP WORKSHOP ON NAT. LANGUAGE PROCESSING & COMPUTATIONAL SOC. SCI. 49; John J. Nay, *Predicting and Understanding Law-Making with Word Vectors and an Ensemble Model*, PLOS ONE (May 10, 2017), <https://doi.org/10.1371/journal.pone.0176999> [<https://perma.cc/6EE6-PXQW>].

¹³³ See Harry Surden, *Machine Learning and Law*, 89 WASH. L. REV. 87, 96 (2014).

¹³⁴ *Id.* at 89 (footnotes omitted).

Machine learning has two primary groups of methods: unsupervised learning and supervised learning.¹³⁵ Supervised learning works by improving the predictive power of a model over time with respect to a specified outcome by adjusting parameters to make more accurate predictions.¹³⁶ This adjustment process necessarily involves human intervention, such as by instructing the program when it has made accurate or inaccurate predictions, to “train” the system.¹³⁷ This is a common approach to e-discovery in modern litigation.¹³⁸ By contrast, for unsupervised learning, observations only include their measured variables and no particular variable has the special status of the outcome variable to be predicted. The goal of supervised learning is to make accurate predictions for new observations and the goal of unsupervised learning is to provide useful compact representations of underlying data that can be used to summarize, cluster, and describe the data.¹³⁹ Topic modeling is a form of unsupervised learning that provides an overview of the various topics (themes) across a large number of documents, and how much each document is devoted to each topic.¹⁴⁰

Before going further, it is important to clarify exactly what is meant by a “topic” in the computational method. Say we were to gather 20,000 recipes from around the world. If we asked a human chef to sort them into a topic model, the chef might construct the model based on cuisines (Mexican, Ethiopian, Indian, etc.), or courses (appetizers, soups, entrees, etc.), or proteins (beef, chicken, soy, etc.), then would go about sorting the recipes. By contrast, in computational topic modeling the “topics” are statistical abstractions. The researcher does not specify the themes, but rather uses the program to extract them based on the algorithms’ search for semantic patterns within the corpus content. One could specify themes if using *supervised* machine learning, but that defeats the point of allowing the unsupervised learning to possibly unearth themes that would not have been evident to a human. Going back to the recipes example, it may very well be the semantic structure of the recipes varies based on cooking method (baking, braising, roasting, etc.) because the instructions for each method follow a pattern distinct from other methods, a feature that might escape the attention of a human classifier.

135 *See id.* at 90–95.

136 *See id.* at 90–92.

137 *See id.* at 90.

138 *See id.* at 112–13.

139 *See id.* at 113–15.

140 *See id.* at 113–14.

This distinction is not as important if one is using computational topic modeling to detect differences in *semantic structure*. For example, Livermore et al. used computational topic modeling to explore whether the writing style of the U.S. Supreme Court has over time become semantically distinct from the style of the lower federal courts.¹⁴¹ This did not require them to assign *substantive content* to the topics. By contrast, using computational topic modeling to construct a substantive-content model faces the challenge that patterns of semantic structure may not correspond to patterns of substantive content. To be sure, word and text patterns contribute to substance, and the computational method pays close attention to those, but that is no guarantee. For example, if a single cook wrote 300 of the recipes in our hypothetical recipe topic model exercise and used the same template form for each, those recipes could contribute to forming a distinct topic in the model notwithstanding that they range across the board with respect to cuisine and protein (although, that would be useful for someone interested in chefs). This is why human intervention ultimately is needed in such cases, to determine the viability of assigning coherent substantive content labels to the semantic structure topics. Hence, although the human classifier might miss deep semantic patterns that differentiate substantively among documents, the computational method might create semantically distinct clusters of documents that have no relevant substantive distinctions, which is why using both methods in some combination may be more powerful than either alone.

With the understanding that even the concept of what a topic is differs substantially between conventional and computational topic modeling, in the following Sections we provide the basics of computational topic modeling, including further description of how it differs from the conventional method, and then we explain our study design.

A. *Computational Topic Modeling Basics*

A computational topic model generates distributions of words for a collection of documents.¹⁴² The computation process is generative, in that it moves from documents to topics and back progressively. The first step is creating topics for an entire corpus of documents based on word distributions. Then the program identifies a topic distribution for each document by pairing each word in a document to a topic from

¹⁴¹ See generally Livermore et al., *supra* note 21.

¹⁴² David M. Blei et al., *Latent Dirichlet Allocation*, 3 J. MACHINE LEARNING 993, 1001 (2003).

the collection-wide distribution of topics. Based on these pairings, the program represents the dominant vocabulary terms for each topic to develop a distribution of the terms, known as the topic's relevant vocabulary.¹⁴³ This process allows documents to be comprised of multiple topics to varying degrees—that is, any document might “load” text into one or more of the collection-wide topics.

For a given number of topics the researcher specifies at the start, estimating the parameters of the model automatically uncovers the topics spanning the corpus, per-document topic distributions, and per-document-per-word topic assignments.¹⁴⁴ A correlated topic model shows how topical prevalence within documents exhibits correlation.¹⁴⁵ For example, a climate change topic can be more likely to co-occur in a given document with a high proportion of words from an energy topic than in a document with a high proportion of words from a financial regulation topic.

The computational topic model method has also been extended to incorporate metadata on time, location, and author.¹⁴⁶ The “structural topic model” flexibly extends the word-correlated topic model to allow topic prevalence to be modeled as a function of document-level variables, such as the year of the document's creation or its author.¹⁴⁷ This allows us to model the relationship between document characteristics and topic prevalences—that is, which document features correlate with which topics. The distribution over words (the content of the topics) is also adapted so that it is a combination of topics, covariates (the explanatory variables for correlations), and interactions between topics and covariates. In this way, both the *prevalence* and the *word content* of topics can be modeled as a function of document metadata, allowing the researcher to test hypotheses about the effects of time and author on topics.¹⁴⁸

To ground this technological explanation in the two different research methods, consider how researchers like Mayer, Ragsdale, and

¹⁴³ See *id.* at 998.

¹⁴⁴ See David M. Blei, *Probabilistic Topic Models*, 55 COMM. ACM, no. 4, Apr. 2012, at 77, 78.

¹⁴⁵ See David M. Blei & John D. Lafferty, *A Correlated Topic Model of Science*, 1 ANNALS APPLIED STAT., no. 1, June 2007, 17, 18.

¹⁴⁶ See David M. Blei & John D. Lafferty, *Dynamic Topic Models*, 23 INT'L CONF. ON MACHINE LEARNING 113, 113–20 (2006).

¹⁴⁷ See Margaret E. Roberts et al., *The Structural Topic Model and Applied Social Science* 1, 1–2 (2013), <https://scholar.harvard.edu/files/dtingley/files/stmnips2013.pdf> [<https://perma.cc/2Z8M-T5TA>].

¹⁴⁸ *Id.* at 3.

West and Sussman use conventional topic modeling methods. First, they must gather the relevant text corpus. For a large corpus, researchers must decide whether to random sample or work through the entirety. Either way, they must read all or a portion of each document to classify its topic. The topic model itself could be predetermined based on a researcher's theoretical construct, such as how one might expect Presidents to use direct actions, or a researcher could allow the topic model to develop organically as the documents are reviewed, or one could start with a model and tweak it along the way. In any of these approaches, each document must be coded for relevant information (date, type, etc.) and the topic determined. Once the documents are sorted into topics, researchers can begin to perform conventional statistical analyses, such as percentage distribution of the documents, numbers over time, and so on. Depending on how extensively a researcher coded the documents, more advanced empirical methods, such as linear regressions, could be performed to test various correlations.

In computational topic modeling, researchers also start with gathering the documents, but from there the process is quite different. The documents must be converted into a form the program requires, including specifying any fields such as the date or author, that will be modeled along with the text. Once the documents are in the appropriate form, the program begins by reducing the desired field of each document, usually the text field, to its collection of words (known as a "bag of words") or some other construct (e.g., numbers).¹⁴⁹ The methods we used are representative: the first step divides the document into its individual words; then the process converts all letters to lowercase, removes numbers, punctuation, and common words that would be found across topics and documents and therefore add little value in creating distinct topics (e.g., "the"); then it removes the endings of many words (e.g., consolidate, consolidated, and consolidating would all be converted to "consolid"); then, as a final preprocessing step, it converts each document to a "one-hot-encoded bag-of-words representation," which is a list of frequencies of terms.

¹⁴⁹ "Bag-of-words" is one approach (albeit the dominant one), but there are other approaches to topic modeling. See, e.g., Mark Andrews & Gabriella Vigliocco, *The Hidden Markov Topic Model: A Probabilistic Model of Semantic Representation*, 2 TOPICS COGNITIVE SCI. 101, 104 (2010) (describing the hidden Markov topic modeling approach); Thomas L. Griffiths et al., *Integrating Topics and Syntax*, 17 ADVANCES NEURAL INFO. PROCESSING SYS. PROC. 537, 537–38 (2005).

A researcher using the computational method then specifies how many topics to generate and, as described above, the program searches the documents, both as a corpus and each one individually, to estimate which words are likely to occur in which topics. As with conventional methods, there are important tradeoffs when specifying the number of topics. Going back to our 20,000 recipes example, the chef could choose a very coarse two-topic model, such as “serve hot” and “serve cold,” but this would not be a very useful product, as a cook looking for “serve hot” Indian cuisine dishes would likely have to sort through thousands of recipes. Or the chef could adopt a finely grained model, such as one comprising all of the regional cuisines of every nation in the world. Although this would improve the ability of cooks with a specialized cuisine in mind to find a recipe, there could be very few or no recipes in many of the cuisine topics.

Similarly, the larger the number of topics specified for a computational topic model, the less likely it is to find documents highly associated with any topic. This is because each document can contribute to more than one semantic pattern, and the program will generate its “loading” proportion for each such topic. Thus, one document might highly load into one topic, whereas another, perhaps because it addresses several themes or uses several different textual approaches, might load into many different topics, each at moderate proportions. As the number of topics specified grows, each document is more susceptible to being split into more and more topics, potentially making the topics themselves less coherent to a human observer. In the recipes example, for instance, specifying 1,000 topics could parse documents so finely that the chef could not translate the semantic topics into any meaningful substantive topics.

Once a researcher has settled on an appropriate topic number specification, the computational process delivers a topic model defining each topic according to its dominant words and ranking each topic according to its proportional content prevalence. This is an important distinction from the conventional one-document/one-topic method that Mayer, Ragsdale, and West and Sussman used. The percentages they assigned to their topics were of total *documents* in the text corpus, and they assigned each document to only one topic. The percentages that the computational method assigns to a topic are of total *content* in the text corpus, with each document loading percentages of its content into one or more topics. We identify where and how this distinction matters in Part III.

At this point, the human researcher must intervene to assign a substantive label to the topic. For a relatively large corpus, such as direct actions, the dominant words for a topic might or might not offer much of a clue as to the best way to label the topic. If the dominant words do not make the topic obvious, one way to gain a deeper understanding of the topic is to read the full version of the documents that loaded most highly into the topic, as these are the most representative of the distinct semantic patterns of the topic. For example, if the twenty highest-proportion documents of a particular topic in our recipe example are soup recipes, there is good reason to believe it is a soup-recipe topic. Another way is to identify documents the researcher is confident address known themes and examine the topics into which the documents loaded and in what proportions. If twenty recipes known to be French cuisine load their highest proportion of text into a topic under consideration for a French cuisine label, this strengthens the basis for the label.

Lastly, the computational method also can readily identify the degree of relatedness among topics based on their document overlap, which can help guide labeling. For example, if three topics believed to represent recipes featuring chicken based on the previous tests also demonstrate close relatedness, that strengthens the “chicken” label. (They may be separated into distinct topics for other reasons, such as one for soups, one for appetizers, and one for entrées.) The computational topic model also generates a variety of other metrics useful in label assignment and testing. If the researcher creates a field representing the date, the documents could be divided into time periods to generate the prevalence of the topic in each time period, which in turn could help label the topic based on the researcher’s understanding of historical trends in the text corpus—e.g., war powers orders ought to rise in prevalence during wartime.¹⁵⁰ As discussed below, we employed all these methods to assist in defining and testing the descriptive accuracy of our topic labels.

Several essential distinctions between the two methods are apparent even from these brief descriptions. The first has to do with how the topic model is constructed. The conventional method requires the researcher to develop the model either before, during, or after review-

¹⁵⁰ See Avinava Dubey et al., *A Nonparametric Mixture Model for Topic Modeling Over Time*, PROC. SIAM INT’L CONF. ON DATA MINING 530, 536 (2013), <https://epubs.siam.org/doi/abs/10.1137/1.9781611972832.59> [<https://perma.cc/UVP6-N6S8>] (modeling the time evolution of topics); see also Liangjie Hong et al., *A Time-Dependent Topic Model for Multiple Text Streams*, PROC. 17TH ACM SIGKDD CONF. 832, 837 (2011).

ing and coding the documents. The computational method relies entirely on the program to sort the documents at the front end, requiring the researcher to interpret the topics afterwards. The two methods thus “see” the corpus through different lenses. The researcher using conventional methods may have strong intentions or unconscious biases about what he or she is looking for, such as the different aspects of what environmental law comprises. The computational program has no intentions or biases at all at the front end—it is simply applying its algorithms to the text corpus—although the researcher’s subsequent assignment of labels to topics could be biased.

The second major distinction has to do with the documents. Unless the researcher using conventional methods engages in the laborious task of coding each document granularly to define multiple topics and assign weights to each, a one-document/one-topic approach is the default method, as Mayer, Ragsdale, and West and Sussman used. By contrast, the computational program dissects each document into multiple topics based on semantic content. The difference can have profound effects on the topic model. For example, consider an executive order or other direct action that instructs federal agencies on how to improve their energy efficiency, waste efficiency, water efficiency, and so on. The conventional one-document/one-topic method would assign the document to one predetermined (or new) topic—perhaps “federal agency management” or “efficiency”—whereas the computational method could assign it to multiple topics, which, after labeling, could correspond to energy, waste, water, and efficiency, and would identify the weight given to each. Replicating this effect of the one-document/one-topic default over the entire corpus of documents could lead to some topics that are clearly represented in the text not being represented in the topic model because they were not sufficiently dominant in a substantial number of documents to warrant classification as a distinct topic. Our environmental policy model results discussed below plainly revealed this difference in outcomes.

The third major difference has to do with the malleability of the corpus. Say a researcher using conventional methods decides mid-stream, based on having worked through the documents, that Topic A and Topic B should be combined into Topic C, or that Topic D should be spilt into two topics. Although this kind of tweaking will require revising the statistics, the merger of the two topics, or the splitting of one topic into two, is a relatively straightforward process—just combine the piles or spilt up one pile into two. By contrast, the computational method cannot easily perform this kind of reorganization of

topics. To illustrate, if a researcher specified a ten-topic model and then decided to change to a twenty-topic model, it would not necessarily follow that each topic in the twenty-topic model would consist of half of a topic in the ten-topic model. Nor would it necessarily be the case that the twenty-topic model would contain all of the topics in the ten-topic model plus ten new topics. Hence, as we did, the researcher must decide whether and how to consolidate or split topics to develop a final synthesis model. Further work in the corpus, however, must continue to use the original number and organization of topics the researcher originally specified. Using the recipes example again, if the chef specified twenty topics and later decided for purposes of compiling the recipe book to combine the recipes from two topics into one chapter, that might make sense for the book, but the computational model would need to continue its analytics using twenty topics, not nineteen. Respecifying the computational model to use nineteen topics would require generating a new model, which could affect how all of the documents load into the topics.

We are not suggesting that either method is necessarily better—each has its advantages and limitations. What we can say, however, is that one way—the computational method—is markedly faster. This difference goes well beyond producing the initial topic model once the documents are gathered. For example, consider a researcher using conventional methods who decides that the ten-topic model he or she has developed is not granular enough and probably fifteen or twenty topics would be better. This could require that researcher recode and re-sort *all* of the documents, likely demanding as much or more time as was needed to generate the original ten-topic model. Using the computational method, the program could generate a fifteen-topic model *and* twenty-topic model in a matter of hours—all the researcher would need to do from there is label the topics. Similarly, a researcher studying a corpus over time, or one consisting of different kinds of documents, could swiftly generate new data or models including or excluding different time periods and document types. Hence, to the extent that the computational method produces topic models that are useful and thus is worth pursuing, its speed of delivery and capacity for enormous text corpus sizes could prove immensely valuable to legal scholars. Of course, that is the critical question—does using the computational method produce a useful topic model? Testing that proposition was the point of our project, and we outline how we went about it in the next Section.

B. Study Methods and Design

To assemble our corpus of direct action documents, we downloaded all executive orders, presidential memoranda, presidential proclamations, and presidential determinations available on John T. Woolley and Gerhard Peters's *American Presidency Project*,¹⁵¹ which is one of several collections of presidential documents. Their site is a convenient source of digital direct actions, but is incomplete prior to President Truman,¹⁵² which as we explain below affected the representation of Presidents Hoover and Franklin D. Roosevelt when included in our models. Bearing in mind these differences in datasets, we designed our models to provide as close a comparison as we could to the Mayer, Ragsdale, and West and Sussman models, as well as to provide a broad model of presidential use of direct actions in environmental policy.

Once we assembled the documents, we curated them according to standard methods to produce “stemmed” words (e.g., “consolidate,” “consolidated,” and “consolidating” convert to “consolid”) that became the vocabulary set for the corpus,¹⁵³ and then for each vocabulary term that appeared in at least five documents, the program calculated the frequency of each term in each document. For each of our model constructions, our first step was to generate several models at different topic number specifications, so that we could identify the topic number that, in our judgment, best balanced coarseness and granularity of topic distinctions.

The program then produced the number of topics specified, ranked by overall corpus prevalence, and provided the dominant terms for each topic.¹⁵⁴ To generate labels for each topic, we interpreted its dominant words but also read the twenty documents with the highest content load for the topic. Although dominant words can often reveal themes effectively, direct actions cover such a variety of

¹⁵¹ Woolley & Peters, *supra* note 30.

¹⁵² See John T. Woolley & Gerhard Peters, *Executive Orders: J.Q. Adams–Trump*, AM. PRESIDENCY PROJECT, http://www.presidency.ucsb.edu/executive_orders.php [<https://perma.cc/F3NG-9W6E>] (“[O]ur collection of executive order texts is complete beginning with the administration of Harry S. Truman through the present.”). Assembling a complete set of executive orders, much less all direct actions, prior to President Truman has been a vexing problem for political scientists. See WARBER, *supra* note 11, at 135–38.

¹⁵³ See M.F. Porter, *An Algorithm for Suffix Stripping*, 14 PROGRAM, no. 3, July 1980, 130, 130–31, <https://www.emeraldinsight.com/doi/pdfplus/10.1108/eb046814> [<https://perma.cc/U9CS-H5BY>].

¹⁵⁴ For topic modeling, we used the open source R package *stm*. See Molly Roberts et al., *stm: An R Package for the Structural Topic Model*, GITHUB, <https://github.com/bstewart/stm> [<https://perma.cc/2ECH-8QEY>].

times, formats, and styles that we found reading the top-twenty documents was essential to assigning accurate and useful labels. We also performed document loading tests on documents we knew to be representative of certain concepts and topic relatedness tests to guide labeling, as described in more detail below.

As we reviewed the top-twenty documents, we also assigned a subjective measure of each topic's "coherence," from very high to low, based on how strongly and uniformly the documents supported our label decision. Even within the top-twenty documents, there were substantive outliers for many of the topics. As explained above, this is to be expected given that topics are constructed around semantic structures. We designed our coherence ratings of very high, high, medium, and low as a metric for that effect.

Occasionally a topic had such low coherence that it defied labeling. As reported below, however, the majority of topics demonstrated high or very high coherence throughout the top-twenty documents.¹⁵⁵ Indeed, in some cases the documents indicated a template-type format that has been used by Presidents for the intended purpose (e.g., to declare national monuments or set tariffs) for many decades. Also, in a few cases the topic, to our eyes, contained two or more coherent substantive themes, even though the program lumped them together. We treated these "hybrid" topics as containing two or more distinct topics and labeled them accordingly.

Lastly, we synthesized the computational model into the most compact set we could by combining closely associated topics, splitting the few that were "hybrids" and assigning proportionate shares of prevalence. We used these synthesis models as the final comparators to the conventional models.

We performed these steps to construct three separate models, two of which were designed to allow us to make apples-to-apples comparisons to the conventional models and one of which was designed to explore how different "top down" researcher assumptions can vastly affect the topic model results:

Meta-Topic Model. We designed this model to come as close as possible to an apples-to-apples comparison to Mayer's and Ragsdale's models. It includes all executive orders from 1936 through 1999. Starting with eight topics, we specified incrementally larger topic numbers until settling on a twenty-topic model, then we applied the steps outlined above.

¹⁵⁵ See *infra* Table 5.

Environmental Case Study—All Direct Actions Model. We designed this model to develop a broad sense of the modern American Presidency and the environment. We started the model with President Hoover, not because of his environmental policy distinction but because his term coincided with a turning point in American history and the Presidency—the Great Depression. Hoover also used direct actions to manage federal agencies more aggressively than his predecessors, and he was the first to run into stiff and persistent congressional pushback, thus marking a turning point in direct action history.¹⁵⁶ Unlike West and Sussman, we made the “top down” decision to include all four forms of direct actions in our dataset, rather than just executive orders, to determine whether that approach more fully captured the scope of presidential attention. For example, Presidents traditionally have used proclamations to make Antiquities Act designations,¹⁵⁷ and presidential memorandums have become as frequent and influential as executive orders;¹⁵⁸ thus, limiting the model to executive orders would skew the profile.

We did not manually select which direct actions were “environmental.” Rather, using the methods described above, we specified increasingly larger topic numbers until, at thirty-five topics, we obtained a model containing a high-coherence topic that we could confidently label “environmental and energy policy.” Various topics associated with public lands, including several distinctly associated with the Antiquities Act, had emerged at lower topic number specifications, but no topic at lower specifications robustly defined any other theme of environmental or energy policy based on our top-twenty documents review. As discussed in more detail in Part III, when we identified such a topic in the thirty-five-topic model, we externally confirmed the validity of labeling it “environmental and energy policy” by examining the loading distribution across all thirty-five topics of over a dozen well-known environmental and energy policy direct actions.¹⁵⁹ The environmental and energy policy topic scored the highest average loading for this cohort of documents.

We then extracted all the documents that had significant proportions (over one-tenth of their content) devoted to the topics that we

¹⁵⁶ See DODDS, *supra* note 2, at 179–83.

¹⁵⁷ Presidents choose this form because the statute states that they may “declare by public proclamation” areas of “land owned or controlled by the Federal Government to be national monuments.” 54 U.S.C. § 320301(a) (2012).

¹⁵⁸ COOPER, *supra* note 2, at 115–21.

¹⁵⁹ We drew these from the list Rodrigues compiled. See RODRIGUES, *supra* note 6, at 274–75 tbl.13.1.

labeled as having to do with public lands, the Antiquities Act, and environment and energy, combining them to create a new “environmental documents” subset, and then estimated a new topic model on those documents.¹⁶⁰ We selected a twenty-topic specification and then applied the steps outlined above to arrive at a final All Direct Actions Model for environmental policy.

Environmental Case Study—Executive Orders Model. We designed this model to provide a closer to apples-to-apples comparison to West and Sussman, given that our Environmental—All Direct Actions Model spans a broader timeframe than theirs and includes all direct-action types, whereas theirs included only executive orders. We used only the executive orders included in the environmental documents subset for FDR’s first term through President Clinton’s second term, and then applied the same 10% loading threshold.¹⁶¹ We specified a twenty-topic model and performed the steps outlined above.

III. FINDINGS AND ASSESSMENTS

A. *The Meta-Topic Model*

Our twenty-topic model of executive orders appearing in our dataset from 1936 through 1999, a span matching Mayer’s time period and encompassing Ragsdale’s, produced a synthesis model bearing close correspondence to Mayer’s model, and less so to Ragsdale’s. Table 4 shows the distribution of executive orders in the model by President. Our total is lower than Mayer’s and Ragsdale’s due to the American Presidency Project’s incomplete digital records prior to President Truman.¹⁶² With that caveat, our model is as close to an apples-to-apples comparison to Mayer’s and Ragsdale’s models as we could construct.

¹⁶⁰ The content proportion filter was designed to reduce “noise,” as it is possible for documents having nothing (to human eyes) to do with environmental or energy policy to have loaded into one or more of the topics at low levels based on the semantic patterns.

¹⁶¹ Our model thus is not a true apples-to-apples comparison in two respects. First, as noted previously, our dataset does not contain all of FDR’s executive orders. Also, we included executive orders through Clinton’s second term, whereas West and Sussman ended their study with 1995. See West & Sussman, *supra* note 7.

¹⁶² See Woolley & Peters, *supra* note 152.

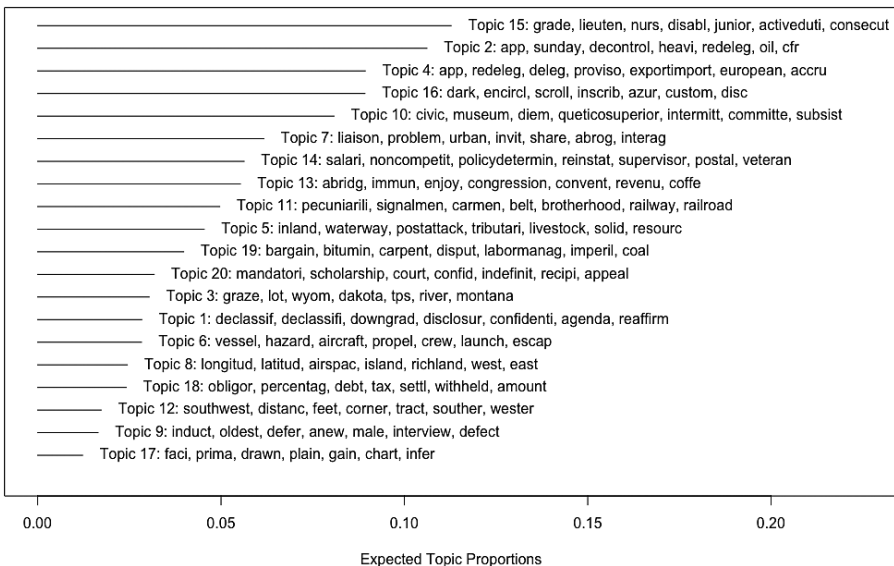
TABLE 4. META-TOPIC MODEL EXECUTIVE ORDERS DATASET—
NUMBER OF ORDERS BY PRESIDENT

President	# EOs
Franklin D. Roosevelt ¹⁶³	401
Harry S. Truman	876
Dwight D. Eisenhower	482
John F. Kennedy	214
Lyndon B. Johnson	325
Richard Nixon	346
Gerald R. Ford	169
Jimmy Carter	320
Ronald Reagan	381
George Bush	166
William J. Clinton	313
TOTAL	3,993

1. Findings

As described in Part II, the computational method’s first deliverable of interest to us is the distribution of topics showing dominant words and expected proportion for each, as shown in Figure 2.

FIGURE 2. META-TOPIC MODEL TOPICS—RANKED BY PREVALENCE
ACROSS THE CORPUS



163 As noted previously, the dataset is incomplete for FDR.

We assigned substantive content labels and coherence ratings to each topic after interpreting the dominant words and top-twenty documents for each. Table 5 shows the labels, ranked in proportion (expressed as a percentage), with our subjective judgment of topic coherence.

TABLE 5. META-TOPIC MODEL TOPIC LABELS—
RANKED BY PROPORTION

Topic #	Topic Label	%	Coherence
15	Military positions, succession, titles	11.3	High
2	International emergencies and sanctions	10.6	High
4	Delegation of statutory authority and functions	9.0	High
16	Foreign affairs/Public land withdrawals	8.9	High
10	Federal Advisory Commission Act commissions	8.0	Very high
7	Creation of positions, offices, councils, commissions, etc.	6.2	High
14	Civil Service	5.6	Very high
13	Declaring tax returns subject to inspection	5.5	Medium
11	Railroad labor disputes	5.0	Very high
5	Emergency planning and response	4.5	High
19	Other labor disputes	4.0	High
20	Military orders, appointments, awards, etc.	3.2	High
3	Public lands—acquisitions	3.0	High
1	Information security and intelligence	2.9	High
6	Military justice and compensation	2.8	High
8	Public lands—controlling access to land and airspace	2.5	Medium
18	No coherent topic	2.4	Low
12	Public lands—interagency and intergovernmental transfers	1.7	High
9	Military Selective Service	1.6	Very high
17	No coherent topic	1.2	Low

The model demonstrated strong correlation to very-high and high coherence substantive content topics. Only two topics (17 and 18) were so scattered in terms of substantive content of the top-twenty documents—possibly acting as “catch-alls”—to deserve what we considered a low coherence rating. We rated two others (8 and 13) as medium coherence, and the rest warranted a high or very high coherence rating.

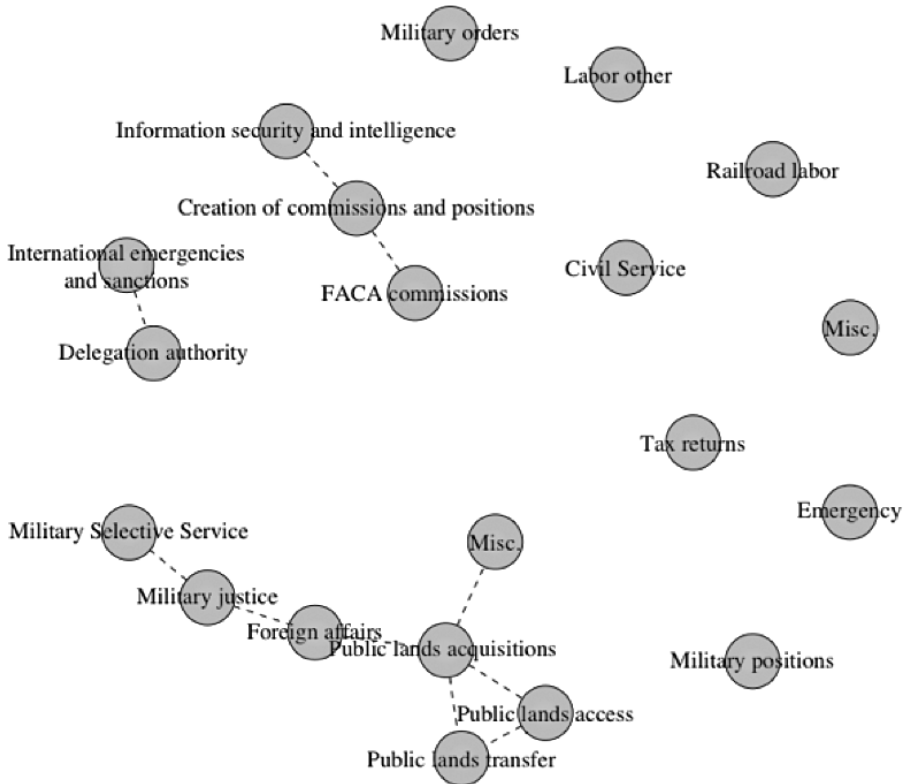
For some topics, the strong coherence likely was the product of Presidents for many decades using the same fill-in-the-blanks tem-

plate for routine actions, such as for national monument designations as discussed above. Nevertheless, even if the template format of the direct action contributed strongly to the formation of the topic, to the extent the template is associated with a particular substantive theme, the model captured the substantive attributes. Other topics, although not associated with a template, are no doubt influenced by the statutory authorities recited and the similarity of introductory words, such as the recitation of statutes under which the direct action asserted authority. In short, for any topic we rated as high or very high coherence, assigning a label required little thought—the substantive content was clear.

Another model output of interest is the network representation of topic prevalence correlation, shown for our Meta-Topic Model in Figure 3. As explained above, the computational method does not use a one-document/one-topic classification approach; rather, a document can load into several topics. This allows the model to account for topics that often appear together within documents, represented in Figure 3 as linked in a network.¹⁶⁴ In some cases, the strength of these *semantic topic* relationships makes *substantive policy* sense. For example, the direct actions in Topics 1 (information security and intelligence), 7 (creation of positions, offices, councils, commissions, etc.), and 10 (federal advisory commissions) are all very much about the President charging a group of agencies or appointees to go do something. Had we specified a lower number of topics, they may very well have been collapsed into one topic. On the other hand, in some cases the substantive connection is not as clear, such as for Topic 2 (international emergencies and sanctions) and Topic 4 (delegation of statutory authority and functions). Other topics sitting “all alone” in the network usually can be explained by their specialized substantive content associated with a distinct textual structure, such as Topic 11 covering railroad labor disputes.


¹⁶⁴ The absolute location of a topic in the network figure is of no relevance; rather, what is important is whether a topic is shown as linked to others.

FIGURE 3. META-TOPIC MODEL—TOPIC CORRELATION NETWORK



Another model output that could be useful in the study of direct actions is a representation of the proportion of a randomly selected document from a given timeframe that would be expected to belong to the topic. Whereas the conventional method reports the distribution of topics by average proportion of totals on a one-document/one-topic basis, this computation metric offers insight into how much attention the topic received from a particular President, or in a given time period, across all topics in all documents. In this example, we modeled the effect of the decade on the prevalence of each topic across the documents. This allowed us to then estimate the effect of a document being from a particular decade on the likelihood it would have a high or low proportion of a topic. For example, Figure 4 shows that, as a topic represented in all executive orders issued by decade, emergency planning and response (Topic 5) soaked up its highest degree of presidential attention relative to other topics in the 1930s (albeit with a higher error bar likely because there is less data in our model during that time) and then in the 1940s and 1970s (with narrower error bars).

FIGURE 4. META-TOPIC MODEL—PROPORTION OF EMERGENCY PLANNING & RESPONSE TOPIC BY DECADE



To generate a more compact synthesis model, we relegated Topics 17 and 18 to “other” for lack of coherence, and we split hybrid topic 16 into two topics—foreign affairs and public lands—and assigned each one half of the Topic 17 proportion score. We then combined similar topics using labels intended to approximate Mayer’s and Ragsdale’s labels as much as reasonably plausible. Table 6 provides the resulting topic model, showing the components from the twenty-topic model comprising each synthesis topic and its respective total prevalence.

TABLE 6. META-TOPIC MODEL—SYNTHESIS MODEL

Final Topic	%	Constituent Topics
Executive branch administration	23.2	4 + 7 + 10
Defense and military policy	18.9	6 + 9 + 15 + 20
Public lands	15.6	3 + 8 + 12 + part of 16
Foreign affairs	10.5	2 + part of 16
Labor policy	9.0	11 + 19
Information security	8.4	1 + 13
Civil Service	5.6	14
War and emergency powers	4.5	5
Other	3.6	17 + 18

Table 7 directly compares our computational synthesis model to Mayer's and Ragsdale's.

TABLE 7. COMPARISON OF META-TOPIC SYNTHESIS MODEL TO MAYER & RAGSDALE MODELS

MAYER		COMPUTATIONAL SYNTHESIS MODEL		RAGSDALE	
Topic	%	Topic	%	Topic	%
Executive branch administration	25.5	Executive branch administration	23.2	Personnel/agency requests	25.3
Civil service	19.6	Defense and military policy	18.9	Governance/economic management	15.3
Public lands	15.6	Public lands	15.6	Defense	15
Defense and military policy	11.9	Foreign affairs	10.5	Foreign trade and diplomacy	14.9
Foreign affairs	11.3	Labor policy	9.0	Natural resources/environmental	12.7
War and emergency powers	7.1	Information security	8.4	Social welfare/civil rights	9.2
Labor policy	5.4	Civil Service	5.6	Ceremonial/cultural	4.4
Domestic policy	3.8	War and emergency powers	4.5	Federalism	4.4
		Other	3.6	Agriculture	1.3
				Foreign aid	1.1

2. *Assessment*

All three models rank executive branch management and defense high in the list, with foreign affairs and public lands also receiving significant shares. Overall, however, our model bears a much closer cor-

respondence to Mayer's model than to Ragsdale's. Neither of their models includes the distinct information security topic our model produces, but otherwise the topics in Mayer's model and ours correspond closely and many rank in the same spots or close thereto. As a consequence of the close fit with Mayer's model, our model thus shares the differences Mayer's model has with Ragsdale's.

Does this mean Mayer's model is better than Ragsdale's? Or, given that we were able to produce a topic model very close to Mayer's with far less time and effort, does this mean the conventional methodology is obsolete and researchers should adopt computational topic modeling exclusively? We make no such claims. Rather, our results suggest computational topic modeling has great potential as a research tool for legal scholars. Had the method been available to Mayer and Ragsdale, they could have used it at the front end of their respective projects to inform how they constructed their topic models, or at the back end to validate their models, as we effectively did. They may have determined, for example, that a distinct information security topic was justified.

Moreover, the computational method can quickly provide deep insight into the corpus that could be replicated using conventional methods only through laborious and time-consuming efforts. The computational method's rapid generation of word dominance, topic correlation networks, over-time proportions, and other metrics provides the researcher a sandbox for exploration of the text corpus. Not to suggest that our study was effortless, but once we assembled the document database, we were able to "play around" with these metrics with relative ease, in a way that for all practical purposes would be unattainable using conventional methods. In short, based on our results we strongly advise in favor of using computation topic modeling methods for any project involving classification of a large legal-text corpus.

Nevertheless, we also would advise equally as strongly against turning over such research entirely to "the machine." For one thing, the computational method is incapable of making the semantic-to-substance translation required for labeling the topics. Recall, moreover, that two of our topics lacked substantive coherence to the point of defying a label.¹⁶⁵ Even for topics demonstrating high coherence based on our top-twenty documents review, as one moves further down the loading scale ranking in a topic set to documents contributing lower percentages of content to the topic, the substantive fit becomes less

¹⁶⁵ See *supra* Section III.A.1.

coherent. This is necessarily the case, as any one document could load content into multiple topics. By no means, therefore, would we suggest that *all* the documents that loaded to a topic coherently fit the label we assigned that topic. For that level of accuracy, we would trust the human over the machine. For constructing the topic *model*, however, we would trust the two of them working together more than we would either working alone.

Overall, the closeness of fit between Mayer's and our models suggests a validation story more than discordance. Mayer might have tweaked his model had he employed the computational method at the front end, but from there his study would have produced results similar to those he reached. By contrast, our environmental case study presented in the next Section more plainly reveals the highs and the lows of computational topic modeling.

B. Environmental Topic Models

As explained in Section II.B, our initial environmental topic model case study was an effort to capture a “big picture” assessment of environmental policy direct actions in the modern Presidency as well as to compare to the West and Sussman model. To be sure, there is plenty of room for debate over when the “modern Presidency” began. For practical reasons, we included FDR because Ragsdale and West and Sussman do. We added Hoover to bookend FDR with a national turning point, the Great Depression. In both cases our dataset was incomplete prior to President Truman as a consequence of the APP source-material gaps.

We developed an initial topic model using all direct actions in the time span from Hoover through June 30, 2017.¹⁶⁶ As explained in Section II.B, robustly coherent topics covering public lands and the Antiquities Act emerged at low-topic-number specifications—these are unmistakably in the environmental policy space if defined to include public lands policy—but it required moving to a thirty-five-topic model before we detected a distinct “environment and energy” topic (“E&E Topic”). The E&E Topic was an interesting hybrid demonstrating how computational text modeling develops topics by statistical abstractions rather than by expert substantive sorting. The topic's top-twenty documents contained direct actions that any environmental or energy lawyer would identify as falling in those domains, such as

¹⁶⁶ We do not discuss this model in full detail, as it was developed primarily to populate our environmental topics subset.

President Obama’s memorandum on greenhouse gas emissions¹⁶⁷ and President Trump’s executive order on energy policy that rescinded several of President Obama’s environmental and energy direct actions.¹⁶⁸ But the top-twenty documents also contained several executive orders from the series President Kennedy issued during the Cuban missile crisis of February 1962. Although these would likely not strike environmental or energy lawyers as falling in their domains, the orders make frequent reference to energy, minerals, and resources security, suggesting the computational algorithms detected common patterns that a human researcher would likely not associate as related, given the larger context of the document.

Given the hybrid content of the E&E Topic, we tested the coherence of the topic as the “destination” for environmental-and-energy-policy themed direct actions by examining the topic loading proportions of several direct actions widely acknowledged as addressing those domains.¹⁶⁹ The average loading score of the environmental direct actions was highest for the E&E Topic, with topics we labeled as “federal commissions” and “civil service” close behind.

Having identified a distinct environment and energy policy topic, we combined direct actions loading at least 10% of content to the topic with those from the public lands and Antiquities Act topics (also applying the loading filter) to form a new environmental topics subset of direct actions. Table 8 shows the distribution of these direct actions by President for the relevant timeframe. We used this subset to develop the two environmental case-study models.

¹⁶⁷ See Press Release, Office of the Press Sec’y, The White House, Presidential Memorandum—Power Sector Carbon Solution Standards (Jun. 25, 2013), <https://obamawhitehouse.archives.gov/the-press-office/2013/06/25/presidential-memorandum-power-sector-carbon-pollution-standards> [<https://perma.cc/LNN4-9VPU>].

¹⁶⁸ See Exec. Order No. 13,783, § 3, 82 Fed. Reg. 16,093, 16,094 (Mar. 28, 2017).

¹⁶⁹ We drew the documents for this test from the list Rodrigues compiled. See RODRIGUES, *supra* note 6, at 274–75 tbl.13.1.

TABLE 8. ENVIRONMENTAL—ALL DIRECT ACTIONS MODEL DATASET—NUMBER OF DIRECT ACTIONS BY PRESIDENT

President	Environmental Direct Actions
Herbert Hoover	14
Franklin D. Roosevelt	214
Harry S. Truman	164
Dwight D. Eisenhower	214
John F. Kennedy	53
Lyndon B. Johnson	69
Richard Nixon	47
Gerald R. Ford	25
Jimmy Carter	85
Ronald Reagan	231
George H. Bush	100
William J. Clinton	154
George W. Bush	156
Barack Obama	152
Donald J. Trump	16
TOTAL	1,561

1. All Direct Actions Model

a. Findings

Figure 5 shows the topic model results for our Environmental—All Direct Actions Model, and Table 9 shows our final topic labels, proportion, and coherence rating.

FIGURE 5. ENVIRONMENTAL—ALL DIRECT ACTIONS MODEL TOPICS—RANKED BY PROPORTION

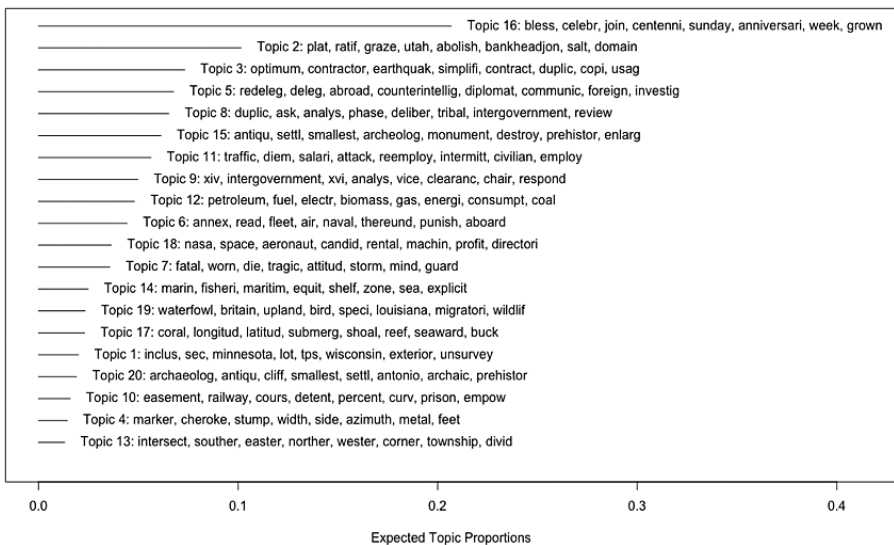


TABLE 9. ENVIRONMENTAL—ALL DIRECT ACTIONS MODEL—
TOPIC LABELS RANKED BY PROPORTION

Topic #	Topic Label	%	Coherence
16	Symbolic proclamations (national day, week, month, etc.)	20.6	Very high
2	Public lands—revoking prior withdrawals	10.2	Very high
3	Agency environmental management directives	7.3	High
5	National security (nonenvironmental)	6.7	Very high
8	Infrastructure & permitting/nonenvironmental budget	6.6	Medium
15	Antiquities Act proclamations (land)	6.3	Very high
11	Emergency preparedness—nonenvironmental	5.6	Very high
9	Emergency preparedness—environmental	4.9	Very high
12	Energy policy	4.8	Very high
6	Floods/resilience/naval reserves	4.5	Low
18	“Space” terms (office space, NASA, etc.)	3.7	Medium
7	Minor proclamations on energy, food, etc./boating week	3.6	Low
14	Marine environment	2.5	Very high
19	Migratory birds	2.4	Very high
17	Marine environment/Tsongas forest	2.3	Very high
1	Public lands—acquisition, withdrawal, transfer	2.0	Very high
20	Antiquities Act proclamations (land)	1.9	Very high
10	Public lands—minor orders on forests, naval reserves, etc.	1.6	Very high
4	Public lands—transfers (TVA, military lands, Hawaii)	1.5	Very high
13	Public lands—designations and expansions/ Antiquities Act land)	1.4	Very high

As with the Meta-Topic Model, most topics in this model produced highly coherent substantive themes based on our top-twenty documents assessment. Indeed, even more so than the Meta-Topic Model, topics in this model often were characterized by direct action “fill in the blanks” templates Presidents have recycled for decades, such as for declaring national boating week and other symbolic declarations (Topic 16), revoking prior withdrawals of public lands from access (Topic 2), and Antiquities Act designations (Topics 15 and 20). Every one of the top-twenty documents in Topic 15, for example, implements Antiquities Act authority using a scripted form including lines such as “Now, Therefore, I, [name], President of the United States of America, by the authority vested in me by Section 2 of the Act of June 8, 1906 (34 Stat. 225, 16 U.S.C. 431), do proclaim that

there are hereby set apart and reserved as the [name] National Monument”¹⁷⁰ Similarly, all of the top-twenty documents in Topic 1 included extensive survey coordinates in connection with public land acquisitions and transfers. No guesswork was required to label topics like these. And although text structure varied more for documents in other topics, the combination of the recitation of associated statutory authorities and word similarity and proximity in the body of the documents no doubt contributed to strong topic formation. The agency environmental-management direct actions in Topic 3, for example, frequently recited statute names distinctive of environmental law.

Several results in the model output, however, suggest that factors beyond form templates and statutory terms explain the topics. For example, the model rather sharply divided emergency-planning direct actions into those having something to do with the environment (Topic 9) and those not (Topic 11). Representative of Topic 9, for example, is President Obama’s executive order on preparing for the consequences of climate change,¹⁷¹ whereas none of the top-twenty documents in Topic 11 had any relation to an environmental theme.¹⁷² Similarly, although the documents in Topic 12 follow varied formats and invoke different authorities, all had a strong association with energy policy. To be sure, the model produced some humorous oddball topics, such as the collection of direct actions in Topic 18 having something to do with different meanings of the term “space,” including office space and outer space. Overall, however, the model produced a robust *substantive* classification of environmental policy direct actions.

The network representation of topic correlations, shown in Figure 6, reinforces the conclusion that the model has policy-substantive as well as thematic-semantic coherence. For example, the two substantive topics that stand out from the others in the model, energy policy (Topic 12) and agency environmental management (Topic 3), show no relation to other topics. Also, the lower coherence topics (Topics 7

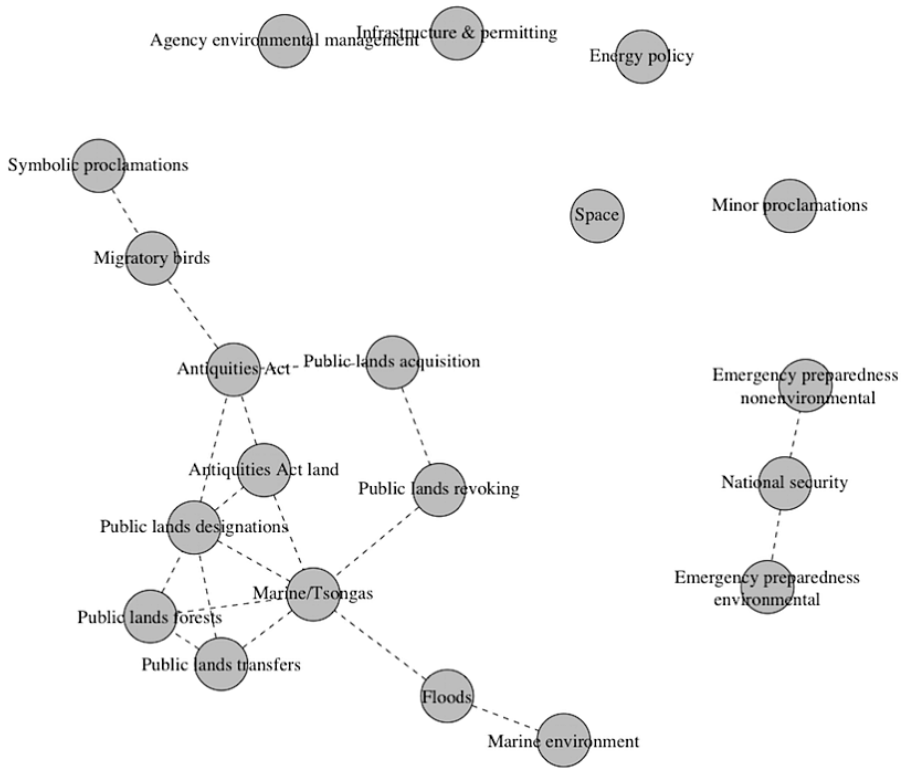
¹⁷⁰ See, e.g., Proclamation No. 6920, 61 Fed. Reg. 37,635 (Sept. 18, 1996) (establishing the Grand Staircase-Escalante National Monument).

¹⁷¹ Exec. Order No. 13,653, 78 Fed. Reg. 66,819 (Nov. 6, 2013).

¹⁷² It was not surprising that Topic 11, as well as Topic 5 (national security—nonenvironmental) and part of hybrid Topic 8 (nonenvironmental budget), showed no signs of environmental or energy policy; indeed, it was encouraging. As explained in Part II, it will often be the case that a document loads semantic content into a topic that, once labeled, does not bear a close substantive relationship to the document. When we selected the topics from the thirty-five-topic model to comprise our environmental subset case study, those kinds of nonenvironmental documents came along for the ride, so to speak. The second iteration model, however, sorted these three sets of nonenvironmental direct actions into distinct clusters, allowing us to remove them from the model when developing the final synthesis model.

and 8) and two fully or partially nonenvironmental topics (Topics 8 and 18) sit alone. By contrast, the four most tightly related topics, Topics 4, 10, 13, and 17, all have to do with public lands and resources, as do their closely linked Topics 1 and 2, covering public lands, and 15 and 20, covering the Antiquities Act.

FIGURE 6. ENVIRONMENTAL—ALL DIRECT ACTIONS MODEL—
TOPIC CORRELATION NETWORK




As an example of an “over-time” representation, Figure 7 shows the proportion of Topic 4, which includes agency environmental directives, for groupings of Presidents that follow Lazarus’s historical account. Lazarus does not discuss Presidents before President Kennedy in any detail. He describes Kennedy and Johnson as governing environmental law in its precursor form to the revolutionary 1970s,¹⁷³ and the Nixon-Ford-Carter trio as overseeing the ramp up and build out of modern statutory environmental law.¹⁷⁴ Presidents Reagan and Ford initiated a hard pushback, and from there, Lazarus argues, it has been

¹⁷³ See LAZARUS, *supra* note 90, at 52–53.

¹⁷⁴ See *id.* at 74.

a “pathological cycle” back and forth from President to President.¹⁷⁵ Although the proportion metric does not measure a substantive policy vector (any more than a statistical average would in conventional methods), our model suggests that Presidents have consistently used direct actions to steer agencies, presumably from one policy direction to the other if Lazarus’s account is accurate.¹⁷⁶

FIGURE 7. ENVIRONMENTAL—ALL DIRECT ACTIONS MODEL—
PROPORTION OF AGENCY ENVIRONMENTAL
MANAGEMENT DIRECTIVES TOPIC BY
PRESIDENT GROUPINGS



As we did for the Meta-Topic Model, we also developed a synthesis model for the All Direct Actions environmental topic model. We relegated Topics 5, 11, and 18, which accounted for a total of 19.3% of

¹⁷⁵ See *id.* at 89.

¹⁷⁶ President Trump’s results are likely distorted by the low number of direct actions he had issued at the time of our study, allowing a small number of environmental policy actions to swamp the proportions.

content, to “nonenvironmental” status and omitted them, adjusting final percentages for other topics accordingly. We split hybrid Topics 6, 8, and 17 and assigned each subtopic a proportionate share of prevalence. Table 10 shows our final model, with synthesized topics ranked by total adjusted proportion.

TABLE 10. ENVIRONMENTAL—ALL DIRECT ACTIONS MODEL—
SYNTHESIS MODEL

Final Environmental Topic	adj %	Topic Components from 20-topic Environmental Model
Symbolic	30.0	16 + 7
Public Lands (excluding Antiquities Act)	23.0	2 + 1 + 10 + 4 + part of 6 + part of 13 + part of 17
Antiquities Act (land)	10.9	15 + 20 + part of 13
Agency Environmental Management	9.0	3
Emergency Preparedness	7.8	9 + part of 6
Marine Environment	6.3	14 + part of 6 + part of 17
Energy Policy	5.9	12
Infrastructure & Permitting	4.1	part of 8
Migratory Birds	3.0	19

Notably, when combining similar topics, we nonetheless decided to keep the Antiquities Act (land) and the marine environment topics distinct from the broader umbrella of public lands, under which they reasonably could have been grouped. We held out the Antiquities Act (land) topic to illustrate the importance of including all four direct action types in any topic model study of direct actions—as noted previously, Antiquities Act designations have historically been made by proclamation, not by executive order. We held out the marine environment topic to highlight the potential for computational topic models to reveal alternative model structure possibilities. A comparison of our final synthesis model to the West and Sussman model, shown in Table 11, drives home these points.

TABLE 11. COMPARISON OF ENVIRONMENTAL—
ALL DIRECT ACTIONS SYNTHESIS MODEL
TO WEST & SUSSMAN MODEL

WEST & SUSSMAN 1933–1995		COMPUTATIONAL SYNTHESIS MODEL 1930–2017	
Topic	%	Topic	Adj %
Land Use	25.5	Symbolic	30.0
General	17.0	Public Lands (excluding Antiquities Act)	23.0
Parks/Forests	11.0	Antiquities Act (land)	10.9
Oil	10.0	Agency Environmental Management	9.0
Energy	9.0	Emergency Preparedness	7.8
Water	9.0	Marine Environment	6.3
Radioactivity/Nuclear	8.0	Energy Policy	5.9
Animal/Plant	4.0	Infrastructure & Permitting	4.1
Mineral/Coal	4.0	Migratory Birds	3.0
Air	0.1		
Waste	0.1		
Preservation	0.1		

b. Assessment

The two models have strong correspondence on public lands, showing that the primary emphasis of presidential use of direct actions in the environment and energy policy spheres, whether focusing just on executive orders or more broadly on all direct actions, is public lands. West and Sussman attribute a total of 36.5% to that theme, and our model, including marine environment in the cluster, puts it at nearly 40%. Several topics overlap (e.g., our migratory birds topic can fit into their animal/plant topic), suggesting substantive coherence. From there, however, the two models show some striking differences.

For example, although the specific “energy” topics have both substantive coherence and roughly the same prevalence, adding (as seems reasonable) their oil, nuclear/radioactive, and mineral/coal topics to the mix brings their energy topics prevalence to 31%, which is far above our result of 5.9%. More overtly, our distinct topics of agency environmental management, emergency preparedness, marine environment, and infrastructure and permitting do not find corollaries of any kind in the West and Sussman model, suggesting that computational topic modeling can provide insights into how to design the ultimate topic model. Conversely, we do not find distinct water or air topics, which is surprising given the importance of clean-air and clean-

water themes in federal environmental policy. This could represent a weakness in our computational modeling approach, or it could indicate a strength in avoiding mismatches between preconceived topic frameworks and the corpus, which give rise to sub-percent prevalences for topics such as air.

An advantage of multitopic sorting of one document in computational topic modeling versus the one-document/one-topic approach commonly taken in conventional research methods is its potential to reveal these different theme possibilities. For example, most documents in the marine environment topic had to do with public resources and many were Antiquities Act proclamations. The computational method could sort one such document to all three topics and thus reveal a major direct-action theme of marine environment. To be sure, public lands, Antiquities Act, and marine environment (because its direct actions usually involved federal waters, the Antiquities Act, or both) could be combined into one “federal public resources” topic, which could be exactly how a human researcher might have started the model. But it should mean something—at least worthy of a hypothesis—that the computational model kept the three topics separate. The Antiquities Act, after all, is a specialized and controversial statute—not a run of the mill public lands law.¹⁷⁷ And the marine environment presents distinct resource management issues compared to public lands.¹⁷⁸

Also, by using only executive orders, West and Sussman necessarily excluded identifying the Antiquities Act (land) and marine environment topics as distinct subsets of public lands, as well as the symbolic component of presidential proclamations concerning environmental themes. Using all four direct action types, our model reveals that the predominant role of direct actions in environmental policy is symbolic (26.7%), that Antiquities Act direct actions have been a major source of presidential influence on environmental policy (10.6%), and that the marine environment stands out as deserving distinct attention as a subset of the public lands/Antiquities Act authori-

¹⁷⁷ Major controversies exist over the size of some national monuments and whether a President can shrink or abolish existing monuments. See *Summary of CONGRESSIONAL RESEARCH SERV., R41330, NATIONAL MONUMENTS AND THE ANTIQUITIES ACT* (2016). Compare Seamon, *supra* note 5, at 51, with Squillace et al., *supra* note 5, at 65 (arguing that monuments cannot be abolished).

¹⁷⁸ See Robin Kundis Craig, *Treating Offshore Submerged Lands as Public Lands: A Historical Perspective*, 34 *PUB. LANDS & RES. L. REV.* 51, 52–53 (2013); Robin Kundis Craig, *Protecting International Marine Biodiversity: International Treaties and National Systems of Marine Protected Areas*, 20 *J. LAND USE & ENVTL. L.* 333, 359–60 (2005).

ties (6.4%). This is further support for Cooper's proposition that all types of direct actions—executive orders, presidential memoranda, proclamations, and determinations—are important to include for a comprehensive profile of topics of presidential direct-action attention.¹⁷⁹

Indeed, our model suggests several strong concerns flowing from these attributes of the West and Sussman model. First, as noted above, they include several topics that account for trivial numbers of executive orders within their sixty-three-year study period. Their preservation, air, and waste topics when combined account for just 1% of the executive orders, and adding in their mineral/coal and radioactivity/nuclear puts the combined percentage at 6%.¹⁸⁰ Devoting five topics of a twelve-topic model to 6% of the corpus provides little analytical leverage beyond noting that those topics have not received much presidential attention. Our model confirmed it is difficult to tease out these themes as distinct within the corpus of environmental policy direct actions, so why force the matter by using a preconceived environmental law casebook “table of contents” approach to constructing the topics? On the other hand, if we combine the four distinct energy-related topics in West and Sussman's model (oil, energy policy, radioactive/nuclear, and mineral/coal) into a single energy topic, the total prevalence is 31%, which is much greater than the 5.9% in our model.¹⁸¹ As we discuss below, this may reflect the different sets of documents (executive orders versus all direct actions) that the two models analyzed.

Second, using that “top down” approach led West and Sussman to miss distinct topics, such as not only the marine environment topic but also the agency environmental management, emergency preparedness, and infrastructure and permitting topics. This could have important implications for how we think of the President as a player in environmental policy. Granted, the marine environment topic is likely not one many legal scholars would include in a “top down” model at the front end, but that is the point of testing the preconceived “table of contents” approach with computational methods. If one examines the record, however, Presidents recently have used direct actions to shape marine environmental policy notwithstanding a relatively inert Congress on the topic.¹⁸² Legal scholars have begun to pay attention

179 See COOPER, *supra* note 2, at 2.

180 See *supra* Table 2.

181 See *supra* Table 11.

182 See, e.g., Exec. Order No. 13,547, 75 Fed. Reg. 43,023 (July 22, 2010) (Obama order

to this trend;¹⁸³ our model results suggest they are onto an important theme.

The agency environmental management, emergency preparedness, and infrastructure and permitting topics present a different concern. The West and Sussman model focuses exclusively on environmental media and resource types. Our agency environmental management, emergency preparedness, and infrastructure and permitting topics suggest a strong *functional* role for direct actions as well, as do Mayer's and Ragsdale's models, and our Meta-Topic Model. Some of the executive orders in our topic have no particular environmental media or resource types in mind; rather, they command practices for generic environmental performance of agencies, coordinate emergency response, or outline a vision and practice for putting infrastructure on the ground. For example, the top loading document in our agency environmental management topic, President Clinton's executive order on agency environmental performance, asserted that "[t]he head of each Federal agency is responsible for ensuring that all necessary actions are taken to integrate environmental accountability into agency day-to-day decisionmaking and long-term planning processes, across all agency missions, activities, and functions," and went on to cover a broad swath of environmental realms and practices.¹⁸⁴ These kinds of functional direct actions find no correspondence to any topic in the West and Sussman model.

We are not suggesting the West and Sussman model is wrong—it seems to sort executive orders into a coherent model—but rather that it is not the only way of constructing the model and likely is not the most useful way for many purposes. Using computational text modeling can assist legal scholars in breaking out of the “top down” approach and possibly construct more useful topic models for their research. Far more so than for our Meta-Topic Model exercise, our Environmental—All Direct Actions model suggests a vastly different array of topics compared with the West and Sussman model. Given the attention given in the literature to the presidential use of direct actions as part of the overall “imperial President” narrative, our functional topics strike us as at least as important, if not more important,

regarding oceans management policy); Exec. Order No. 13,158, 65 Fed. Reg. 34,909 (May 31, 2000) (Clinton order expanding the system of marine protected areas).

¹⁸³ See, e.g., Robin Kundis Craig, *Ocean Governance for the 21st Century: Making Marine Zoning Climate Change Adaptable*, 36 HARV. ENVTL. L. REV. 305, 307 (2012).

¹⁸⁴ Exec. Order No. 13,148, 65 Fed. Reg. 24,595 (Apr. 26, 2000).

to include in the topic model as the resource-specific topics that exclusively inform the West and Sussman model.

Stepping back from there to the bigger picture, our model also suggests, consistent with Lazarus's historical account, that West and Sussman may overstate the attention Presidents have given the environment through their direct actions. Consider that they characterize 22% of all executive orders issued in their time frame as "environmental," or 11% if FDR is excluded.¹⁸⁵ This does not square well with how the environment plays in our model. Recall that it was not until we specified a thirty-five-topic model that we identified a distinct "environmental and energy" topic other than those involving public lands and the Antiquities Act. This does not mean that environmental direct actions were not there; rather, a distinct topic capturing core environmental and energy policy themes as opposed to public lands did not emerge *across the entire corpus* until we reached a fairly granular scale of modeling. It is important in this respect to remember again that West and Sussman counted by document and used a one-document/one-topic method,¹⁸⁶ whereas our model identifies topics through prevalence within and across documents. Their percentages are of total *documents*, whereas our percentages are of total *content*. Using the semantic structure representation of the theme, our model indicates that outside of public lands, the environment has played a minor role in direct action *content* over the study time frame. As Figure 8 shows, our E&E Topic ranked twentieth out of thirty-five in topic prevalence of the initial thirty-five-topic All Direct Actions model from which we extracted the "environmental direct actions" subset. And this low priority has been the norm over time. As Figure 9 shows, the proportion of a direct action randomly selected from the full corpus expected to be devoted to the environmental and energy policy topic has consistently been below 2.5%. Only President Trump departs from that norm, likely due to our time frame capturing only the first six months of his term, during which he paid special attention to reversing President Obama's environmental and energy policy initiatives.¹⁸⁷

¹⁸⁵ West & Sussman, *supra* note 7, at 80 tbl.4.1.

¹⁸⁶ See *supra* Section II.A.

¹⁸⁷ See, e.g., Exec. Order No. 13,783, 82 Fed. Reg. 16,093 (Mar. 31, 2017) (rescinding several Obama climate and environment policy direct actions).

FIGURE 8. 35-TOPIC MODEL—TOP WORDS FOR EACH TOPIC¹⁸⁸

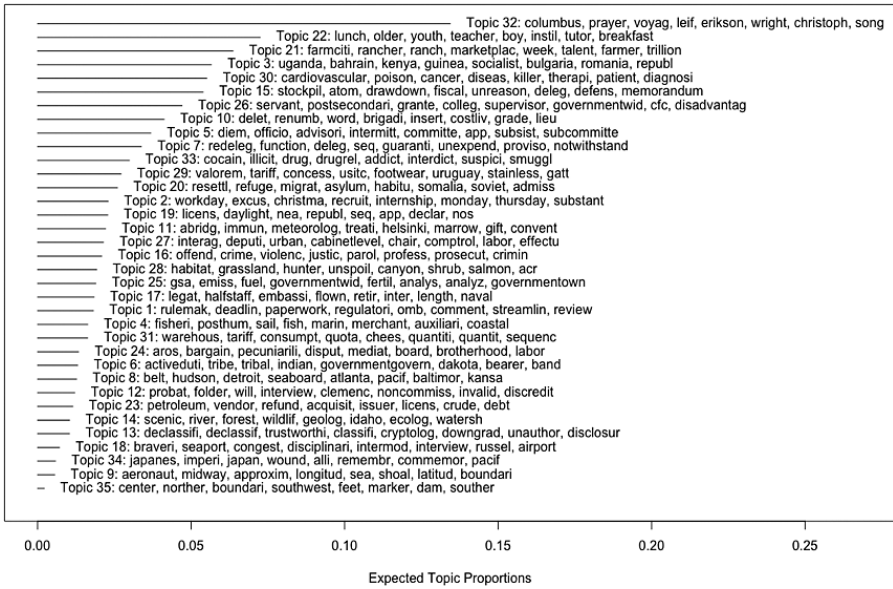



FIGURE 9. 35-TOPIC MODEL—PROPORTION OF ENVIRONMENT & ENERGY TOPIC BY PRESIDENT GROUPINGS



188 The size of the line and the vertical positioning ranks the topics by their prevalence across the corpus.

Several factors could account for the differences among the models in this respect. First, by including only executive orders, which Warber classified as predominantly routine or policy in scope, the West and Sussman model reduces the content influence of symbolic direct actions, which are primarily made through proclamations.¹⁸⁹ Indeed, the top three topics in our thirty-five-topic model—Topics 32, 22, 21 in Figure 8—we labeled as symbolic proclamations. Removing this content from the corpus would necessarily increase the percentage share of the E&E Topic (Topic 25 in Figure 8). Moreover, our environmental subset used for comparison to the West and Sussman model also included the topics from our thirty-five-topic model we labeled as associated with public lands and the Antiquities Act (Topics 4, 9, 14, 25, and 35 in Figure 8). Topic 25 thus is not the only source of direct actions in our environmental subset models. Even so, the combined share of those topics in our thirty-five-topic model was below 8%, far short of the 22% West and Sussman assign to the environment with FDR included, albeit not out of line for the 11% figure they reach (based on our computations) for President Truman forward. These differences suggest that one front-end decision—which direct action types to include in the analysis—can lead to substantially different results. Indeed, in the next Section we illustrate how influential that choice was for our purposes.

2. *Executive Orders Model*

In fairness to West and Sussman, comparing our Environment—All Direct Actions Model to the West and Sussman Model is arguably too much of an apples-to-oranges proposition given the differences in timeframe (theirs ends with Clinton in 1995; ours runs to June 2017) and direct action types (theirs includes only executive orders; ours includes four direct action types). As noted above, we developed a third model to move closer to the West and Sussman parameters, with the caveats that our dataset is incomplete for FDR and includes more of President Clinton's years in office than theirs did.

Table 12 shows the distribution of executive orders by President in our model compared to the West and Sussman distribution. Our two models appear to disagree markedly regarding FDR, but recall that the APP database is incomplete prior to President Truman. Indeed, if anything, our model supports West and Sussman's depiction of FDR as active in the environment and energy space. Our full

¹⁸⁹ See Rudalevige, *supra* note 2, at 146.

dataset includes 401 FDR executive orders (Table 4), and our environmental subset contains 214 FDR direct actions (Table 8), 201 of which were executive orders (Table 9). Because our dataset is incomplete for FDR and we do not know whether the APP database is systematically biased on content,¹⁹⁰ we cannot say that half of FDR’s total executive orders addressed environmental or energy policy, but the fact that our model produced that result for our partial set makes a statement about FDR’s environmental and energy policy focus.¹⁹¹

TABLE 12. ENVIRONMENTAL—EXECUTIVE ORDERS MODEL DATASET—NUMBER OF ORDERS BY PRESIDENT

President	Environmental EOs Our Model	Environmental EOs West & Sussman’s Model
Franklin D. Roosevelt	204	1,147
Harry S. Truman	114	119
Dwight D. Eisenhower	46	56
John F. Kennedy	39	19
Lyndon B. Johnson	34	28
Richard Nixon	31	46
Gerald R. Ford	10	21
Jimmy Carter	29	48
Ronald Reagan	46	26
George Bush	16	14
William J. Clinton	49	17
TOTAL	618	1,541

Discrepancies between the two models for other Presidents are relatively minor, not out of line with differences between Rodrigues and West and Sussman, for example.¹⁹² Our content loading threshold of 10% may have omitted some true environmental policy orders, and our computational topic model and the West and Sussman model may disagree over what constitutes environmental or energy policy. As explained above, computational topic modeling uses word occurrence patterns to identify statistical abstractions called “topics,” whereas conventional methods rely on the human researcher’s expert judgment to define topics and sort documents into them.¹⁹³ Our review of

¹⁹⁰ See *Comparing the Pace of President Trump’s Executive Orders & Memoranda to Other Recent Presidents*, AM. PRESIDENCY PROJECT, <http://www.presidency.ucsb.edu/index.php> [<https://perma.cc/7KKL-VMYJ>].

¹⁹¹ See *supra* Table 4, Table 8, Table 9.

¹⁹² See *supra* Section I.A.2.

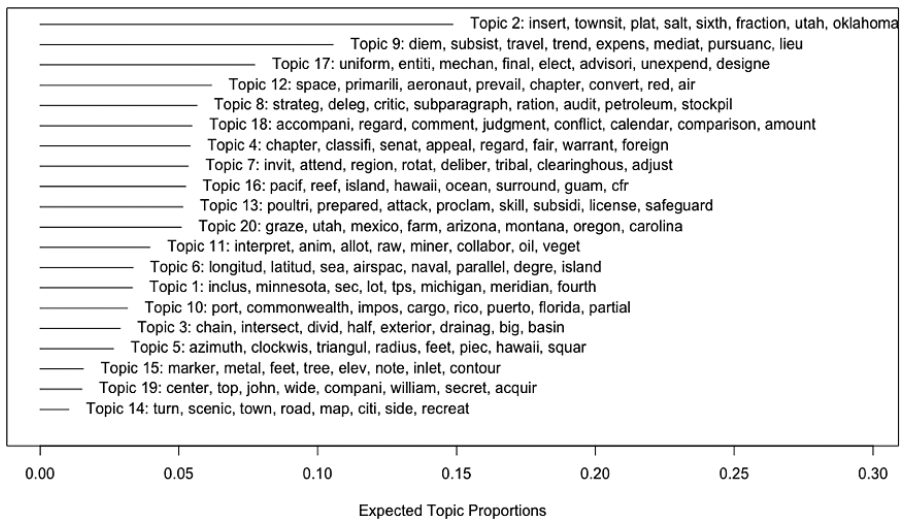
¹⁹³ See *supra* Part II.

dominant words and top-twenty documents supported what we believe could confidently be labeled as “environmental and energy” topics, including the public lands and Antiquities Act topics along with our distinct E&E Topic from the thirty-five-topic model. However, a human expert may not characterize *all* of the documents in those topics as “environmental” or “energy,” and our model might have missed some on the other side.¹⁹⁴ For example, as explained above, President Kennedy’s Cuban missile crisis executive orders are contained in our E&E Topic.¹⁹⁵

a. Findings

Taking those differences in dataset size and direct-action characterization into account, our most striking finding from this modeling exercise was the failure of the computational method to produce a model that would be useful to further understanding of how Presidents have used executive orders in the environmental and energy domains, likely due to the low number of documents. Figure 10 shows the topic model results, and Table 13 shows the topics with our final labels, the proportion scores, and our coherence ratings.

FIGURE 10. ENVIRONMENTAL—EXECUTIVE ORDERS MODEL
TOPICS—RANKED BY PROPORTION



¹⁹⁴ See *supra* Section III.B.1.b.

¹⁹⁵ See *supra* Section III.B.

TABLE 13. ENVIRONMENTAL—EXECUTIVE ORDERS MODEL—
TOPIC LABELS RANKED BY PROPORTION

Topic #	Topic Label	%	Coherence
2	Public lands—withdrawals	14.9	High
9	Infrastructure (some environmental)	10.6	Medium
17	Agency directives (some energy)	7.7	Medium
12	No coherent topic	6.2	Low
8	Agency directives (some energy)	5.6	Medium
18	Agency directives (substantial environmental)	5.6	High
4	National security/other (nonenvironmental)	5.4	Medium
7	Federal commissions, etc. (some environmental)	5.3	High
16	No coherent topic	5.2	Low
13	Emergency preparedness (nonenvironmental)	5.1	High
20	Public lands (some environmental)	5.1	High
11	Emergency preparedness/wildlife refuges	3.9	Medium
6	Public lands/Tongass/military airspace	3.3	Medium
1	Public lands—designations and transfers	3.3	High
10	Federal facilities (buildings, ports, etc.)	3.1	High
3	Public lands—forests and commissions	2.9	High
5	Public lands—Hawaii; other transfers	2.6	Medium
15	Public lands—designations and transfers	1.6	High
19	Public lands—designations and transfers	1.5	High
14	Public lands—designations and management	1.0	Medium

b. Assessment

Other than public lands topics, which predominate, as they also do in the West and Sussman model, the computational model produced only two distinct, robust subtopics of environmental or energy policy. Environmental and energy orders were scattered thinly throughout top-twenty documents for many of the topics, with the notable exceptions of a concentration of environmental and energy policy orders in Topics 7 (federal commissions) and 18 (agency directives). To test whether those two presidential function topics could justifiably be labeled “environmental and energy,” we examined the loading shares of several well-known environmental and energy policy executive orders in the topic model, and indeed Topics 7 and 18 had the two highest average loading scores. Although that could plausibly justify combining the two and labeling the synthesis topic something like “agency and commission directives,” that is as far as we could support adopting an environmental and energy policy topic from our results, and it is certainly not as granular as the topics West

and Sussman developed. Additional topics in our model would be difficult to label with anywhere near a distinct environmental or energy policy theme.

A compact synthesis model therefore could include three topics: public lands, agency and commission directives, and “other,” which would not be useful for drawing meaningful comparisons to West and Sussman or for guiding further research hypotheses. Public lands could plausibly be unpacked to anywhere from four to six topics rather than a single umbrella topic, given the different thrusts of the seven topics falling within the field, but is not clear that doing so would add any utility to the topic model other than for studying the nuances and typology of public lands executive orders. In short, the computational method, hamstrung by the low number of documents in the text corpus, did not produce a model that could provide the basis for meaningful comparison to the West and Sussman model.

In one important respect, however, our Environmental—Executive Orders model reinforces the major finding of the Environmental—All Direct Actions model—a significant focus of presidential direct action on environmental and energy policy is functional rather than resource specific. The two topics we determined were sufficiently representative of environmental and energy policy content were both functional in focus, as were several other topics we did not classify as environmental or energy but were unmistakably functional in focus (Topics 9, 8, 13, and 18). Overall, therefore, public lands and presidential functions drove the topic model far more than did environmental media types.

IV. RESEARCH USE CASES

Our project grew out of a mutual interest in the Presidency, environmental and energy policy,¹⁹⁶ and the use case for deploying artificial intelligence in legal and policy contexts. We bit off what we considered a manageable and accessible intersection of those three themes—presidential direct actions, and in particular, environmental and energy policy direct actions. The substantial number of direct actions makes any study of them a “large number” challenge for conventional research methods, yet the relatively small number of environmental and energy policy executive orders proved a challenge for the computational method.

¹⁹⁶ At the inception of this project, we were all affiliated with Vanderbilt University’s Institute for Energy and the Environment, a broadly interdisciplinary community of researchers focused on those themes. Co-author John Nay has since moved to his present positions.

We conclude that computational topic modeling demonstrates substantial value to legal research—in that it could move the needle in understanding of the content structure of a legal-text corpus—but also with the sober appreciation that the computational method would be of little value without the human researcher guiding its applications and interpreting its results. The two working together are better than either working alone. Advancing that core theme, in this Part we outline some broad use case applications for legal scholars.

A. *Front End: Generating Model Hypotheses*

One of the leading e-discovery providers, Ringtail, markets an early case assessment “concept clustering & searching” function that “clusters documents based on conceptual similarity. Fully searchable and interactive, these maps allow reviewers to shape, shift, filter and sift documents to reveal key facts and key fact patterns.”¹⁹⁷ In short, Ringtail delivers computational text modeling at the front end of litigation to assist the user in defining types and themes of documents in the discovery document corpus, which as any litigator knows can include millions of documents. It is up to the user to interpret and label the conceptual document clusters to determine their relevance to the litigation, but the conceptual clustering work is done at the front end in far less time than even a large team of lawyers could hope for.

This front-end application of computational text modeling has just as much potential value to legal scholars as it does to litigators. Computational topic modeling lets the documents speak for themselves. When approaching any study involving classification of a large corpus of legal text into substantive topics, using computational methods to assist in the design of the topic model can generate hypotheses about the scope and themes of the model, which later can be tested using conventional methods. Mayer and Ragsdale, for example, may have tested the viability of the information security topic our model produced as they reviewed executive orders. Even more so, West and Sussman may have reconsidered their resource-specific model had they been presented with the evidence our two environmental case study models produced of significant direct-action content addressing presidential functions, such as agency directives and emergency preparedness, with no single resource in mind.

¹⁹⁷ See *Visual Analytics + Your Expertise = Better Early Case Assessment, Investigations and Document Review*, RINGTAIL, <https://www.ringtail.com/ringtail-ediscovery-software/early-case-assessment> [<https://perma.cc/CNC6-5R98>].

The scale and speed at which computational topic modeling works also can allow researchers to ask questions they would not have when armed only with conventional methods. Using conventional methods to model the entire 15,000-plus document corpus of direct actions is a daunting proposition; using them to model the entire corpus of federal statutes, federal regulations, or public company SEC 10-K filings, would be a ludicrous proposition. With computational methods, legal scholars can set their targets that high.

B. Back End: Validating Models

Our application of computational topic modeling was not to generate hypotheses for moving forward, but to test the validity of existing topic models derived from conventional methods. To be blunt, judging by our results, Mayer's model looks solid; Ragsdale's could use some rethinking; and the West and Sussman model, while useful if one is interested only in which specific resource types executive orders have addressed, strikes us as missing a substantial part of the environmental and energy direct action story. Our Environmental—All Direct Actions model also calls into question the decision in all three cases to limit the topic model to executive orders. Consistent with Cooper's assessment,¹⁹⁸ our model points strongly in the direction of advising any legal scholar studying presidential direct actions to include at least presidential memoranda, proclamations, and directives in addition to executive orders.

Of course, back end model validation is not limited to direct actions. Many existing topic models of a large legal-text corpus, such as the C.F.R. or West's Topic and Key Number System, are the result of incremental growth of new topics and some path dependence. There may also be resistance to adding new topics to the model as the corpus grows, leading in some cases to poor fits between documents and topics. With computational topic modeling, it is not audacious to ask, does the C.F.R. title structure make sense for the C.F.R. content? Is there another way to structure the titles that would be more usefully arranged based on the actual content?

C. Meta-Scale: Building and Challenging Theory

Carrying that back-end validation theme further, consider the question we raised when first introducing the environmental direct actions topic in Part I—what is environmental law? This question has

¹⁹⁸ See COOPER, *supra* note 2, at 114, 172–74.

both practical and theoretical dimensions. As Todd Aagaard has observed, the answer usually identifies the canonical sources in the form of statutes, regulations, and cases practitioners and scholars typically considered to sit at the core of the field.¹⁹⁹ But not all environmental law is within the canon. As Aagaard suggests, at the outer reaches some is embedded in nonenvironmental programs:

Embedded environmental laws, a subspecies of noncanonical environmental law, are contained within a statute or program that is not primarily aimed at regulating environmental impacts and usually are administered by an agency that does not specialize in environmental issues. Essentially, embedded environmental laws are environmental laws organized with other, non-environmental laws. Embedded environmental laws thus lie within overlapping legal fields—both environmental law and whatever field they are embedded within.²⁰⁰

So, how does one find embedded environmental law? Is there really as much of it as Aagaard believes there could be? Is it really environmental law? The conventional method would answer these questions by having the researcher read the law—*all* the law—code the law for topics, and measure the weight given to environmental topics. Any volunteers?

A much faster way to test the theory of embedded environmental law would be to run the relevant legal text through a computational topic model. The U.S. Code is a large text corpus—very large—but is not too large for computational topic modeling. One could, for example, take it one title at a time and, much as we did, specify increasing numbers of topics to observe whether an environmental topic emerges. Because “the machine” has no conception that the Tax Code is supposed to be about tax law, it has no reason not to “see” embedded environmental law in the semantic structure. Perhaps it is not there—at least not in a way that produces a distinct topic in the computational model—but perhaps it is.

CONCLUSION

Gone are the days when teams of young law firm associates pored over piles of litigation discovery documents, sorting and searching for important documents and highlighting key passages. Today the docu-

¹⁹⁹ See Todd S. Aagaard, *Environmental Law Outside the Canon*, 89 IND. L.J. 1239, 1243–44 (2014).

²⁰⁰ *Id.* at 1264.

ments are loaded into computational topic model programs and plumbed with supervised machine learning algorithms, in a fraction of the time it took the lawyers of the recent past (though perhaps not a fraction of the cost). And this is just the beginning of the artificial intelligence revolution in legal practice.

There is good reason to believe that legal scholars can also employ these and other artificial intelligence applications in their work. As Livermore et al. put it with more technical panache, “The ability of topic models to quantitatively capture semantic features of very large corpora of legal documents has substantial potential to aid the work of empirical legal scholars in many domains.”²⁰¹ We could not agree more.

The question, though, is how best to leverage computational topic modeling and the other new tools of empirical legal studies. In their present capacities, they are by no means ready to replace human ingenuity. Rather, they enhance it, by leagues. Not only do they do more, faster, they also open up windows for legal scholars to find new insights that may never have been possible to see. Whether it is to start a topic model from scratch, test and refine one a researcher has developed from theory and experience, or rebuild conceptions of legal texts as large as the U.S. Code, legal scholars today have tools available which place previously unimaginable research undertakings within reach.

Consider our study of the Presidency through direct action documents.²⁰² We asked a question that other scholars examined through conventional methods: what policy realms have Presidents attempted to influence through direct actions?²⁰³ The work Mayer, Ragsdale, and West and Sussman put into their empirical analyses was impressive. Our computational case studies suggest Mayer’s topic model was spot on, whereas we derived a very different model of environmental-policy direct actions compared to West and Sussman. But in both the conventional and computational studies, a corpus of roughly 3,000 documents was at stake.²⁰⁴ Consider if there were *three million* direct action documents, which likely is not far off the number of federal judicial opinions on record.²⁰⁵ For Mayer and his fellow researchers of

²⁰¹ Livermore et al., *supra* note 21, at 863.

²⁰² See *supra* text accompanying notes 126–29; see also *supra* Part III.

²⁰³ See *supra* text accompanying notes 55–108.

²⁰⁴ See *supra* Table 1 and accompanying text.

²⁰⁵ There are over 750,000 opinions in a public database that the U.S. Government Printing Office maintains that covers most of the federal lower courts and dates back only to 2004. See *64 Federal Courts Now Publish Opinions on FDsys*, U.S. COURTS (Nov. 13, 2013), <http://www.us>

direct actions, that would have made their work excruciatingly difficult, if even possible. For our computational topic model, that would have added an extra day at most to run the program.²⁰⁶ From there, we would have done exactly what we did to label and test the topics, in the same amount of time.²⁰⁷ In short, bring on the legal text—no set is too large!

Legal scholars must proceed with caution, however, rather than with hype. As our study of presidential direct actions demonstrates, plenty of work remains for humans when using computational topic modeling, and legal scholars must understand the limits and idiosyncrasies of the technology so as to interpret what “the machine” hands us. On the other hand, the questions researchers in all disciplines ask are limited by what we know is possible to test empirically. Computational topic modeling, machine learning, and other artificial intelligence applications move that frontier exponentially outward in many directions. There may be questions no legal scholar thought to ask simply because there was no imaginable way to examine them empirically. Perhaps in that respect—in dreaming up new questions to ask—legal scholars should throw caution to the wind.

courts.gov/news/2013/11/13/64-federal-courts-now-publish-opinions-fdsys [https://perma.cc/MR5Z-VBM3].

²⁰⁶ See *supra* text accompanying note 18.

²⁰⁷ See *id.*