COMPUTATIONALLY ASSISTED REGULATORY PARTICIPATION

*Michael A. Livermore, Vladimir Eidelman & Brian Grom**

With the increased politicization of agency rulemaking and the reduced cost of participating in the notice-and-comment rulemaking process, administrative agencies have, in recent years, found themselves deluged in a flood of public comments. In this Article, we argue that this deluge presents both challenges and opportunities, and we explore how advances in natural language processing technologies can help agencies address the challenges and take advantage of the opportunities created by the recent growth of public participation in the regulatory process. We also examine how scholars of public bureaucracies can use this important new publicly available data to better understand how agencies interact with the public. To illustrate the value of these new tools, we carry out computational text analysis of nearly three million public comments that were received by administrative agencies over the course of the Obama administration. Our findings indicate that advances in natural language processing technology show great promise for both researchers and policymakers who are interested in understanding, and improving, regulatory decision-making.

INTRODUCTION

The public comment process is one of the hallmarks of the American administrative state.[1] As the informal notice-and-comment rulemaking procedure has grown into one of the most important national policy-making

* Michael Livermore is an Associate Professor at the University of Virginia School of Law. Vladimir Eidelman is the Vice President of Research at FiscalNote. Brian Grom is a senior principal scientist at FiscalNote. Thanks for comments from participants at a faculty workshop at the University of Virginia School of Law and the 2016 Political Economy and Public Law conference at Cornell Law School. Special thanks to Kevin Quinn and Jed Stiglitz. The views expressed here do not necessarily reflect those of FiscalNote.

1 As interpreted over the past several decades, the Administrative Procedure Act creates a right for all persons who are interested in a regulatory proposal to submit official comments to the issuing agency that must be duly accepted and considered. Administrative Procedure Act § 4, Pub. L. No. 79-404, 60 Stat. 237 (1946) (codified at 5 U.S.C. § 553); United States v. N.S. Food Prods. Corp., 568 F.2d 240, 252 (2d Cir. 1977) ("It is not in keeping with the rational process to leave vital questions, raised by comments which are of cogent materiality, completely unanswered.").

venues, the public comments process has become a forum for both organized interest groups and ordinary individuals to engage in public deliberation and political debate.[2] In recent years, as both the ease of participation and interest in rulemaking have grown, there has been an explosion of public participation, and agencies now receive millions of comments from the public each year concerning proposed agency actions.[3] These comments are voluntarily generated by individuals and organizations representing a vast diversity of interests—from large industrial trade associations representing businesses with billions of dollars at stake to individual citizens who have an interest in a particular regulatory outcome.

There is a substantial academic literature that studies, critiques, defends, and proposes reforms to the notice-and-comment process and informal rulemaking.[4] The advent of electronic rulemaking (commonly referred to as "e-rulemaking"), which involves the use of digital technologies by agencies to distribute information and collect comments, initiated a wave of scholarly commentary on this new phenomenon.[5] Opinions on the potential value of

2 For some defenders of the American administrative state, the open and deliberative character of the informal rulemaking process is vital to its legitimacy. *See, e.g.*, Mark Seidenfeld, *A Civic Republican Justification for the Bureaucratic State*, 105 HARV. L. REV. 1511 (1992).

3 *See infra* Section I.B.

4 *See, e.g.*, Cary Coglianese et al., *Transparency and Public Participation in the Federal Rulemaking Process: Recommendations for the New Administration*, 77 GEO. WASH. L. REV. 924 (2009) (summarizing recommendations of nonpartisan group interested in reforming rulemaking); Thomas O. McGarity, *Some Thoughts on "Deossifying" the Rulemaking Process*, 41 DUKE L.J. 1385 (1992); Wendy Wagner et al., *Rulemaking in the Shade: An Empirical Study of EPA's Air Toxic Emission Standards*, 63 ADMIN. L. REV. 99 (2011); Susan Webb Yackee, *Sweet-Talking the Fourth Branch: The Influence of Interest Group Comments on Federal Agency Rulemaking*, 16 J. PUB. ADMIN. RES. & THEORY 103 (2005). *See generally* William West, *Administrative Rulemaking: An Old and Emerging Literature*, 65 PUB. ADMIN. REV. 655 (2005) (providing review of rulemaking literature in social sciences).

5 *See, e.g.*, Cary Coglianese, *Citizen Participation in Rulemaking: Past, Present, and Future*, 55 DUKE L.J. 943 (2006) (arguing that digital technologies reduce costs of substitutes to participation and limited the total effect of e-rulemaking); John M. de Figueiredo, *E-Rulemaking: Bringing Data to Theory at the Federal Communications Commission*, 55 DUKE L.J. 969 (2006) (finding that e-rulemaking had not substantially changed the nature of the public commenting process); Cynthia R. Farina et al., *Rulemaking 2.0*, 65 U. MIAMI L. REV. 395 (2011) (discussing challenges in facilitating genuine participation through online rulemaking); Stephen M. Johnson, *The Internet Changes Everything: Revolutionizing Public Participation and Access to Government Information Through the Internet*, 50 ADMIN. L. REV. 277 (1998) (arguing that new technologies have transformative potential); Lauren Moxley, *E-Rulemaking and Democracy*, 68 ADMIN. L. REV. 661 (2016); Beth Simone Noveck, *The Electronic Revolution in Rulemaking*, 53 EMORY L.J. 433 (2004) (proposing an optimistic vision of the potential for e-rulemaking); David Schlosberg et al., *Democracy and E-Rulemaking: Web-Based Technologies, Participation, and the Potential for Deliberation*, 4 J. INFO. TECH. & POL. 37 (2007) (comparing electronic and nonelectronic comments); Stuart W. Shulman et al., *Electronic Rulemaking: A Public Participation Research Agenda for the Social Sciences*, 21 SOC. SCI. COMPUTER REV. 162 (2003) (calling for "social science reflection" on the development of e-rulemaking).

e-rulemaking vary from reasonably optimistic to highly skeptical, with some scholars remaining hopeful that new technologies can facilitate more inclusive and participatory decisionmaking, while others point to limited successes so far and question whether any increase in participation that has occurred has added anything of value.[6]

In this Article, we argue that at least some of the disappointment and pessimism surrounding e-rulemaking is premature and reflects inadequate attention to the potential to deploy digital technologies not only to solicit comments, but to analyze and understand them as well. Recent advances in machine learning and natural language processing have made powerful text analysis tools more broadly available. Both commercial enterprises and academic researchers have recently begun to put these tools to use in a variety of settings, from tracking employee morale based on email communications to testing the relationship between online blogging and political opinions.[7] Computational text analysis of public comments, however, is relatively rare, leaving largely untapped a substantial resource for both scholars and policymakers.[8] This Article explores how these new tools can be used by researchers, agencies, and oversight institutions to understand and improve agency decisionmaking.

From a positive, descriptive perspective, public comments are a valuable source of data that can be used to empirically examine how bureaucratic institutions interact with the public. As a form of political participation that is unique to the bureaucratic setting, commenting behavior is an interesting and important phenomenon in its own right and provides information on how agencies and their actions shape, and are shaped, by the publicly expressed views of individuals and groups. In recent years, a small number of political scientists and others interested in bureaucratic behavior have begun to take advantage of public comments to study agencies.[9]

In this Article, we contribute to this nascent field by conducting the first large-scale sentiment analysis of public comments to examine how word choices in nearly three million public comments are related to measures of agency ideology. Sentiment analysis has become a widespread tool used in a variety of settings to estimate how texts reflect the attitudes of their authors. Applying a basic, replicable procedure of sentiment analysis to public comments received for all nonminor rulemakings over the course of the Obama administration, we find that agencies with more moderate ideological lean-

6 *See* Stephen M. Johnson, *Beyond the Usual Suspects: ACUS, Rulemaking 2.0, and a Vision for Broader, More Informed, and More Transparent Rulemaking*, 65 ADMIN. L. REV. 77 (2013) (summarizing competing views).

7 *See* Diana Maynard & Adam Funk, *Automatic Detection of Political Opinions in Tweets*, *in* THE SEMANTIC WEB: ESWC 2011 WORKSHOPS 88 (Raúl García-Castro et al. eds., 2012) (estimating political leaning from pre-election tweets); Michael Schrage, *Sentiment Analysis Can Do More than Prevent Fraud and Turnover*, HARV. BUS. REV. (Jan. 5, 2016), https://hbr.org/2016/01/sentiment-analysis-can-do-more-than-prevent-fraud-and-turnover.

8 *See infra* Section I.C.

9 *See infra* notes 124–28 and accompanying text.

ings tend to receive comments that contain more positive language. This analysis indicates, as a threshold matter, that political characteristics of agencies are correlated with comment characteristics. Future work can build on this insight to inform subsequent research into the relationship between agencies' behavior and the public comments they receive.

Moving from the descriptive to the normative, we examine how agencies and agency oversight institutions can use computational text analysis of public comments to improve agency decisionmaking and accountability. In the era of mass commenting, agencies face both a "needle-in-the-haystack" problem (i.e., identifying the most substantive comments) and a "forest-for-the-trees" problem (i.e., extracting overall trends or themes in large, unstructured collections of documents). To examine the usefulness of text analysis techniques to address these challenges, we carry out a case study of the comments received by the Environmental Protection Agency (EPA) in response to its proposed rule to limit greenhouse gas emissions from the electricity-generating sector, the Clean Power Plan. We find that, although not perfect, existing techniques already have value for agencies and can be further refined to improve their current performance.

The Clean Power Plan case study illustrates how text analysis can help agencies, but there is an open question of whether agencies will put these tools to use. Doctrinally, existing interpretation of agencies' duties under the Administrative Procedure Act (APA) will likely provide sufficient incentive for agencies to take reasonable steps to address the haystack problem—agencies may soon deploy more sophisticated text analysis tools for this purpose. There is currently very little incentive, however, for agencies to address the forest problem. We will argue that this is a challenge that is worth taking up, and will discuss steps that could be taken by oversight institutions, including courts and the White House, to encourage agencies to do so.

The remainder of the Article will proceed as follows. Part I examines the practice of public commenting and how it has changed in recent years. We first address the various normative goals served by the public comment process, which include gathering technocratic and political information and serving expressive and procedural justice functions. We then turn to difficulties created by the recent participation explosion. For public comments to serve their social function, someone needs to read them. Reading a few comments presents no particular technical difficulty, but for many important rulemakings, agencies are presented with tens of thousands of comments—and sometimes upward of a million. These situations present a very severe resource and cognitive dilemma. These challenges are in part a problem created by technology, but just as technology has created new challenges, technological innovation can offer new solutions. When agencies receive a flood of comments, they are in essence facing a "big data" problem.[10] The

10 The term "big data" is often used in conflicting and confusing ways in both academic discourse and popular media. The general definition used in this Article is that the term "big data" applies to datasets for which "traditional data management and analysis technologies become inadequate" because of the sheer volume of data or some other char-

tools of data science—and in particular sophisticated techniques that have been developed in recent years to analyze large textual datasets—are meant to respond to exactly these big data challenges. Part I concludes with an introduction to the computational text analysis tools that can help translate the dilemmas posed by the era of mass public commenting into an opportunity to enhance understanding of, and participation in, the rulemaking process.

Part II focuses on the potential contribution of computational text analysis to better understand agencies and their relationships to the public. We first provide a brief overview of the substantial literature in law and political science that examines bureaucratic politics and agency decisionmaking. We then introduce the natural language processing technique referred to as *sentiment analysis*, which is a tool used to extract the attitudes of authors from the word choices within a text. A variety of new sentiment analysis techniques developed in recent years have found uses in both commercial settings and academic research. We apply a relatively straightforward form of sentiment analysis to nearly three million comments received by federal agencies during the Obama administration to test the relationship between agency ideology and comment sentiment, finding a significant relationship: agencies that occupy the center are more likely to receive comments with relatively more positive sentiment; agencies closer to either ideological pole tend to receive comments with more negative sentiment. This analysis is a useful contribution to the existing literature on bureaucratic policies and illustrates the potential of computational text analysis of public comments to inform scholarship in this field.

Part III discusses the value of computational text analysis in improving public participation in the regulatory process. We begin by characterizing two challenges that agencies face when responding to a very large set of public comments: the *haystack problem* and the *forest problem*. The haystack problem occurs when comments of high substantive value are hidden within a very large set of documents of lower substantive value, creating the risk that agencies will fail to locate and appropriately consider high-value comments. The forest problem occurs when agencies are unable to identify aggregate patterns within a large collection of comments because they treat each comment in isolation rather than drawing connections between them. We then test whether current text analysis techniques can be of use, using EPA's controversial and high-comment-volume climate change rule as a case study. We first examine how a basic text analysis tool can separate more substantive comments from less substantive comments and discuss how agencies, given their greater internal knowledge, could improve on our basic approach. We then deploy a text analysis tool referred to as a "topic model" to examine the semantic content of comments, extracting general, aggregate trends in the themes contained within the public comments on the rule. For each tool, we

---

acteristic, such as complexity or lack of structure. Amir Gandomi & Murtaza Haider, *Beyond the Hype: Big Data Concepts, Methods, and Analytics*, 35 INT'L J. INFO. MGMT. 137, 139 (2015); *see also infra* Section I.C.

examine its current efficacy at addressing the relevant agency challenge and discuss potential improvements and refinements to better capture their promise.

Part IV concludes with a discussion of how oversight institutions—specifically the courts and the Office of Information and Regulatory Affairs (OIRA) in the White House—can facilitate the broader adoption of computational text analysis tools by agencies. Courts exercising review under the APA already require agencies to identify and respond to highly substantive comments. Similar incentives could be created so that agencies use reasonably available technologies to extract, publish, and consider general trends within comments. OIRA can also play a role by encouraging agencies to process public comments to facilitate intra–executive branch coordination and deliberation on rulemaking proposals and, potentially, to build public support for contested policy choices.

I. THE ERA OF MASS COMMENTS

During the 1970s, there was an explosion of growth in notice-and-comment rulemaking.[11] In the pre– and post–New Deal periods, agency decisionmaking typically took the form of adjudications of individual decisions or, when more general rules were set, formal, trial-like proceedings. With the substantial expansion of federal regulatory authority over an increasingly large set of issues—including environmental protection and workplace safety—the "informal" notice-and-comment rulemaking created by the APA became correspondingly more important. As a consequence, public comments came to take a more prominent place in many administrative proceedings, a trend that has been amplified with recent technological and political developments.[12] This Part examines the justification for a robust system of public participation in agency rulemakings and discusses the challenges and opportunities created by recent trends toward very high volume commenting on controversial rulemakings.

A. The Value of Public Comments

Agencies spend a considerable amount of effort soliciting, collecting, and responding to public comments that are received during the regulatory process. In part, agencies engage in these efforts because they are legally required to do so. The APA sets out specific procedural requirements that agencies must follow prior to issuing a rulemaking. Among other requirements, agencies must issue a notice of proposed rulemaking and then "give interested persons an opportunity to participate in the rule making through

11 *See* Christopher DeMuth, *Can the Administrative State Be Tamed?*, 8 J. LEGAL ANALYSIS 121, 126–27 (2016).

12 Informal rulemaking has become a primary mechanism for the executive branch to take on a greater explicit policy-making function. *See* Elena Kagan, *Presidential Administration*, 114 HARV. L. REV. 2245 (2001) (discussing the importance of rulemaking in promoting presidential preferences through executive action).

submission of written data, views, or arguments."[13] After providing that opportunity for comment, agencies may only issue a final rule "[a]fter consideration of the relevant matter presented."[14] This procedural notice-and-comment requirement has been given teeth (and arguably claws) through the hard look review that courts undertake under the APA's "arbitrary" or "capricious" standard.[15] Agencies anticipate that major rulemakings will be litigated, and if they fail to acknowledge and respond to substantive concerns that are raised during the rulemaking process, they can expect a difficult time in court.[16]

There are several potential normative justifications for providing the public with a meaningful opportunity to comment on proposed regulations.[17] Some of the benefits of the notice-and-comment process flow from the value of information for an agency's decision-making process. When public comments bring unforeseen consequences of a rule to light, or call attention to new scientific and technical information, agency personnel can use this information to improve regulatory outcomes. This class of benefits will be referred to as the *output value* of the notice-and-comment process: the benefit of collecting information that results in improved regulatory design.

From the agency's perspective, some of the information that is generated by the public comment process will be useful because it bears on how the rule will affect social welfare. This is the kind of information that would be put to use by a disinterested and public-spirited social decisionmaker to make improvements to a regulation.[18] A scientific study on the relationship between exposure to a toxic substance and a health end point (e.g., cancer risk) is an example of this type of technocratic information that expert-driven agencies will seek out. Although agency personnel are frequently knowledgeable in their substantive areas of expertise, and agencies can draw from the stock of published scholarly literature, no one is omniscient. Especially for complex regulations, there are good reasons to believe that stakeholders

13 5 U.S.C. § 553(c) (2012).

14 *Id.*

15 5 U.S.C. § 706(2)(A).

16 *See, e.g.*, Bus. Roundtable v. SEC, 647 F.3d 1144, 1152 (D.C. Cir. 2011) (discussing agency failure to address cost issues raised in comments).

17 *See* United States v. Reynolds, 710 F.3d 498, 517 (3d Cir. 2013) ("[A]mong the purposes of the APA's notice and comment requirements are (1) to ensure that agency regulations are tested via exposure to diverse public comment, (2) to ensure fairness to affected parties, and (3) to give affected parties an opportunity to develop evidence in the record to support their objections to the rule and thereby enhance the quality of judicial review." (quoting Prometheus Radio Project v. FCC, 652 F.3d 431, 449 (3d Cir. 2011)) (internal quotation marks omitted)).

18 Agencies often state that changes in regulatory text are due to comments. *See, e.g.*, Federal Implementation Plans: Interstate Transport of Fine Particulate Matter and Ozone and Correction of SIP Approvals, 76 Fed. Reg. 48,208, 48,216 (Aug. 8, 2011) (to be codified at 40 C.F.R. pts. 51, 52, 72, 78, & 97) (noting several changes that were made "[a]s a result of updated analyses and in response to public comments"; changes included revision to state budgets, updated modeling and analysis tools, and a revised allocation methodology for emitters).

outside of agencies will have relevant information. The regulated community will likely know better than the agency how it would go about complying with new requirements; experts employed in academia, the private sector, and civil society organizations may be able to add to the pool of technical know-how that is brought to bear on a regulatory problem. In addition to information of a technical nature, agencies are often called on to make decisions that implicate important questions of values and political norms.[19] Public comments can also generate information on public responses that provide insight into the value-laden decision that agencies make.

Agencies might also be interested in (and respond to) information that is political in nature.[20] Agencies and agency personnel seek to maximize their goals within a system of political constraints and incentives and may therefore attend to the reactions of powerful interest groups to regulatory decisions, or how blocs of voters respond to how an issue is framed. Career civil servants must deal with internal bureaucratic politics and ultimately report to the political appointees who serve in the main managerial positions at most agencies. These senior managers have their own partisan and institutional loyalties and have attained positions of trust within an administration by responding adroitly to political rewards and punishments. The White House oversees the actions of administrative agencies, most importantly through the process of regulatory review by OIRA but also through other White House offices such as the Domestic Policy Council and the Council of Economic Advisers.[21] Agencies must obtain their budgets through the appropriations processes, negotiated by Congress and the President.[22] Congress has other oversight tools including hearings and information requests, the appointments process, and taking advantage of the mass media to criticize an agency.[23] Success in this environment requires close attention to the political context and consequences of regulatory actions, and a portion of the output value of public comments lies in helping agencies craft politically savvy (not only technocratically proficient) rules.[24]

A second benefit arises from the value that people place on their ability to weigh in on potential regulations. This benefit will be referred to as the *input value* of the notice-and-comment process. This input value is associated with ideas such as due process, procedural justice, and participatory democracy. Government agencies make important decisions that have profound

19 *See* Nina A. Mendelson, *Foreword: Rulemaking, Democracy, and Torrents of E-Mail*, 79 GEO. WASH. L. REV. 1343, 1350–51 (2011).

20 *Cf.* Michael A. Livermore, *The Perils of Experimentation*, 126 YALE L.J. 636 (2017) (discussing types of information that can be used by policymakers in the context of experimentation in decentralized policy regimes); Moxley, *supra* note 5.

21 *See generally* Kagan, *supra* note 12 (discussing the role of White House offices).

22 *See id.* at 2259 n.38.

23 *See* Michael A. Livermore, *Political Parties and Presidential Oversight*, 67 ALA. L. REV. 45, 97–100 (2015) (discussing Congress's oversight role).

24 *Cf.* Mathew D. McCubbins et al., *Structure and Process, Politics and Policy: Administrative Arrangements and the Political Control of Agencies*, 75 VA. L. REV. 431 (1989) (arguing that process is used by Congress to promote the preferences of the operative majority).

effects on people's lives, from the cancer risks from pesticides in food to the cost of borrowing for a first home. When the government makes individual decisions that bear on life, liberty, or property, courts have long held that affected parties have participatory rights that at least include the opportunity for a public hearing at which they can present oral testimony.[25] When these decisions are made for a group, the consequences are no less important. Although the form of the procedural right may differ,[26] democratic norms require some mechanism for broad participation in decisionmaking, if only through the process of electoral accountability.[27] Given the attenuated connection between regulatory decisions and the political branches, an opportunity for comment and other procedural rights helps lend some degree of democratic legitimacy to agency decisionmaking.

Social scientists and organizational psychologists have studied the potential for procedure to enhance the perceived legitimacy of decision-making processes.[28] This research, which has been undertaken in contexts as varied as criminal courts and corporate cubicles, examines how decision-making procedures, and in particular the mechanisms through which stakeholders are able to participate in and voice their perspectives, affects the acceptability of the ultimate decisions.[29] There are many controversies and subquestions within this literature, but the general thrust of the research is that people are generally more inclined to accept a decision when they feel that the decisionmaker has listened, considered, and responded to their concerns.[30] In this way, the public comment process may help increase the perceived legitimacy of regulatory decisionmaking, especially when contrasted to a model in which agencies explicitly refused to consider the perspective of stakeholders.

An additional category of input value derives from the place the public comments process occupies within the broader system of political contesta-

25 *See* Londoner v. City & County of Denver, 210 U.S. 373 (1908) (holding that due process rights applied to tax assessment decision for small number of property owners).

26 *See* Bi-Metallic Inv. Co. v. State Bd. of Equalization, 239 U.S. 441 (1915) (holding that city-wide tax decision need not be preceded by opportunity for hearing because a large number of individuals were implicated).

27 There is a substantial debate within the European Union over how to best incorporate democratic and participatory norms into EU-level regulatory decisionmaking. *See generally* Edoardo Chiti, *European Agencies' Rulemaking: Powers, Procedures and Assessment*, 19 EUR. L.J. 93 (2013) (mapping and evaluating the current state of procedural rules within European administrative institutions).

28 *See* E. Allan Lind & Christiane Arndt, *Perceived Fairness and Regulatory Policy: A Behavioral Science Perspective on Government-Citizen Interaction* (OECD Regulatory Policy, Working Paper No. 6, 2016) (summarizing research on procedure and perception of government decisions).

29 *See generally* D. Ramona Bobocel & Leanne Gosse, *Procedural Justice: A Historical Review and Critical Analysis*, *in* THE OXFORD HANDBOOK OF JUSTICE IN THE WORKPLACE 51 (Russell S. Cropanzano & Maureen L. Ambrose eds., 2015).

30 *See generally* Kees van den Bos et al., *On the Role of Perceived Procedural Justice in Citizens' Reactions to Government Decisions and the Handling of Conflicts*, 10 UTRECHT L. REV. 1 (2014) (examining the Dutch process of citizen engagement when receiving unfavorable administrative decisions).

tion, primarily structured as rivalry between the main two political parties.[31] The same arguments that are developed in public comments can be deployed in other forums, including various media sources and during rounds of hearings and debates in Congress. The public comment process augments these alternative venues by providing incentives for the contestants to invest in more thorough, well-researched, and sophisticated analysis, which helps to elevate the quality of public policy debates that can easily become mired in symbolism rather than substance.[32] In addition, the open nature of the public comments process provides a means for every interested person to gain some measure of voice in the conversation, no matter their ability to generate press or summon lawmakers to their cause.

A final category is *observer value*. Agencies have an interest in improving the quality of their regulatory outputs, and the public has an interest in having some input into regulatory decisions, but outside observers also have an interest in understanding the relationship between the public and agencies. These observers include oversight institutions such as White House or congressional committees, who have political incentives to care about public reactions to rulemakings and may also rely, to some extent, on information generated in comments to identify salient issues that require their attention.[33] Other outsiders include social scientists, who have studied public bureaucracies and their interactions with the public for decades.[34] As is discussed in more detail below, public comments can provide researchers with important insights into the relationship between agencies, interest groups, and the public.[35]

Although the public comments can have multiple sources of value, they can also present agencies with considerable challenges when they arrive in overwhelming amounts. As barriers to participation have fallen and interest in the regulatory process has increased, there has been substantial growth in the number of comments received by agencies. This development is the subject of the following Section.

---

31 *Cf.* Livermore, *supra* note 23, at 100–04.

32 *See* Thomas O. McGarity, *Administrative Law as Blood Sport: Policy Erosion in a Highly Partisan Age*, 61 DUKE L.J. 1671 (2012) (discussing the decline in quality of public debate over regulatory decisionmaking).

33 *See* Mathew D. McCubbins & Thomas Schwartz, *Congressional Oversight Overlooked: Police Patrols Versus Fire Alarms*, 28 AM. J. POL. SCI. 165 (1984) (discussing the role of interest groups in congressional oversight).

34 *See infra* Section II.A.

35 *See* Kimberly D. Krawiec, *Don't "Screw Joe the Plummer": The Sausage-Making of Financial Reform*, 55 ARIZ. L. REV. 53, 58 (2013) ("Though scholars may debate the extent to which comment letters can—and should—reveal information to agencies, comments can reveal a great deal of information to the interested researcher . . . ." (footnote omitted)).

B. The Participation Explosion

Over the past several decades, informal rulemaking has become one of the most important policy-making forums in American politics.[36] Partisan rancor and divided government have often inhibited the ability of Congress to pass meaningful legislation.[37] Legislative gridlock on the major issues of the day, including immigration and climate change, has led to an ever more active executive branch.[38] Informal rulemaking is perhaps the Executive's preeminent tool for setting domestic policy, and administrations of both political parties have wielded it to great effect. This state of affairs has had several effects on the public comment process. Comments have taken on greater importance as a means of influencing major policy decisions and, at least occasionally, serve as a preliminary step in litigation over high stakes rulemakings. In addition, the public comment process is sometimes incorporated into broader advocacy efforts to influence public opinion and politicians. Advocacy campaigns around rulemakings increase public attention, leading to a higher volume of public comments.

At the same time as the stakes of the public comment process grew, information technology lowered the costs of participation. In the past, there were fairly substantial barriers to learning about a rule, engaging in research on the public policy choices involved, and submitting comments.[39] Now, rather than attempting a trip to the local library for a copy of the Federal Register, interested individuals can quickly access a diverse array of information about proposed regulations online. From websites of individual agencies or the comprehensive government-wide regulatory portal at "regulations.gov," interested persons can now easily and inexpensively identify ongoing regulatory proceedings, access relevant documents, and submit comments. In addition to extensive explanatory regulatory preambles, agencies typically include a great deal of additional substantive information on their rulemakings on agency websites, and any official supporting documents

36 *See* Kagan, *supra* note 12; Livermore, *supra* note 23.

37 *See* John J. Coleman, *Unified Government, Divided Government, and Party Responsiveness*, 93 AM. POL. SCI. REV. 821, 832–33 (1999) (finding greater legislative productivity and responsiveness to public opinion during periods of united government). Some earlier work had found that there was little difference between periods of divided and united government. DAVID R. MAYHEW, DIVIDED WE GOVERN: PARTY CONTROL, LAWMAKING, AND INVESTIGATIONS, 1946–2002, at 4 (2d ed. 2005) (finding that periods of divided government are not less productive). Subsequent work has challenged these conclusions. Sean Q. Kelly, *Divided We Govern? A Reassessment*, 25 POLITY 475, 482–83 (1993) (reexamining Mayhew's data and finding significant differences between divided and united governments). The increase in party polarization also appears to have exacerbated the threat of gridlock during periods of divided government, a phenomenon well illustrated by the vast fall in legislative productivity when the Republican party took control of both houses early in President Obama's first term.

38 *See* Kagan, *supra* note 12, at 2344.

39 For an overview of the preinternet difficulties of accessing the relevant information, see generally JUDITH SCHIEK ROBINSON, TAPPING THE GOVERNMENT GRAPEVINE: THE USER-FRIENDLY GUIDE TO U.S. GOVERNMENT INFORMATION SOURCES (1st ed. 1988).

are also made available on "regulations.gov." Advocacy organizations also publish their own analyses of proposed rules, and journalists and other content authors (i.e., bloggers, opinion writers, academics, etc.) often provide additional information for free. Within the time it once took to drive to the library and find a proposed rule's text, a relatively well-informed and conscientious researcher can amass a substantial amount of information about any rule of interest.

With lower costs and higher stakes, participation in the notice-and-comment process has ballooned. Several recent high profile rulemakings have generated what might be called *megaparticipation*, with comments numbering well over a million.[40] The State Department's Keystone XL oil pipeline decision received more than 2.5 million comments; the Federal Communications Commission received over 1.25 million comments on its net neutrality rules; and the EPA received over 4 million comments on its proposed Clean Power Plan.[41] These rates of participation begin to approach a nontrivial percentage of individuals who participate in elections, with the Clean Power Plan receiving nearly twenty percent as many comments as the total votes cast in House races in the 2014 midterm elections.[42] Within these voluminous submissions are form comments that have been circulated by advocacy groups; detailed, well-researched submissions by nongovernmental organizations, industry, and academics; and comments from other interested groups and individuals, including local organizations, states and municipalities, and members of the general public.

By way of comparison, in 1996, the EPA issued a high profile and controversial proposal to increase the stringency of the National Ambient Air Quality Standards (NAAQS) for particulate matter and ozone.[43] That proposal was the subject of significant discussion in the media, as well as congressional hearings and a substantial lobbying push from both opponents and supporters.[44] The EPA established an email address (which was cutting edge technology at the time) to receive comments and staffed a telephone hotline to

40 The prefix "mega" denotes one million in the metric system, as in "megawatt."

41 Gautham Nagesh, *Federal Agencies Are Flooded by Comments on New Rules*, WALL ST. J. (Sept. 3, 2014), http://www.wsj.com/articles/federal-agencies-are-flooded-by-comments-on-new-rules-1409786480; *see also* Carbon Pollution Emission Guidelines for Existing Stationary Sources: Electric Utility Generating Units, 80 Fed. Reg. 64,662, 64,704 (Oct. 23, 2015) (to be codified at 40 C.F.R. pt. 60).

42 Roughly seventy-five million votes were cast for major party House candidates in the 2014 election. *See* KAREN L. HAAS, CLERK OF THE HOUSE OF REPRESENTATIVES, STATISTICS OF THE CONGRESSIONAL ELECTION OF NOV. 4, 2014, at 54 (2015).

43 National Ambient Air Quality Standards for Ozone: Proposed Decision, 61 Fed. Reg. 65,716 (Dec. 13, 1996) (to be codified at 40 C.F.R. pt. 50).

44 *See, e.g.*, *Clean Air Act: Ozone and Particulate Matter Standards: Hearings Before the Subcomm. on Clean Air, Wetlands, Private Prop. & Nuclear Safety and the S. Comm. on Env't & Pub. Works*, 105th Cong. (1997); FARA CHURGIN, PAST, PRESENT AND FUTURE: THE NATIONAL AMBIENT AIR QUALITY STANDARDS FOR OZONE AND PARTICULATE MATTER (1997), http://www.wise-intern.org/journal/1997/.

take oral comments as well.[45] At the end of the commenting process, the agency had received roughly 50,000 comments, which was a substantial number at the time.[46] In 2014, the Obama EPA again proposed a measure to increase the standards for ozone alone, which generated similar levels of opposition from industry and support from environmental and public health groups.[47] That proposal received nearly an order of magnitude higher number of public comments.[48] The massive increase in participation on these two very similar rules provides a clear indication of the degree to which the notice-and-comment process has changed in the intervening years.

In terms of simple administrative manageability, opening the floodgates of participation in the rulemaking process presents clear difficulties. As discussed above, agencies are obligated to consider and respond to substantive comments; having to review many millions of comments to even determine their substance is an extraordinary burden. Even if many of the comments are repeats of form submissions, agencies must still separate out the unique comments and give them at least some cursory examination. Given the lower cost of acquiring information, agencies may also face a higher volume of substantively meaningful comments. Even if these comments contain valuable information, processing them can require substantial commitment of agency resources.

From the perspective of expanding citizen involvement in administrative decisionmaking, recent innovations create obvious opportunities. From an output value perspective, a large number of comments can mean more information. This information can take many forms, and may include technical information about the interaction of a rulemaking with technology or business practices, information about public preferences or desires, forecasts about how individuals or organizations may change their behavior in light of a rulemaking, or regulatory alternatives. By considering this information, agencies can improve their decisions along both technocratic and political dimensions. From an input value perspective, when more people participate in the comment process, they can (at least theoretically) enjoy procedural justice and due process benefits and contribute to the broader democratic conversation over regulatory decisionmaking. If submitting comments increases an individual's willingness to accept the legitimacy of an agency's final decision, the larger the number of commenters, the greater the legitimacy-conferring benefits of the process. Just as higher voter turnout is often

45 Press Release, EPA, EPA Administrator Carol Browner to Sign Ozone (Smog) and Particulate (Soot) Air Quality Standards (July 16, 1997), http://bit.ly/2ic1cO4.

46 *Id.*

47 *See* INST. FOR 21ST CENTURY ENERGY, U.S. CHAMBER OF COMMERCE, GRINDING TO A HALT: EXAMINING THE IMPACTS OF NEW OZONE REGULATIONS ON KEY TRANSPORTATION PROJECTS (2015), https://www.globalenergyinstitute.org/sites/. . ./Grinding-to-a-halt-9-18.compressed.pdf.

48 *See* EPA, RESPONSES TO SIGNIFICANT COMMENTS ON THE 2014 PROPOSED RULE ON THE NATIONAL AMBIENT AIR QUALITY STANDARDS FOR OZONE 1 (2015) ("More than 430,000 written comments were received from various commenters during the public comment period on the proposal . . . .").

interpreted as a sign of a more robust democracy, a larger number of public comments indicates a more inclusive and participatory administrative process.

But translating the promise of mass participation into a public comment process with enhanced output or input value has proven to be no easy challenge.[49] Indeed, there is some controversy over whether there *is* any value in many of the comments that are now submitted to agencies. Many of the comments received by agencies amount to very short statements of approval or disapproval, with very little in the way of substantive support.[50] Prominent administrative law scholars have disputed whether agencies ought to take these comments into consideration and, if so, how.[51] In one camp are Cynthia Farina and scholars associated with the Cornell eRulemaking Initiative (CeRI), who argue that certain preferences "may be good enough for electoral democracy, but . . . are not good enough for rulemaking."[52] They come to this conclusion based on their view about the nature of the administrative process, which is meant to be more than the mere aggregation of preferences, and instead involves "[r]easoned decisionmaking," which requires both "technocratic rationality" as well as a considered effort to "make choices among interests and values deliberatively."[53] For Farina and her collaborators, the deliberative nature of the rulemaking process excludes consideration of bare preferences that are not supported by public-regarding reasons. Others have raised similar concerns, focusing especially on the low substantive quality of many comments.[54] An additional concern is that mass public

49 *See generally* David Schlosberg et al., *Deliberation in E-Rulemaking? The Problem of Mass Participation*, *in* ONLINE DELIBERATION: DESIGN, RESEARCH, AND PRACTICE 133 (Todd Davies & Seeta Peña Gangadharan eds., 2009). The Administrative Conference of the United States has taken a particular interest in this issue, commissioning multiple reports from independent academics on the subject in recent years. *See* STEVEN J. BALLA, PUBLIC COMMENTING ON FEDERAL AGENCY REGULATIONS: RESEARCH ON CURRENT PRACTICES AND RECOMMENDATIONS TO THE ADMINISTRATIVE CONFERENCE OF THE UNITED STATES (2011); MICHAEL HERZ, USING SOCIAL MEDIA IN RULEMAKING: POSSIBILITIES AND BARRIERS (2013).

50 *See* Nina A. Mendelson, Response, *Should Mass Comments Count?*, 2 MICH. J. ENVTL. & ADMIN. L. 173, 173–74 (2012).

51 *Compare id.*, *with* Cynthia R. Farina et al., *Rulemaking vs. Democracy: Judging and Nudging Public Participation that Counts*, 2 MICH. J. ENVTL. & ADMIN. L. 123, 137 (2012).

52 Farina et al., *supra* note 51, at 137.

53 *Id.* at 139.

54 *See* Krawiec, *supra* note 35, at 58 (contrasting industry comments on the SEC's Volcker Rule, which were "meticulously drafted, argued, and researched" with "citizen letters [which were] short and provide little evidence that citizen commenters even understand, or care, what proprietary or fund investment is, much less the ways in which agency interpretation of the Volcker Rule's complex and ambiguous provisions might govern such activities"); Stuart W. Shulman, *The Case Against Mass E-mails: Perverse Incentives and Low Quality Public Participation in U.S. Federal Rulemaking*, 1 POL'Y & INTERNET 23 (2009); *see also* Elise Hu, *1 Million Net Neutrality Comments Filed, but Will They Matter?*, NPR (July 21, 2014), http://www.npr.org/sections/alltechconsidered/2014/07/21/332678802/one-million-net-neutrality-comments-filed-but-will-they-matter (quoting Richard Pierce stating that "[t]he vast majority of . . . comments are utterly worthless").

comments can be manipulated by "astroturf campaigns" run by special interest groups that seek to create the impression of broad public support for their positions.[55]

In the opposite camp is Nina Mendelson, who has argued that even relatively unsophisticated comments should still be considered by agencies because, in many instances, agencies are called to "decide values and policy questions left unresolved by their authorizing statutes" and, to some degree, the administrative process may appropriately be considered a site of pluralist contestation and not only impartial deliberation.[56] Mendelson also emphasizes that even comments that appear information-poor may still provide useful insights for agency decisionmakers.[57]

Michael Herz has staked out something of a middle position, acknowledging the theoretical desirability of considering even unsophisticated comments, but noting the difficulty of extracting valuable content from them and figuring out how exactly that content should figure into agency decisionmaking.[58]

Although Mendelson sees some use in analyzing and considering unsophisticated mass comments, she stays well within a broad expert consensus that the notice-and-comment process should not be treated as a plebiscite.[59] Scholars, agencies, and courts appear to largely agree that bare preferences have little place in the rulemaking process, which should be oriented toward collecting and aggregating information (including, perhaps, about public values) rather than deferring to the judgments expressed in comments about the desirability of a policy.[60] The advice on "regulations.gov" emphasizes that "[t]he comment process is not a vote,"[61] and courts have also repeatedly emphasized that agencies "need not respond at all to comments that are

55 Victoria Peng, Note, *Astroturf Campaigns: Transparency in Telecom Merger Review*, 49 U. MICH. J.L. REFORM 521 (2016).

56 Mendelson, *supra* note 19, at 1350–51 (citing Richard B. Stewart, *The Reformation of American Administrative Law*, 88 HARV. L. REV. 1669, 1683 (1975)).

57 As an example, Professor Mendelson offers the case of a Coast Guard decision to establish "live-fire zones" on the Great Lakes for purposes of training exercises. That decision had a truncated public comment process, in part because the Coast Guard apparently expected little resistance. As it turned out, there was a substantial public outcry when the practice began, and the Coast Guard ended up withdrawing its plan. Mendelson, *supra* note 50, at 177. In that case, it appears, a more robust comment process, along with attention to informal comments that may have been submitted, may have saved the agency some wasted time and resources on what was ultimately a failed decision.

58 Michael Herz, *"Data, Views, or Arguments": A Rumination*, 22 WM. & MARY BILL RTS. J. 351, 374–78 (2013).

59 Mendelson, *supra* note 19, at 1374 (stating that agencies should not "tally up" comments to "serve as a referendum or a dispositive vote of some sort").

60 *See* Bridget C.E. Dooling, *Legal Issues in E-Rulemaking*, 63 ADMIN. L. REV. 893, 901 n.31 (2011) (stating that the antiplebiscite position is "settled").

61 *Tips for Submitting Effective Comments*, REGULATIONS.GOV, https://www.regulations.gov/docs/Tips_For_Submitting_Effective_Comments.pdf (last visited Nov. 28, 2017).

'purely speculative and do not disclose the factual or policy basis on which they rest.'"[62]

This expert consensus, however, appears to cut against widely held views among participating individuals, advocacy groups, and journalists that the public expression of preferences should and does carry some weight, entirely apart from whatever substantive justification for those preferences is offered.[63] Farina and colleagues point to two "powerful cultural patterns" that reinforce the "plebiscite assumption" that participation in notice-and-comment proceedings are equivalent to a vote in which the majority preference will prevail.[64] "The first is the popular equation in the United States of democratic voice with casting a vote," and the second is a feature of "online culture" in which "web content lives or dies by whether the crowd gives thumbs up, or down."[65] Together, Farina and her colleagues at CeRI argue that "these two patterns may create . . . a powerful 'voting instinct'" that reinforces the idea among participants that rulemaking is akin to a plebiscite.[66]

The current state of affairs in which megaparticipation exists alongside deeply contested views about how mass comments should be treated creates a serious dilemma. Given popular perceptions of the public comment process, it seems unlikely that agencies will, as a political matter, be able to both disregard mass comments altogether, and be entirely candid about doing so, without risking substantial public backlash. At the same time, paying little heed to mass comments, while also soliciting them and ostensibly taking them in to consideration may, over time, "cause citizens justifiably to lose faith in the prospect of genuine opportunities to participate in government."[67] Simply discouraging comments or attempting to raise barriers to commenting also seems both unlikely and risky.

There are two potential strategies to addressing the mass-commenting dilemma. One way to respond to a large number of unsophisticated comments is to improve their sophistication.[68] There are both internal government efforts as well as external initiatives to promote this goal. Within the government, the General Services Administration's digital service agency, referred to as 18F, has undertaken a substantial effort to develop a new public comment platform, "eRegulations," that attempts to provide an enhanced user experience for commenters that ultimately results in comments that are

62 Pub. Citizen, Inc. v. FAA, 988 F.2d 186, 197 (D.C. Cir. 1993) (quoting Home Box Office, Inc. v. FCC, 567 F.2d 9, 35 n.58 (D.C. Cir. 1977) (per curiam)).

63 *See* Farina et al., *supra* note 5, at 431–32.

64 *Id.* at 431–32.

65 *Id.* at 432.

66 *Id.*

67 Mendelson, *supra* note 50, at 175–76.

68 *See* Fred Emery & Andrew Emery, *A Modest Proposal: Improve E-Rulemaking by Improving Comments*, 31 ADMIN. & REG. L. NEWS, Fall 2005, at 8.

"more useful, and more focused."[69] The eRegulations platform was initially developed in collaboration with the Consumer Financial Protection Bureau, but is open source and meant to be expanded on and improved by what 18F refers to as "intrepid civic hackers both within and outside of the government."[70] The hallmark of the eRegulations platform is the ability for commenters to target specific sections of a regulatory text, which is meant to facilitate more focused responses and allow agencies to more easily categorize comments. The 18F platform has been used by several additional agencies, including the EPA and the Bureau of Alcohol, Tobacco, Firearms and Explosives.[71]

Outside of government, CeRI has tested the value of facilitation of policy deliberation as part of the rulemaking process through their Regulation Room.[72] This highly ambitious project, which was chosen by the Department of Transportation (DOT) as its open-government flagship initiative and received a Leading Practices Award by the White House, attempted to increase the value of public comments by hosting an online deliberative portal that enabled the interaction of potential commenters with human facilitators who were trained to engage with participants in ways that probe for the kinds of information that are more likely to be of value to agency decisionmaking.

A second approach to addressing mass participation attempts to extract more meaning from comments in their current form. Rather than trying to improve the quality of public comments at the source, this method uses advanced information-processing and text analysis techniques to extract as much possible usable information from the type of public comments that are already submitted to agencies. The second approach is not a substitute for the first and can be used in tandem with efforts to increase the sophistication of comments—as the quality of comments improves, there will be more information to extract. In the balance of this Article, we will focus on how the data that already exists in the form of agency comments can be put to use, first to better understand how agencies interact with the public, and then to improve those interactions to ultimately increase the value of the public comment process.

---

69 Andre Francisco & Will Sullivan, *New Pilot Aims to Streamline Notice and Comment Process*, 18F BLOG (July 26, 2016), https://18f.gsa.gov/2016/07/26/new-pilot-aims-to-streamline-notice-and-comment-process/.

70 CM Lubinski, *An Open Source Government Is a Faster, More Efficient Government*, 18F BLOG (Dec. 9, 2015), https://18f.gsa.gov/2015/12/09/an-open-source-government-is-a-faster-more-efficient-government/. For more information, see *18F/eregs-platform Wiki*, GITHUB, https://github.com/18F/eregs-platform/wiki (last updated May 8, 2017).

71 *See* Francisco & Sullivan, *supra* note 69; Lubinski, *supra* note 70.

72 *See* Dmitry Epstein & Gilly Leshed, *The Magic Sauce: Practices of Facilitation in Online Policy Deliberation*, 12 J. PUB. DELIBERATION, no. 1, 2016; Farina et al., *supra* note 5; Cynthia R. Farina et al., *Rulemaking in 140 Characters or Less: Social Networking and Public Participation in Rulemaking*, 31 PACE L. REV. 382 (2011); *History*, REG. ROOM, http://regulationroom.org/about/history (last visited Nov. 28, 2017).

C. *Public Comments as Big Data*

In April 2010, the Library of Congress and Twitter announced a deal to allow every publicly available microblog on the service to be archived and made publicly available.[73] Two and a half years later, the Library reported that it had received over 170 billion "tweets," which required 133.2 terabytes of storage and that it was continuing to receive and archive data from Twitter.[74] Although the Library has collected this massive dataset, it has not provided any access for researchers; the data sits unused on Library servers, taking up storage space and providing no intellectual value. The problem is that the Library faces a host of difficult logistical and policy questions concerning how to provide access, to whom to provide access, and on what terms. The institution has found itself unable to make progress on answering these questions, and there are no immediate plans to allow access to the data.[75]

The Library is not alone in struggling with what has become known as "big data." That term has gained substantial popularity in recent years and is now used in a variety of settings, from popular news stories to National Science Foundation solicitations.[76] By definition, big data is associated with a set of difficulties and challenges.[77] As the name suggests, the primary defining characteristic of big data is its volume—the sheer amount of information that is stored and potentially subject to analysis, as measured in the number of bytes of storage required to house the data.[78] Several other characteristic features of big data that have been proposed include the "velocity" at which data accumulates, problems with "veracity," and relatively low "value density."[79] Big data is also commonly unstructured, meaning that it does not follow a standardized tabular format.[80] Collections of texts or images are often unstructured and are simply presented as unorganized sets of data,

73 Michael Zimmer, *The Twitter Archive at the Library of Congress: Challenges for Information Practice and Information Policy*, 20 FIRST MONDAY, July 2015.

74 *Id.*

75 Andrew McGill, *Can Twitter Fit Inside the Library of Congress?*, ATLANTIC (Aug. 4, 2016), http://www.theatlantic.com/technology/archive/2016/08/can-twitter-fit-inside-the-library-of-congress/494339/.

76 *Critical Techniques, Technologies and Methodologies for Advancing Foundations and Applications of Big Data Sciences and Engineering*, NAT'L SCI. FOUND., https://www.nsf.gov/funding/pgm_summ.jsp?pims_id=504767 (last visited Oct. 23, 2017); Steve Lohr, *How Big Data Became So Big*, N.Y. TIMES (Aug. 11, 2012), http://www.nytimes.com/2012/08/12/business/how-big-data-became-so-big-unboxed.html.

77 The University of Wisconsin-Madison statistics department defines "big data" as "a collection of data sets so large and complex that it becomes difficult to process using on-hand database management tools or traditional data processing applications." Dep't of Statistics, *Big Data*, U. WIS.-MADISON, https://www.stat.wisc.edu/bigdata (last visited Nov. 28, 2017).

78 *See* Gandomi & Haider, *supra* note 10, at 138.

79 *Id.* at 138–39.

80 Data that was organized according to several variables linked together as unique observations would be a common example of a standard tabular format.

demarcated at the document or image level, often with little in the way of identifying information.[81] As more of this type of data is collected, researchers in several disciplines have developed new techniques of data management and analysis to deal with the challenges that it presents.[82] These tools are now deployed in both academic and industry settings on a wide range of data sources, ranging from social media content to financial market transactions.[83]

Agencies collect a great deal of data of many different varieties, from greenhouse gas emissions data collected by the EPA to the financial filings of publicly traded companies collected by the Securities and Exchange Commission (SEC).[84] Some of this data is relatively structured (such as the EPA's emission data, which are organized by firm and year and easily interpretable); other data is highly unstructured (as is the SEC's 10-K filings, which are text documents with a limited amount of metadata). Some recent efforts have been made to bring big data analytic tools to bear on the information collected by agencies—for example, a recent collaboration between the EPA and researchers at the University of Chicago used the agency's inspection data under its Resource Conservation and Recovery Act enforcement program to predict facilities that were most likely to be out of compliance with environmental requirements.[85] These early efforts are promising, but there remains a great deal of untapped potential to bring sophisticated analysis techniques to government data.

With the move to e-rulemaking and the shift from paper to digital submissions of public comments, agencies began accumulating a near archetypical example of big data. Public comments make up a large volume dataset—with many millions of documents—that is textual and relatively unstructured. Given its size and unstructured nature, traditional tools cannot fully exploit this data source; on the one hand, simply reading the documents is enormously cumbersome and time consuming and may miss important large scale trends; on the other hand, standard statistical tools are ill-suited to large and unstructured textual corpora.

Approaching public comments as big data allows for the use of natural language processing techniques that generate interpretable and meaningful analyses of large textual corpora. Natural language processing is a broader field within computer science and artificial intelligence that is focused on computational approaches for processing natural languages, and more spe-

81 *See* Gandomi & Haider, *supra* note 10, at 138.

82 *See id.* at 140–42.

83 *See generally* VIKTOR MAYER-SCHÖNBERGER & KENNETH CUKIER, BIG DATA: A REVOLUTION THAT WILL TRANSFORM HOW WE LIVE, WORK, AND THINK (2013) (discussing a variety of applications for big data techniques).

84 *See Greenhouse Gas Reporting Program*, EPA, https://www.epa.gov/ghgreporting (last updated Oct. 5, 2017); *Annual Report*, SEC, https://www.sec.gov/fast-answers/answers-annrephtm.html (last modified Oct. 15, 2014).

85 *See* Jane Zanzig, *EPA: Data-Driven Hazardous Waste Detection*, https://dssg.uchicago.edu/2015/10/16/epa-data-driven-hazardous-waste-detection/ (last visited Oct. 22, 2017).

cifically, on designing software that can derive information from natural language artifacts (either written or spoken) as well as generate language that is recognizable and meaningful to human users. The aim, at a high level of abstraction, is to extract, structure, and analyze quantitative information from streams of words, as embodied textually in documents or verbally in speech.[86]

Computational text analysis focuses on extracting or creating information from written texts and can include *syntactic parsing* (which uses automatically extracted grammatical rules to tag words within a text according to parts of speech and grammatical categories),[87] *sentiment analysis* (which extracts information on authors' attitudes),[88] *topic modeling* (which is a form of clustering that identifies subject matter categories within unstructured corpora),[89] *lexical semantics* (which is a mechanism to detect the meanings of words from context),[90] and *stylometry* (which focuses on author attribution and writing style).[91] In recent years, there has been considerable growth in computational social science: the use of computational text analysis in the humanities and social sciences.[92] These "distant reading" approaches combine quantitative analysis with qualitative interpretation to generate new insights into a wide range of fields.[93] Within legal scholarship, researchers are beginning to apply computational text analysis techniques to legal docu-

86 Problems within natural language processing include machine translation, speech recognition, automatic summarization, optical character recognition, and developing natural language user interfaces. For general background on the field, see Daniel Jurafsky & James H. Martin, Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition (Aug. 28, 2017) (unpublished manuscript), https://web.stanford.edu/~jurafsky/slp3/ed3book.pdf.

87 *See* Alton F. Sanders & Ruth H. Sanders, *Syntactic Parsing: A Survey*, 23 COMPUTERS & HUMAN. 13, 13–14 (1989). Both Google and the Natural Language Processing group at Stanford have publicly available parsers. *See* Slav Petrov, *Announcing SyntaxNet: The World's Most Accurate Parser Goes Open Source*, GOOGLE RES. BLOG (May 12, 2016), https://research.googleblog.com/2016/05/announcing-syntaxnet-worlds-most.html; STANFORD PARSER, http://nlp.stanford.edu:8080/parser/ (last updated Sept. 12, 2016).

88 *See infra* Section II.B.

89 *See infra* Section III.C.

90 *See* Tomas Mikolov et al., *Efficient Estimation of Word Representations in Vector Space*, ARXIV, https://arxiv.org/abs/1301.3781 (last updated Sept. 7, 2013).

91 *See* Keith Carlson et al., *A Quantitative Analysis of Writing Style on the U.S. Supreme Court*, 93 WASH. U. L. REV. 1461 (2016).

92 *See* Matthew G. Kirschenbaum, *What Is Digital Humanities and What's It Doing in English Departments?*, 150 ADE BULL. 55 (2010) (discussing growing role of computational text analysis in a new generation of literary scholarship); David Lazer et al., *Computational Social Science*, 323 SCIENCE 721 (2009) (discussing role of natural language processing in facilitating social science research).

93 *See* FRANCO MORETTI, DISTANT READING (2013); James A. Evans & Pedro Aceves, *Machine Translation: Mining Text for Social Theory*, 42 ANN. REV. SOC. 21 (2016).

ments, including Supreme Court opinions,[94] legal briefs,[95] national constitutions,[96] Federal Register notices,[97] and bankruptcy opinions.[98]

There have been some initial experiments with natural language processing approaches for improving the quality of comment submissions, either as a complement to, or as a substitute for, the human facilitation techniques pioneered at the Regulation Room. In one implementation, a team at Cornell University constructed an interactive natural language tool that helped commenters construct more "evaluable" arguments that explicitly related their claims to supporting premises or facts.[99] This tool can help improve the quality of comments by avoiding long chains of unsupported conclusions that provide little in the way of information for regulators.[100] A second implementation constructed a "Discussion Facilitation Agent" that was powered by several natural language processing tools (including one designed to help participants answer substantive questions), a word cloud generator, and email reminders with links to additional content.[101] This automated agent was intended to empower participants to understand highly complex rules for the purpose of facilitating high-value comments.

These approaches to improving comment quality may yield important techniques that facilitate higher value comments that contain more information of value to the agency. Even these higher value comments, however, will pose big data challenges given the large volume of comments and the reality that they are likely to remain relatively unstructured. Techniques in computational text analysis have substantial potential to extract information from comments that are submitted—however sophisticated or unsophisticated they may be. In fact, many big data analysis techniques are designed to deal with datasets that have low "value density"—which is to say that there is valuable information, but it is widely diffused. In data where there is a low signal-to-noise ratio, only a high total volume of data along with sophisticated analysis techniques makes it possible to recover the weak signal. Used in tandem with more sophisticated natural language techniques that facilitate higher

94 RYAN C. BLACK ET AL., U.S. SUPREME COURT OPINIONS AND THEIR AUDIENCES (2016); Carlson et al., *supra* note 91.

95 Ryan C. Black et al., *The Role of Emotional Language in Briefs Before the US Supreme Court*, 4 J.L. & CTS. 377 (2016).

96 David S. Law, *Constitutional Archetypes*, 95 TEX. L. REV. 153 (2016).

97 Edward H. Stiglitz, *Unaccountable Midnight Rulemaking? A Normatively Informative Assessment*, 17 N.Y.U. J. LEGIS. & PUB. POL'Y 137 (2014).

98 Jonathan Macey & Joshua Mitts, *Finding Order in the Morass: The Three Real Justifications for Piercing the Corporate Veil*, 100 CORNELL L. REV. 99 (2014).

99 *See* Joonsuk Park et al., *Toward Machine-Assisted Participation in eRulemaking: An Argumentation Model of Evaluability*, *in* PROCEEDINGS OF THE 15TH INTERNATIONAL CONFERENCE ON ARTIFICIAL INTELLIGENCE AND LAW 206, 206 (2015).

100 *See id.* at 207–08.

101 *See* Peter Muhlberger et al., *An Experiment in E-Rulemaking with Natural Language Processing and Democratic Deliberation*, *in* CITIZEN 2.0: PUBLIC AND GOVERNMENTAL INTERACTION THROUGH WEB 2.0 TECHNOLOGIES 23 (Kathryn Kloby & Maria J. D'Agostino eds., 2012).

quality comments, computational text analysis can provide a means to turn what is now a relatively low-value corpus of unstructured public comments into a powerful tool for understanding and communication.

II. A New Lens on Agencies

Public comments express opinions, make arguments, and offer perspective about proposed agency actions. These comments can be thought of as having an attitude, valence, or more commonly in the literature, *sentiment*. This Part discusses the potential application of computational analysis of such sentiment to the study of agency-public interactions. We begin by discussing the literature on bureaucratic politics and the potential for public comments to shed light on important questions within that literature. We then provide an overview of the field of sentiment analysis in computational linguistics and report the results of an analysis that examines sentiment in nearly three million public comments received by agencies during the Obama Administration.

A. Agencies, Politics, and the Public

The relationship between public bureaucracies and the public they serve, and the appropriate place of politics in that relationship, is one of the most enduring questions in the social scientific study of government.[102] In the late nineteenth century, Woodrow Wilson, in one of the founding texts of the field of public administration, pondered how public bureaucracies could cure the ills of the spoils system.[103] At the turn of the century, Frank Goodnow argued forcefully for a separation of politics and administration to avoid the incompetence and ineffectiveness that resulted from the influence of partisanship and patronage on government bureaucracies.[104] Writing roughly contemporaneously, but from a European perspective, Max Weber came to the opposite concern that the growth of a powerful and independent professional bureaucracy would crowd out politics, resulting in rule by administrators rather than politically responsive politicians.[105]

From these origins, scholars have taken the study of bureaucracies in a wide variety of empirical and normative directions, but the interplay of administration and politics has remained a core concern. The notion of a strict politics/administration dichotomy promoted by Goodnow, in which politics could be separated from the business of public administration, exper-

102 *See* Dwight Waldo, *Politics and Administration: On Thinking About a Complex Relationship*, *in* A Centennial History of the American Administrative State 89 (Ralph Clark Chandler ed., 1987).

103 *See* Woodrow Wilson, *The Study of Administration*, 56 Pol. Sci. Q. 481 (1941).

104 *See generally* Frank J. Goodnow, Politics and Administration: A Study in Government (1900).

105 Max Weber, *Bureaucracy*, *reprinted in* From Max Weber: Essays in Sociology 196, 232 (H. H. Gerth & C. Wright Mills eds., 1946).

ienced a heyday in the early to mid-twentieth century,[106] but eventually fell out of favor as scholars and public perception shifted to see the work of administrative agencies as inevitably involving decisions with controversial political consequences.[107] During the 1960s, skepticism about the regulatory state emerged, with writers such as William Niskanen and Theodore Lowi arguing that agencies consistently undermined the public interest by pursuing their own ends or those of favored clients.[108] Capture Theory expanded on this skepticism and introduced more explicit economic models and empirical methods to attempt to tease out the effect of interest group politics on agency decisionmaking.[109] More recent generations of scholars have applied concepts such as transaction costs and principal-agent theory to the study of public bureaucracies,[110] with a particular emphasis on the relationship between administrative agencies and other government institutions, particularly the presidency, the White House, Congress, and the courts.[111]

106 *See* Gerald E. Caiden, *In Search of an Apolitical Science of American Public Administration*, *in* POLITICS AND ADMINISTRATION: WOODROW WILSON AND AMERICAN PUBLIC ADMINISTRATION 51, 51 (Jack Rabin & James S. Bowman eds., 1984).

107 *See, e.g.*, PAUL H. APPLEBY, POLICY AND ADMINISTRATION (1949); DWIGHT WALDO, ADMINISTRATIVE STATE: A STUDY OF THE POLITICAL THEORY OF AMERICAN PUBLIC ADMINISTRATION (2d ed. 1984); DWIGHT WALDO, ADMINISTRATIVE STATE: A STUDY OF THE POLITICAL THEORY OF AMERICAN PUBLIC ADMINISTRATION (1st ed. 1948); Stewart, *supra* note 56.

108 THEODORE J. LOWI, IDEOLOGY, POLICY, AND THE CRISIS OF PUBLIC AUTHORITY (1st ed. 1969); THEODORE J. LOWI, THE END OF LIBERALISM: THE SECOND REPUBLIC OF THE UNITED STATES (2d. ed. 1979); WILLIAM A. NISKANEN, JR., BUREAUCRACY AND REPRESENTATIVE GOVERNMENT (1971).

109 *See* Gary S. Becker, *A Theory of Competition Among Pressure Groups for Political Influence*, 98 Q.J. ECONOMICS 371, 372 (1983); William T. Gormley Jr., *A Test of the Revolving Door Hypothesis at the FCC*, 23 AM. J. POL. SCI. 665 (1979); Sam Peltzman, *Toward a More General Theory of Regulation*, 19 J.L. & ECON. 211, 212 (1976); Richard A. Posner, *Theories of Economic Regulation*, 5 BELL J. ECON. & MGMT. SCI. 335, 336 (1974); George J. Stigler, *The Theory of Economic Regulation*, 2 BELL J. ECON. & MGMT. SCI. 3, 3 (1971). *See generally* William D. Berry, *An Alternative to the Capture Theory of Regulation: The Case of State Public Utility Commissions*, 28 MIDWEST J. POL. SCI. 524, 544–45 (1984) (describing Capture Theory).

110 *See* Kenneth J. Meier & George A. Krause, *The Scientific Study of Bureaucracy: An Overview*, *in* POLITICS, POLICY, AND ORGANIZATIONS: FRONTIERS IN THE SCIENTIFIC STUDY OF BUREAUCRACY 1, 7–14 (George A. Krause & Kenneth J. Meier eds., 2003) (providing an overview of the field).

111 *See* John Ferejohn & Charles Shipan, *Congressional Influence on Bureaucracy*, 6 J.L. ECON. & ORG. 1 (1990) (investigating the effect of Congress on bureaucracy); Thomas H. Hammond & Jack H. Knott, *Who Controls the Bureaucracy?: Presidential Power, Congressional Dominance, Legal Constraints, and Bureaucratic Autonomy in a Model of Multi-Institutional Policy-Making*, 12 J.L. ECON. & ORG. 119 (1996); McCubbins & Schwartz, *supra* note 33; Barry R. Weingast & Mark J. Moran, *Bureaucratic Discretion or Congressional Control? Regulatory Policymaking by the Federal Trade Commission*, 91 J. POL. ECON. 765 (1983); William F. West & Joseph Cooper, *Legislative Influence v. Presidential Dominance: Competing Models of Bureaucratic Control*, 104 POL. SCI. Q. 581 (1989).

There is now a considerable and methodologically eclectic empirical and theoretical literature on bureaucratic politics.[112] Important contributions include emphasizing the heterogeneous nature of agencies and the importance of institutional design and administrative procedure in affecting relationships between agencies and political institutions.[113] Even narrowing focus to U.S. federal agencies, there is great variation in agencies' missions, the interest group environments in which they operate, their institutional structures and procedures, and their internal cultures.[114] Taken together, the collection of these differences can be thought to contribute to distinct profiles that remain at least somewhat consistent over time, inclining some agencies toward certain behaviors while inclining other agencies toward other behaviors. Broadly, for example, independent agencies, which are governed by bipartisan commissions with appointment protections, may be less responsive to political principals in the White House.[115] An agency that is narrowly focused on regulating a single industry may be more subject to clientelism and capture than a more generalist agency with a broadly balanced constituency.[116]

Among the differences that appear to matter for agency decisionmaking is the makeup of the career personnel and specifically their policy prefer-

112 *See, e.g.*, JOEL D. ABERBACH ET AL., BUREAUCRATS AND POLITICIANS IN WESTERN DEMOCRACIES (1981) (conducting interviews of over 1000 professional bureaucrats and politicians in seven countries to examine the politics/administration distinction across political cultures); JOHN BREHM & SCOTT GATES, WORKING, SHIRKING, AND SABOTAGE: BUREAUCRATIC RESPONSE TO A DEMOCRATIC PUBLIC (1999); Martin Baekgaard et al., *When Politics Matters: The Impact of Politicians' and Bureaucrats' Preferences on Salient and Nonsalient Policy Areas*, 28 GOVERNANCE 459 (2015); Jonathan Bendor & Adam Meirowitz, *Spatial Models of Delegation*, 98 AM. POL. SCI. REV. 293 (2004); Morris P. Fiorina & Roger G. Noll, *Voters, Bureaucrats and Legislators: A Rational Choice Perspective on the Growth of Bureaucracy*, 9 J. PUB. ECON. 239 (1978) (presenting a blame-shifting theory of agencies).

113 JAMES Q. WILSON, BUREAUCRACY: WHAT GOVERNMENT AGENCIES DO AND WHY THEY DO IT (1989) (discussing the great deal of variation within the general form of administrative agencies); Steven J. Balla, *Administrative Procedures and Political Control of the Bureaucracy*, 92 AM. POL. SCI. REV. 663 (1998); Neal Devins & David E. Lewis, *Not-So Independent Agencies: Party Polarization and the Limits of Institutional Design*, 88 B.U. L. REV. 459 (2008) (examining the effects of institutional design on political control); David Epstein & Sharyn O'Halloran, *Divided Government and the Design of Administrative Procedures: A Formal Model and Empirical Test*, 58 J. POLITICS 373 (1996); Mathew D. McCubbins et al., *Administrative Procedures as Instruments of Political Control*, 3 J.L. ECON. & ORG. 243 (1987).

114 *See* WILSON, *supra* note 113 (discussing agency variety); *see also* Rachel E. Barkow, *Insulating Agencies: Avoiding Capture Through Institutional Design*, 89 TEX. L. REV. 15, 30–37 (2010) (discussing structural differences between agencies); Kirti Datla & Richard L. Revesz, *Deconstructing Independent Agencies (and Executive Agencies)*, 98 CORNELL L. REV. 769 (2013) (demonstrating that "independence" arises from a set of characteristics).

115 *See* Devins & Lewis, *supra* note 113, at 460 (noting delayed responsiveness in independent agencies).

116 *See* Michael A. Livermore & Richard L. Revesz, *Regulator Review, Capture, and Agency Inaction*, 101 GEO. L.J. 1337, 1362–67 (2013).

ences.[117] Although policy is at least sometimes responsive to the desires of political principals,[118] the substantial policy discretion given to career personnel, and their role in structuring and informing the decisions made by principals, gives them substantial ability to shape outcomes toward their preferred policies.[119] This fact does not imply bad faith on the part of career staff—they may genuinely intend to serve the public interest, as they understand it, and have little at stake in terms of personal satisfaction for the policy decisions made by agencies.[120] Nevertheless, the values, perspectives, and beliefs of career personnel can—perhaps appropriately—influence the choices that agencies make.

There is a subfield within the bureaucratic politics literature that attempts to estimate the policy tendencies of agency personnel. This work builds on earlier efforts within political science to examine the role of "ideology"—understood as a consistent set of preferences over policy outcomes—in other decision-making contexts such as Congress and the courts.[121] These ideal point models enable a data-driven method for estimating preferences via construction of a low dimensional latent space that captures similarity among individuals. This same notion of ideology has been applied to agencies and agency personnel, and a variety of methods have been used to estimate this "latent variable," including prior agency decisions, the opinions of outside experts, the political moment of an agency's formation, and the cam-

117 *See* Baekgaard et al., *supra* note 112; John Brehm & Scott Gates, *Donut Shops and Speed Traps: Evaluating Models of Supervision on Police Behavior*, 37 AM. J. POL. SCI. 555 (1993); Sean Gailmard & John W. Patty, *Slackers and Zealots: Civil Service, Policy Discretion, and Bureaucratic Expertise*, 51 AM. J. POL. SCI. 873 (2007); Canice Prendergast, *The Motivation and Bias of Bureaucrats*, 97 AM. ECON. REV. 180 (2007).

118 *See* Louis M. Imbeau et al., *Left-Right Party Ideology and Government Policies: A Meta-Analysis*, 40 EUR. J. POL. RES. 1 (2001) (finding mixed results when examining multiple studies of party influence); Christoph Knill et al., *Do Parties Matter in Internationalised Policy Areas? The Impact of Political Parties on Environmental Policy Outputs in 18 OECD Countries, 1970–2000*, 49 EUR. J. POL. RES. 301 (2010) (finding effects from party turnover); B. Dan Wood & Richard W. Waterman, *The Dynamics of Political Control of the Bureaucracy*, 85 AM. POL. SCI. REV. 801 (1991) (finding that agencies are often responsive to political principals).

119 *See generally* DANIEL CARPENTER, REPUTATION AND POWER: ORGANIZATIONAL IMAGE AND PHARMACEUTICAL REGULATION AT THE FDA (2010) (detailed study of gradual accumulation of power by FDA personnel); DANIEL P. CARPENTER, THE FORGING OF BUREAUCRATIC AUTONOMY: REPUTATIONS, NETWORKS, AND POLICY INNOVATION IN EXECUTIVE AGENCIES, 1862–1928 (2001) (examining origins of modern administrative independence); *see also* Michael A. Livermore, *Cost-Benefit Analysis and Agency Independence*, 81 U. CHI. L. REV. 609 (2014) (examining the role played by agencies in developing the methodology of cost-benefit analysis, the primary substantive standard used to inform executive political oversight).

120 *See* Michael E. Levine & Jennifer L. Forrence, *Regulatory Capture, Public Interest, and the Public Agenda: Toward a Synthesis*, 6 J.L. ECON. & ORG. 167, 172 (1990) (noting that the definition of "public interest" is often contested).

121 *See* KEITH T. POOLE, SPATIAL MODELS OF PARLIAMENTARY VOTING (2005); Andrew D. Martin & Kevin M. Quinn, *Dynamic Ideal Point Estimation via Markov Chain Monte Carlo for the U.S. Supreme Court, 1953–1999*, 10 POL. ANALYSIS 134 (2002).

paign contributions and survey responses of agency personnel.[122] These studies tend to come to relatively consistent results, confirming that agencies have something like an ideological profile that persists over time. Agency ideology has been found to have several interesting consequences for the relationship between political and administrative decisionmaking. For example, Presidents treat agencies differently in their political appointment decision in light of agency ideology; some agencies are targeted for patronage, while other agencies are targeted for more intensive policy supervision, depending on how well agency ideology tends to align with the governing philosophy of the White House.[123]

An important but understudied question in the bureaucratic politics literature is how agency ideology affects interactions between agencies and the broader public. In addition to making official decisions—promulgating rules, issuing licenses, initiating enforcement actions—agencies engage in a variety of public engagement activities, which include not only those that are required by law (such as the notice-and-comment process) but also through voluntary initiatives that include public meetings, publications, media relations, and social networking. These actions speak to the importance of managing public perception for agencies, not only for direct reputational benefits, but also as part of a broader effort to influence the oversight activities of actors that are more directly accountable to the public.

One of the difficulties of studying agency-public interactions is that data and methods have not been as fully developed as in other areas of the bureaucratic politics literature. Some studies have focused on the identity of commenters as the primary explanatory variable to test whether interests that commenters have submitted in comments tend to influence the regulatory process.[124] Survey techniques have also been used to examine how participants perceive their role in the regulatory process.[125] Recently, researchers

122 *See* DAVID E. LEWIS, PRESIDENTS AND THE POLITICS OF AGENCY DESIGN: POLITICAL INSULATION IN THE UNITED STATES GOVERNMENT, 1946–1997 (2003) (political history); Joshua D. Clinton & David E. Lewis, *Expert Opinion, Agency Characteristics, and Agency Preferences*, 16 POL. ANALYSIS 3 (2008) (expert opinion); Joshua D. Clinton et al., *Separated Powers in the United States: The Ideology of Agencies, Presidents, and Congress*, 56 AM. J. POL. SCI. 341 (2012) (employee surveys); David E. Lewis, *Testing Pendleton's Premise: Do Political Appointees Make Worse Bureaucrats?*, 69 J. POL. 1073 (2007) (same); David C. Nixon, *Separation of Powers and Appointee Ideology*, 20 J.L. ECON. & ORG. 438 (2004) (prior decisions); Adam Bonica, Jowei Chen, and Tim Johnson, Automated Methods for Estimating the Political Ideology of Individual Public Bureaucrats Across Time and in a Common Ideological Space (unpublished manuscript) (campaign contributions).

123 *See* DAVID E. LEWIS, THE POLITICS OF PRESIDENTIAL APPOINTMENTS: POLITICAL CONTROL AND BUREAUCRATIC PERFORMANCE (2008); Alex Bolton, Creating Capacity: Presidential Control and the Allocation of Career SES Employees (Sept. 9, 2014) (unpublished manuscript), https://static1.squarespace.com/static/53557985e4b049723f6e96f3/t/540f7d95e4b07fe1be587aeb/1410301333063/ses_july_draft.pdf.

124 *See, e.g.*, Balla, *supra* note 113.

125 *See* Scott R. Furlong & Cornelius M. Kerwin, *Interest Group Participation in Rule Making: A Decade of Change*, 15 J. PUB. ADMIN. RES. & THEORY 353 (2005) (updating an earlier survey-based approach on interest group participants); Susan Webb Yackee, *Participant*

in political science, public administration, and law have begun to exploit public comments to study agency-public interactions.[126] Susan Webb Yackee, in particular, brought attention to the value of public comments in understanding agency-public interactions with several studies that relied on hand coding a large number of comments.[127] Two recent papers published in law journals also engage in human coded analysis of public comments to gain insight into how the public perceives highly salient agency rulemakings.[128]

One underexplored feature of comments is their sentiment. Similar to consumer reviews on websites such as Amazon and Yelp, which express satisfaction or dissatisfaction with a product or service, public comments express agreement (satisfaction) or disagreement (dissatisfaction) with a rulemaking or provisions thereof. The sentiment carried within public comments provides a valuable lens into attitudes and perceptions held by the public on agency decisions. Of course, comments are not submitted by a random sample of the population and should not be thought of as replacing public opinion surveys, which are carefully designed to provide insight into the general public. But the self-selection process itself conveys information—public comments express the views of the interested public, a number that can exceed over a million individuals for mass-comment-volume rules. This large group of interested individuals may be more likely to vote, contribute to or volunteer for political campaigns, or act as opinion leaders within local social networks. Whatever their status vis-à-vis these other political activities, the group of public commenters is sufficiently large that it is a useful object of study in its own right, so extracting and analyzing the sentiment in their comments has the potential to provide worthwhile information for the study of how agencies and the public interact. Although the sentiment within public comments may be valuable, extracting information about that sentiment from the very large and unstructured textual corpus that comprises the responses to a rulemaking proposal is not a trivial task. The next Section discusses the computational text analysis technique of "sentiment analysis" that addresses this difficult challenge.

B. Sentiment Analysis

Sentiment analysis is an application of natural language processing "that analyzes people's opinions, sentiments, appraisals, attitudes, and emotions

---

Voice in the Bureaucratic Policymaking Process, 25 J. PUB. ADMIN. RES. & THEORY 427 (2014) (using a survey to test for whether participants in rulemakings feels sense of "efficacy").

126 Anne E. Boustead & Karlyn D. Stanley, *The Legal and Policy Road Ahead: An Analysis of Public Comments in NHTSA's Vehicle-to-Vehicle Advance Notice of Proposed Rulemaking*, 16 MINN. J.L. SCI. & TECH. 693 (2015); Marissa Martino Golden, *Interest Groups in the Rule-Making Process: Who Participates? Whose Voices Get Heard?*, 8 J. PUB. ADMIN. RES. & THEORY 245 (1998); Krawiec, *supra* note 35; Amy McKay & Susan Webb Yackee, *Interest Group Competition on Federal Agency Rules*, 35 AM. POL. RES. 336 (2007); Yackee, *supra* note 4.

127 *See* McKay & Yackee, *supra* note 126; Yackee, *supra* note 4.

128 *See* Boustead & Stanley, *supra* note 126; Krawiec, *supra* note 35.

toward entities and their attributes expressed in written text."[129] At the heart of sentiment analysis is the concept of *sentiment*, which is a relation between a person and a target.[130] Sentiment is meant to capture a range of related concepts that are in the neighborhood of affect, attitude, or evaluation. Simply and intuitively, the sentiment of *A* toward target *X* is whether *A* likes or dislikes *X*, *A*'s positive or negative evaluation of *X*, *A*'s attitude toward *X*, and the affect brought about in *A* by contact with or contemplation of *X*. Although, in theory, sentiment analysis could distinguish between the many nuanced flavors of these mental and psychological states (say distinguishing between anger and fear), the general tendency in the field (to date at least) has most often been to reduce sentiment to a single dimension between *positive* and *negative* poles. In this tendency toward unidimensionality, there is resonance between sentiment analysis and hedonic utilitarianism, in the sense that the many diverse and plural forms of human experience are ultimately understood along a single dimension of good/bad, pleasurable/painful, or positive/negative.

But sentiment analysis is not a psychological or moral theory; it is a methodology of text analysis—a means of algorithmically extracting meaning from language. Sentiment is, naturally, only one element of the meaning of any string of words. The statements, "this cake is delicious," and "I had a hard time using this can opener," have content beyond their sentiments; but sentiment is a part of the meaning of these statements. Absent other context, a reader can fairly understand these sentences to imply a generally positive sentiment toward "this cake" and a generally negative sentiment toward "this can opener." The abstract question of whether this sentiment resides in the speaker or in the text is generally beside the point. A sentence in a fictional novel can have sentiment, although there is no human being with a corresponding mental or psychological state.[131] Sentiments sometimes matter, however, exactly because there is some ability to map the sentiment expressed by a text to an actual human evaluation, attitude, or affect. For this reason, a variety of private industries are interested in extracting consumer sentiment to inform product development and marketing.[132] Work has also been done to use sentiment analysis to predict box office returns,

129 BING LIU, SENTIMENT ANALYSIS: MINING OPINIONS, SENTIMENTS, AND EMOTIONS 1 (2015).

130 *See id.*

131 *See* Saif Mohammad, *From Once Upon a Time to Happily Ever After: Tracking Emotions in Novels and Fairy Tales*, *in* PROCEEDINGS OF THE 5TH ACL-HLT WORKSHOP ON LANGUAGE TECHNOLOGY FOR CULTURAL HERITAGE, SOCIAL SCIENCES, AND HUMANITIES 105, 105 (2011) (exploring sentiment in novels and fairy tales).

132 One researcher states that, in private consultancy work, he had conducted sentiment analysis for clients in a large variety of consumer product areas, including "automobile, mobile phone, earphone, printer, fridge, washing machine, stove, Blu-ray, laptop, home theater, television, e-book, GPS, LCD monitor, dieting, hair care product, coffee maker, mattress, paint, cruise, restaurant, hotel, cosmetics, [and] fashion." LIU, *supra* note 129, at 6.

the outcomes of political elections, and, increasingly, as a key alternative dataset in financial services for stock market returns.[133]

The difficulty of the sentiment analysis problem and the appropriate technique are closely determined by the degree to which a researcher attempts to disaggregate the targets of sentiment within a single document. The simplest approach examines a single positive/negative dimension and considers a text (or set of texts) jointly, either assuming a single target or examining sentiment toward all targets. In contrast, multi-aspect sentiment analysis tries to capture the various aspects, or targets, within a single document.[134] Both approaches are typically undertaken at the document level. The former analysis reduces all of the text within the document to a single sentiment score, which can be understood as a point in a single dimension sentiment space. Because there is no attempt to determine whether different targets are referenced, this type of analysis might be thought to best capture sentiment as the overall mood of the author.[135]

Although simple word frequencies can perform reasonably well at capturing general mood, further analysis can provide more information. Take, for example, this sentence: "I had a lousy day because my commute was blocked up by a terrible accident and I had to wait around all morning for a boring meeting with my boss." Words like "lousy," "terrible," "accident," and "boring," all have a negative valiance and are not counterbalanced by positive words, and so the generally sour mood can be fairly easily captured with simple word frequencies. More sophisticated forms of sentiment analysis would attempt a more fine-grained analysis by developing enhanced measures of sentiment, extracting multiple targets of sentiments within a document or both. Imagine that the "lousy day" sentence appeared on Twitter with geotagged information. If a traffic monitoring and predicting service wanted to use real time social media information to improve its performance, it would be important to extract sentiment concerning some of the targets in the text (commute, accident) while ignoring the other irrelevant sentiments.

An important component for sentiment analysis at any level is a "sentiment lexicon" that categorizes words according to the sentiment that they

133 *See id.* at 5–7.

134 Bin Lu et al., *Multi-Aspect Sentiment Analysis with Topic Models*, *in* PROCEEDINGS OF THE 2011 IEEE 11TH INTERNATIONAL CONFERENCE ON DATA MINING WORKSHOPS (2011).

135 In a basic informal analysis by Livermore, Riddell, and Rockmore, the authors treat the writings of Justices on the U.S. Supreme Court as documents and examine the overall sentiments of the authors—no attempt is made to extract all of the targets of those sentiments (which might range from facts in individual cases to prior precedent to their fellow Justices). That analysis examines roughly 25,000 documents over the two-hundred-year history of the Court. Each Justice is assigned a sentiment score based on the percentage of positive and negative words of the total sentiment-laden words. Each Justice is placed in time according to his or her median year on the Court. The analysis finds a strong time trend toward increasingly negative sentiment. This highly aggregated form of analysis hides a great deal of detail, but there is no need to tease out the specific targets of sentiments within the opinions to identify the overall time trend. *See* Carlson et al., *supra* note 91.

convey.[136] In a sentiment lexicon, some words will be classified as positive (e.g., wonderful, intelligent, great), while others will be classified as negative (e.g., terrible, stupid, bad). There are two general approaches to developing a sentiment lexicon: the *thesaurus* approach and the *natural corpus* approach. For the thesaurus approach, the researcher relies on existing reference documents, such as dictionaries and thesauruses, and begins with some set of human-identified "seed" words that are obviously positive or negative.[137] Based on those seed words, the thesaurus approach identifies a set of words that are closely linked together in a network of synonyms. A thesaurus approach is typically augmented with some human supervision.[138]

The natural corpus approach begins with some set of documents produced naturally through some generative process, such as Amazon reviews or judicial opinions.[139] When metadata is associated with documents that facilitate classification, supervised learning algorithms provide a simple and straightforward approach to develop a lexicon. For example, for text data that is derived from consumer reviews, there may be numerical rating information provided by consumers (i.e., 1–4 "stars") that can be used as classifications of the documents along a sentiment axis. The ratings can be used as training data for a machine learning approach. The main textual features relied on are word frequencies that can be augmented by parts of speech tags or recognition of other grammar or usage constructs (such as the negation in "not great"). Lacking metadata, a human generated list of seed words can be used to identify other candidate sentiment-laden words by their close co-occurrence with the seed words in documents in the corpus.

There is a great deal of additional complexity to the problem of sentiment analysis, and research continues to focus on the wide range of interesting and difficult computer science and computational linguistic challenges embedded within the field.[140] One obvious example is that sentiment analysis models typically have a difficult time correctly classifying sarcastic statements, a clear problem when analyzing certain corpora, such as social media posts.[141] Although significant progress has been made since its inception in the early 2000s, there are likely to be substantial strides in the future. Nevertheless, existing tools can be deployed right now to useful effect.

---

136 *See generally* Maite Taboada et al., *Lexicon-Based Methods for Sentiment Analysis*, 37 COMPUTATIONAL LINGUISTICS 267, 267–72 (2011) (explaining the methodology for building a sentiment lexicon).

137 *See id.* at 271.

138 *See id.*

139 *See id.*

140 Liu characterizes sentiment analysis as a "mini" version of the entire natural language processing problem, such that "every subproblem of [natural language processing] is also a subproblem of sentiment analysis, and vice versa." LIU, *supra* note 129, at 14.

141 Diana Maynard & Mark A. Greenwood, Who Cares About Sarcastic Tweets? Investigating the Impact of Sarcasm on Sentiment Analysis (unpublished manuscript), https://gate.ac.uk/sale/lrec2014/arcomem/sarcasm.pdf.

C. Ideology and Comment Sentiment

This Section describes the results of an analysis of the content of nearly three million public comments received by federal agencies during the first seven years of the Obama administration. Using widely available sentiment analysis tools, we examine whether the content of public comments changes with the ideology estimates of agencies from the political science literature on bureaucratic politics. Specifically, we hypothesize that agencies with relatively more polarized ideologies will tend to receive comments that are relatively less favorable. As discussed below, we find that there is a relationship between ideological polarity and comment sentiment: the further agencies sit toward the ideological poles, the more they can expect negative sentiment in the comments that they receive.

In a simple model of participation in the public comment process, a proposed regulation can be understood as representing a point within a one-dimensional, left-right ideological space. A liberal proposed regulation (for example, one that increases regulation of financial products) would occupy a position to the left of a conservative proposed regulation (for example, deregulating financial products). Both agencies and commenters in this model have "ideal points" within this ideological space that represent their preferred policy outcomes.[142] Agencies will tend to propose regulations that are close to their ideal point. The distance between a commenter's ideal point and the proposed regulation affects the sentiment of the comment. A proposed regulation that is relatively close to a commenter's ideal point will spur positive sentiment, and proposals that are relatively distant will generate more negative sentiment. The likelihood of commenting may also be related to the distance between a potential commenter's ideal points and proposed regulations, as well as other factors, such as the salience of the underlying issues affected by a rulemaking. Individuals are more likely to comment on proposals they strongly favor or disfavor and on rulemakings that involve high-salience issues.

This model of commenting involves strong assumptions about the nature of ideological space and the motivations of commenters. In reality, people's ideological dispositions may vary on multiple dimensions and be poorly represented by a simple left-right spectrum. Even on a single dimension, we are unable to directly observe the distribution of ideal points of potential commenters, which may be relatively normally distributed and clustered toward the center or, alternatively, may be bimodal, with relatively few moderates and separate groupings oriented toward the extremes or even uniformly distributed across the space. We also have little information on how individuals learn about regulatory proposals and decide to comment, and what mix of salience, satisfaction with the policy choices made by an agency, and other factors influence commenting behavior. As discussed in the prior

142 For background on ideal point estimation as an empirical undertaking, see David A. Bateman & John Lapinski, *Ideal Points and American Political Development: Beyond DW-NOMINATE*, 30 STUD. AM. POL. DEV. 147 (2016).

Section, sentiment itself is complex and may involve more than simple positive or negative affect.

What simple models lack in nuance, sophistication, and complexity, they can sometimes offset with parsimony. Despite the limitations in the micro-level account of commenting behavior offered above, it can nonetheless help motivate some macro-level predictions about the relationship of sentiment and the ideological characteristics of agencies. At the aggregate level, certain features of comments may be predictable, even if the behavior of individual commenters is highly stochastic, unchanging, or affected by a wide range of unobserved variables.[143] The macro-level prediction that we test is that agencies with more moderate ideological tendencies will receive comments with relatively more positive sentiment. This result is consistent with the simplified model in which personnel ideology is correlated with the likely policy choices made by agencies, the desire to comment is relatively uniform across the ideological spectrum, and there is some centralized tendency to the distribution of ideal points within the pool of potential commenters.

Given the number of potential variables in play, there are other, more micro-level accounts that would also be consistent with a positive correlation between ideological moderation and sentiment. For example, under a bimodal distribution, moderate proposals might be distant from all potential commenters' ideal points, but many unhappy potential commenters may be demotivated by milquetoast proposals and decide not to comment at all. There may be other causal stories as well, such as the possibility that negative sentiment in comments causes ideologically moderate individuals to avoid seeking work at an agency. If this is the case, the location of regulatory proposals in ideological space does not cause negative sentiment in comments; rather, some other variable causes negative sentiment (e.g., an unpopular mission), which then affects the type of personnel attracted to an agency, which in turn affects ideology estimates. Or there may be unobserved variables—such as the influence of a special interest group or congressional oversight committee—that affect both agency ideology and comment sentiment.

The current analysis will not be able to distinguish between these alternatives. Rather, we offer a simple account, investigate whether there are empirical associations in the data that are consistent with that account, and open the door for additional research to develop and test alternative theories of commenter behavior. Given the relative newness of the computational techniques deployed here and their potential to grow in sophistication in coming years, the value of this analysis is its ability to offer observations and prompt questions that form the basis for future work. After discussing our data and results, we will return at the end of this Section to potential paths that this work may take.

143 *Cf.* ROBERT S. ERIKSON ET AL., THE MACRO POLITY 6 (2002) ("When we aggregate over time, those who act as if at random cancel out. Those who act always the same produce no variance. The aggregate 'signal' arises almost wholly from those who are orderly in their behavior.").

For our analysis, we exploit the extensive dataset of public comments compiled by FiscalNote, a Washington D.C.-based government analytics firm. FiscalNote scraped all publicly available comments for all agencies during the study period from "regulations.gov" whenever possible and from the individual agencies' websites otherwise. It is worth noting that not all comments are publicly released; for example, agencies sometimes release only unique comments.[144] Every comment that was available was analyzed. Our study period is the first seven years of the Obama administration. During that time, several million public comments were received and released in response to solicitations by U.S. federal administrative agencies concerning pending actions (primarily, but not exclusively, rulemakings). Our analysis is based on 1445 rules adopted during the study period for which one hundred or more comments were received. The data covers 99 administrative agencies. The three rules with the largest number of comments during the study period were a rule by the Department of Health and Human Services implementing provisions of the Affordable Care Act, an Internal Revenue Service rule concerning candidate-related political activities by tax-exempt organizations, and a State Department permit application concerning the TransCanada Keystone XL Pipeline.

Our predictor variable of interest is agency ideology. We will test the relationship between ideology and sentiment with the basic hypothesis that agencies with more "extreme" ideologies will receive public comments with lower mean sentiment. To construct the ideology variable, we rely on a measure based on responses to a 2007–2008 survey conducted as part of the *Survey on the Future of Government Service* project.[145] That survey was sent to over seven thousand senior level officials within the federal government, both career and political appointees. Completed surveys were received from over two thousand respondents (with an overall response rate of thirty-four percent). To develop an ideology measure that was commensurable with Congress, the survey included several questions on how the respondent would have voted on fourteen questions that were subject to House and Senate votes in 2006 (including the confirmation of Justice Alito and a bill to increase the minimum wage). Using congressional actions allowed the researchers to compare the hypothetical votes of survey respondents to the actual votes of Senators and Representatives, placing them in common ideological space.

The *Survey on the Future of Government Service* researchers reported separate ideology scores for career officials and appointees; for our analysis we will use the ideologies for the careerists only, disregarding the political appointees. Given the relatively low level of turnover within the ranks of career civil servants, a greater degree of interadministration ideological consistency can be expected. Use of only the career respondents may also allow

144 For example, despite stating that it has received more than four million comments in response to the proposed Clean Power Plan, EPA released fewer than 50,000 comments.

145 Clinton et al., *supra* note 122. We use data from an earlier working paper version that reported more detail on the survey responses.

us to examine more persistent agency-level effects that are less likely to be subject to variation as the White House changes between Presidents and parties.

The FiscalNote data includes fairly granular specification of the issuing agency, which is more finely grained than the estimates of agency ideology. In those cases, we identify the "parent" agency and tag an ideology score accordingly. For example, the Fish and Wildlife Service is a bureau within the Department of the Interior—since we have no specific ideology score for the Fish and Wildlife Service, all of its rules are attributed to Interior. This lack of granularity adds a fair amount of noise to our analysis; it is quite possible that the ideology of personnel associated with the Fish and Wildlife Service (which is charged with administering the Endangered Species Act) is different than personnel at the Bureau of Ocean Energy Management (which administers offshore oil drilling), even though both fall within the Department of the Interior. The coarseness of our ideology measure can be expected to attenuate our results, rather than lead to spurious correlations.

Our ideology variable is not simply a location on the left-right spectrum. Rather, we estimate the degree to which an agency departs from the middle of the ideological distribution. To do this, we calculate an ideological midpoint by taking the mean ideology for agencies in our sample and then using the distance between that midpoint and an agency's score as an estimate of the degree of ideological polarity.

We analyzed the sentiment of all of the comments in our dataset using the publicly available sentiment analysis model in the pattern Python library.[146] The text was preprocessed using TextBlob and NLTK to tokenize, stem, and remove stop words.[147] Every comment is given a sentiment score based on the occurrence of lexical items indicating positive and negative polarity within the document; these comment sentiment scores are then categorized by the corresponding rule to create a per rule distribution, and a mean sentiment value is determined for every rule. This mean sentiment is used as the dependent variable in our analysis. For all rules in our dataset, the distribution of mean sentiment is roughly normal.

There are eighteen agencies in our dataset for which we have both ideological scores and rules. Many of these agencies have bureaus, departments, or offices that are not represented separately. Figure 1 presents the relationship between agency ideological polarity and mean sentiment. Recall that each agency's ideological score is not a location in left-right ideological space, but rather a measure of the distance between that agency's ideology and the midpoint in the space. Most of the agencies have a unique score, but

146 Tom De Smedt & Walter Daelemans, *Pattern for Python*, 13 J. MACHINE LEARNING RES. 2063 (2012).

147 We evaluated the hypothesis using results obtained from a more sophisticated proprietary sentiment analysis model, and they were consistent with those obtained from the simpler model offered by pattern; thus we use the latter in further discussion for reproducibility.

four of the agencies share ideological scores with each other (there are sixteen unique scores for eighteen agencies).

FIGURE 1. Agency polarity and comment sentiment

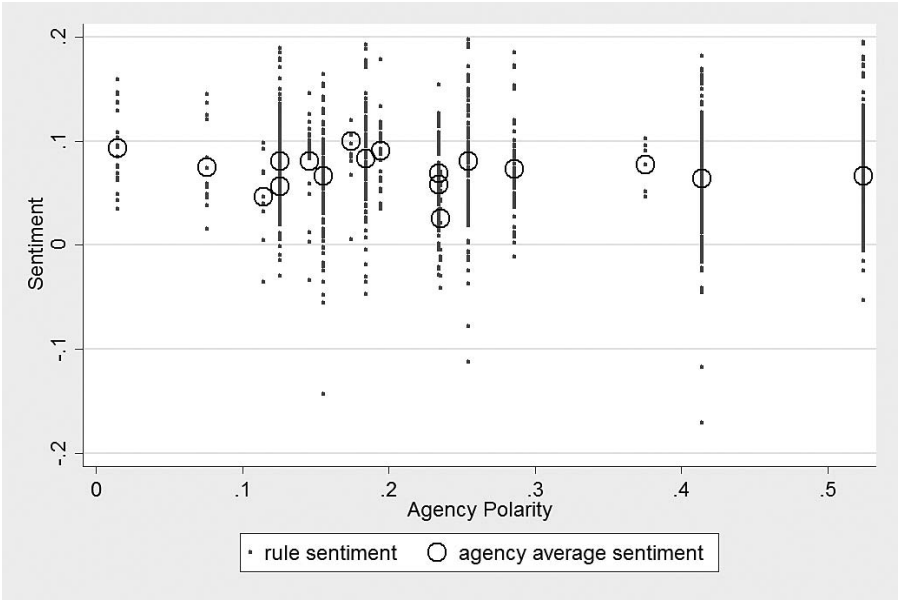

Note: The horizontal axis presents a measure of agency polarity that is calculated by taking the absolute difference between an agency's ideology and the mean ideology. Sentiment falls on the vertical axis, with the scores representing the percentage of positive words minus the percentage of negative words. Each dot represents a single rulemaking. Circles are the sentiment scores aggregated by agency.[148]

Visually, it is clear from Figure 1 that there is a substantial amount of variation at the agency level; the same agencies issue rules that generate comments with very different levels of positive or negative sentiment. Because ideology is assigned at the agency level, its usefulness in predicting sentiment at the rule level is relatively low and at first glance it is not clear that there is a relationship at all. The large number of rules in our sample, however, allows us to detect even fairly weak signals in this noisy data. When sentiment is averaged at the agency level, the relationship between ideology and sentiment becomes somewhat more visually apparent.

To examine the size of the relationship and test for significance, we estimate a linear model in which sentiment is treated as a dependent variable and agency ideological polarity is treated as a predictor variable, with controls for the number of comments (which may be correlated with both agency ideology and sentiment) and year-fixed effects. An ordinary least

148 Seventeen rules with sentiment scores that fell outside the range -0.2–0.2 were excluded from Figure 1 for purposes of presentation; they were included in the statistical analysis.

squares regression shows that agency ideological polarity is significant at the 1% threshold with a coefficient of -0.03 and an adjusted R-squared of 2%.[149] This coefficient can be interpreted to mean that agencies with highly moderate ideologies can expect, other things being equal, comments that include roughly 1.5% more net positive words compared to agencies at the extreme end of our sample.[150] The relatively low R-squared value is consistent with the limited predictive power of ideology for sentiment at the rule level. We also estimate a separate model in which agency average sentiment replaces rule sentiment with the same controls for total comments received and year-fixed effects. In that model, ideological polarity is again significant at the 1% threshold and R-squared increases to 21%.

To check whether any specific agency accounts for the relationship, we conducted an additional round of analysis. We first analyzed whether EPA rules were particularly controversial, adding an indicator for whether a rule was issued by the EPA; there was no change in the results.[151] We then constructed a new model with the issuing agencies as indicator variables and identified seven agencies with significant correlations with mean sentiment (Commerce, Education, Interior, Transportation, Housing and Urban Development, Nuclear Regulatory Commission, and State). We gave each of these agencies the same treatment as the EPA, adding them individually to the third model. In none of these specifications did the significance of agency polarity fall below the 1% level.[152]

The upshot of this analysis is that there is a relationship between agency ideology and the sentiment that is expressed in comments to the agency. More ideologically moderate agencies tend to receive comments that are, on average, more positive. It is worth emphasizing that these results are detectable even though the measures of both sentiment and agency ideology are quite noisy. The sentiment analysis scores are based on publicly available technology that is not tailored to the regulatory context, and the document level aggregation obscures the content of the sentiment within the comments. Agency ideology is assigned at a very high level on the hierarchal ladder within the executive branch, which misses out on the potentially substantial ideological differences between bureaus and offices within agencies. Nonetheless, in the aggregate, there appears to be an identifiable relationship.

---

149 The total number of comments received is also significant at the 1% threshold, with an increase of 0.0005 in sentiment per 1000 comments. We also estimate alternative models without year-fixed effects and with neither year-fixed effects nor total number of comments—the coefficient size for ideological polarity and significance remain the same in these models, with the R-squared values declining somewhat.

150 For this explanation, the extreme end of the sample is taken to be an agency with a polarity score of 0.5.

151 The EPA indicator variable was not significant. Agency polarity remained significant at the $p<0.01$ level.

152 The agencies with a significant effect ($p<0.10$) were: Education (positive coefficient); Interior (positive coefficient); and the Nuclear Regulatory Commission (negative coefficient).

This result is consistent with the simple model of commenting introduced above, but as discussed, there are also other plausible stories that might account for this relationship—a fuller reckoning with the underlying causal mechanisms will require further research. Our limited understanding of how agency ideology interacts with the comment process also reduces our ability to draw normative conclusions. Under a very strong Goodnowian understanding of agencies, in which politics should be kept entirely separate from administration, our result would raise red flags.[153] But on a more nuanced understanding that allows for some interaction of politics and administration, it is at least potentially unproblematic for agencies to occupy different places along the ideological spectrum and, accordingly, generate differing public responses to their rules.

Certainly, it would be overly hasty to conclude from our analysis that the more ideologically distinctive agencies are poorly performing their regulatory tasks. It is not obvious that agencies should seek to avoid negative sentiment and encourage positive sentiment—unlike in the case of consumer products, immediate satisfaction may simply not be the best interpretation of agency performance. Furthermore, there is no reason to think that the comments that are submitted to agencies are representative of the views of the American public in general; they may well better reflect the interests of organized groups or ideological activists, in which case negative sentiment from commenters may be associated with rules that better protect the public interest.

Although both the empirical and normative conclusions that can be drawn from our results are murky at this stage, the potential directions for future research are clearer. We propose three research areas that are raised by the preceding analysis, that could be addressed by leveraging existing data sources, and that have the potential to generate a richer understanding of the public comment process and interactions between public bureaucracies and the public more generally.

The first research area would be to develop better measures of comment characteristics as well as measures of rule characteristics based on regulatory texts. The basic sentiment measure used in our analysis is very rough and is not tailored to the corpus or question of interest. Better measures could capture more relevant features of comments that are estimated with greater precision. In addition, we do not analyze the content of regulatory texts and relate that content to agency ideology or comment content. Some simple analyses of regulatory texts have been carried out,[154] but more sophisticated measures that reveal the policy content or ideological valence of regulatory text could serve as the basis for more nuanced analysis of the relationship between proposed regulations and the comments they generate.

153 *See* GOODNOW, *supra* note 104, at 83–84.

154 *See* Omar Al-Ubaydli & Patrick A. McLaughlin, *RegData: A Numerical Database on Industry-Specific Regulations for all United States Industries and Federal Regulations, 1997–2012*, 11 REG. & GOVERNANCE 109 (2017); Stiglitz, *supra* note 97.

The second research area concerns commenter ideology. In the analysis above, identifying information was not used and commenters were treated as an undifferentiated mass. As with other public filings, however, commenter identity is not hidden, and can, in theory, be extracted from comment texts. Generally speaking, there are two types of commenters: organizations and individuals. Once commenter identity is revealed, that information can be used to categorize comments by the authoring organization or individual and can be linked via that information to other sources of data, such as campaign donation information for individuals, or lobbying filings for organizations. This information can provide insight into the distribution of ideological tendencies within the pool of commenters and how commenter ideology relates to agency ideology and comment characteristics.

A third research area would involve interactions of comment characteristics with political oversight. Our analysis is limited to the Obama presidency, and so does not capture variation in the party that occupies the White House. As the Trump administration unfolds, this new data will become available. Nor do we mine information on congressional oversight activities or the institutional design characteristics of agencies (such as whether they submit rules to OIRA for clearance) to examine how oversight from political bodies might affect regulatory text or the content of comments. Data from media sources might also be used to estimate the public salience of regulatory matters, which may interact with political oversight or affect the propensity of interested individuals to comment.

Overall, the initial exploratory analysis presented in this Section indicates that there is substantial potential to make use of public comments to understand the relationship between agencies and the public. Of course, as with any source of data, public comments do not provide a complete picture, but they nonetheless are a valuable lens into the thoughts, concerns, and reactions of interested parties to the proposed actions of administrative agencies. As new generations of computational text analysis tools become available, they have the potential to provide substantial new insights that challenge existing understandings and create opportunities for new avenues of research that enrich and deepen our understanding of questions at the heart of the social scientific study of government.

III. Enhancing Participation

The explosion of public participation in the rulemaking process and the growth of computational text analysis tools present opportunities not only for researchers, but also for agencies. As discussed in Section I.B, high-volume-comment rules create challenges for agencies, and there is some controversy about how, and even whether, the large mass of unsophisticated comments should be taken into account when setting regulatory policy. Those controversies notwithstanding, the fact of the matter is that megaparticipation, where agencies receive more than a million comments on their proposed actions, is now commonplace. This deluge of comments creates practical difficulties, but just as technology has helped create these difficulties, techno-

logical tools may help agencies address them. This Part begins by discussing in more detail the challenges faced by agencies in the era of mass commenting and then examines how two computational text analysis tools can be used to address those challenges.

A. Haystacks and Forests

The most salient risk posed to agencies when faced with a large number of comments is the failure to identify and respond to substantive comments that are subsequently used as the basis for litigation. Courts enforce the APA's requirement for agencies to collect and consider comments, occasionally striking down rules when agencies fail to adequately respond to comments on a proposed rule.[155] Agencies are not under a general obligation to "discuss every item of fact or opinion" offered in a comment,[156] but they are required to include in the record sufficient response to comments to show a reviewing court that the "major issues of policy were ventilated" during the rule making deliberations.[157]

Regardless of controversies over how to respond to less sophisticated mass comments, there is broad consensus that *highly substantive comments* that contain information of obvious output value for agencies ought to be appropriately considered. Courts have emphasized the importance for agencies to respond to comments that raise relevant scientific or technical information, illuminate undesirable consequences of the proposed rule, or offer alternative courses of action for the agency to consider.[158] A judicial reversal imposes major costs. Agencies spend considerable resources on rule making, including personnel hours and funds to hire consultants.[159] The most significant category of rulemaking costs, however, may be political. Each agency's rulemaking agenda is extremely constrained, with the potential to issue only a small number of rulemakings each presidential term. The opportunity cost for any rulemaking includes all of the other policies that could have been pursued. Agencies also expend general political capital during rulemakings, inevitably courting controversy and opposition even in contexts where regulation is broadly popular. Given the costly investment that a major rulemaking represents, when a rule is struck down it is an enormous disappointment for the agency. When agencies fail to account for highly substantive comments, not only do they undermine their own decision-making process, they also open themselves up to these litigation risks.

---

155 *See, e.g.*, Del. Dep't of Nat. Res. & Envtl. Control v. EPA, 785 F.3d 1, 15–16 (D.C. Cir. 2015) (striking down an EPA rule as arbitrary and capricious for failing to appropriately grapple with commenter objections).

156 Auto. Parts & Accessories Ass'n v. Boyd, 407 F.2d 330, 338 (D.C. Cir. 1968).

157 *Id.*

158 *See, e.g.*, Pub. Citizen, Inc. v. FAA, 988 F.2d 186, 197 (D.C. Cir. 1993).

159 *See Background on Rulemaking Process*, CTR. FOR EFFECTIVE GOV'T, https://www.foref fectivegov.org/sites/default/files/regs/rcenter/backgroundpdfs/IV.Rulemaking.pdf (last visited Oct. 30, 2017).

Although agencies have both internal reasons and external incentives to identify highly substantive comments, the low cost of submitting comments means that more substantive submissions can be buried under a mountain of less substantive comments. We refer to this challenge as the *haystack problem*—agencies must find the proverbial needle in the haystack. Substantive comments must be unearthed from within a large number of comments that are highly unlikely to pose any litigation risk. Failure means that the agency loses the opportunity to improve its rulemaking or head off a judicial challenge. When there are a large number of unique comments of a less substantive nature, and a relatively small number of sophisticated comments, the task becomes more difficult. Agencies that find themselves inundated with tens of thousands, or even millions, of comments have an especially difficult chore.

There is some similarity between the haystack problem faced by agencies and the difficulties associated with discovery in large scale litigation.[160] The recipient of a massive discovery disclosure can face a vast amount of information, and the goal is to comb through and identify material that may be relevant to the ongoing litigation. On the other side of the case, the party subject to the discovery request must identify material that must be disclosed while screening for documents that do not, or may not, be disclosed. In large litigation, discovery can impose substantial costs as plaintiffs and defendants employ extremely expensive techniques—including paying large numbers of junior attorneys—to comply with the discovery process. The move to digital documents has transformed the process of discovery, leading to the burgeoning field of electronic discovery in which computational algorithms now play a substantial role.[161]

Agencies have developed various practices to screen for sophisticated comments. They hire consultants who are charged with combing through the comments, identifying salient arguments, and constructing summaries for policymakers. These efforts, which involve human readers, are resource intensive. They are also potentially error prone because it would be inordinately expensive to hire seasoned experts to carry out this task; there is little choice but to rely on the work of relatively inexperienced individuals who lack substantial professional training. Because the current process is expensive and error prone, there are opportunities for technological tools to reduce costs and improve performance.

The second challenge agencies face is how best to respond to the less sophisticated comments offered on behalf of small stakeholders or interested individuals. These are not the technocratic, jargon-laden comments that are

160 *See generally* SHIRA A. SCHEINDLIN ET AL., ELECTRONIC DISCOVERY AND DIGITAL EVIDENCE (2d ed. 2012).

161 *See id.* The analogies to e-discovery should not be overstated. Perhaps most importantly, agency officials and commenters have very different incentives when examining comments than do litigants: commenters are not trying to hide their views or intentionally bury agencies, and, although defendants must use reasonable means to comply with discovery demands, they seek mere compliance, not a perfect retrieval rate.

submitted by large law firms on behalf of major industry, but are nonetheless genuine expressions of concern or support. The welfare stakes can be quite real for the small business owner who faces a rule that would cut into already slim profits or the parent commenting on an air quality rule out of concern for a child with asthma. Although there is some disagreement among experts about the information value of these *informal comments*, the reality is that they are regularly submitted, and they may, if properly analyzed, provide insights that are of use to regulators.

This second challenge is much greater than the haystack problem because it is not just a matter of screening out the informal comments. Since the information in these comments is diffused, there has to be some mechanism to distill and aggregate meaning that might be spread over many tens of thousands of comments. It is relatively easy to sit down and extract the content from a set of well-researched, persuasively written comments that have been prepared by professionals whose expertise is exactly in communicating to agency officials. It is an altogether different task to face perhaps thirty thousand comments submitted by individuals from a huge range of backgrounds—many of whom have no experience communicating to agencies, lack familiarity with the governing jargon, and have little time to devote to researching the relevant issues—and attempt to extract any kind of collective meaning from those.

This difficult task we refer to as the *forest problem*. For a major rulemaking with tens of thousands of unique comments, it is impossible for any person to gain a sufficient vantage point from which to view the entirety. Any single person can read only a small share, and attempt to summarize as best as possible; others may read the summaries and try to detect broader trends, but as the level of resolution decreases, important details can blur.[162] The limits of human cognition (not to mention hours in the day) require a trade-off, and the process of communication between the fine-grained level of individual comments and higher-order meaning (trends within the group of comments) is highly imperfect.

The technique of sentiment analysis, introduced in Part II, provides an example of how computational text analysis tools can be brought to bear on the haystack and forest problems. With respect to the first, agencies have information on which provision within a rulemaking might be particularly vulnerable to challenge. With more sophisticated parsing of comments, the agency could attempt to identify all of the comments that include that provision as a target of sentiment, and then search the positive-sentiment comments for supporting information and probe the negative-sentiment comments for arguments that may reappear in litigation. An agency might also attempt to use intracomment parsing to identify whether there is an

162 For example, the analysis of the public comments on the Volcker Rule in Krawiec, *supra* note 35, is based on a hand-coding exercise, which typically must start with a set of preexisting variables for coders to identify, and has less room for spontaneous observation from the comments themselves. *See* Krawiec, *supra* note 35, at 72.

alternative to the provision that is proposed with sufficient clarity that the agency might want to consider or be called to respond to.

The forest problem is perhaps an even more promising application of sentiment analysis. Here, the goal for the agency is not to attempt to identify individual documents that require additional scrutiny, but to find overall trends within the comments. For this task, sentiment analysis could be done at the document level or, with greater sophistication, within individual comments. The document-level approach would assign a single sentiment score to determine the overall attitude that the commentator has taken toward the rule. Based on an aggregation of this data, the agency could examine the distribution of sentiment toward the rule for various features: Is the central tendency toward positive or negative sentiment? Does sentiment have a high or low variance, or display skew in one direction or another, or display fat tails at either extreme? These aggregate measures of sentiment would provide agencies with a new lens into how the public is responding to its rulemaking, including whether the weight of sentiment is for or against the agency, and the degree of uniformity of responses.

With more sophisticated analysis, agencies could delve deeper into the higher order content within the comments. Agencies could develop distributions of sentiments over many features and provisions of a rule, some of which the agency may be able to pre-identify, and others that might be identified in the first instance by commentators. Every element of a rulemaking, which could include not only the substantive requirements, but also the agency's justification and general characteristics of the rulemaking (such as the length or level of complexity) could potentially serve as a target of sentiment in the comments. Through this aggregated analysis, the agency may be able to identify discussion of specific regulatory effects that it might have missed. For example, issues that are not at the center of any comment, but which appear across a wide range of comments, may be difficult to identify by human evaluators because the degree of shared concern may only become apparent when a very large number of comments are digested.

Moving beyond sentiment analysis, there are a range of other computational text analysis techniques that could help agencies address the haystack and forest problems. The following two Sections explore some possibilities.

B. Gravitas

At the heart of the haystack problem is the need to identify the most useful comments within a large, unstructured corpus of documents. In essence, agencies would like to be able to separate comments that have substantive weight, or *gravitas*, from the mass of less substantial informal comments. Under Farina et al.'s account, agencies can focus exclusively on the comments with greater gravitas,[163] while under Mendelson's, agencies should analyze even the less substantial comments.[164] Regardless, the treat-

163 *See* Farina et al., *supra* note 5.

164 *See* Mendelson, *supra* note 50.

ment given to comments will vary depending on their gravitas—even if informal comments are not ignored altogether, the types of analysis that they will be subject to, and the information that will be extracted, will be very different. When a major law firm representing a multinational company submits comments to an agency on a proposed rule, agency lawyers must comb through the document to identify potential legal lines of attack, and economic, engineering, and policy personnel will examine the comments for data or arguments that inform potential changes in the rule. The information in informal comments will often be of a very different sort, and is more likely to bear on issues such as risk communication, framing, or the political interpretation of the proposal.

There are some obvious approaches to addressing the haystack problem that do not require advanced computational techniques. For example, comments that are filed by major companies or organizations, or their representatives, are likely to receive attention. The comments of established experts may also be more likely to be flagged for special treatment. Other relevant features of comments that can be identified fairly easily using basic computational tools include the length of the submission, the sophistication of the vocabulary, the relevance of the comment to the rule (perhaps based on shared vocabulary or phrases), the amount of citation, and basic identifying information on the author (such as whether the comment was from an institution or an individual, or whether the commenter participated in other rulemakings).

We conduct an initial assessment of the viability of this technique using comments from a recent rulemaking by the EPA to reduce greenhouse gas emissions from existing electricity-generating units. This rule is one of the most consequential environmental regulations in the nation's history, and the number of public comments received by the agency reflected this historical status of the rule: the EPA received over 4.3 million comments. Given this massive volume of comments, it is pretty clear that the agency cannot give each one a great deal of individual attention, and so some automated means of directing the agencies focus would be particularly helpful in this context.

A first step is simply to remove duplicate comments. EPA appears to have done so; although the agency reports over four million comments received, less than one percent (34,388) are posted by the agency online.[165] Even with the duplicates removed, however, there are still tens of thousands of comments for the agency to review, and so additional steps to prioritize responses and categorize comments remain helpful.

165 Regulations.gov specifically states that "some agencies may choose to redact, or withhold, certain submissions (or portions thereof) such as those containing private or proprietary information, inappropriate language, or duplicate/near duplicate examples of a mass-mail campaign." *See, e.g.*, *Agency Information Collection; Proposals, Submissions, and Approvals*, REGULATIONS.GOV (Oct. 2, 2017), https://test.regulations.gov/document?D=NHTSA_FRDOC_0001-1837. The latter category is by far the most likely, although it wouldn't be a huge surprise if some of the comments sprinkled some occasional inappropriate language.

The following analysis relies on a gravitas measure developed by FiscalNote. This measure ranks comments based on several identifying features including comment length, attachment count, the complexity (or coarseness) of the language that is used, whether the author is an organization, key person, or ordinary individual, the number of cogent arguments expressed, and other cues that together serve as a proxy for sophistication. According to the FiscalNote gravitas measure, the single most sophisticated comment was submitted by the environmental organizations Sierra Club and Earthjustice. This comment was truly behemoth, clocking in at 274 pages with an additional seventy-three exhibits (an example, "The History of Energy Efficiency") and thirteen appendices (including "Literature Survey, Efficiency Improvements through Upgrades of Existing Plants"). Other gravitas leaders include the environmental group Center for Biological Diversity and the Ameren Corporation (an energy holding company that serves 2.4 million electric and nearly 1 million natural gas customers).[166]

Some other points on the gravitas spectrum give a sense of the wide diversity of comments received by the agency. At the top of the bottom quintile of the FiscalNote measure are two short comments sent by individuals, one of which appears to be a form comment:

> Please reduce carbon pollution from existing power plants to protect public health. Set strict limits on carbon pollution from power plants. (forwarded via the American Lung Association)
>
> Thank you, thank you for your courage and foresight. Change must happen immediately to save the planet.

At the top of the second quintile there are two short comments from individuals that remain fairly unsophisticated:

> This sounds good. Up front at least. It's been a while since this country has passed any major environmental movements, and I definitely agree with one like this.
>
> [S]upport american energy and not the marxist agenda.

The comment with the median gravitas score was submitted by an individual and reads as though it might be a mixture of a form comment along with some additional personalized commentary:

> I strongly support the EPA's effort to limit industrial carbon pollution from existing power plants. These new clean air standards will protect public health, fight climate change, and create jobs through innovation in cleaner, safer energy technology. It's our obligation to protect our children and future generations from the effects of climate change—and that means moving forward with these clean air protections now. I have been screaming about this for 50 years! What the hell are you waiting for, the end of life on earth?

At the top of the third quintile is the following:

166 *Ameren Facts and Figures*, AMEREN, https://www.ameren.com/about/facts (last visited Oct. 27, 2017).

> It is absolutely imperative that this plan be enacted. As an American citizen, I am ashamed by the near-total lack of action on climate change in this country. Climate change is a very real threat to the health and well-being of my children and their future children. If we do not take decisive action now, there will be no world for my grandchildren to inherit. This plan makes economic sense and promotes public health and national security. It would be a complete disgrace if this plan was not enacted because of the short-sighted special interests of a small group of industry executives and politicians that are beholden to the coal industry. It is time for the United States to begin taking actions to reduce our disproportionate share of global greenhouse gas emissions.

At the top of the fourth quintile is a comment that may be, or have elements drawn from, a form comment:

> Please keep in mind all of us who support clean renewable energy, which is very do-able in the USA. We vote and we spread information. Along with all 2.4 million Sierra Club members and supporters, I want to see strong positive climate action. I'm encouraged with the framework put forward by EPA's proposed safeguards against carbon pollution from existing power plants but want to see the standard strengthened. We need states to be encouraged to choose clean energy and energy efficiency. It's not the appropriate time to build more dangerous nuclear power plants[ ] and invade our fragile water supplies with fracking. We must dedicate resources to create more productive jobs in an industry that doesn't pollute our air and water or disrupt the climate. (forwarded via the Sierra Club)

It is, of course, the most sophisticated comments that are likely to receive the most agency attention. Even focusing at the top, there are still a fairly large number to consider. For example, the comment with the one hundredth highest gravitas score was from a group called the Small Business & Entrepreneurship Council, a northern Virginia–based "nonpartisan advocacy and research organization."[167] These comments are fairly long at 2300 words and include comparatively sophisticated arguments in support of the three criticisms of the rule summarized by the group as:

> After a thorough review of the CPP, [we] believe the wisest course of action is for EPA to withdraw the proposed rule and abandon its costly agenda to regulate carbon dioxide under the Clean Air Act.
>
> The reasons for our position are straightforward. EPA's proposed rule:
> • Is illegal, stretching far beyond the narrow boundaries of Section 111(d) of the Clean Air Act;
> • Imposes high costs for no meaningful benefits. Electricity will be more expensive for small business owners and entrepreneurs, which will slow economic growth, harm competitiveness, and destroy jobs. There will be no effect on global temperatures and climate change; and
> • Threatens the reliability of the nation's bulk electric power system, which raises the prospect of blackouts and brownouts, which can in turn increase operating expenses and uncertainty, as well as reduce output and revenues.

167 *Our Mission and History*, SMALL BUS. & ENTREPRENEURSHIP COUNCIL, http://sbecouncil.org/about-us/mission/ (last visited Oct. 30, 2017).

The one thousandth most sophisticated comment was submitted by the West Virginia Community Action Partnership, a "statewide membership association" of local community groups that serve over 40,000 "low-income families annually in all of our state's disadvantaged communities, providing [them] with multiple services to promote their economic security."[168] These comments clocked in at 1600 words and offered support as well as detailed criticism for the rule. The group summarized its view follows:

> We are in strong support of the general goal and ultimate result of the 111[d] rule. We are enthusiastic about a framework that allows a building block of efficiency and clean energy policies as part of the solution. However, we believe EPA has failed to establish:
> 1) The needed framework for participatory planning;
> 2) The criteria for establishing that each state's plan is equitable with respect to access to clean energy;
> 3) Protection from environmental degradation; and,
> 4) Protection of human health.

Any automated measure will sometime fail, and one form comment broke the top five most sophisticated because it was delivered in bulk as a set of images (rather than individually as text) creating a very large file. The comment, which appears to have been individually hand signed by over 30,000 electricity customers in South Carolina, urged the agency to give "equitable treatment for [South Carolina's] nuclear units under construction." But even if imperfect, the ability to conduct a rough triage using automated tools can have substantial savings. Assuming that it would take an employee ten seconds per comment to group the comments into rough categories, the human resource cost to conduct this simple operation on all of the comments received on the Clean Power Plan would be over ten thousand person hours—a considerable return.

Agencies also have access to information that could build on the simple approach used above. FiscalNote's gravitas measure is based on a handful of very easily identifiable characteristics that are used to generate an intuitive measure: longer comments are more substantial; comments with more sophisticated language are more substantial. These relationships make sense, but they are not independently verified. Agencies, on the other hand, have access to metadata that could be used to validate a constructed measure, or perhaps even more promising, to train a supervised model for classification. Specifically, agencies have data from past rulemakings on which comments were considered serious enough to warrant a substantive response—either in terms of revision to a proposed rule, a response in a preamble, or response-to-comments document. In essence, the agencies have "tagged" certain documents as worthy of deeper consideration. This information could be used to generate a training set for a supervised machine-learning experiment. It is possible that a properly constructed and trained algorithm could predict, with a fair degree of accuracy, those comments that are most

168 *About Us*, W. VA. COMMUNITY ACTION PARTNERSHIP, http://www.wvcommunityactionpartnership.org/wvcap.html (last visited Oct. 30, 2017).

likely to warrant additional attention. At the very least, such an approach could take a first cut to identify high-substance comments, with the remainder subjected to the existing human evaluation procedure.

The preceding discussion focused on tasks that are now done by human reviewers that could be made less resource intensive or error prone through the use of computational tools. The following Section discusses how big data techniques could create entirely new ways to understand and engage in public deliberation over agency decisionmaking.

C. Identifying Emergent Meaning

At the heart of the forest problem is the potential for emergent meaning that arises at the level of the corpus (i.e., the collection of comments) that is not apparent when the component documents (i.e., the comments) are analyzed in isolation.[169] This emergent meaning exists at the aggregate level in the form of patterns and regularities that are determined, at least in part, through the relationship of documents to each other, as well as the internal relationships between concepts, arguments, and ideas within documents. Addressing the forest problem requires aggregated representations of a large number of comments that can be interpreted by agencies and released to the public.

As it stands, in major rulemakings, agencies release comments in an entirely unstructured fashion.[170] In response-to-comments documents or in regulatory preambles, agencies will detail how they have considered substantive issues that were raised during the public comment process.[171] Sometimes, agencies will include some sweeping language concerning informal comments, characterizing the grounds for general support for or opposition to the rule. When they identify common themes that are the subject of many comments, or when they flag some comments as duplicates that have been generated by mass email campaigns, agencies already engage in some rudimentary efforts to identify emergent meaning. Both of these observations require agencies to compare comments to each other, in addition to taking each on its own terms. But frequently, the response-to-comments document is not indexed at all to the pool of comments, so it is impossible to track which comments match which revisions or responses (other than, theoreti-

169 In recent decades, the concepts of emergence and emergent phenomenon have become important in several scientific disciplines and have spawned a considerable philosophical literature. *See generally* PAUL HUMPHREYS, EMERGENCE: A PHILOSOPHICAL ACCOUNT (2016); STUART A. KAUFFMAN, THE ORIGINS OF ORDER: SELF-ORGANIZATION AND SELECTION IN EVOLUTION (1993); Jeffrey Goldstein, *Emergence as a Construct: History and Issues*, 1 EMERGENCE 49 (1999). We are using the term here in the context of "meaning" to imply the possibility that at least some useful interpretations can only be constructed when a corpus of public comments is treated as a collection in which meaning can be found in the relationships between comments, as well as in their isolated content.

170 *See Regulatory Timeline*, REGULATIONS.GOV, https://www.regulations.gov/docs/FactSheet_Regulatory_Timeline.pdf (last visited Oct. 30, 2017).

171 *See id.*

cally, reading all of the comments themselves). Ultimately, although these rudimentary efforts are useful, there is very little nuance, and the diversity of the individual voices that participated in the commenting process and their relationship to each other is lost.

The volume of comments, their low information density, and the relative resource scarcity of agencies make more sophisticated attempts to develop higher order meaning difficult without the aid of computational tools. Topic modeling is one possible approach. Topic models are a family of computational tools that have become popular in many academic settings in recent years for discovering latent thematic structure, typically in a completely unsupervised fashion, within a document collection.[172] Introduced in the early 2000s,[173] researchers in a variety of disciplines, including political science and many humanities, now take advantage of topic models to analyze large textual datasets as diverse as congressional floor speeches and twentieth-century German studies journals.[174] Topic models work by identifying patterns in the co-occurrence of words within documents, which can be used to infer subject matter categories within an unstructured corpus.[175] A major advantage of topic models is their ability to radically reduce the number of dimensions for analysis, from a vocabulary of tens of thousands of words down to a small number of *topics* (set by the user, typically to between 10 and 100). With this information, researchers can examine how documents relate to each other and how semantic content varies over time.[176]

Intuitively, a topic is a subject matter classification. In one application of a topic model to 17,000 articles from the journal *Science*, David Blei reports topics that include genetics, evolution, disease, and computers.[177] These topics match intuitive subject matter categories that might be discussed in a scientific periodical. In the technical topic modeling sense, each document is characterized by a probability distribution over topics, where each topic is a probability distribution over words (i.e., the vocabulary). If there are 20,000 unique words in a corpus, then a topic is a weighted distribution where every word is given a nonzero weight, and all of the weights sum to one. The most heavily weighted words in a topic, sometimes referred to as "top words," give a sense of the subject matter categories that it is associated with. For exam-

172 *See* David M. Blei, *Probabilistic Topic Models*, 55 COMM. ACM, Apr. 2012, at 77.

173 *See* David M. Blei et al., *Latent Dirichlet Allocation*, 3 J. MACHINE LEARNING RES. 993 (2003).

174 *See* Kevin M. Quinn et al., *How to Analyze Political Attention with Minimal Assumptions and Costs*, 54 AM. J. POL. SCI. 209 (2010); Allen Beye Riddell, *How to Read 22,198 Journal Articles: Studying the History of German Studies with Topic Models*, *in* DISTANT READINGS: TOPOLOGIES OF GERMAN CULTURE IN THE LONG NINETEENTH CENTURY 91 (Matt Erlin & Lynne Tatlock eds., 2014).

175 Blei, *supra* note 172, at 77.

176 Two implementation of topic models to examine relationships between documents are Michael A. Livermore et al., *The Supreme Court and the Judicial Genre*, 59 ARIZ. L. REV. 837 (2017); and Greg Leibon et al., Bending the Law (unpublished manuscript).

177 Blei, *supra* note 172, at 79 fig.2; *see also* David M. Blei & John D. Lafferty, *A Correlated Topic Model of Science*, 1 ANNALS APPLIED STAT. 17 (2007).

ple, the top words in one of the *Science* topics are "stars," "astronomers," "universe," and "galaxies"—these words match an intuitive subject matter category that corresponds to the field of astronomy.[178]

The mathematical details of topic models are quite technical, but in essence, computational algorithms work on a corpus of documents to infer a set of topics that best explain the observed data, which are the words within the actual documents in the corpus.[179] Topic models have grown in usage because they tend to generate results that conform to intuitive expectations about the subject matter within a corpus, which provides some confidence that they are able to capture something meaningful about the documents they are used to analyze.[180] The *Science* topics match up to the disciplinary categories of anyone familiar with the departmental structure of a contemporary university; the top topics the model applied to U.S. Senate floor speeches in 1997 to 2004 were a topic on debt, the deficit, and social security, and a topic on jobs, wages, and the general state of the economy, which are unsurprising subjects for political debate during that time period.[181] Topic models that have been applied to legal corpora, which include U.S. Supreme Court opinions, bankruptcy opinions, and national constitutions, have found topics that accord well with expert understandings of the relevant fields.[182] Topic models also create a representation of a textual corpus that is highly amenable to additional analysis that can provide further insight into a corpus. For example, topics can be used to construct a network that is based on the document in a corpus that illustrates how topics relate to each other, showing, in essence, how often subject matter categories are mixed within a document.[183] Topics can also be used as a way to estimate the semantic distinctiveness between documents or groups of documents, which allows analysts to determine, for example, the degree of relatedness between authoring institutions.[184]

From the perspective of an agency, topic modeling comments and carrying out associated analyses could help identify relationships between portions of the rule that are not obviously apparent from the face of the regulation. These emergent relations could provide the agency with insights into how revisions could be made that mitigate a shared concern through a pathway that the agency did not anticipate. This type of representation might also help the agency target its revisions toward areas of the rule that are of concern to many diverse stakeholders by addressing issues that tend to be central

178 Blei & Lafferty, *supra* note 177, at 27 fig.2.

179 For a reasonably nontechnical overview, see Blei, *supra* note 172.

180 *Cf.* Patrik Ehrencrona Kjellin & Yan Liu, *A Survey on Interactivity in Topic Models*, 7 INT'L J. ADVANCED COMPUTER SCI. & APPLICATIONS 456 (2016).

181 *See* Quinn et al., *supra* note 174, at 217 tbl.2.

182 *See* Law, *supra* note 96 (applying topic models to constitutions); Livermore et al., *supra* note 176 (applying topic models to Supreme Court decisions); Macey & Mitts, *supra* note 98 (applying topic models to bankruptcy decisions).

183 *See, e.g.*, Blei & Lafferty, *supra* note 177, at 19–20.

184 Livermore et al., *supra* note 176 (comparing Supreme Court opinions with circuit court opinions).

within a topical network (i.e., are linked with many different clusters of issues) rather than dealing with concerns that are on the outskirts of the network (and therefore have few relations with other issues).

By combining topic modeling and sentiment analysis, agencies could potentially delve even deeper into the emergent content of individual rules. For example, an agency could develop distributions of sentiments over many features and provisions of a rule, some of which the agency may be able to preidentify, and others that might be identified in the first instance by commenters. Every element of a rulemaking, which could include substantive requirements and the agency's justification, as well as general characteristics of the rulemaking (such as the length or level of complexity), could potentially serve as a target of sentiment in the comments to create a multiaspect sentiment analysis wherein the aspects are automatically defined through topics. Through this aggregated analysis, the agency may be able to identify discussion of specific regulatory effects that it might have missed. For example, issues that are not at the center of any comment, but that appear across a wide range of comments, may be difficult to identify by human evaluators because the degree of shared concern may only become apparent when a very large number of comments are digested.

A combination of sentiment analysis and topic modeling could also contribute substantially to the input value of the public comments process. If it is difficult for agencies to find general trends within a large group of comments, for interested individuals (or even organizations) the task is essentially impossible. Commenters can have little sense of how the concerns and recommendations they have offered relate to all of the other comments. It would be difficult for a commenter to know, for example, whether he or she was an outlier, offering a set of ideas that are far outside the norm, or instead was comfortably within the majority. Commenters also have no way to know whether their concerns are addressed by other comments, or whether there are perspectives, experiences, or viewpoints raised in other comments that he or she failed to anticipate or consider.

Actors outside of government have already experimented with applying advanced computational text analysis tools to a large volume rulemaking for the purpose of identifying emergent content.[185] In 2014, the Federal Communications Commission (FCC) commenced a rulemaking on net neutrality

185 One important nongovernmental initiative, Docket Wrench, which was hosted by the Sunlight Foundation, is no longer available. Docket Wrench provided a graphical interface for rulemaking dockets for a large number of rulemakings. The tools also provided text analysis tools and research tools. Nicko Margolies, *Docket Wrench: Exposing Trends in Regulatory Comments*, SUNLIGHT FOUND. (Jan. 31, 2013), https://sunlightfoundation.com/2013/01/31/docket-wrench-exposing-trends-regulatory-comments/. Unfortunately, funding concerns led to the shutdown of this tool, although there may be plans to revive it elsewhere. *See* Kat Duffy, *Sunlight Labs Update: Nonprofits Step Up to Preserve Tools for Transparency*, SUNLIGHT FOUND. (Nov. 1, 2016), https://sunlightfoundation.com/2016/11/01/sunlight-labs-update-nonprofits-step-up-to-preserve-tools-for-transparency/.

that received over a million public comments.[186] Two different actors deployed computational text analysis tools to examine the large number of comments received by the FCC. One of the analyses was conducted by the Sunlight Foundation, a not-for-profit interested in open government and campaign finance reform.[187] Analysts for the foundation, Bob Lannon and Andrew Pendleton, applied an open source topic model to over 800,000 public comments that were released by the FCC. Their analysis extracted topics and then clustered comments by topics, which were then used to create an interactive web interface that allows users to pull out individual comments according to their location within a set of nested topical clusters.[188]

An example of a topic cluster is one that relates to classifying internet service providers as common carriers. The topics in the cluster are:

> carrier, request, respectfully, classify, common
> broadband, carrier, common, classify, thank
> ban, tiered, carrier, mobile, classify
> phone, support, carrier, common, classify[189]

Comments in this cluster include:

> I'm requesting that Internet Service Providers be classified as common carriers. Fast lanes ruin equality on the internet, and will smother and kill innovation. . . .
>
> I would like to ask the FCC to classify the Internet as a Type II common carrier. I support a free and open internet. . . .
>
> Classify ISPs as common carriers. Defend net neutrality.[190]

The Lannon and Pendleton tools facilitate navigation through the comments according to their content. This step moves beyond generating aggregate level information by creating the ability to draw out and examine individual comments and see their relationship to each other. This tool fundamentally shifts the interaction between a user and what would otherwise be an undifferentiated and indigestible mass of comments.

A second analysis was carried out by a private firm, Quid, which describes itself as an "artificial intelligence company that accelerates research and insights to address the world's most complex issues."[191] Quid was hired by the Knight Foundation to undertake an analysis of the public discourse on

186 Mike Snider, *Record 9 Million Comments Flood FCC on Net Neutrality*, USA TODAY (July 19, 2017), https://www.usatoday.com/story/tech/news/2017/07/19/record-9-million-comments-flood-fcc-net-neutrality/488042001/.

187 Bob Lannon & Andrew Pendleton, *What Can We Learn from 800,000 Public Comments on the FCC's Net Neutrality Plan?*, SUNLIGHT FOUND. (Sept. 2, 2014), https://sunlightfoundation.com/2014/09/02/what-can-we-learn-from-800000-public-comments-on-the-fccs-net-neutrality-plan/.

188 *Id.*

189 *Id.*

190 *Id.*

191 Anthony Ha, *Quid Raises $39M More to Visualize Complex Ideas*, TECHCRUNCH (Mar. 19, 2015), http://techcrunch.com/2015/03/19/quid-series-d/.

the FCC rule, including an analysis of the public comments.[192] Although the details of the methodology used by Quid are somewhat unclear, they appear to have undertaken some form of topic modeling of the comments and developed a similar clustering analysis as done by the Sunlight Foundation, although with less detail and transparency. They also appear to have taken some liberties in interpreting the topics. Some examples of their topic categories include:

> Unfair to charge for preferential treatment
>
> Concerns over the pay-to-play internet
>
> ISPs already have monopolistic power[193]

These categories may accurately reflect the content of the public comments, but it is difficult to tell from the information released to the public as part of the Quid analysis. One particular advantage of the Sunlight Foundation's interactive web interface is the ability to connect the topic clusters to individual comments; the paper (i.e., PDF) report released by Quid and the Knight Foundation does not allow for that kind of analysis.

Comparing the Quid and Sunlight Foundation reports helps illuminate one of the dangers in the enterprise of using topic models to analyze public comments. Although some of the tools that we describe are quite straightforward to understand, others are more sophisticated, and it is often unclear, even to those with considerable experience, how best to interpret the outputs of complex analyses like topic models. This provides considerable discretion for an analyst that can be used to manipulate results in a favored direction. This problem is especially grave when government actors are carrying out this analysis, and there is likely a role for oversight institutions to develop and enforce norms and best practices concerning how analyses are carried out and results communicated.

For purposes of illustration, we now report the results of a standard topic model applied to the nearly 40,000 publicly released comments received by the EPA in regards to the Clean Power Plan. We use the Latent Dirichlet Allocation implementation provided by Genism for building the topic model.[194] The text was tokenized and stemmed, stop words were removed, and the top 10k unigrams (single-word expressions) and 5k bigrams (two-word expressions) were retained. As discussed in detail above, there was a broad diversity in the sophistication of comments, with many simply offering words of support or opposition, and others engaging in detailed analysis. Because topic models are based, at a very general level, on word co-occurrence, sophisticated and unsophisticated comments are given the same

192 KNIGHT FOUND., DECODING THE NET: AN ANALYSIS OF MEDIA, PUBLIC COMMENT AND ADVOCACY ON OPEN INTERNET (2014), http://www.knightfoundation.org/features/netneutrality/.

193 *Id.*

194 *See models.ldamodel – Latent Dirichlet Allocation*, GENSIM, https://radimrehurek.com/gensim/models/ldamodel.html (last updated Nov. 6, 2017).

weight. That said, if more sophisticated comments tend to use similar language, that may be reflected in the topics that are generated by the model.

This actually seems to be the case. In a twenty-topic model of the Clean Power Plan comments, one of the topics includes a number of two-level phrases that are very specific to the regulatory context at issue and seem quite likely to be associated with fairly sophisticated players. The top words (or words indicated by word stems) for this topic include "stationary source," "electric utility," "generating unit," and "emissions guidelines." All of these are terms of art specific to the Clean Air Act context, and so are likely a flag for sophisticated comments.

In that same twenty-topic model, there is a fair amount of overlap in the vocabulary, with words like "EPA," "energy," carbon," "clean," and "rule" showing up across multiple topics. Looking at the top words that are not common across topics, some interesting results emerge. For example, one of the topics includes "carbon tax" as a top unit when bigrams were included. That same topic included a relatively high distribution of the words "fee," "dividend," and "citizen." To an observer of climate policy, this topic conforms to arguments forwarded by those who favor an approach to greenhouse gas reductions that favors a revenue neutral price on carbon that is accompanied by a per capita "dividend" of the collected funds.[195]

Other topics in the same model include a topic that appears oriented toward the gravity of the threat from climate change, which includes several bigrams within the topic words; these include "children enact," "today threaten," "problem greatest," "challenge face," and "undeniable problem." Two topics appear to focus on the question of how nuclear power should fit into climate change policy: one with top words like "nuclear power," "waste," "reactor," "sustain," and "radioactive"; the other with top words such as "nuclear," "clean," "current," "construct," "calculate," "credit," and "baseline." The role of environmental groups in organizing commentators also shows up, with one comment substantially devoted to words associated with the Sierra Club (e.g., "sierra," "sierra club," "sierraclub," "sierraclub_org") and another with words associated with the Environmental Defense Fund (e.g., "edf," "edf_org").

From a content perspective, the analysis of the net neutrality and Clean Power Plan comments provide some interesting insights, but the primary utility of these early efforts is in illustrating the ability to extract emergent content from a large unstructured corpus of comments. The Sunlight Foundation analysis of the net neutrality comments provides an extremely useful template for agencies going forward. By providing well-organized, structured representations of aggregated comments, agencies can give voice to commenters by allowing them to communicate in a meaningful way to another audience, even if as part of a whole. At the same time, there are even more ambitious possibilities for a new kind of public deliberation that is

195 *See How Does It Work?*, CAP & DIVIDEND, http://climateandprosperity.org/ (last visited Nov. 28, 2017).

created when agencies provide a feedback mechanism within the group of commenters, rather than a closed one-way flow of information from commenters to agency.

IV. A Role for Oversight Institutions

Computational text analysis of public comments is not, to put it mildly, a cure for all that ails American democracy. Nevertheless, as the preceding Parts show, there is a powerful set of new technological tools that, if used well, can lead to enhanced understanding of agency-public relationships and more robust and useful participation in agency rulemaking by the public. In this final discussion, we briefly lay out a normative vision of the potential value of machine reading public comments for participatory rulemaking, discuss agencies' incentives to bring that vision about (or lack of such incentives), and examine how oversight institutions could prod agencies in the right direction.

Piecing together some of the different techniques and applications discussed throughout this Article, one can imagine a technologically enhanced notice-and-comment process that is superior to the current approach in many ways. Natural language processing tools could be embedded directly into the public comment interface and allow commenters easier access to information on a rulemaking that could include automated summaries of rulemaking text or links to relevant background documents, which themselves could be automatically summarized and distilled for ease of consumption.[196] Prompts could be used to help commenters understand the types of information that they might want to include. And automated argument analysis could help facilitate more persuasive and well supported comments.[197]

These tools could help improve the baseline quality of comments, but it is in facilitating interaction between commenters that there is even greater potential value. For example, an agency could engage in a two-step public comment process that allowed for an initial round of comments that closed by a certain date, followed by a pause and a second public comment window. During the pause, a specifically designed content analysis engine that combined elements of topic modeling, sentiment analysis, and perhaps other computational text analysis tools could work on the first round of comments to develop an interactive representation that builds on the basic structure of the Sunlight Foundation approach to the net neutrality comments discussed above. During the second-round comment period, automated tools could guide participants through the first-round comments and flag areas of interest and points of agreement or disagreement—even, perhaps, facilitating upvoting or down-voting particular comments or ideas.[198] An automated

196 *See supra* notes 99–101 and accompanying text.

197 *See supra* notes 99–101 and accompanying text.

198 Similar ideas have been floated by others. *See, e.g.*, Stuart Minor Benjamin, *Evaluating E-Rulemaking: Public Participation and Political Institutions*, 55 DUKE L.J. 893 (2006); Farina et al., *supra* note 5; Noveck, *supra* note 5.

approach could be used to identify higher quality comments that could be presented to commenters to inform their own understanding of the rule, or could be integrated into an annotation engine that would allow commenters to voice agreement or disagreement with the arguments presented by others.[199]

Once the final set of comments were collected, the agency (or a third-party agency that might have a more neutral attitude) could analyze the digital traces of the conversation, examining points of agreement and disagreement, collating opinions or sentiment with commenter characteristics (such as industry affiliation or region), scraping through for new bits of technical information, and developing a final aggregate representation that captures the state of deliberation on the rulemaking. Because it would be designed to be highly automated, this more interactive notice-and-comment process could be used for rulemaking at many different levels of broader public salience, from highly charged rulemakings that generate millions of comments to the more run-of-the-mill rulemakings in which only several hundred comments are submitted.

This technologically enhanced public comment process would likely have substantial input value, as participants would be given a much richer environment in which to engage in public deliberation, with the associated legitimacy-enhancing effects. There would also be greater output value, as higher quality comments are generated and more of the information that is embedded within them is extracted. The final aggregated representation of the comments could then serve as the foundation for revisions to the proposed rule. The agency could also use the comments as a starting place for a final substantive discussion that grapples with both the technical decisions and value-laden choices that it made. Higher order aggregation at the multirule level is also possible. Oversight bodies, such as OIRA or congressional committee staff could compare comments on a proposed rule to other rules of similar scope to determine whether closer-than-average scrutiny is warranted. An agency that consistently issues rules that tend to be highly controversial could also warrant additional oversight. And, of course, researchers could take advantage of this information for purposes of studying how agencies interact with a variety of different external actors.

Although this may be an attractive normative vision, and the technology to bring it about exists or is rapidly approaching, it is not clear that agencies have obvious legal incentives to embrace a socially desirable amount of technological sophistication in the notice-and-comment process. Courts have given agencies adequate reason to address the haystack problem by identifying and responding to the most substantive comments, but otherwise, the legally required threshold amounts to a willingness to accept and acknowledge receipt of comments and review them to determine if they offer novel

199 For an example of a sophisticated annotation engine applied to the lyrics of the musical *Hamilton*, see *Hamilton (Original Broadway Cast Recording)*, GENIUS, https://genius.com/albums/Lin-manuel-miranda/Hamilton-original-broadway-cast-recording (last visited Nov. 28, 2017).

insights or new information.[200] Indeed, courts have frequently disclaimed that there is any responsibility on the part of agencies to respond to the weight of opinion within the group of commenters.[201]

One possible route would be for courts to emphasize agencies' procedural obligation to appropriately consider commenters' views rather than a purported substantive obligation to follow the expressed preferences of commenters. This procedural obligation could be understood to change with the times: as technology evolves in a direction that enhances the ability of agencies to engage in more probing analysis of the comments, the obligation of nonarbitrariness ultimately creates a requirement to do so. There is some precedent that fits, at least loosely, this style of argument. An example can be found in *Center for Sustainable Economy v. Jewell*, where the D.C. Circuit rejected petitioners' argument that the agency had acted arbitrarily by failing to engage in a quantitative assessment of the information value of delaying the sale of offshore oil exploration leases.[202] The court specifically linked this holding, however, to the state of development of the underlying methodology, stating that although the agency did not have an obligation "to employ methods . . . at the 'frontiers of scientific knowledge.' Had the path been well worn, it might have been irrational for Interior not to follow it."[203] A similar logic could be applied to computational text analysis of comments—as the tools in this area become more well developed and deployed by a broader range of actors, both inside and outside of government, agencies may find themselves obligated to at least show some semblance of keeping up with best practices.

Although there may be some role for courts, it is unlikely that they will drive agencies toward a new technological horizon. The White House, perhaps operating through OIRA, is a more likely institutional candidate to encourage genuine innovation. Beth Simone Noveck has laid out perhaps the most ambitious vision for technology driven expansion in participatory government, informed by her own experience in the White House Office of Science and Technology Policy during the Obama administration.[204] Certainly a presidential executive order, enforced through OIRA, could be effective at prompting agencies to adopt new technologies. An analogy could be made to the use of regulatory impact assessment by federal agencies, which was a nascent technique until agencies were required to perform them by an

200 Several commentators have argued that the existing legal structure fails to require updated approaches to public participation. *See, e.g.*, Lisa Blomgren Bingham, *The Next Generation of Administrative Law: Building the Legal Infrastructure for Collaborative Governance*, 2010 WIS. L. REV. 297; Noveck, *supra* note 5.

201 *See, e.g.*, U.S. Cellular Corp. v. FCC, 254 F.3d 78, 87 (D.C. Cir. 2001) (stating that an agency is under no obligation to follow the approach endorsed by the majority of commenters).

202 Ctr. for Sustainable Econ. v. Jewell, 779 F.3d 588, 611–12 (D.C. Cir. 2015) (citation omitted).

203 *Id.* at 612 (quoting California v. Watt, 712 F.2d 584, 600 (D.C. Cir. 1983)).

204 *See* Noveck, *supra* note 5.

executive order by President Reagan in 1981.[205] When the order first took effect, agencies lacked the technical know-how to engage in sophisticated cost-benefit analyses of their regulatory proposals, but with time, they developed the needed methodologies and mastered the technique.[206]

The more pressing question is not whether the President could require more of agencies, but whether there are incentives to do so. Technologically enhanced rulemaking could provide information that enhances the efficacy of political review, akin to the role that cost-benefit analysis is thought to play by some commentators.[207] But greater public participation is also unruly and unpredictable; once unleashed, the force of public opinion may be difficult to control. In addition, it is not clear that the kind of information that would be generated would be well-suited to facilitating political control. Cost-benefit analysis, by contrast, aggregates technical information that is relevant to almost any consideration of regulatory effects; it is by its nature jargon-filled and dense, making negative spillover into public discourse less likely; and a great deal of communication goes on between OIRA and agencies outside of the public eye, allowing the White House to have greater control over the process. As a mechanism to gain the kind of information that is most useful for political oversight at low cost, enhanced public participation is not an obvious choice.

An alternative ground for Presidents to want to promote enhanced participation would be to increase the legitimacy of regulatory decisionmaking. As policy-making prerogative has migrated to the White House from Congress and the structure of partisan contestation has become primarily organized around the President, periods of divided government have seen especially heated interbranch clashes over the appropriate scope of executive power. The ability to marshal and demonstrate public support for specific rulemakings may become a powerful weapon for the President in these battles—to the extent that agencies adopt popular rules, the public comment process can be a means of translating latent support into publicly expressed opinions, and potentially even political action. If Presidents believe that technologically enhanced notice-and-comment rulemaking can generate these political returns, it would be surprising if they failed to take advantage of it.

Conclusion

In this Article, we report the results of the first large-scale computational text analysis of public comments received by all agencies over the course of a presidential administration. The purpose of this undertaking is twofold. First, we hope to have illustrated how new tools in natural language processing, combined with the massive (and publicly available) corpus that is gener-

205 Exec. Order No. 12,291, 46 Fed. Reg. 13,193, 13,194 (Feb. 19, 1981).

206 *See* Livermore, *supra* note 119.

207 *See* Eric A. Posner, *Controlling Agencies with Cost-Benefit Analysis: A Positive Political Theory Perspective*, 68 U. Chi. L. Rev. 1137 (2001).

ated by the notice-and-comment rulemaking process, create research opportunities for scholars of public bureaucracies, who can use the information embedded within public comments to test theories on the interaction of agencies and the public they serve. Second, we argue that agencies can use advanced computational techniques to respond to the challenges of the era of mass commenting, specifically by identifying the most substantive comments that require more sustained attention and by aggregating and analyzing comments to identify emergent content that is only apparent when comments are understood in relationship to each other and not simply read as individual, atomized responses to a regulatory proposal. Although agencies already have some incentives to adopt these tools, there may be a role for oversight institutions, such as OIRA or the courts, to prod agencies to adopt the next generation of information technology tools to facilitate deliberative and participatory regulatory governance.