Measuring Discretion and Delegation in Legislative Texts: Methods and Application to US States

Matia Vannoni[1], Elliott Ash[2] and Massimo Morelli[3]

[1] King's College London, WC2B 4BG London, UK. Email: matia.vannoni@kcl.ac.uk
[2] ETH Zurich, 8092 Zürich, Switzerland. Email: ashe@ethz.ch
[3] Bocconi University and CEPR, 20136 Milan, Italy. Email: massimo.morelli@unibocconi.it

Abstract

Bureaucratic discretion and executive delegation are central topics in political economy and political science. The previous empirical literature has measured discretion and delegation by manually coding large bodies of legislation. Drawing from computational linguistics, we provide an automated procedure for measuring discretion and delegation in legal texts to facilitate large-scale empirical analysis. The method uses information in syntactic parse trees to identify legally relevant provisions, as well as agents and delegated actions. We undertake two applications. First, we produce a measure of bureaucratic discretion by looking at the level of legislative detail for US states and find that this measure increases after reforms giving agencies more independence. This effect is consistent with an agency cost model, where a more independent bureaucracy requires more specific instructions (less discretion) to avoid bureaucratic drift. Second, we construct measures of delegation to governors in state legislation. Consistent with previous estimates using non-text metrics, we find that executive delegation increases under unified government.

Keywords: natural language processing, text analysis, executive delegation, bureaucratic independence, US politics

1 Introduction

The use of text data in political science has expanded rapidly in recent years (Gentzkow and Shapiro 2010; Grimmer and Stewart 2013; Roberts *et al.* 2014; Lucas *et al.* 2015), with notable examples including the detection of legislative agendas or topics and estimating the ideological positions of parties (Laver and Garry 2000) or single legislators (Lauderdale and Herzog 2016). The standard approach is to break down the syntactic structure of the text and represent it as a sequence of tokens or phrases, thereby losing the potentially vital information encoded in syntax and grammar. This paper shows how to extract this syntactic information and bring it back into the analysis, paving the way for richer text representations in political science.

With some exceptions, the mainstream approach to political text analysis is a bag-of-words (or bag-of-phrases) representation. First, the text is split up into tokens (single words or groups of words, which relate to a concept) and the set of informative tokens is filtered (Monroe, Colaresi, and Quinn 2008). Second, tokens are assigned a probability distribution to analyze associations with a speaker, party, topic, or another covariate. In a nutshell, this approach starts from text as unstructured data and transforms it into a frequency distribution over tokens (Klebanov, Diermeier, and Beigman 2008).

This mainstream approach potentially misses essential information in the text. Any piece of written text comes with a language structure that conveys a potentially large amount of lexical, syntactic, and semantic information.[1] For example, we would want to know whether mentions of the "governor" in state legislation have the governor as a subject (undertaking an action) or an object (the target or recipient of an action). Here, we explore how political science research

1 This list is not exhaustive; the text contains pragmatic and potentially other information that might also be empirically important.

***Political Analysis* (2020)**

DOI: 10.1017/pan.2020.9

Corresponding author
Matia Vannoni

Edited by
Jeff Gill

could benefit from taking this language structure of texts into consideration, building on natural language processing (NLP) techniques.

By looking at the lexical and syntactic features of a sentence, NLP techniques serve to retrieve richer information than a list of tokens. Our rule-based labeling approach (called "information extraction") starts by automatically parsing the lexical and syntactic structure of a sentence, extracting information on what is the subject, what type of verb is present, and so on. The structure is matched against *frames*, templates that determine what different provision types look like lexically and syntactically. For example, sentences with "governor" as subject and a strict modal verb (e.g. "governor shall enforce regulations") can be understood as a delegation of authority to the role of governor. Our role labeling rules follow dependency relations between words in a sentence and therefore are not constrained by word order (as is the case with N-grams or dictionary matching). The result is a classification of sentences according to their meaning, with information on the agents involved.

This paper makes two contributions. First, we validate the use of the parser-based method for extracting measures of delegation from legislative texts. Second, we demonstrate the usefulness of the method in two empirical applications using historical statute data from US states.

To validate the method, we apply it to hand-annotated language features from Franchino (2004), as well as our own additional hand coding of legislative text. We document that the information extraction method replicates human annotations more accurately than a simpler lexicon-based method. The more basic strategy of counting modals tends to generate more false positives. Still, the error rate for our method is not negligible, and in the Supplementary Appendix we provide a discussion of cases where machine annotation disagrees with human annotation.

Our first empirical application looks at bureaucratic discretion in US states. Motivation comes from Huber and Shipan (2002), who find using manual coding of statutes (the traditional method) that an independent bureaucracy may result in agency drift. As such, legislators would want to put into place a series of control mechanisms to restrain the bureaucracy, such as writing more detailed laws. To get at this question, we apply our information extraction method to a unique corpus, which consists of the full text of US state session laws from the 20th century. We find that the introduction of merit systems, namely independent bureaucracies, across US states is associated with statutes containing more legal provisions. This trend is consistent with the discretion model in the literature: legislators introduce stronger ex ante control mechanisms to discipline the more independent bureaucracy.

The second application analyzes delegation of powers from the legislature to the governor. The previous literature has used standard datasets to produce robust evidence that under unified government (governor and legislature controlled by the same party), the delegation of powers to the executive is more likely to take place (Epstein and O'Halloran 1999; Franchino 2004). Using a new measure of delegation constructed from the syntactic parse, we find confirming evidence for this empirical regularity. In line with the previous literature, we find that the number of statements delegating powers to the governor, discounted by statements constraining the governor, increases in unified government situations.

In both of these applications, therefore, we find that previous results using standard methods generalize to larger-scale text datasets using our information extraction method. These consistent applications, along with our other validations, support the use of legislative information extraction in future work. We hope that our information extraction approach can expand the use of text analysis to a broader range of topics in political science.

2 Legislative Information Extraction

This section summarizes the method of legislative information extraction. The approach relies on computational linguistics tools to produce parse data—statistical representations of the syntactic and lexical content in legal clauses. For example, it will identify the subject and verb of a sentence, the adjectives that describe the subject, and the objects of the verb. Meanwhile, we construct role labeling rules—a set of tags or rules that identify relevant clauses from the linguistics data—which, in our applications, provide measures of discretion or delegation. For example, an extraction rule could be "governor subject with permissive modal verb (e.g. may)", which would indicate a permission for the governor. We apply these types of extraction rules to the parse data to construct datasets for empirical analysis. The method can be understood as a form of rule-based semantic role labeling (SRL) using the domain-specific structure of legal language.

Automated methods to extract relevant information from legislative texts have recently been used for both federal laws (Al-Ubaydli and McLaughlin 2017) and state laws (Vakilifathi 2019). Vakilifathi (2019), the closest paper to ours, measures the level of statutory discretion in statutes regulating charter schools by counting the number of mandatory and optional statements, which are based on dictionaries of words and phrases. The author identifies these statements mainly by looking at modal verbs, associating "shall" to mandatory sentences and "may" to optional ones. She also includes in the analysis some alternative optional and mandatory phrases. Our method has some advantages over this approach. Using parse information and extraction rules (based on ontologies) allows us to filter out false positives: the modal counting method would treat "shall not be expected" as mandatory, while our extraction rules would not.[2]

2.1 Syntactic Dependency Parsing

Automated legislative information extraction is possible because computers can now quickly and reliably extract detailed lexical and syntactic information from large corpora. A key technology in this area is syntactic dependency parsing, developed in computational linguistics. Dependency parsing produces annotations on the syntactic structure of a sentence—the words and the grammatical relations between them (Jurafsky and James 2000).

First, parsers tag the parts of speech (POS)—verb, noun, adjective, and so on—of each word in a sentence. This identifies the function of each word. Second, parsers tag dependencies—the function relations between each word in the sentence. A dependency relation consists of a headword and a dependent word, related to each other through a functional dependency. Examples of functional dependencies are nominal subject (linking a subject and a verb), direct object (linking a verb and a direct object), attribute (linking an adjective and the noun it describes), and so on.

The dependency parser tells us whether a noun is the subject or the object of the sentence. It tells us rich information about the verb—whether it is the main verb or just an auxiliary, whether it is active or passive, and so on. A key category of verb in statutes is the modal verb, which in legal language assigns responsibilities and grant permissions. These annotations provide the ingredients from which our extraction rules build measures of delegation.

In the demonstrations reported below, our dependencies are produced using the Python package spaCy (Choi, Tetreault, and Stent 2015; Honnibal and Johnson 2015). The spaCy parser obtains state-of-the-art performance on the standard computational linguistics metrics. Like most parsers, it is trained on corpora of hand-parsed sentences (Goldberg and Nivre 2012). We inspected many samples and were happy with its performance on statute language. More detail is provided in the Supplementary Appendix.

2.2 Extraction Rules

A key step in legislative information extraction is to consider what information is available from the syntactic parser and then to define a set of provision types that are relevant to the research

2 Some work in international relations (O'Connor, Stewart, and Smith 2013) and political communication (Van Atteveldt, Kleinnijenhuis, and Ruigrok 2008) has used syntactic parsing.

question (Saias and Quaresma 2004; Soria *et al.* 2007). For example, one might be interested in statements that expand the governor's powers, versus statements that constrain them. With this goal in mind, one can identify a set of lexical units that could serve as tags or rules for identifying relevant provisions (Lame 2003; van Engers, van Gog, and Sayah 2004). These extraction rules can then be applied to the syntactic parser output to create the dataset for use in the analysis.

In most research, constructing extraction rules can be done using large-scale repositories of coded ontologies. These are dictionaries of words and dependencies that have been annotated to serve a theme, such as making a promise. An example of these ontology dictionaries is FrameNet (Baker, Fillmore, and Lowe 1998; Ruppenhofer *et al.* 2006). Lexicons of synonyms and categories, such as WordNet (Miller 1995), can be useful for constructing ontologies. Other work that has engaged with legal provision types using syntactic features includes Lame (2003), Saias and Quaresma (2004), Ceci *et al.* (2011), and Ash, MacLeod, and Naidu (2017).

Thanks to the linguistic regularities in legal language, the syntactic markers obtained from dependency parsing can be used to label semantic roles. From an extensive examination of example statements, we know (for example) that a subject attached to an active verb is the agent. A (direct or indirect) object, in turn, is the patient. The use of modal verbs "shall," "will," "must," "can," and "may" in legal language are universally deontic, whereas in common language they would often refer to nondeontic cases such as conditional or future tense. From these semantic labels, we construct the following categories: delegation, prohibition, permission, and entitlement (see Table 1). In defining these legal provisions, we start by deciding which modal and special verbs are associated with them. For instance, legal provisions that delegate authority, such as "The Governor shall act." These "delegations" contain strict modals, such as "shall" (unlike permissions, which would take a permissive modal such as "may"). Unlike prohibitions (which are negative—e.g. "shall not"), delegations are positive. Besides, delegations could be articulated through several "delegation verbs," such as "require," "expect" and so on. An example of this would be "The Governor is expected to."

A detailed and reproducible articulation of the tags and rules underlying our extraction rules may be found in Table 1. As enumerated in the table, a delegation is characterized by one of two structures: (1) a nonnegated strict modal followed by an active verb ("Governor shall act"), or (2) a nonnegated nonpermissive modal (either a nonmodal or a strict modal) followed by a delegation verb ("Governor is expected to"). Constraints are characterized by (1) a negated modal ("Governor shall not"), (2) a negated permission verb ("Governor is not allowed"), or (3) a nonnegated strict modal followed by a constraint verb ("Governor shall be prohibited from"). Permissions are characterized by a 1) nonnegated permission verb ("Governor is allowed to"), (2) a nonnegated permissive modal followed by a nonspecial verb ("The Governor may act"),[3] or a (3) negated constraint verb ("Governor is not prohibited from"). Finally, entitlements are characterized by (1) a nonnegated entitlement verb ("Governor retains the power to"), (2) a nonnegated strict modal followed by a passive verb ("Governor shall be considered"), or (3) a negated delegation verb ("Governor is not obligated to").

A key feature of our approach, relative to lexicon-based approaches that for example count modal verbs, is that the subject of any given legal provision is identified by the parser. A potential issue in this regard is coreferencing: namely, the use of a pronoun as a subject of a sentence which refers to a subject of a previous sentence. While coreference resolution is a major problem in most language domains, such as newspaper articles (Van Atteveldt *et al.* 2008), legislation uses relatively few pronouns, making the identification of the subject of each sentence easier. In our case, we found in samples of the data that our measures of delegation were unaffected

3 "Special verbs" are those listed in the top half of Table 1.

Table 1. Lexical units and pseudocode for extraction rules.

Lexical Units	
Strict modals	"shall", "must", "will"
Permissive modals	"may", "can"
Delegation verbs	"require", "expect", "compel", "oblige" , "obligate", "have to", "ought to"
Constraint verbs	"prohibit", "forbid", "ban", "bar", "restrict", "proscribe"
Permission verbs	"allow", "permit", "authorize"
Extraction Rules	
Delegation	strict modal + active verb + not negation OR not permissive modal + delegation verb + not negation
Constraint	modal + not delegation verb + negation OR strict modal + constraint verb + not negation OR permission verb + negation
Permission	permission verb + not negation OR permissive modal + not special verb + not negation OR constraint verb + negation
Entitlement	entitlement verb + not negation OR strict modal + passive + not negation OR delegation verb + negation

by the use of coreference resolution. Therefore, we chose not to run coreference resolution on the whole corpus (which would have been computationally demanding) for this analysis.

As mentioned, this process is similar to SRL. Semantic role labeling software, such as AllenNLP's implementation of PropBank, would serve to identify "who does what to whom" by labeling agents, patients, and associated verbs. The information from SRL, along with the modality modifier, could in principle deliver equivalent information for use in extracting legal provisions. But in our experiments comparing an SRL approach to the dependency-parse approach, we got better results with the latter for legal language. Our sense is that SRL annotates subtler relations in sentences, which are less transparent and rely more on the specific features of the training corpus. The training corpus for SRL is nonlegal language, and we have not fully assessed the performance of off-the-shelf SRL on legal language. In contrast, we have analyzed many samples of dependency parsing on legal language and were pleased with the results. It is necessary to note that our legal ontology would not work well on nonlegal language. We expect that techniques such as SRL will be needed to extend these methods to broader language domains.

2.3 Validation

In this section, we provide some validation for our method in the context of identifying delegations and constraints in texts. First, we compare our machine-annotated counts to hand-annotated counts from a previous paper (Franchino 2004). Second, we compare it to the lexicon-based strategy of counting modals.

To compare machine annotations to hand annotations, we apply our information extraction technique to the corpus from Franchino (2004). This dataset contains more than 150 European Communities legislative acts, hand-coded with the number of delegations and constraints. Our machine coding identifies delegations and constraints by counting the number of matches to the respective rules articulated in Table 1.

The panel on the top of Figure 1 shows the binned scatterplot of the relationship between our machine-annotated counts (horizontal axis) and Franchino's hand-annotated counts (vertical axis) for delegations. The measures are strongly correlated, with an R^2 of 0.44. We can see that the machine-coded measure identifies about twice as many delegations as the hand annotations, probably because the human annotators treated related/redundant statements as a single delegation.

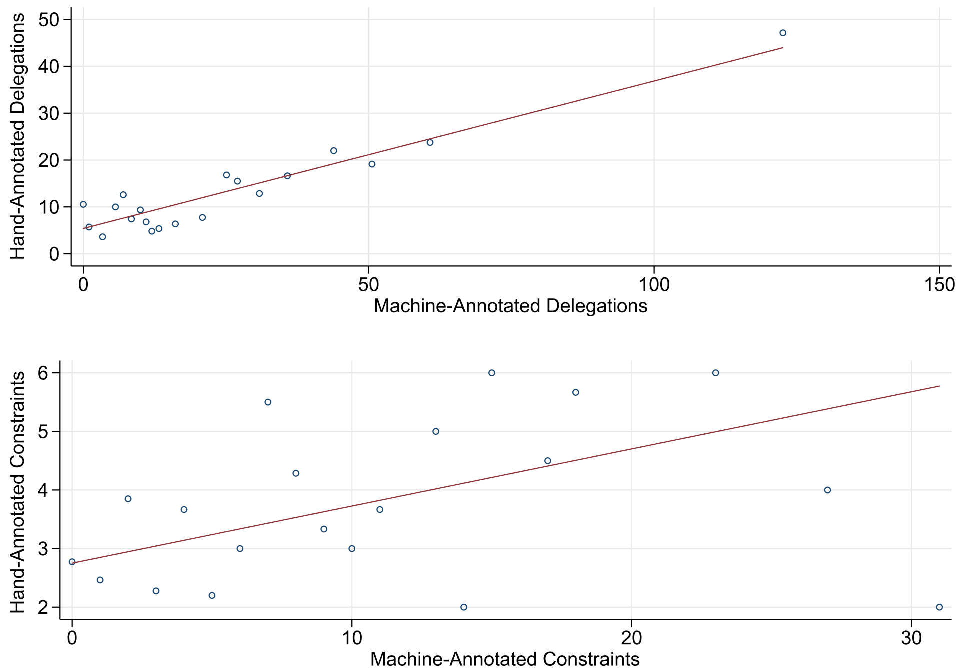

Figure 1. Validation with Franchino (2004): delegation and constraint counts.

The panel on the bottom of Figure 1 shows the same figure for constraints. While the measures are correlated, the performance is much lower, with $R^2 = 0.06$. Again, the machine coding measure identifies more constraints than hand coding. The low R^2 for constraints may be due to the subjective nature of coding constraints in the European Union (EU) data (Franchino 2004). In the future, we should work further on validating the constraint measure in the US state context.

Next, we compare our method for measuring delegations to a more standard lexicon-based approach based on counting modal verbs. For this validation exercise and the empirical demonstrations below, we use a unique dataset consisting of the full text of US state session laws from the 19th century to the 21st century. This corpus, introduced by Ash (2016), consists of all the new statutes enacted by a legislature during a session, which are published annually or biennially. We process this raw data by removing all nonstatute material from the texts and merging them.

For the validation check, we follow our method and compute the number of delegations with "governor" as the subject on the US state session laws corpus. This gives a count of sentences matching our extraction rule for delegations for each state and each biennium for the years 1900–2010. The lexicon-based comparison is the count of the bigrams "governor shall" and "governor will."

These measures are highly correlated, as one would expect from the similarity of the definition. However, we find that they result in different time series in our corpus. Figure 2 shows the ratio of the lexicon-based measure to the parser-based measure along with 95% confidence intervals. The figure shows that (although decreasing over time) the ratio is always statistically greater than one, suggesting that simply counting modals tends to generate false positives.

In the Supplementary Appendix, we provide further validation of our method. We blindly hand-coded a sample of delegations and find that our parser agrees with the human annotations over 80% of the time. The parse measure produces fewer false positives and false negatives than a lexicon-based measure counting modals. We also provide some examples of false positives and negatives generated by our method and we acknowledge some limitations which would require further work on the extraction rules.

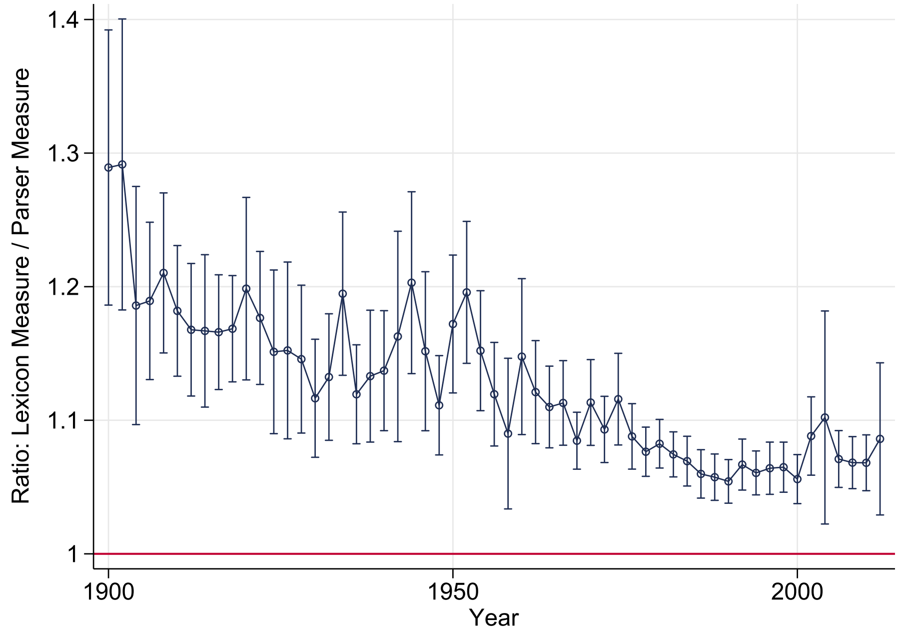

Figure 2. Modal counts tend to generate false positives.

3 Bureaucratic Discretion in US States

In recent decades, the literature on bureaucracy has focused on whether and how politicians delegate tasks to bureaucrats. In particular, they look at what control instruments legislators put in place to manage policy implementation (McCubbins and Schwartz 1984; McCubbins, Noll, and Weingast 1987; Levine and Forrence 1990; Epstein and O'Halloran 1994; Martin 1997; Gailmard and Patty 2012). On a leading framework for this process, legislators can use either ex ante or ex post control mechanisms (Martin 1997). Ex post control mechanisms refer to backward-looking incentives, such as firing bureaucrats who fail to implement a policy correctly. Ex ante mechanisms are more forward-looking and try to structure the bureaucracy to maintain the desired policy. These include administrative procedures (McCubbins *et al.* 1987), for example, and the level of detail of legislation. Detailed laws can be used to micromanage policy implementation (Huber and Shipan 2002). The delegation literature studies whether these two types are substitutes or complements (Huber and Shipan 2008).

We build on these ideas to analyze the introduction of an independent bureaucracy. These reforms weaken the legislators' capacity to control bureaucrats ex post, so legislators might write more detailed legislation as a form of ex ante control. As a set of natural experiments, we study the introduction of merit systems in the civil service in US states (Volden 2002; Wood and Bohte 2004). Note that an alternative expertise model of civil service reform would predict that legislation might become less detailed, if increased professionalism among bureaucrats means they need less legislative guidance.

The first step in this analysis is to measure legislative detail, which is central in analyzing bureaucratic discretion. A leading analysis in this area is Huber and Shipan (2002), who examine variation in detail of the statutes implementing the federal Medicaid program across US states. First, they select the relevant statutes for Medicaid by searching legal databases. Second, they use manual annotation to distinguish between procedural and policy language in the statutes. They argue that procedural language is less constraining than policy language because

> a bureaucrat can comply with the need to write a report or to consult particular groups or to conclude his or her work in a specified time period without being sharply constrained with

respect to the policy implemented. But if the statute says to do X, the bureaucrat cannot do Y (at least without some risks) (Huber and Shipan 2002, p. 48).

They then measure discretion quantitatively. As a baseline, they use a simple length-based measure of legislation as a proxy for the discretion left to bureaucrats: the longer the statutes, the greater the effort to reduce discretion. In addition, they look at the share of policy language, which gives less discretion.

The approach in our paper is a compromise between a length-based baseline and a hand-annotated measure like policy-versus-procedure share. On the one hand, the length of legislation alone is missing a lot of linguistic detail and treats legally relevant statements identically to boilerplate and other irrelevant text. On the other hand, the distinction between procedural and policy language is costly to annotate, somewhat subjective, and cannot be easily applied to other cases. We build at this intersection by looking for *legally* (rather than *policy*) relevant information from texts. Applying the information extraction techniques described above, we count the most common types of legal provisions listed in Table 1 (delegations, constraints, permissions, entitlements).

Formally, our outcome is $\log(\text{LegalProvisions}_{st})$, the logged number of legal provisions in the statutes of state s for each biennium t. We test the effect of the introduction of an independent bureaucracy on this outcome, where more provisions means less discretion. We analyze 50 US states from 1900 to 2000. The Supplementary Appendix reports some results using alternative text measures of discretion.

The estimating equation is

$$\log(\text{LegalProvisions}_{st}) = \alpha \text{Merit}_{st} + \beta X_{st} + \gamma_s + \delta_t + \phi_s t + \varepsilon_{st}, \tag{1}$$

where Merit_{st} is the variable that measures the introduction of a comprehensive merit system, X_{st} is a vector of time-varying state characteristics, γ_s and δ_t are state and time (biennium) fixed effects, and $\phi_s t$ represents state–time trends. The state fixed effects control for time-invariant state characteristics, while time fixed effects address any factors that change over time, but not across states, such as influence from the federal level.[4] The state trends allow for confounding trends at the state level. The equation is estimated using the reghdfe Stata package (Correia 2016) and standard errors are clustered to allow serial correlation within state.

Table 2 shows the results for the fixed effects estimates of Equation (1). The introduction of the civil service is statistically associated with higher levels of detail in legislation. The coefficient and standard errors are robust across specifications, including state trends and time-varying controls (Column 1). There is no change from adding the lagged dependent variable (Column 2), addressing the issues of long-term serial correlation in state panel data documented by Caughey, Xu, and Warshaw (2017). Adding a separate dummy variable for the year of the reform (Column 3) does not change the results either, meaning that the effect happens after the introduction of the merit system and not contemporaneously with it. The results do not change when interacting the treatment with Divided Government (Column 4), meaning that our results are not driven by the correlated changes in government structure. Finally, in Column 5 we also include in the treatment variable repeals of the merit system (which occurred in 15 states from 1996), finding similar results.

The dynamics of this effect are illustrated in Figure 3. The event study graph plots the log provision count (residualized on state/time fixed effects, state trends, and time-varying controls corresponding to Column 1 of Table 2), binned by biennium, for the two bienniums before and two

4 In particular, we can rule out influences from vertical delegation of powers from the federal to the state level. Assuming that the delegation of competences from the federal to the state level occurs at the same time for all the states, time fixed effects control for this.

Table 2. Civil service reform and legislative detail.

	(1) Leg Detail	(2) Leg Detail	(3) Leg Detail	(4) Leg Detail	(5) Leg Detail-repeal
Civil service	0.137	0.112	0.157	0.147	0.131
	(0.0625)	(0.0643)	(0.0646)	(0.0705)	(0.0588)
Introduction of drafting system	0.0755	0.111	0.0775	0.0764	0.0820
	(0.0807)	(0.0766)	(0.0804)	(0.0804)	(0.0783)
Divided government	−0.0256	−0.0153	−0.0255	−0.0359	−0.0255
	(0.0294)	(0.0289)	(0.0288)	(0.0308)	(0.0285)
Observations	1,438	1,382	1,438	1,438	1,485
R-squared	0.838	0.814	0.838	0.838	0.838
State FE	X	X	X	X	X
Time FE	X	X	X	X	X
State trends	X		X	X	X
Lagged DV		X			
Interaction				X	
Reform year			X		

Notes: Column 1 shows the results for the ordinary least squares (OLS) regression model with state and biennium fixed effects, time-varying controls (introduction of drafting system and divided government) and state-specific time trends. Column 2 adds the lagged dependent variable (without state-specific time trends). Columns 3 and 4 use the same specification of Column 1, but respectively add a dummy variable for the reform year and the interaction between divided government and the introduction of the merit system. Column 5 uses as treatment variable the introduction and the repeal of merit system. In all models standard errors are clustered by state.

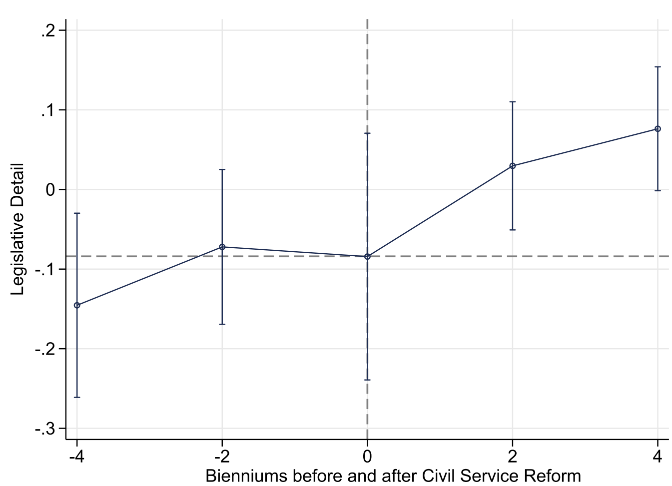

Figure 3. Event study graph. *Note:* Event study graph for effect of civil service reform on legislative detail. Dots give the binned mean residuals of log provision counts (vertical axis) from a regression on state fixed effects, biennium fixed effects, state time trends, and time-varying controls (Column 1 of Table 2), binned by the bienniums before and after the reform (horizontal axis). Error spikes give 90% confidence intervals from standard errors of the mean.

bienniums after civil service reform. The plot suggests no pretrend, with an increase in legislative detail taking place the next biennium after the introduction of an independent bureaucracy.

After the establishment of an independent bureaucracy, legislators start writing more detailed statutes. This finding is consistent with the idea that more independent bureaucrats are prone to agency drift, so legislators tend to micromanage policy implementation. Without ex post control mechanisms (such as firing bureaucrats at will), legislators start putting in place ex ante control mechanisms (more detailed legislation). The data do not support the alternative professionalism model, where expert bureaucrats would require less legislative guidance.

An additional set of model specifications and robustness checks are reported in the Supplementary Appendix, which shows the results for the regression models with different types of provisions as dependent variables. Results are robust across types, suggesting an increase in entitlements, permissions, constraints, and delegations associated with the introduction of an independent civil service. In addition, we test whether divided government affects legislative complexity in those years where the merit system was not in place. Results show that in those years there is no effect of divided government on legislative complexity, providing further evidence that divided government is not driving the results.

4 Executive Delegation in US States

A consistent prediction from delegation models is that when preferences between principal and agent converge, more delegation will take place (e.g. Huber and Shipan 2002, 2008). Empirical support for this prediction includes Volden (2002), who studies welfare boards in US states. He finds that, when the preferences of the legislature and the governor are aligned (that is, they come from the same party), legislators tend to give governors more appointment power over welfare boards.

The work on delegation is part of the broader literature on the powers of governors, such as appointment powers, control over the budget, term limits, and so on (Beyle 1990, 2007; Krupnikov and Shipan 2012; Kousser and Phillips 2012).

Another way of analyzing delegation to governors is to look at the content of legislation that delegates powers (Huber and Shipan 2008). Epstein and O'Halloran (1999) introduce a measure of statutory executive delegation which considers two components.[5] First, the degree of authority delegated to the executive branch, measured by the proportion of provisions in a legislative act delegating policy authority. Second, the degree of constraints imposed on the executive branch, measured by the number of constraints imposed in legislation. The total measure of statutory executive delegation is given by the share of provisions delegating powers in an act, weighted by the constraints imposed on executive action.

Epstein and O'Halloran (1999) apply this measure to the delegation of powers from US Congress to the president. They find less delegation under divided government. Franchino (2004) extends this analysis to delegation of powers in the EU. He looks at the Council of Ministers (the EU's equivalent to a second legislative chamber) and finds they delegate more to the Commission (the equivalent of the executive) where Member States' preferences converge.

This previous work has computed delegation through a combination of qualitative and quantitative methods. First, they identified relevant pieces of legislation, according to some guidelines, such as previous research (Epstein and O'Halloran 1999) or the relevant jurisprudence (Franchino 2004). Second, they manually code provisions according to whether they grant policy discretion or not. Finally, they identify potential categories of procedural constraints and manually count their frequency in the documents. This approach has some limitations. Perhaps most importantly, it is time and resource-intensive. Manual coding requires expert knowledge of the legal documents and associated legal system. The coders must go through hundreds of documents and preferably cross-validate results. In addition, manual coding requires subjective judgments on a series of important factors: which documents to sample, which statements are

5 In the original work this is referred to "statutory executive discretion" and not "statutory executive delegation", but in this work we use the latter to avoid confusion with the measure of discretion used in the first analysis.

relevant, what the potential categories of procedural constraints look like, and so on. The method is necessarily domain-specific, which limits opportunities for clean replication.

The time and resource requirements of hand coding legislative clauses can be ameliorated by machine learning from labeled documents. O'Halloran *et al.* (2016) is a promising example of this approach. However, machine classification does not address the issue of subjective judgments in labeling the documents. Besides, there is still the problem that documents labeled in one legal context would not be valid for machine classification in other legal contexts. We view the rule-based information extraction method and the machine learning method as complementary approaches.

In this section, we aim to address some of these issues using legislative information extraction. The empirical context is legislation in US states, and our outcome of interest is delegation to the governor. Following Epstein and O'Halloran (1999) and Franchino (2004), Delegation_{st} to the governor of state s at biennium t is computed as

$$\text{Delegation}_{st} = \frac{D_{st}}{M_{st}} - \frac{C_{st}}{M_{st}} \cdot \frac{D_{st}}{M_{st}}, \tag{2}$$

where D_{st} is the number of delegation statements with governor as subject, M_{st} is the total number of statements in that session's legislation, and C_{st} is the number of constraint statements with governor as subject.[6] This is the delegation ratio minus the constraint ratio (weighted by the delegation ratio). In the Supplementary Appendix, we report similar results for alternative outcome specifications that ignore constraints and/or use the number of provisions with governor as subject (rather than all provisions) as the denominator M_{st}.

Figure 4 illustrates how these factors have evolved over time in the US state session laws for the years 1900–2000. The left panel shows that the delegation ratio had a mostly flat trend roughly until World War II, then an increase in delegation until the 1980s, and then again a decreasing trend starting in the 1990s. These trends for governors are similar to the delegation trends at the federal level documented by Epstein and O'Halloran (1999, Figure 5.10, p. 138). The right panel shows the evolution of the constraint ratio, which was flat until the 1950s but then began a positive trend. Again, this is similar to trends at the federal level Epstein and O'Halloran (1999, Figure 5.11, p. 139). Moreover, these trends are broadly in line with anecdotal evidence on the powers of governors provided by the literature. Ruhil and Camões (2003) argues that the powers of governors increased after the Great Depression, while Rosenthal (1982) argues that powers became more balanced starting in the 1980s.

These descriptive statistics are promising initial support for our method. But our main inquiry is whether the previous evidence on unified government and delegation to the governor can be replicated using the new text-based measure. If our measure is valid, we would expect a positive relationship between government unity and statutory executive delegation.

To measure unified government, we use data from Klarner (2003) for the years 1935 through 2010. While we experiment with different specifications in the Supplementary Appendix, our preferred measure Unified_{st} takes value one when a single party (Democrat or Republican) controls the governorship and both chambers of the legislature in state s during biennium t. If at least one of the three government bodies is controlled by a different party, it takes value zero.

Our estimating equation is

$$\text{Delegation}_{st} = \alpha \text{Unified}_{st} + \beta X_{st} + \gamma_s + \delta_t + \phi_s t + \varepsilon_{st}, \tag{3}$$

6 Note that this formula is slightly modified from that used by Epstein and O'Halloran (1999) and Franchino (2004). They compute delegation as $Y = \frac{D}{M} - \frac{C}{K} \cdot \frac{D}{M}$, where K is the number of possible constraints. The choice of K requires expert knowledge of the possible set of constraints and is not feasible to do in our diverse context (50 states, 100 years).

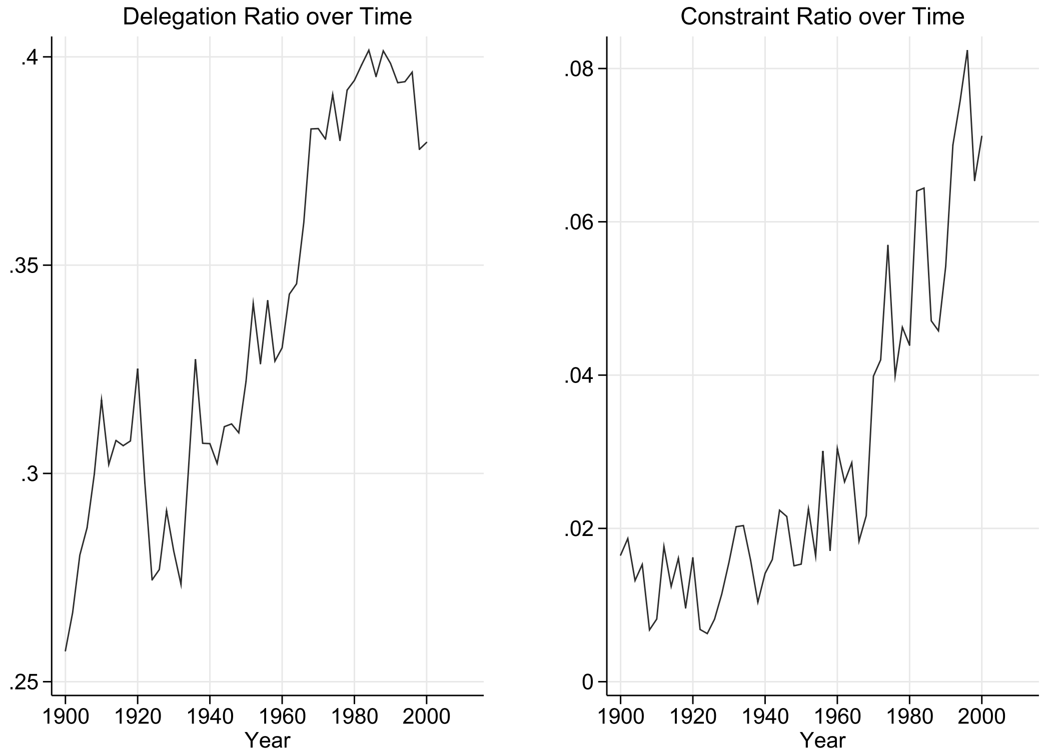

Figure 4. Average delegation and constraint ratios in state session laws, 1900–2000.

where as before, X_{st} is a vector of time-varying state characteristics, γ_s and δ_t are state and time (biennium) fixed effects, and $\phi_s t$ represents state–time trends. Controls include the introduction of the civil service because, as seen above, it affects the number of provisions in the statutes. As before, standard errors are clustered by state.

Table 3 shows the results of the fixed effects regression from Equation (3). A positive relationship is present between unified government and executive delegation, which suggests that where a single party controls the legislature and the executive, legislators tend to delegate more powers to the executive. Results are robust to different specifications, including the inclusion of state–time trends (Column 2), the lagged dependent variable (Column 3), and controls for civil service reform (Column 4). The preferred specification is robust to specifying the outcome as just the delegation ratio (Column 5), as well as using just governor statements (rather than all statements) as the denominator (Column 6).

In conclusion, we find evidence for a positive relationship between unified government and the statutory executive delegation to the governor. In other words, when the legislators and the governor are from the same party and hence they converge in their policy preferences, the former delegate more powers to the latter. This is in line with the findings of an extensive set of previous delegation studies and hence lends support to our information extraction approach to measure executive delegation.

5 Conclusion

In this work, we introduce a new approach to political text analysis—instead of a bag-of-words text representation, we look at richer language representations. By looking at the lexical and syntactic features of texts, we can classify statements according to more refined meaning. We show how to retrieve some legal provisions, namely delegations, entitlements, and constraints, from legal texts.

Table 3. Effect of unified government on executive delegation to the governor.

Variables	(1) Exec Del	(2) Exec Del	(3) Exec Del	(4) Exec Del	(5) Del Ratio	(6) Del Ratio Gov
Unified govt	0.0054	0.0046	0.0045	0.005	0.00678	0.008
	(0.003)	(0.0027)	(0.0025)	(0.0027)	(0.0031)	(0.004)
Observations	2,270	2,270	2,185	2,223	2,223	2,221
R-squared	0.396	0.464	0.434	0.463	0.529	0.328
State FE	X	X	X	X	X	X
Time FE	X	X	X	X	X	X
State trends		X		X	X	X
Lagged DV			X			
Civil service				X	X	X

Notes: Column 1 shows the results for the OLS regression model with state and biennium fixed effects. Column 2 adds state-specific time trends and Column 3 adds the lagged dependent variable. Column 4 adds the introduction of an independent civil service as control. Column 5 and Column 6 use "Delegation Ratio" and "Delegation Ratio Gov" as dependent variable, respectively. In all models standard errors are clustered by state.

We illustrate the validity of this approach by analyzing two predictions in the literature. First, the introduction of a merit system in the civil services of US states is associated with an increase in the number of legal provisions contained in statutes. Second, the number of provisions delegating powers to the governor in US state session laws is associated with government unity.

This is only one of the many potential contributions computational linguistics can make to social research. In future research, we shall use lexical and syntactic information in legal texts to distinguish contingent clauses, namely those clauses which specify different realizations of states of the world, from noncontingent clauses (or spot clauses) and test the differential effects of these types of clauses on productivity and economic growth. In this way, we will empirically answer a key question in political economy, namely whether more regulation is good or bad for the economy and when.

Our approach can also be used to extract information about exceptions, loopholes, or suspensions from legal texts. Recent work in legal studies uses an approach similar to the one discussed above to extract suspension norms (Ceci *et al.* 2011; Palmirani *et al.* 2011). Other work has tried to retrieve exceptions, which are another subcategory of efficacy provision and represent a modification of the norm, where the rules are restricted with respect to the original scope (Palmirani *et al.* 2011). Loopholes have also been recently studied in tax legislation from a computational linguistic perspective. This focus can be interesting for political scientists studying the effect of gridlock and vetoes on decision-making, a growing area of scholarship.

Acknowledgments

We thank Jake Brounstein and Tiziano Toniolo for research assistance and Gabriele Gratton, Russell Mills, and Moritz Osnabruegge for helpful comments. We also wish to thank seminar participants at Bocconi University, the Fifth Columbia Political Economy Conference, the Charting New Directions in Public Management Research at SDA Bocconi School of Management, and MPSA 2018 for helpful comments. Finally, we are grateful to Fabio Franchino for sharing his data.

Funding

This work was supported by the European Research Council Advanced Grant [694583].

Data Availability Statement

The replication materials for this paper can be found at Vannoni, Ash, and Morelli (2020).

Supplementary material

For supplementary material accompanying this paper, please visit https://doi.org/10.1017/pan.2020.9.

References

Al-Ubaydli, O., and P. A. McLaughlin. 2017. "RegData: A Numerical Database on Industry-Specific Regulations for All United States Industries and Federal Regulations, 1997–2012." *Regulation & Governance* 11(1):109–123.

Ash, E. 2016. "The Political Economy of Tax Laws in the US States." Working Paper.

Ash, E., B. MacLeod, and S. Naidu. "The Language of Contract: Promises and Power in Union Collective Bargaining Agreements." Working Paper.

Baker, C. F., C. J. Fillmore, and J. B. Lowe. 1998. "The Berkeley Framenet Project." In *Proceedings of the 36th Annual Meeting of the Association for Computational Linguistics and 17th International Conference on Computational Linguistics-Volume 1*, 86–90. Stroudsburg, PA: Association for Computational Linguistics.

Beyle, T. 1990. "The Powers of the Governor in North Carolina: Where the Weak Grow Strong-Except for the Governor." *North Carolina Insight* 12:27–45.

Beyle, T. 2007. "Gubernatorial Power: The Institutional Power Ratings for the 50 Governors of the United States." University of North Carolina at Chapel Hill.

Caughey, D., Y. Xu, and C. Warshaw. 2017. "Incremental Democracy: The Policy Effects of Partisan Control of State Government." *The Journal of Politics* 79(4):1342–1358.

Ceci, M., L. Lesmo, A. Mazzei, M. Palmirani, and D. P. Radicioni. 2011. "Semantic Annotation of Legal Texts Through A Framenet-Based Approach." In *International Workshop on AI Approaches to the Complexity of Legal Systems*, 245–255. Heidelberg, Germany: Springer.

Choi, J. D., J. Tetreault, and A. Stent. 2015. "It Depends: Dependency Parser Comparison Using a Web-based Evaluation Tool." In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing*, 387–396. Stroudsburg, PA: Association for Computational Linguistics.

Correia, Sergio. 2016. "Linear Models with High-Dimensional Fixed Effects: An Efficient and Feasible Estimator." Technical Report. Working Paper.

Epstein, D., and S. O'Halloran. 1994. "Administrative Procedures, Information, and Agency Discretion." *American Journal of Political Science* 38(3):697–722.

Epstein, D., and S. O'Halloran. 1999. *Delegating Powers*. Cambridge: Cambridge University Press.

Franchino, F. 2004. "Delegating Powers in the European Community." *British Journal of Political Science* 34(2):269–293.

Gailmard, S., and J. Patty. 2012. "Formal Models of Bureaucracy." *Annual Review of Political Science* 15(1):353–377.

Gentzkow, M., and J. Shapiro. 2010. "What Drives Media Slant? Evidence from US Daily Newspapers." *Econometrica* 78(1):35–71.

Goldberg, Y., and J. Nivre. 2012. "A Dynamic Oracle for Arc-Eager Dependency Parsing." In *Proceedings of COLING 2012*, 959–976.

Grimmer, J., and B. M. Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* 21(3):267–297.

Honnibal, M., and M. Johnson. 2015. "An Improved Non-monotonic Transition System for Dependency Parsing." In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 1373–1378. Stroudsburg, PA: Association for Computational Linguistics.

Huber, J. D., and C. R. Shipan. 2002. *Deliberate Discretion?: The Institutional Foundations of Bureaucratic Autonomy*. Cambridge: Cambridge University Press.

Huber, J. D., and C. R. Shipan. 2008. *Politics, Delegation, and Bureaucracy*. Oxford: Oxford University Press.

Jurafsky, D., and H. James. 2000. *Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech*. London: Pearson Education.

Klarner, C. 2003. "Measurement of the Partisan Balance of State Government." *State Politics & Policy Quarterly* 3(3):309–319.

Klebanov, B. B., D. Diermeier, and E. Beigman. 2008. "Lexical Cohesion Analysis of Political Speech." *Political Analysis* 16(4):447–463.

Kousser, T., and J. H. Phillips. 2012. *The Power of American Governors: Winning on Budgets and Losing on Policy*. Cambridge: Cambridge University Press.

Krupnikov, Y., and C. Shipan. 2012. "Measuring Gubernatorial Budgetary Power: A New Approach." *State Politics & Policy Quarterly* 12(4):438–455.

Lame, G. 2003. "Using Text Analysis Techniques to Identify Legal Ontologies' Components." In *ICAIL 2003 Workshop on Legal Ontologies & Web Based Legal Information Management*, 50–61.

Lauderdale, B., and A. Herzog. 2016. "Measuring Political Positions from Legislative Speech." *Political Analysis* 24(3):374–394.

Laver, M., and J. Garry. 2000. "Estimating Policy Positions from Political Texts." *American Journal of Political Science* 3(44):619–634.

Levine, M. E., and J. L. Forrence. 1990. "Regulatory Capture, Public Interest, and the Public Agenda: Toward a Synthesis." *Journal of Law Economics and Organization* 6:167–198.

Lucas, C., R. A. Nielsen, M. E. Roberts, B. M. Stewart, A. Storer, and D. Tingley. 2015. "Computer-Assisted Text Analysis for Comparative Politics." *Political Analysis* 23(2):254–277.

Martin, E. M. 1997. "An Informational Theory of the Legislative Veto." *Journal of Law Economics and Organization* 13(2):319–343.

McCubbins, M. D., R. G. Noll, and B. R. Weingast. 1987. "Administrative Procedures as Instruments of Political Control." *Journal of Law Economics and Organization* 3(2):243–277.

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

Miller, G. A. 1995. "WordNet: A Lexical Database for English." *Communications of the ACM* 38(11):39–41.

Monroe, B. L., M. P. Colaresi, and K. M. Quinn. 2008. "Fightin' Words: Lexical Feature Selection and Evaluation for Identifying the Content of Political Conflict." *Political Analysis* 16(4):372–403.

O'Connor, B., B. M. Stewart, and N. A. Smith. 2013. "Learning to Extract International Relations from Political Context." *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics, Volume 1: Long Papers*, 1094–1104.

O'Halloran, S., S. Maskey, G. McAllister, D. K. Park, and K. Chen. 2016. "Data Science and Political Economy: Application to Financial Regulatory Structure." *RSF: The Russell Sage Foundation Journal of the Social Sciences* 2(7):87–109.

Palmirani, M., M. Ceci, D. Radicioni, and A. Mazzei. 2011. "FrameNet Model of the Suspension of Norms." In *Proceedings of the 13th International Conference on Artificial Intelligence and Law*, 189–193.

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

Rosenthal, A. 1982. "The State of State Legislatures: An Overview." *Hofstra Law Review* 11:1185–1204.

Ruhil, A. V. S., and P. J. Camões. 2003. "What Lies Beneath: The Political Roots of State Merit Systems." *Journal of Public Administration Research and Theory* 13(1):27–42.

Ruppenhofer, J., M. Ellsworth, M. R. L. Petruck, C. R. Johnson, and J. Scheffczyk. 2006. *FrameNet II: Extended Theory and Practice*. Berkeley, CA: International Computer Science Institute.

Saias, J., and P. Quaresma. 2004. "Using NLP Techniques to Create Legal Ontologies in A Logic Programming Based Web Information Retrieval System." In *Workshop on Legal Ontologies and Web Based Legal Information Management of the 9th International Conference on Artificial Intelligence and Law*.

Soria, C., R. Bartolini, A. Lenci, S. Montemagni, and V. Pirrelli. 2007. "Automatic Extraction of Semantics in Law Documents." In *Proceedings of the V Legislative XML Workshop*, edited by C. Biagioli, E. Francesconi, and G. Sartor, 253–266. European Press Academic Publishing.

Vakilifathi, M. 2019. "Constraining Bureaucrats Today Knowing You'll Be Gone Tomorrow: The Effect of Legislative Term Limits on Statutory Discretion." *Policy Studies Journal* 47(4):978–1001.

Van Atteveldt, W., J. Kleinnijenhuis, and N. Ruigrok. 2008. "Parsing, Semantic Networks, and Political Authority using Syntactic Analysis to Extract Semantic Relations from Dutch Newspaper Articles." *Political Analysis* 16(4):428–446.

van Engers, T. M., R. van Gog, and K. Sayah. 2004. "A Case Study on Automated Norm Extraction." In *Legal Knowledge and Information Systems. Jurix 2004: The Seventeenth Annual Conference*, edited by T. Gordon, 49–58. Amsterdam, Netherlands: IOS Press.

Vannoni, M., E. Ash, and M. Morelli. 2020. "Replication Data for: Measuring Discretion and Delegation in Legislative Texts: Methods and Application to U.S. States." https://doi.org/10.7910/DVN/FQTA4A, Harvard Dataverse, V1, UNF:6:jXzgM/eYOZkItuUH2Zdarw== [fileUNF].

Volden, C. 2002. "A Formal Model of the Politics of Delegation in a Separation of Powers System." *American Journal of Political Science* 46(1):111–133.

Wood, B. D., and J. Bohte. 2004. "Political Transaction Costs and The Politics of Administrative Design." *Journal of Politics* 66(1):176–202.