Legislative Studies Quarterly

ORIGINAL ARTICLE

The Language of Delegation: An NLP Analysis of Congressional Bill Text

Austin Bussing[1] | Joshua Y. Lerner[2] | Gregory P. Spell[3]

[1]Department of Political Science, Trinity University, San Antonio, Texas, USA | [2]NORC at the University of Chicago, Chicago, Illinois, USA | [3]Independent Scholar, USA

Correspondence: Austin Bussing (abussing@trinity.edu)

Received: 18 June 2024 | **Revised:** 21 April 2025 | **Accepted:** 24 April 2025

Keywords: administrative agencies | bureaucracy | deep learning | delegation | interbranch relations | large language models | separation of powers

ABSTRACT

Delegation of powers from the legislature to the executive branch is a nearly ubiquitous feature of modern lawmaking. However, much of what scholars know about delegation is gleaned from an exclusive focus on landmark legislation. We introduce a method to identify delegating language across a larger universe of legislation. Using an active learning convolutional neural network on bill text, we classify bill sections by their delegation content, applying an iteratively improving coding scheme that enhances existing supervised learning approaches. We develop a novel dataset that allows us to answer important questions about interbranch relations. First, we find that legislator ideology, partisanship, and institutional position affect the delegatory content of introduced legislation. We then explore the role of delegation in the advancement of bills through the legislative process. Finally, we evaluate the ally principle, finding that variation in delegation is driven by cross-agency differences in ideology and structural independence.

Delegation of powers from the legislature to the executive branch remains a central challenge of modern governance. The mere existence of delegated policymaking authority can blur the boundaries between the powers and responsibilities of the three branches of the American government. This fact is further emphasized by the Supreme Court's recent decision in *Loper Bright Enterprises v. Raimondo*, which can be seen as an attempt by the judiciary to reassert itself as the ultimate arbiter of the statutory boundaries of legislative delegations of authority to the executive branch. Whereas *Chevron* deference would have required courts to defer to any reasonable agency interpretation of a statutory ambiguity, the Supreme Court now endorses Sunstein's assertion that an "ambiguity is simply not a delegation of law-interpreting power" (Sunstein 1989, 445). Ultimately, this shift from *Chevron* deference to what Vermeule (2025) calls "Loper Bright delegation," makes the identification and measurement of explicit textual delegation even more important to understanding inter-branch relationships and policymaking.

While the study of delegation in political science has been theoretically rich, large-scale empirical work has been hindered by measurement problems. Delegation is fundamentally a textual act: the legislature restricts or enables an agency's authority through written law and legislative history. Consequently, even the most rudimentary measurement of delegation has to grapple with the written word, either directly or indirectly. In a single piece of legislation, this task is straightforward—one can identify the agency and any change to its authority through close reading. However, this approach is prohibitively time-intensive for large-scale applications.

We would like to thank the following people for their helpful comments on earlier versions: Mathew McCubbins, John Aldrich, Kristen Renberg, Kiran Auerbach, Hannah Ridge, Robert Shaffer, Justin Grimmer, Michael Crespin, Sarah Bouchat, Roger Noll, and attendees at PolMeth 2018, APSA 2018, and the Mathew McCubbins Memorial Conference 2022.

Innovations in machine learning simplify scaling up these connections. In particular, these methods can predict qualitative labels from quantifiable features. After being trained on a set of hand-labeled examples, a learning algorithm can be deployed to predict labels for a much larger data set, greatly reducing the up-front costs of a human-run labeling project. However, most machine learning approaches separate the hand-coding from the evaluation process, treating it largely as a single attempt to achieve high accuracy. Complex concepts can be difficult to capture in a simple coding scheme, and thus, traditional machine learning approaches may exacerbate certain measurement challenges.

In this paper, we argue that researchers can leverage their expertise through an interactive machine-learning framework that effectively learns on the go—commonly referred to as "active learning." We argue that this framework best combines the portability and flexibility of machine learning with the expert evaluation practices typically found in smaller-scale projects or case studies.

Scaling up the identification of delegation in statutory language to a broader universe of legislation allows new leverage on important questions pertaining to legislatures in separation of powers systems. In this paper, we focus on three specific applications. First, we leverage the text of introduced legislation to discern patterns in delegation as a function of sponsor-level characteristics, such as ideology and committee membership. Second, we examine the role of delegation throughout the legislative process, finding that bills with higher delegation ratios are more likely to make it past the committee stage, but less likely to be passed on the floor of either chamber or to become law. Finally, we examine the workings of the ally principle during a transition from divided government (110th Congress) to unified government (111th Congress), demonstrating that variation in delegation is driven by cross-agency differences in ideology and structural independence.

This paper proceeds as follows: we first explain why delegation poses a unique measurement challenge. We then present a deep and active learning model that addresses these challenges, discussing how it differs from other machine-learning approaches. Next, we assess classification accuracy and discuss the performance of our active learning model and how it improves upon existing models. Finally, we illustrate the value of our newly labeled dataset through three substantive applications related to delegation and interbranch relations before offering concluding observations.

1 | Approaches to Measuring Delegation

The empirical literature on delegation is characterized by a number of conceptual and methodological divides. We present a summary of these conceptualizations in Table 1. While these studies generally share a broad definition of delegation as a grant of policy-relevant authority from the legislature to some other governmental body, they vary on how to operationalize this concept, where to look for it, and the appropriate unit of analysis at which to measure it. Some studies, such as Huber and Shipan (2002), evade the measurement of delegation altogether by only studying statutes that delegate and moving directly to the measurement of constraint and discretion.

Of the studies that do explicitly measure delegation, most identify delegation at the legislative provision level and then aggregate up to the statute level (e.g., Epstein and O'Halloran 1999). One recent study (Bolton and Thrower 2019) measures delegation at the agency-year level rather than the statute level. Studies generally rely on text to identify and measure delegation, but there is variation in the types of text used—from full legislative text (Huber and Shipan 2002; Farhang and Yaver 2016; Anastasopoulos and Bertelli 2020; Vannoni et al. 2021) to legislative summaries (Epstein and O'Halloran 1999) to subcommittee reports (Bolton and Thrower 2019).

Studies also vary in the extent to which they rely on human coders to identify and code delegation; this methodological variation has important implications for the scalability of each approach. Furthermore, studies differ in the units they assign to measures of delegation, ranging from dollars to counts of delegating provisions. Finally, while most studies search for evidence of delegation in legislatively produced texts, McCann et al. (2022) primarily rely on executive agency actions as indicators that authority has been delegated.

Our work here is the only application of automated methods to code delegation in U.S. legislation at the federal level. We also expand our focus beyond bills that become law, measuring delegation in each version of every introduced bill. Our measurement strategy is tied directly to legislative text, and we identify delegation and the recipients of delegated authority at the bill section level before aggregating up to the bill level.

1.1 | Delegation, Constraints, and Discretion

Except for McCann et al. (2022), all of the work discussed above treats delegation primarily as an intermediary step towards the measurement of another concept: discretion. Before describing our measurement strategy and methodological approach in detail, we differentiate between the interrelated concepts of delegation and discretion and justify our decision to focus on delegation. Delegation is the decision made by a principal (e.g., Congress) to assign a task to an agent (e.g., an executive branch actor). Discretion, on the other hand, refers to the leeway an agent is given to carry out the delegated task.

In some form or another, discretion is usually operationalized as the amount of authority that is left to an agency after taking into account any constraints the legislature includes in its delegation. These constraints typically come in the form of administrative procedures imposed on recipients of delegated authority (McCubbins et al. 1987, 1989; Epstein and O'Halloran 1999) or in policy details specified in legislative language (Huber and Shipan 2002). Discretion is delegation net of constraints.

In this paper, we set discretion aside and focus directly on delegation. We conceptualize the legislative decision-making process in three sequential steps—the decision of whether or not to delegate to an executive branch actor, the decision (conditional on delegating) of which executive branch actor(s) will receive delegated authority, and the decision of how much discretion to give the recipient(s) of delegated authority. Identifying and measuring delegation—which can be directly observed in legislative

TABLE 1 | How has delegation been measured?

Citation	Sample	Method	Treatment of delegation
Epstein and O'Halloran (1999)	257 Mayhew (1991) important laws	Close reading of law summaries	Legislative provisions that grant substantive policy discretion to some other governmental body
Huber and Shipan (2002)	67 state laws governing Medicare managed care programs	Word count for legislation	Not explicitly measured, as the sample selects on laws that delegate
Farhang and Yaver (2016)	218 Mayhew (1991) important regulatory laws—through 2008	Close reading of legislative text	Legislative provisions giving agencies the authority to perform regulatory functions
Bolton and Thrower (2019)	All executive agencies from FYs 1976–2015	Reading Appropriations subcommittee reports	New budget authority recommended for an agency in a given year
Anastasopoulos and Bertelli (2020)	59,423 EU laws (1958–2017)	Machine learning framework with existing coded data	Same as Epstein and O'Halloran (1999)
Vannoni et al. (2021)	2550 state laws (1900–2000)	Natural language processing approach using syntactic parsing trees	Legal provisions with strictly defined syntax granting authority to some government actor
McCann et al. (2022)	443 Mayhew (1991) important laws—through 2016	Use of ProQuest's Regulatory Insight database, supplemented by close reading of legislative text	Evident in agency citations of law when promulgating rules/regulations
Bussing et al. (2025)	28,906 bill-versions (2007–2010)	Use of convolutional neural network to identify delegating language in bill sections based on hand-coded data	Written mandate or permission for a federal agency or actor to exercise public authority in some way

text—allows us to gain leverage on the first two steps of this process.

We agree with McCann et al. (2022) that, while differentiating between the degrees of importance, or the amount of discretion, associated with different delegations of authority constitutes an important question, we must first be able to accurately identify delegation in legislation. In many ways, the shift towards explicitly textual delegation in jurisprudence following the decline of *Chevron* deference directly justifies focusing on the first two aspects of legislative decision-making.[1]

The following section describes our methodology for identifying delegation in legislative texts. As mentioned above, delegation is conceptually straightforward and relatively easily identifiable for a reader familiar with legislative texts. However, the difficulty is that close reading does not scale up for large-N empirical research. Our method provides a scalable solution to coding legislative texts for delegatory language.

2 | An Active Learning Convolutional Neural Network for Classifying Text

Classification is a common objective in text-as-data research, typically identified as supervised learning, where inputs and labeled outputs are used to discover a mapping between them. While supervised methods often require thousands of training examples, there are creative ways to reduce the effort required. We see an example related to our own in Anastasopoulos and Bertelli (2020), who use existing classifications of agency delegation in EU legislation to perform supervised labeling of other years not covered in the dataset.

Active learning is a machine learning approach to bolster classification performance by selecting (or "querying") further examples for model training. This querying allows models to obtain higher classification accuracies as training data are added than if new examples were instead randomly chosen. An active learner typically poses its queries over unlabeled data, which are then labeled by a human annotator and added to the dataset. Active learning is particularly useful when unlabeled data may be abundant, but labels are complicated, time-consuming, or expensive to obtain (Settles 2009; Miller et al. 2020).

Active learning also supports dynamic model updates as researchers uncover new patterns. For example, identifying administrative delegation requires recognizing agencies and tasks in text. The general framework for delegating authority is relatively straightforward: some agency is given a task—often told it "shall" or "must" do something. A standard learner would identify the agency and the verbs specifying the task and then use that information to identify delegation. However, there are hundreds of currently active agencies and programs, some of which have unusual names (the "Corporation for National and Community Service" as an example), that would make delegation difficult to identify.[2]

We encountered this problem early on when discussions of the "Attorney General" were frequently mislabeled because, up until that point, there were no observations in our training data with a cabinet-level secretary referred to as anything but "Secretary." Because the specific words "Attorney General" have few analogous positions in other departments, it would have to be hand-coded explicitly for the model to learn what it is. Generally, this is not a complicated fix: hand-labeling some of these aberrant observations should solve the problem. However, what we could not know a priori was what exact issues were going to appear: moving to an interactive labeling and machine-learning framework alleviated these concerns because the classifier would be able to tell us where it was having difficulties discriminating between classes.[3]

Traditionally, in supervised learning, models are iteratively trained, evaluated, and refined until performance is sufficient for deployment. Given the impossibility of anticipating all classification issues or verifying every output, we argue that active learning offers a systematic approach to address potential uncertainties. After querying these uncertainly labeled documents from the model, they may be manually labeled by the researchers and appended to the training dataset. The process of training, evaluation, and querying then repeats until the model performs well enough to be deployed and label all remaining documents.[4]

Our criterion for actively querying new data examples is the model's uncertainty in label prediction—"uncertainty sampling." In the case of classification, this strategy amounts to identifying the observations closest to the classification boundary, such as examples with the smallest margin in max-margin models (e.g., SVM) or examples with logits closest to 0.5 for logistic regression. We use a convolutional neural network (CNN) with a multi-layer perceptron (MLP) as our primary classifier.[5]

2.1 | Convolutional Neural Networks for Text Classification

We follow the example of Kim (2014) who defined a CNN for text classification, as well as Zhang and Wallace (2017) who provide practical guidance on using such models. The primary advantages of using deep learning for our text analysis—rather than "bags of words" approaches—are to model the similarity between words and to account for word order in text sequences. Both of these advantages involve considering words in context.

Underpinning the use of CNNs for text analysis is the distributed representation of words, wherein each word in a vocabulary is associated with a real-valued feature vector (Mikolov et al. 2013). These expressive vector representations encode many linguistic regularities and patterns, such as the relationships between synonymous words, and their use has been shown to improve the accuracy of text classification (Turian et al. 2010). Unsupervised training of word embeddings is typically accomplished by predicting the incidence of words given local context words, but practitioners may prefer to randomly initialize word embeddings to be learned as hyperparameters of their specific task (Rodriguez and Spirling 2022). While we opt for the latter strategy, Figure S7 depicts a visualization of word embeddings trained on the corpus of bills from the 110th Congress (2007–2008).

Figure 1 visually illustrates how our CNN takes our text data and generates our predictive labels.[6] This example is included for illustrative purposes only; for further details, please view Appendix Section 2, which details exactly how our CNNs were fit.

3 | Data

We utilize data for all versions of all bills (both successful and unsuccessful) from the 110th (2007–2008) and 111th (2009–2010) Congresses. We separate each version of each bill into sections and analyze them at the bill section level. We do this for three reasons: first, because each section in each bill deals with a particular agency or activity and likely contains an entire delegatory phrase, keeping the task more straightforward. Second, any given bill could delegate authority to multiple agencies in multiple sections, so to avoid missing any additional delegations, we wanted to reduce the task to units that are about a single delegation. The final reason we chose bill sections is that bill sections are the smallest comparable distinct units of a bill: sections are more comparable to one another than either a sentence or an entire bill would be.

Our complete dataset has several components, divided into labeled and unlabeled sub-datasets. The labeled component comprises 2098 bill sections from the 110th Congress that were read by human annotators and assigned a binary label with respect to delegation; the unlabeled component contains the remaining 137,616 bill sections from the 110th Congress as well as all sections from the 111th Congress. As is standard practice within machine learning, we divided our labeled data into subsets for training, validation, and testing. In performing this split, we ensured that all sections/versions from the same bill were apportioned to the same subset. We randomly divided the bills in our labeled data into training, validation, and test using proportions of 65%, 15%, and 20%, respectively. The remaining bills from the 110th Congress were also apportioned into training, validation, and testing using the same proportions. The numbers of bills for each subset are presented in Appendix Section 1.

3.1 | Delegation Coding

Essential to our project is a consistent definition of delegation to administrative agencies. An act of delegation is a mandate or permission for a federal agency or program (including the president) to exercise public authority in some way (see McCubbins et al. 1987; Kiewiet and McCubbins 1991; Huber and Shipan 2002; McCann et al. 2022, for a discussion of this

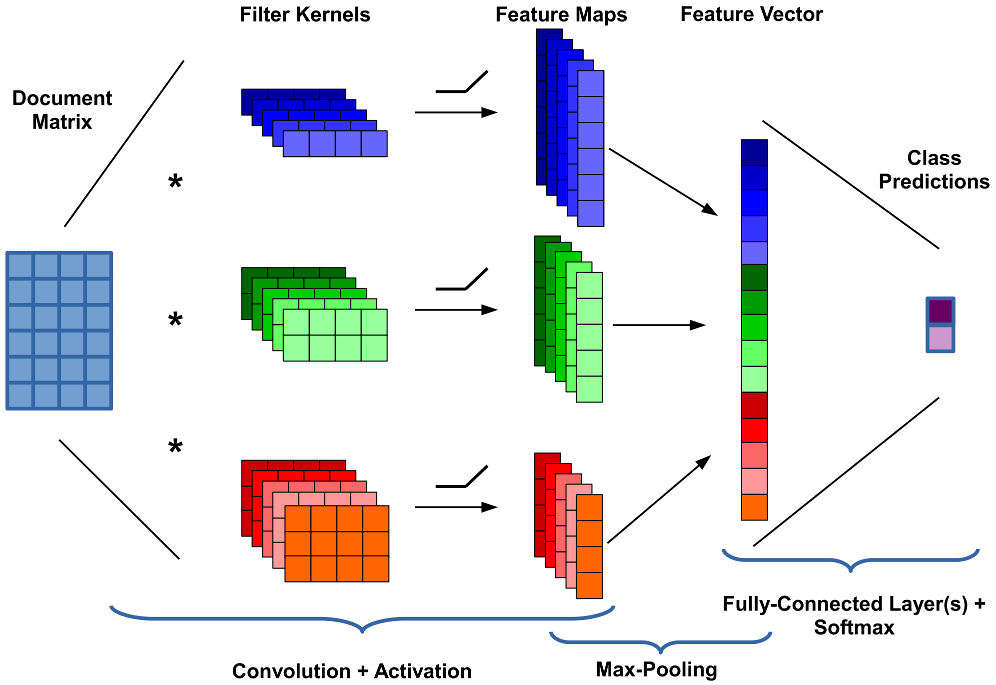

FIGURE 1 | Illustration of CNN for text classification. For this example, $d = 4$, there are five filters of lengths {1,2,3}, and there are two classes.

point). For our task, we stated that allocating money for federal agencies to spend, instructing agencies to promulgate rules, granting agencies the ability to exempt themselves from preexisting rules, requiring agencies to compile reports or commission pilot studies, and charging agencies with the enforcement of specific policies are all delegation (Kiewiet and McCubbins 1991).

For the hand-coding, we gave straightforward instructions on identifying delegation. First, is Congress acting upon an administrative agency? This will include all references to both the agency itself and the person in charge of that agency. We operated with a list of administrative agencies and matched each instance of delegation to one of those agencies.[7]

Second, what is the section asking the agency to do? In general, Congress delegates authority by asking an agency to perform a specific task, collect information, write new regulations, hire people, write a report to Congress on their activities, and distribute an award, among many other things. A bill is not delegating authority if it only appropriates money, if it references actions already taken, or if Congress is writing new rules or regulations. Keeping these actions separate allows us to track statutorily derived authority for the agencies, not merely what funds they have been allotted.

Below are example bill sections the active learner selected in early runs as uncertain and how they were coded.

- *Section.2402*. Energy conservation projects. using amounts appropriated pursuant to the authorization of appropriations in section.2403.a.6, the secretary of defense may carry out energy conservation projects under chapter 173 of title 10, united states code, in the amount of 800,000 (delegates authority to an administrative agency)
- *Section.2*. Reemployment of foreign service annuitants ... the authority of the secretary to waive the application of subsections a through d for an annuitant pursuant to subparagraph c of paragraph 1 shall terminate on September 30, 2008. the authority of the secretary to waive the application of subsections a through d for an annuitant pursuant to subparagraph c ii of paragraph 1 shall terminate on September 30, 2009 (does NOT delegate authority to an administrative agency)

In the above examples, it is clear why the algorithm would have selected them as ambiguous classifications and why human readers would classify them correctly. Take the top section, dealing with "Energy Conservation Projects." A read of the section makes it clear that Congress is delegating authority to Defense (through the Secretary of Defense) to spend $800,000 on energy conservation projects. The classifier may have had issues with the added verbiage of the task ("may carry out") and the addition of US code language in between. In the second section, where Congress is setting up "foreign service annuitants", it is clear that the agency is not being given an extra task or authority but only describes how applications must be processed. This section is an example where context-free language would indicate the possibility of delegating authority, but additional context makes it clear that this is not occurring. These are only two examples pulled from early runs of the active learning module, set up to illustrate the nature of the classification task.[8]

4 | Classification Results

We explore classification modeling results in the Appendix. To evaluate active learning performance, we started with a small training set and iteratively added new examples, either through active learning or random sampling, and then measured validation performance. Our results showed that active learning improved classification accuracy quicker than random sampling, particularly for training sizes of 20–800 examples for the SVM and 200–600 examples for the CNN. We also benchmark overall classification performance against more conventional ML approaches. We found that the CNN outperformed the SVM and other baselines. See Appendix Section 2 for details of the model, Section 3 for comparisons and how classifier performance improved with active learning versus conventional random sampling, and Section 4 for details of different baseline specifications for the ML models. Overall, the best iteration of our CNN with active learning reached a 90.4% classification rate on the validation set, with most other iterations moving between 86% and 90%. This classification rate is excellent and meets or exceeds conventional machine learning benchmarks.

We also look at, in Appendix Section 6, some tests of construct validity and compare our measure of delegation, aggregated up to the bill level, to results from McCann et al. (2022), Farhang and Yaver (2016), and Epstein and O'Halloran (1999). We find a 0.53 correlation between our measure of bill level delegation and McCann et al. (2022)'s measure of the number of agencies delegated to in important statutes based on a database of federal agency regulatory activity, and we find a 0.93 correlation between our measure and Farhang and Yaver (2016)'s measure of the count of regulatory commands in each bill (again, limited only to landmark legislation). Additionally, we rank-order committees based on the average delegation ratio of bills referred to those committees, and compare that rank-ordering to (Epstein and O'Halloran's 1999, 205) ordering of committees based on discretion of their reported bills. The Spearman's ρ for the correlation between these two rankings is 0.81, indicating a strong positive correlation. We find these results to be strong indicators of the validity of our measure.

5 | Delegation and Interbranch Relations Results

With delegation predictions for each version of each bill section from the 110th and 111th Congress, the remainder of the paper will use this novel data to address important substantive questions about interbranch relations. First, we present descriptive data on the distribution of delegations in introduced bills across cabinet-level agencies, broken down by sponsor party. Then, we examine how individual sponsor-level characteristics affect delegation decisions in introduced legislation. We also present an evaluation of the role of delegation in the bill advancement process, made uniquely possible by the fact that we have measurements of delegation from bills as they exist at each stage in the process.

Finally, we use our delegation data to evaluate aspects of the ally principle. Here, we are particularly interested in testing whether congressional delegation decisions are more responsive to stable agency characteristics-like ideology and structural independence or to shifts from divided to unified government. Our results suggest that congressional decisions about delegation are shaped by the stable contours of the administrative state rather than the top-level political leadership that changes with presidential administrations.

5.1 | Delegation in Introduced Legislation

In this section, we attempt to discern the delegatory approach of Congress by observing patterns in introduced legislation from our new dataset. There is a voluminous literature on bill introductions in Congress (e.g., Schiller 1995; Sulkin 2005; Woon 2008; Waggoner 2019), as scholars have viewed bill sponsorship as an arena in which individual legislators have a relatively free hand to express their legislative priorities and take positions that may be popular with key constituencies.[9] While introduced bills have been studied for their ideological content (e.g., Woon 2008) or their specific policy focus (e.g., Sulkin 2005; Waggoner 2019), corresponding attention has not been paid to the recipients of delegated authority in introduced legislation. From an interbranch relations standpoint, Moe (1989) and others emphasize that the bureaucracy is the product of competition among political principals pursuing different, often individualistic, goals, and we may expect legislators' preferences in these competitions to be reflected in their introduced legislation.

While the battle between legislative committees for jurisdictional turf has been well documented (e.g., King 1997; Clinton et al. 2014), the role of individual legislators and bill introductions in these battles is less well understood. Work by Bertelli and Grose (2009) demonstrating that administrative agencies use their discretionary authority to benefit ideological allies in Congress suggests that we should see legislators introducing bills that delegate more authority to ideologically proximate agencies. In most issue areas, legislators have multiple options of implementing agencies to which they may delegate authority (Farhang and Yaver 2016). We expect ideological differences between these agencies to be an important factor influencing legislators' delegation decisions.

Delegation decisions may also be affected by the partisanship of the president and the existence of unified or divided government. The transition from the Bush to Obama administrations provides an opportunity in our data to view the congressional majority party remaining stable (Democratic House and Senate) while the White House changes parties. If there are substantial shifts in delegation moving from divided to unified government, this would be strong evidence that MCs are treating the agencies as the sum of their political appointments; if there are minimal or no changes, this would suggest that MCs treat the agencies as stable reflections of the bureaucracy: more consistent with a semi-permanent administrative state (Lowande 2018).

Figure 2 shows the proportion of bills delegating to cabinet-level agencies by sponsor party for the 110th and 111th Congresses, respectively. Within each Congress, there are a few partisan patterns that could feasibly be driven by ideological considerations—Republicans sponsoring more bills delegating the Department of Defense and the Department of Homeland Security compared with Democrats sponsoring more bills delegating to the Department of Education. However, most of the partisan differences in delegating activity are not particularly large. Additionally, this descriptive look at the data does not suggest large shifts in delegating behavior brought about by a change in presidential administration.

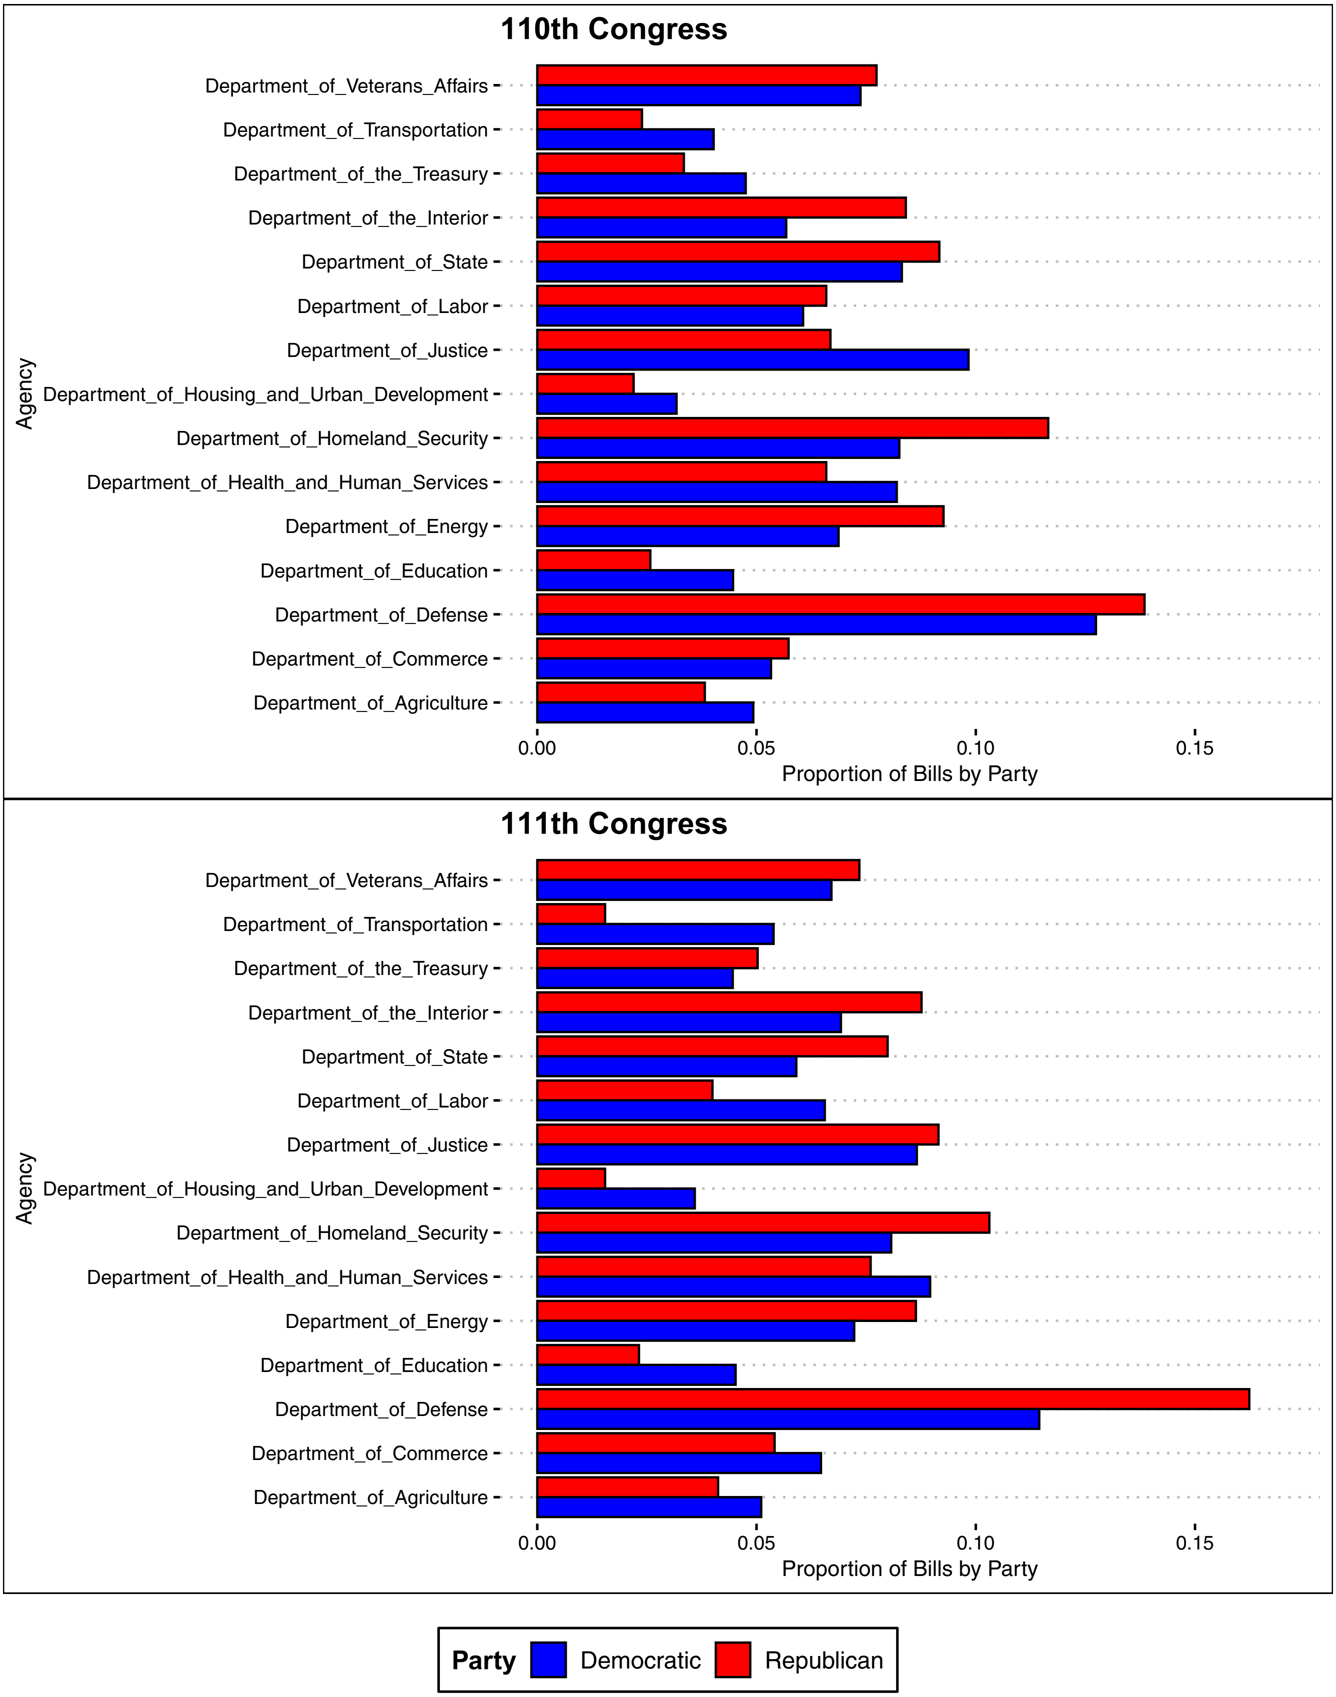

FIGURE 2 | Proportion of bills by party delegating to cabinet level agencies.

In the analyses that follow, we subject these descriptive trends to a number of systematic tests. Specifically, we explore the relationship between individual sponsor-level characteristics—such as partisanship, ideology, and committee assignments—and the ideology of agencies that receive delegated authority in introduced legislation.

These analyses require a measure of agency ideology—or the aggregated policy preferences of each executive branch agency. We use a measure developed by Clinton et al. (2012), who administered a survey to agency appointees and careerists asking them to take positions on salient congressional roll call votes, which generate agency-level ideological scores comparably scaled to those for MCs. To transform these agency-level scores into bill-level scores, we create an average of the ideological scores for each agency that receives delegated authority in a bill, weighted by the prevalence of each agency in that bill.[10] The following equation is used to arrive at these bill-level agency ideology scores:

$$ideo_j = \frac{\sum_{i=1}^{n} \theta_i d_{ij}}{\sum_{i=1}^{n} d_{ij}}$$

in which the subscript j indexes bills, the subscript i indexes agencies, θ_i is an agency-specific ideal point, and d_{ij} is an agency-bill-specific count of delegations.

Our first set of analyses examines the relationship between sponsor partisanship and the ideology of agencies receiving delegated authority in introduced bills. The models presented in Table 2 are linear regressions fit separately on House bills and Senate bills. The dependent variable is $ideo_j$, the bill-level agency ideology score described above. Both models include fixed effects for the bill-level Policy Agendas Project (PAP) topic area code, and a Congress variable.[11] The covariate of primary interest is the partisanship of the bill sponsor.

Like with DW-NOMINATE, the $ideo_j$ scale increases as policy conservatism increases. The positive and significant coefficients on the GOP Sponsor variable in both models demonstrate that bills sponsored by Republicans delegate to more conservative agencies than bills sponsored by Democrats. The effect size is larger in the House than in the Senate—just over one-fifth of a standard deviation of the dependent variable across all introduced HR bills.

The inclusion of fixed effects for PAP topic area codes ensures that the ideological differences found are not driven by Democrats and Republicans sponsoring bills on different topics. Instead, these results demonstrate that the ideological differences apparent across sponsor party hold within policy areas. For example, conditional on introducing an environmental bill, Democrats are more likely to delegate authority to a liberal agency like the Environmental Protection Agency, whereas Republicans are more likely to delegate to a more conservative agency with related jurisdiction, such as the Coast Guard.[12]

TABLE 2 | Ideological delegation and sponsor party.

	Agency ideol. profile	
	House	**Senate**
GOP sponsor	0.181*	0.075*
	(0.023)	(0.032)
Num.obs.	5241	3053
R^2	0.262	0.212
Topic area FEs	✓	✓
Congress FEs	✓	✓

Note: OLS models fit on bill-level data. DV is the bill-level agency ideology score, $ideo_j$. Model 1 is fit on all introduced HR bills, and Model 2 is fit on all introduced S bills.
$^*p < 0.05$.

Figure 3 plots the marginal effect of bill sponsor ideology (measured by DW-NOMINATE) on $ideo_j$, the bill-level agency ideology score. This figure comes from a model like those presented in Table 2, except with a covariate for sponsor DW-NOMINATE instead of sponsor partisanship. The model was fit on House and Senate bills together, and the results demonstrate a positive and statistically significant relationship between sponsor ideology and bill-level agency ideology score.

5.1.1 | Committees and Delegation

In addition to partisanship and ideology, institutional factors such as committee membership likely affect legislators' preferences over delegation. There is a long line of literature that explores the effects of standing committee service on legislator behavior and policy outcomes (e.g., Deering and Smith 1997; Grimmer and Powell 2013), and more specifically, work by Schiller (1995) shows that serving in a committee leadership position affects legislators' bill introductions.

Given the well-documented jurisdictional struggles between standing committees, we expect that membership on a particular committee gives a legislator a stake in growing the portfolio of the executive branch agencies their committee oversees. This should be reflected in bill introduction behavior, with legislators delegating more to agencies their committee oversees. We expect this trend to be especially strong in bills introduced by committee chairs, as chairs may be in a position to benefit even more from augmenting the delegated authority of agencies under their purview (Weingast and Moran 1983).

To test these expectations, we model the number of delegations in introduced bills as a function of the sponsor's service on (or chairmanship of) a committee to which the bill is referred. Just as in the models above, our modeling strategy here holds bill policy area constant. We also opt here for a within-member design, including a sponsor fixed effect. Therefore, the variation we are leveraging in Table 3 comes from whether or not a bill introduced by a given member is referred to a committee on which that member serves.

Table 3 shows the results from four different models: two for each chamber, one looking at the effect of referral committee membership on expected delegation count, and the other looking at the effect of referral committee chairmanship on the expected delegation count. Bills that are referred to a committee on which their sponsor serves have roughly two more delegations in the House and roughly one more delegation in the Senate compared to bills introduced by the same member but referred to different committees. As expected, the effect is stronger for committee chairs.[13] Bills sponsored by the chair of the referral committee have roughly five more delegations in the House and nearly

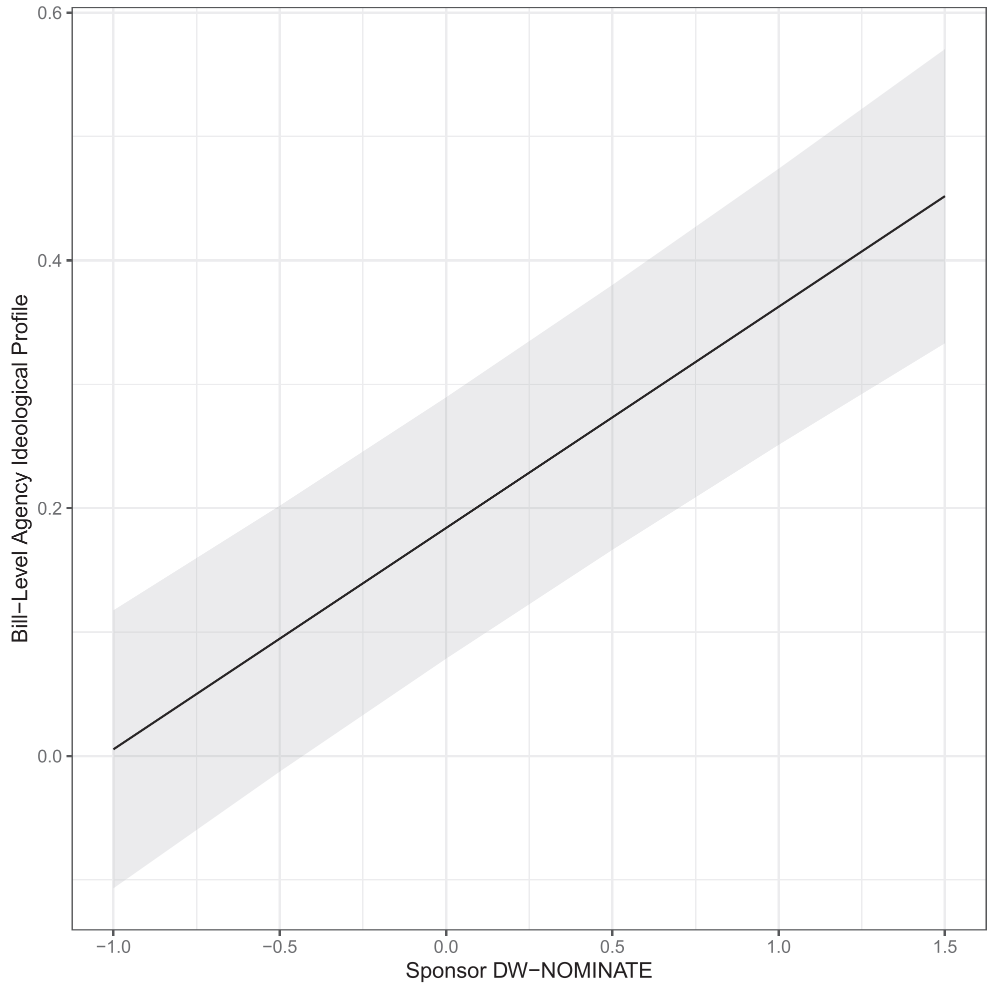

FIGURE 3 | Ideological delegation in introduced bills. Marginal effects plot from regressions of $ideo_j$ on sponsor DW-NOMINATE, conditioning on PAP topic area and Congress.

three more delegations in the Senate compared to bills introduced by the same member but referred to other committees.[14]

These results are in line with expectations developed by Bawn (1997), who posits that legislators serving on the relevant committee of jurisdiction prefer to delegate broad authority to agencies they oversee, whereas off-committee legislators prefer to constrain the authority of those agencies with stricter statutory controls. The findings are important not only for illuminating the effect of committee membership on legislator preferences over delegation, but they also have implications for lawmaking and the shape of the administrative state. Given declining opportunities for members to amend legislation on the floor (Lynch et al. 2020), committee preferences for increased delegation to agencies within their jurisdiction are unlikely to be overturned by off-committee members.[15]

5.1.2 | Delegation Throughout the Legislative Process

To this point, our analyses have shed light on factors that influence delegation in introduced legislation. Because our dataset includes all versions of all bills in two Congresses, we are also able to explore how delegation affects the progress of introduced bills towards becoming law. We are interested here in how the delegation ratio of a bill affects the likelihood of that bill advancing through the legislative process. The delegation ratio (Epstein and O'Halloran 1999) is a ratio of the number of delegating sections in a bill to the total number of sections in that bill.

We model bill advancement separately in the House and Senate, looking at the committee stage, the floor stage in each parent chamber, and the enactment stage.[16] We model each stage separately with a logistic regression, including a covariate for the "Important Bill" indicator from the Congressional Bills Project (Adler and Wilkerson 2017), as well as fixed effects for Congress, PAP topic area code, and bill sponsor. Our main variable of interest, with coefficients plotted in Figure 4, is the delegation ratio of each bill.

In the House, the delegation ratio is positively associated with a bill being reported out of committee, while in the Senate, the relationship is not statistically significant. Moving on to the next stage, among bills that are reported out of committee in both chambers, the delegation ratio is negatively associated with passing the parent chamber. Among the set of bills that pass their chamber, the delegation ratio is also negatively associated with becoming law.

At each stage, we model the probability of advancement as a function of bill-level characteristics as they exist coming into that stage. For example, the results for the committee stage capture the effect of the delegation ratio in introduced bills on the probability of those bills being reported from committee. While these results can tell us something about the characteristics on which bills are being selected for advancement from each stage, these analyses do not speak directly to the question of how advancement changes these bills.

TABLE 3 | Sponsor committee status and the number of delegations.

	Num. of delegations			
	House		Senate	
	(1)	(2)	(3)	(4)
Member of ref. cmte.	0.731* (0.044)		0.322* (0.056)	
Chair of ref. cmte.		1.685* (0.134)		1.014* (0.115)
Num.obs.	13,414	13,414	7420	7420
AIC	40,279.0	40,421.7	24,338.7	24,288.9
Topic area FEs	✓	✓	✓	✓
Sponsor FEs	✓	✓	✓	✓

Note: Negative binomial models fit on bill-level data. DV is the number of delegations in each bill. Models 1 and 2 include the introduced versions of all HR bills. Models 3 and 4 include introduced versions of all S bills.
*$p < 0.05$.

To get leverage on this question, we conduct a series of paired *t*-tests comparing the mean delegation ratio in bills that advance from one stage to the next. For each chamber, we look at the difference in means between introduced and reported versions, reported and passed versions, and passed and enacted versions. The pattern of results is the same in each chamber: committee markup marginally increases the delegation ratio, the delegation ratio then remains unchanged when reported bills pass their parent chamber, and finally, the delegation ratio decreases marginally when these chamber-passed bills go on to become law. These results are available in Appendix Section 12.

The results at the committee stage are consistent with the findings in Table 3, which demonstrate committee members prefer more delegation to the agencies their committee oversees. Not only are committees in the House selecting bills with higher delegation ratios to report out, they are also using committee markups to increase the delegation ratio of these bills further.

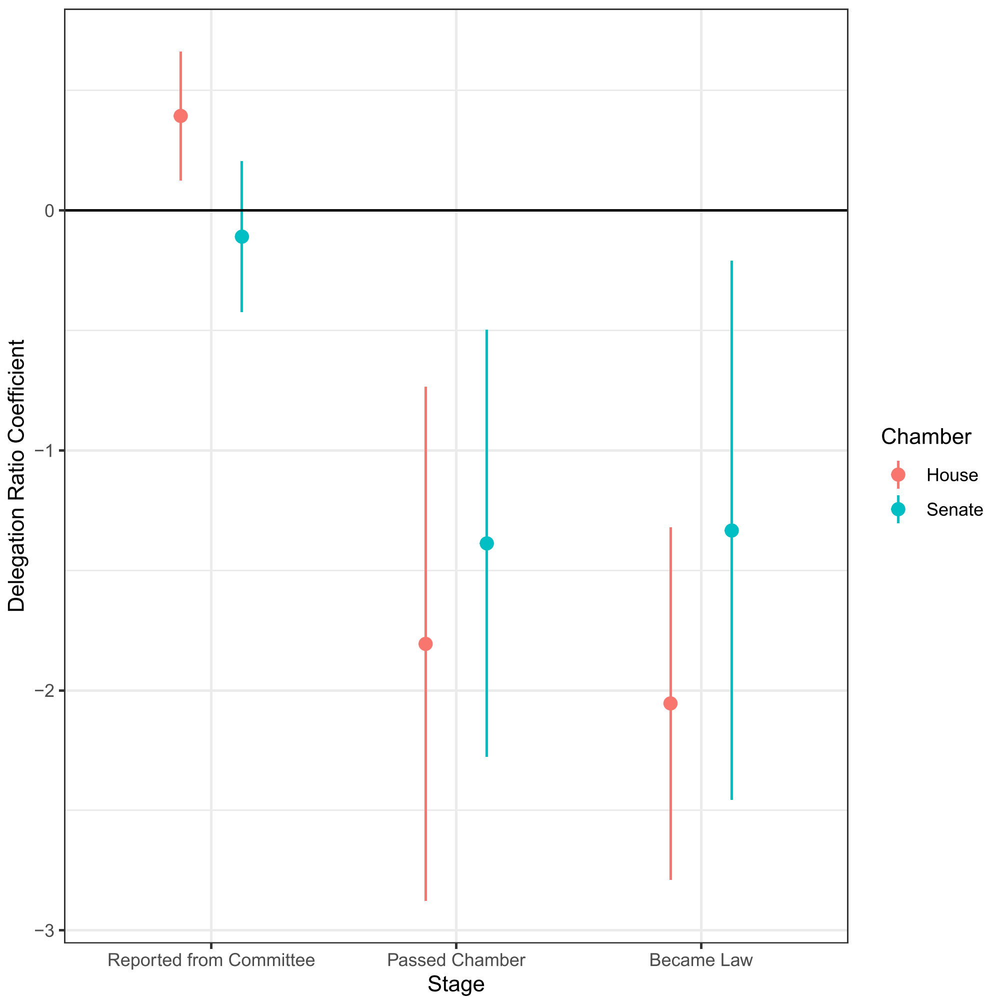

FIGURE 4 | Effect of bill delegation ratio on advancement. Coefficient from a regression of the delegation ratio against each separate stage of bill advancement. We include fixed effects for Congress, PAP topic area code, and bill sponsor.

TABLE 4 | Estimating delegation ratio as a function of regime type and agency independence and ideology.

	Dependent variable					
	Delegation ratio					
		HR bills		S bills		
	HR bills	(Passed house)	S bills	(Passed senate)	Laws	Mayhew
Unified Government	0.010	−0.030	0.026	0.162*	−0.087	−0.312
	(0.022)	(0.031)	(0.031)	(0.076)	(0.080)	(0.331)
|Agency Ideol.—H.Maj. Median|	−0.051*	−0.170*				
	(0.025)	(0.037)				
|Agency Ideol.—S.Maj. Median|			−0.068*	−0.368*		
			(0.032)	(0.080)		
|Agency Ideol.—C.Maj. Median|					−0.455*	−0.545
					(0.100)	(0.320)
Agency Independence Profile D1	−0.340*	−0.352*	−0.325*	−0.135	−0.236*	−0.981
	(0.023)	(0.033)	(0.034)	(0.073)	(0.081)	(0.555)
Agency Independence Profile D2	−0.054*	−0.147*	−0.114*	−0.528*	−0.344*	0.258
	(0.022)	(0.034)	(0.030)	(0.080)	(0.095)	(0.496)
Num.obs.	8520	4074	4301	775	492	24
R^2	0.337	0.512	0.213	0.432	0.586	0.405
Topic area FEs	✓	✓	✓	✓	✓	
Bill version FEs	✓	✓	✓	✓		

Note: The three ideological distance variables in the table are bill-level variables measuring the absolute difference of the DW-NOMINATE score of the median majority party member in each chamber (or in both chambers combined, as is the case for the model fit on bills that became law) and the weighted average of the Clinton et al. (2012) agency ideology scores for agencies receiving delegated authority in each bill.
*$p < 0.05$.

The amending process on the floor of both parent chambers does not reduce the delegation ratios of these bills, and they are only marginally reduced in whatever negotiations are necessary for enactment. It appears that committee preferences for more delegation largely win out in the lawmaking process.

5.1.3 | Agency Characteristics, Presidential Partisanship, and Delegation

Perhaps the most prominent finding from the extensive literature on delegation is the ally principle, which posits that interbranch ideological conflict drives variation in legislative delegation decisions. Specifically, the ally principle states that the scope and amount of delegation decrease as the ideological distance between the legislature and the recipient of delegated authority—usually an executive agency—increases.

In this final section, we construct a set of empirical tests to evaluate scope conditions of the ally principle. Much of the literature on delegation operationalizes the ally principle by comparing instances of unified government with instances of divided government (Epstein and O'Halloran 1996, 1999; Huber and Shipan 2002; Farhang and Yaver 2016), implying that it is the partisanship of the president that drives variation in congressional delegation. However, given the diversity of executive branch agencies—both from an ideological and a structural standpoint—it is likely that the effect of regime type on congressional delegation decisions is impacted by specific agency characteristics. We have already discussed agency ideological measurements in previous sections, and in these final analyses we will also use measures of agency independence to see how these structural characteristics affect delegation decisions.

Agency independence generally refers to the insulation of executive branch agencies from political pressure exerted by Congress or the president. To operationalize this concept, we use measures developed by Selin (2015), who estimated agency ideology on two dimensions—one tapping limits on the appointment of key agency officials and the other tapping the extent to which agency decisions are subject to political review. In order to turn these agency-level measures into bill-level measures, we create average scores at the bill level, weighted by the prevalence of each agency in each bill, in the same manner that we created bill-level agency ideology scores.

To test for the conditions under which the ally principle operates through presidential partisanship, we fit a series of regressions. For our dependent variable, we use a bill's delegation ratio, based on the assumption that bills with high delegation ratios delegate substantively more authority, while bills with low delegation ratios delegate substantively less authority. This fits well with the discussion in Epstein and O'Halloran (1999, 94–95), who find that laws in their dataset with low delegation ratios "delegated very little authority to the executive," while laws with high delegation ratios "granted the executive great leeway to set policy."

We fit a series of regressions predicting a bill's delegation ratio based on the distance between bill-level ideology and the relative majority party medians, two dimensions of agency independence, and a dummy variable for unified government.[17] Because the delegation ratio is constrained between zero and one, we use a beta regression to fit the model.[18]

The models presented in Table 4 below are fit on six different subsets of data—a subset of all HR bills introduced, a subset of all HR bills that passed the House, a subset of all S bills introduced, a subset of all S bills that passed the Senate, a subset of bills that were enacted into law, and finally a subset of important bills by Mayhew's classification (Mayhew 1991).[19] This setup allows us to observe variation in the effects of aggregate agency ideology and independence on delegation ratios across different legislative contexts. Note that the actual agencies' scores will not change between Congresses, so the variation in these ideological variables is based on the constellation of agencies receiving delegated authority in each bill. Models fit on all introduced bills will include messaging bills not intended to become law, so evidence of the ally principle may be diluted in these larger samples. Models fit on bills that passed either chamber or became law yield estimates of the effects of agency characteristics on delegation in actual lawmaking.[20]

Following the logic of Cox and McCubbins (2005), we assume all bills that pass a legislative chamber have the support of the majority party in that chamber, which for our period (110th and 111th Congresses) is the Democratic party in both chambers. Therefore, the coefficients on the agency ideological profile variable can be interpreted as the expected change in delegation ratio based on the ideology of the agency or agencies receiving delegated authority.[21] Similarly, the coefficients on both agency independence profile variables capture the expected effect of structural agency independence on the delegation ratio.

The dichotomous indicator for unified government straightforwardly tests the main expectation of the ally principle as traditionally operationalized. In our specific case, coefficients on this variable can be interpreted as the change in delegating activity from a Democratic Congress attributable to the change in the presidency from George W. Bush to Barack Obama, the same change exploited by both Lowande (2018) and Eldes et al. (2024) to test congressional relations with the changing executive branch.

Because delegation varies across issue area, all models presented below, with the exception of the model fit on the Mayhew bills subset in Table 4, include fixed effects for the PAP topic area of each bill, as coded by the Congressional Bills Project (Adler and Wilkerson 2017). Because the unit of analysis for the data we use is bill-version, we also include a fixed effect for bill version (introduced, reported from committee, passed House, etc.) to account for systematic changes in delegation that occur as bills advance through the legislative process. These fixed effects are included in all models below, except for models that focus exclusively on enacted bills (because these models only include the singular version of bills that became law).

The results presented in Table 4 demonstrate that agency characteristics—specifically agency ideology and independence—retain statistically significant explanatory power even when conditioning on the existence of unified government.[22] In the model fit on bills that passed the Senate—the only model in which the coefficient on the unified government indicator is statistically significant—the coefficient is in the positive direction predicted by the ally principle. However, all other models have a null result for unified government.

The findings on agency ideology and independence are mostly consistent across all models besides the one fit on the subset of bills classified as important by the Mayhew coding system (Mayhew 1991), which is underpowered given the small number of observations. The overall trend demonstrates that the Democratic majorities of the 110th and 111th Congresses wrote less delegation into bills that delegated to more conservative and more independent agencies, regardless of whether these bills were passed during unified or divided government.[23] The results presented in Table 4 suggest that, when making decisions about the scope and extent of delegation, congressional majorities are more attuned to the relatively stable traits of agency ideology and independence than they are to changes in political leadership in the White House and at the top levels of agency hierarchies.[24]

6 | Discussion

In this paper, we have demonstrated how an active learning convolutional neural network for classifying text can be used to study congressional delegation to administrative agencies. We hope that the proposed methods are straightforward and usable for other applications, and that gains in classification accuracy reduce the need for extra documents to be hand-labeled, helping researchers tackle challenging classification problems. This information is part of an ongoing endeavor to learn how Congress uses statutory language to enact its agenda and how the modern legislative process provides oversight and guidance to implement policies. These classifications will help provide the means through which we can test our theories of delegation and congressional oversight on a larger scale and provide the nexus for increased research into the implications of statutory language.

Our substantive analyses in this paper examine the connection between legislator characteristics and delegation decisions, demonstrating that partisanship, ideology, and committee assignments all meaningfully affect the agencies to which legislators prefer to delegate. We also study the role of delegation in bill progression through the legislative process, finding that committees select on bills with high delegation ratios and seek to increase those delegation ratios through their markups. Finally, we examine the workings of the ally principle during a transition from divided government (110th Congress) to unified government (111th Congress), demonstrating that variation in delegation is driven by cross-agency differences in ideology and structural independence.

A notable limitation of the empirical work in this paper is that we focus on only two Congresses. Though we have clear effects, there are some concerns about the generalizability of our results. We acknowledge this limitation, though we note that this particular set of Congresses has been studied extensively for the same reason we do: a change between presidents (and their party) while holding control of Congress constant—see Lowande (2018); Eldes et al. (2024) for two papers studying related questions using the same two Congresses.

An opportunity for future researchers would be to extend our approach—developed for the 109th and 110th Congresses—to additional sessions, either prior or subsequent. Our scope was intentionally limited for several reasons, including the substantial number of observations: once bill texts are broken down by bill version and section, the number of observations per Congress exceeds several hundred thousand. Narrowing the scope was therefore necessary to make data collection and analysis more tractable. Furthermore, we trained our model exclusively on data from the 110th Congress. As a result, applying it to the 111th Congress raises relatively few concerns; statutory language tends not to shift dramatically from one Congress to the next. However, extending the analysis across a broader temporal range—say, from the 103rd to the 118th Congresses—could introduce potential issues related to changes in legislative language over time. In such cases, it would likely be advisable to retrain the model using a broader or differently selected set of training data.

We expect that deep and active learning models of text can be used to further our knowledge of Congressional use of statutory language. While we discuss some concerns about the measurement of discretion in this paper, it would be possible to apply our general method to the systematic identification of the categories of constraints laid out in Epstein and O'Halloran (1999), which would then allow for some measurement of discretion. Recent work by Lowande and Shipan (2021) has made important strides towards the identification of discretion as the president wields it in certain policy areas, but this is a separable task from identifying the *sources* of this discretion. Our methodology could help identify and measure discretion as it exists in legislative texts. As the *Loper Bright* decision makes clear, the emerging legal understanding of discretion requires the identification of explicit textual delegation first. Given this, we believe that our measurement strategy can be particularly helpful in developing a more complete understanding of both delegation and discretion in an interbranch context.

Furthermore, our methodology has multiple potential applications beyond delegation. In particular, we believe that problems involving the classification of complex texts, particularly those requiring expertise for labeling, can benefit from schemes similar to our method. We envision this as the first of many applications of this model to the coding of legislative texts for a number of important features with potential implications for interbranch relations, policymaking, and legislative bargaining.

Data Availability Statement

The data that support the findings of this study are available from the corresponding author upon request. All data and code will be made publicly available conditional on acceptance of submission.

Endnotes

1 While it was ultimately the *Loper Bright* decision that overturned *Chevron* deference, this was the culmination of an evolution in administrative law going back many years. See Merrill (2024); Vermeule (2025) for a discussion of how uses of *Chevron* decreased in the preceding decades and had been unused by the Supreme Court since 2016.

2 Unless that agency also showed up in the training data which, given that there are hundreds of agencies, is likely to miss many.

3 This is a limitation of grammar-based approaches to identifying delegation. Since our method adapts and learns as we hand-code more uncertain observations, it begins to cover the realm of possible mismatches more broadly.

4 This process is depicted in Figure S2.

5 Various other supervised learning methods can accommodate active learning, and the choice of model is the researcher's preference.

6 To prepare our text for the model we have described, we performed the following preprocessing steps: converted all text to lowercase, removed whitespace and punctuation, and replaced all digits with a "digit" token. Then, each word token was assigned a unique integer identifier, creating a dictionary between words and corresponding integers. This dictionary was used to map the words of each bill section to a sequence of integers, and these sequences of integers were used as the inputs to our neural network model.

7 Most often, if Congress refers to a governmental entity, it is an administrative agency. We make an exception for delegating to the courts, states, and local governments because those tasks are defined differently.

8 A potential concern with our approach is that the data we generate from our models is only as good as our human coding. We did manual evaluation and validation of the hand-coding that made up the training set, and we've included the coding instructions in Appendix Section 9. The training set with the hand labels is provided as part of the replication files.

9 But see Ritchie (2023) on how the existence of cross-pressures for many legislators makes it advisable to look to other, less public-facing, behaviors for indications of legislator's policy priorities.

10 Importantly, our measurement strategy here varies from Clinton et al. (2012), who assign a primary recipient of delegated authority for each bill. In contrast, our measure uses information on all agencies receiving delegated authority in a bill. Over 55% of the bills in our data delegate to only one agency, making agency-level information completely harmonious with bill-level information for those observations. To check if our findings are robust to different treatments of the remaining bills—those that delegate to multiple agencies—we present a number of different approaches in Appendix Section 13.2. Results from these analyses are substantively similar to those presented here.

11 The coefficient on the Congress variable, which is just an indicator variable for unified government in the 111th Congress, is not significant in either model, demonstrating that legislators are not systematically changing their ideologically-driven delegation decisions based on the party of the president.

12 Models presented in Appendix Section 11 bear this specific example out.

13 In Appendix Section 13.1, we present results for two alternate specifications. One includes the number of words in each bill as a covariate, which accounts for the fact that committee chairs often sponsor long bills, and another uses the delegation ratio as the dependent variable. The results of these sets of models are substantively the same as those presented in Table 3.

14 The substantive effect sizes in this paragraph are reached by exponentiating the negative binomial coefficients presented in Table 3.

[15] As documented by Curry (2015), party leadership has become more active in making post-committee adjustments to legislation through self-executing amendments or facilitating bypass of the committee consideration process entirely. Therefore, the extent to which committee preferences for increased delegation actually make it into law appears increasingly contingent on party leadership.

[16] These analyses do not account for bills that bypass committee in either chamber, as the committee stage analyses rely on the existence of both an introduced and a reported version of each bill.

[17] Some models also include fixed effects for policy area and for bill version. More details on these measures are provided in the subsequent section.

[18] Similar results are found if we fit our models with OLS. See Appendix Section 13.5 for details.

[19] The models are fit on the subset of bills that delegate authority to actors in the executive branch; bills that do not include any delegation are excluded. What is being analyzed is variation in the *amount* of delegation in each bill (measured by the delegation ratio), as opposed to the initial decision of whether or not to delegate.

[20] Appendix Section 13.7 includes models fit on subsets of the data based on measures of substantive importance. The results from those models are similar to those presented in the paper, demonstrating that the main findings are not driven by the inclusion of legislation of little substantive importance.

[21] The ideology variables are constructed as the absolute difference between the ideological score of the median Democrat in the relevant chamber and the bill-level agency ideological profile score (the weighted average of the ideological score of each agency delegated to in each bill). For models fit on bills that passed the House, the score is constructed using the median House Democrat in each Congress. For models fit on bills that became law, we take the median of House and Senate Democrats pooled together by Congress. Models in Appendix Section 13.6 use two differently constructed ideological measures, and the results from those models demonstrate that our findings are robust to different assumptions. One set of models uses the raw agency ideological scores (without subtracting them from any legislator's ideological score), and another uses the absolute distance between the bill sponsor's DW-NOMINATE score and the agency's ideological score.

[22] In Appendix Section 13.3, we demonstrate that the mostly null results with respect to the unified government variable are robust to a simple model structure that regresses delegation ratio on a dichotomous indicator for unified government, without agency-specific variables.

[23] In Appendix Section 13.4, we present analyses in which we interact the second dimension of agency independence with both our unified government indicator and our agency ideological variables. The results demonstrate that Congress delegates more to politicized agencies during unified government, and that more politically independent agencies receive relatively greater delegation during divided government. The interactions between agency independence and agency ideology show that the classic ally principle result is borne out for highly-politicized agencies, but that the opposite pattern is found for highly independent agencies.

[24] It is worth noting here that the agency-level ideology and independence measures we use to construct our variables are time-invariant, which assumes that these constructs are not meaningfully affected by changes between unified and divided government. Models included in Appendix Section 13.6 relax this assumption by using time-varying agency ideology scores from Chen and Johnson (2015). The results from these models largely support the findings presented here. Most variation in delegating activity seems to be driven by ideological differences across agencies rather than ideological shifts within agencies over time.

References

Adler, E. S., and J. Wilkerson. 2017. Congressional Bills Project NSF 00880066 and 00880061.

Anastasopoulos, L. J., and A. M. Bertelli. 2020. “Understanding Delegation Through Machine Learning: A Method and Application to the European Union.” *American Political Science Review* 114, no. 1: 291–301.

Bawn, K. 1997. “Choosing Strategies to Control the Bureaucracy: Statutory Constraints, Oversight, and the Committee System.” *Journal of Law, Economics, and Organization* 13, no. 1: 101–126.

Bertelli, A. M., and C. R. Grose. 2009. “Secretaries of Pork? A New Theory of Distributive Public Policy.” *Journal of Politics* 71, no. 3: 926–945.

Bolton, A., and S. Thrower. 2019. “The Constraining Power of the Purse: Executive Discretion and Legislative Appropriations.” *Journal of Politics* 81, no. 4: 1266–1281.

Chen, J., and T. Johnson. 2015. “Federal Employee Unionization and Presidential Control of the Bureaucracy: Estimating and Explaining Ideological Change in Executive Agencies.” *Journal of Theoretical Politics* 27, no. 1: 151–174.

Clinton, J. D., A. Bertelli, C. R. Grose, D. E. Lewis, and D. C. Nixon. 2012. “Separated Powers in the United States: The Ideology of Agencies, Presidents, and Congress.” *American Journal of Political Science* 56, no. 2: 341–354.

Clinton, J. D., D. E. Lewis, and J. L. Selin. 2014. “Influencing the Bureaucracy: The Irony of Congressional Oversight.” *American Journal of Political Science* 58, no. 2: 387–401.

Cox, G. W., and M. D. McCubbins. 2005. *Setting the Agenda: Responsible Party Government in the US House of Representatives*. Cambridge University Press.

Curry, J. M. 2015. *Legislating in the Dark: Information and Power in the House of Representatives*. University of Chicago Press.

Deering, C. J., and S. S. Smith. 1997. *Committees in Congress*. Congressional Quarterly Inc.

Eldes, A., C. Fong, and K. Lowande. 2024. “Information and Confrontation in Legislative Oversight.” *Legislative Studies Quarterly* 49, no. 2: 227–256.

Epstein, D., and S. O'Halloran. 1996. “Divided Government and the Design of Administrative Procedures: A Formal Model and Empirical Test.” *Journal of Politics* 58, no. 2: 373–397.

Epstein, D., and S. O'Halloran. 1999. *Delegating Powers: A Transaction Cost Politics Approach to Policy Making Under Separate Powers*. Cambridge University Press.

Farhang, S., and M. Yaver. 2016. “Divided Government and the Fragmentation of American Law.” *American Journal of Political Science* 60, no. 2: 401–417.

Grimmer, J., and E. N. Powell. 2013. “Congressmen in Exile: The Politics and Consequences of Involuntary Committee Removal.” *Journal of Politics* 75, no. 4: 907–920.

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

Kiewiet, D. R., and M. D. McCubbins. 1991. *The Logic of Delegation*. University of Chicago Press.

Kim, Y. 2014. “Convolutional Neural Networks for Sentence Classification.” *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*: 1746–1751.

King, D. C. 1997. *Turf Wars: How Congressional Committees Claim Jurisdiction*. University of Chicago Press.

Lowande, K. 2018. “Who Polices the Administrative State?” *American Political Science Review* 112, no. 4: 874–890.

Lowande, K., and C. R. Shipan. 2021. “Where Is Presidential Power? Measuring Presidential Discretion Using Experts.” *British Journal of Political Science* 52, no. 4: 1–15.

Lynch, M. S., A. J. Madonna, and A. S. Vick. 2020. *The Erosion of “Regular Order” in the Us House: A Historical Examination of Special Rules*, 147–167. New Directions in Congressional Politics. Routledge.

Mayhew, D. R. 1991. *Divided We Govern*. Yale University.

McCann, P. J., &. S. Clouser, and C. R. Shipan. 2022. “How Many Major US Laws Delegate to Federal Agencies?(Almost) All of Them.” *Political Science Research and Methods* 10, no. 2: 438–444.

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, no. 2: 243–277.

McCubbins, M. D., R. G. Noll, and B. R. Weingast. 1989. “Structure and Process, Politics and Policy: Administrative Arrangements and the Political Control of Agencies.” *Virginia Law Review* 75: 431–482.

Merrill, T. W. 2024. “The Demise of Deference-And the Rise of Delegation to Interpret?” *Harvard Law Review* 138: 227.

Mikolov, T., I. Sutskever, K. Chen, G. Corrado, and J. Dean. 2013. “Distributed Representations of Words and Phrases and Their Compositionality.” In *Advances in Neural Information Processing Systems*. Curran Associates, Inc.

Miller, B., F. Linder, and W. R. Mebane. 2020. “Active Learning Approaches for Labeling Text: Review and Assessment of the Performance of Active Learning Approaches.” *Political Analysis* 28, no. 4: 1–20.

Moe, T. 1989. “The Politics of Bureaucratic Structure.” *Can the Government Govern?* Brookings Institution Press. 267–329.

Ritchie, M. N. 2023. *Backdoor Lawmaking: Evading Obstacles in the US Congress*. Oxford University Press.

Rodriguez, P., and A. Spirling. 2022. “Word Embeddings: What Works, What Doesn't, and How to Tell the Difference for Applied Research.” *Journal of Politics* 84, no. 1: 101–115.

Schiller, W. J. 1995. “Senators as Political Entrepreneurs: Using Bill Sponsorship to Shape Legislative Agendas.” *American Journal of Political Science* 39: 186–203.

Selin, J. L. 2015. “What Makes an Agency Independence?” *American Journal of Political Science* 59, no. 4: 971–987.

Settles, B. 2009. *Active Learning Literature Survey. Tech. rept*. University of Wisconsin-Madison Department of Computer Sciences.

Sulkin, T. 2005. *Issue Politics in Congress*. Cambridge University Press.

Sunstein, C. R. 1989. “Interpreting Statutes in the Regulatory State.” *Harvard Law Review* 103: 405.

Turian, J., L. Ratinov, and Y. Bengio. 2010. “Word Representations: A Simple and General Method for Semi-Supervised Learning.” In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, July*. Association for Computational Linguistics.

Vannoni, M., E. Ash, and M. Morelli. 2021. “Measuring Discretion and Delegation in Legislative Texts: Methods and Application to US States.” *Political Analysis* 29, no. 1: 43–57.

Vermeule, A. 2025. “The Old Regime and the Loper Bright Revolution.” *Supreme Court Review*.

Waggoner, P. D. 2019. “Do Constituents Influence Issue-Specific Bill Sponsorship?” *American Politics Research* 47, no. 4: 709–738.

Weingast, B. R., and M. J. Moran. 1983. “Bureaucratic Discretion or Congressional Control? Regulatory Policymaking by the Federal Trade Commission.” *Journal of Political Economy* 91, no. 5: 765–800.

Woon, J. 2008. “Bill Sponsorship in Congress: The Moderating Effect of Agenda Positions on Legislative Proposals.” *Journal of Politics* 70, no. 1: 201–216.

Zhang, Y., and B. C. Wallace. 2017. “A Sensitivity Analysis of (and Practitioners' Guide to) Convolutional Neural Networks for Sentence Classification.” In *Proceedings of the Eighth International Joint Conference on Natural Language Processing*, 253–263. Asian Federation of Natural Language Processing.

Supporting Information

Additional supporting information can be found online in the Supporting Information section.