Networks of Power:

Extracting Measurements of De Jure Power from Constitutional Text

December 16, 2021

Abstract

A topic that has been of perennial interest in political science is that of government actors' political power. However, measurement of political power has thus far been a highly costly process involving hand-coding by experts; consequently, only the executive power of a limited number of polities has been examined in depth. This paper proposes a novel methodology to extract measures of political power directly from the text of US state constitutions using machine learning, text as data, and network methods. With these measures, I find support for conventional wisdom surrounding the changes in state executives' relative power over time and validate my the measure both through comparison of predicted power of government actors with the power estimated using hand-coded data and qualitative examination of constitutional texts. This project contributes to the development of a systematic, objective, and cost-efficient method of measuring de jure political power using constitutional texts.

Keywords: de jure power, constitutions, measurement

Introduction

How powerful is an executive compared to a legislature, or how powerful is the electorate relative to either? More generally, what is the best approach to compare how influential one set of institutional actors are relative to others? These questions have been a fixture of political science and have been closely examined by both the comparative politics and political economy literature. In this paper, I present a method by which I propose de jure power can be directly measured from the text of constitutions that applies the concepts and assumptions of both bodies of literature.

Perhaps the most significant efforts to measure political power have been made by the executive power literature. Shugart and Carey (1992) conceptualized power as being derived from the executive's ability to influence policy and the degree of control that they possessed over congress and their cabinet. Consistent with that definition, they developed a popularly cited measure of executive power in which experts assigned power ratings to executives that were increasing in their ability to influence legislation, appoint subordinates, and influence the legislature. Metcalf (2000) proposed some modifications to Shugart and Carey's methodology to allow it to function on semi-presidential regimes. Similarly, Frye (1997) developed a measure of executive power with 27 measures instead of 10.

One could argue that de jure power, a concept that is common in the political economy literature, is a generalization of the power described by Shugart, Carey, Metcalf, and Frye. Acemoglu and Robinson (2008) describe de jure power as power derived from institutions. In their 2008 paper, Acemoglu and Robinson examine the strategic interactions between a citizenry and elites and how their willingness to invest in de facto power influenced institutional changes, leading to differing levels of de jure power. In a parallel to the executive power theorists, Bertocchi and Dimico (2017) and Debs (2010) describe constitutions as being a source of de jure power of the and cite the example of Alabama legislators pass-

ing poll taxes and literacy tests as examples of de jure power being transferred away from African Americans and to white social and economic elites. In an interesting parallel to the executive power literature, the political scientists who do empirical or theoretical research involving de jure power have described constitutions as being a major determinant of which actors have de jure power and which do not (Ager, 2012; Bertocchi and Dimico, 2017).

While the executive power literature has been focused on how institutions, like constitutions, impact executives' cross-country power, the de facto power literature has examined how constitutions have impacted the relative influence of the public relative to political and economic elites.

The existing literature on executive and de jure power has yet to answer the following questions: 1) how do you compare the powers of the executive with those of other important actors, such as the legislature or the citizenry; 2) how does the power of the executive and other actors depend on the powers of other actors with whom they interact; 3) how can these powers be estimated on a broad set of governing documents cost-effectively? I attempt to address these challenges questions by developing a new methodology that builds off the assumptions and design of preexisting measures and extends them so that it is the structure of networks rather than the judgment of experts that determines the relative power of government actors.

This paper first describes the theory and assumptions that are the basis for my proposed measurement. Second, I examine the data to which this measurement will be applied and describe several existing predictions about how de jure has changed over time. Third, I will describe the methodology by which I convert text to a measure of de jure power. Finally, I analyze the data and assess its validity using both predictions from the literature and comparison with hand-coded data.

Theory

The conceptualization of presidential power in the executive power literature is grounded in hierarchical relationships between government entities (Shugart, 2005; Carpenter et al., 2013; Samuels and Shugart, 2010). Shugart, Metcalf, and Frye describe executive power as increasing when the executive is the principal in principal-agent relationships, can impose costs on or dissolve other important actors - such as Congress or the heads of bureaucracy, or acts as a gatekeeper that is able to veto legislation. Essentially, political actors' power increases with the number and strength of the hierarchical political relationships that they have with other actors (Shugart, 2005). While hierarchy is not strictly defined in the executive power literature, I infer hierarchical relationships as being relationships defined in a constitution in which one of two actors can impose coercion or otherwise influence other actors' composition or behavior. The de jure hierarchical relationships described by Shugart and Carey are the basis for the measurements proposed in this paper.

To measure de jure power, I assume that a measure of the political power of an actor should be increasing in the number of hierarchical ties in which it is the dominant actor and in the power of the other actors that the actor dominates. Substantively, this framework implies that presidents can be seen as influential in part because they can appoint and dismiss other powerful government actors, such as members of their cabinet or other government officials. Likewise, the legislature is powerful insomuch as it can influence other influential actors, such as impeaching the president or approving other government officials' appointments.

One natural implication of this framework is that all political actors exist in a network

of hierarchical relationships, and their power is endogenous to the power of other actors. Moreover, any individual actors' influence is dependent on the structure and distribution of power among actors in the entire network. How might one create a metric of power from this highly endogenous set of actors and relationships? Fortunately, perhaps one of the most commonly used network centrality metrics satisfies each of the above criteria: directed eigenvector centrality.

Eigenvector centrality of a node is proportional to the nodes' centralities to which it is connected (Bonacich, 2007). Moreover, it increases in both the number of connections a node has and the centrality of those nodes with which it is connected. Consequently, eigenvector centrality has been used as a proxy for social influence and social power in several studies Mizruchi (1982); Hossain (2009). There are both directed and non-directed versions of eigenvector centrality. In the case of directed eigenvector centrality, actor i's centrality increases as the number of links to them from other actors increases and as the centrality of those actors with incoming ties increases. Rephrased in this paper's terminology, an actor i's directed centrality increases as they appoint and oversee other actors and as their agents' importance increases. In the case of undirected centrality, all ties are seen as bi-directional; i.e. there is no principal actor. While directed eigenvector centrality is more consistent with our theoretical understanding of de jure power, the analysis performed in the paper's final section utilizes undirected centrality out of practical necessity. [1]

Data

Constitutions are broadly considered by both the comparative and political economy literatures to be a source of political power for government actors. This is because they delineate

[1]Establishing directed links requires a natural language processing algorithm that is trained to identify links of the type of interest to the researcher. To train such a model would require a training data set of thousands of data points, which would be beyond the scope of this paper.

relationships, such as whom can appoint and remove whom, what checks and balances exist, and who has the right to vote if such a right is guaranteed. The powers delineated by constitutions are the 'hierarchical relationships' described by Shugart and Carey. They have been used as the basis for many of the indexes used in the presidential power literature (Frye, 1997; Metcalf, 2000; Shugart and Carey, 1992). Likewise, constitutions have been the primary means of identifying de jure power by political economists (Acemoglu and Robinson, 2008; Ager, 2012; Bertocchi and Dimico, 2017). Accordingly, I follow the precedent of these two bodies of literature when I apply this proposed measure of de jure power to the sets of entities and hierarchical links described in US 122 state constitutions gathered by the Maryland State Constitutions Project (Wallis, 2005).[2]

There are several advantages to using the constitutions of US states to estimate de jure power. First, US states typically follow a standard model - the US constitution - that allows for easier comparisons of entities and relationships between constitutions. Second, there is literature with qualitative assessments of how various state actors' powers have evolved over time. This allows me to validate my findings by comparing them to experts' assessments. Finally, US constitutions have been developed over a wide range of time, geographies, and constituencies; consequently, there should be sufficient diversity in their structure to allow for observable variation in the power of various state actors within and between states over time.

In Figure 1 I present the appearance of new state constitutions over time. There are regional discrepancies in the overall number of constitutions produced and the periods in which they were adopted. First, it is clear that the American Southeast has been the most prolific producer of new constitutions. In particular, they were active in producing new constitutions immediately after the Civil War. This was due to Reconstruction-era requirements from the federal government that forced Southern states to pass new constitutions, giving African Americans the right to vote and run for office. These rights were undermined in

[2]These constitutions include all constitutional versions and amendments implemented since the United States' founding in 1776 until 2000.

the subsequent wave of constitutional reforms, which can be seen in the 1880s-early 1900s in which restrictions on suffrage were instituted to limit African Americans from participating in Southern politics (Bois, 1910). Other regions were less prolific, with Western and Midwestern states mostly implementing one or two constitutions after their formation in the late 19th and early 20th century.

Figure 1: Adoption of New Constitutions over Time by Region

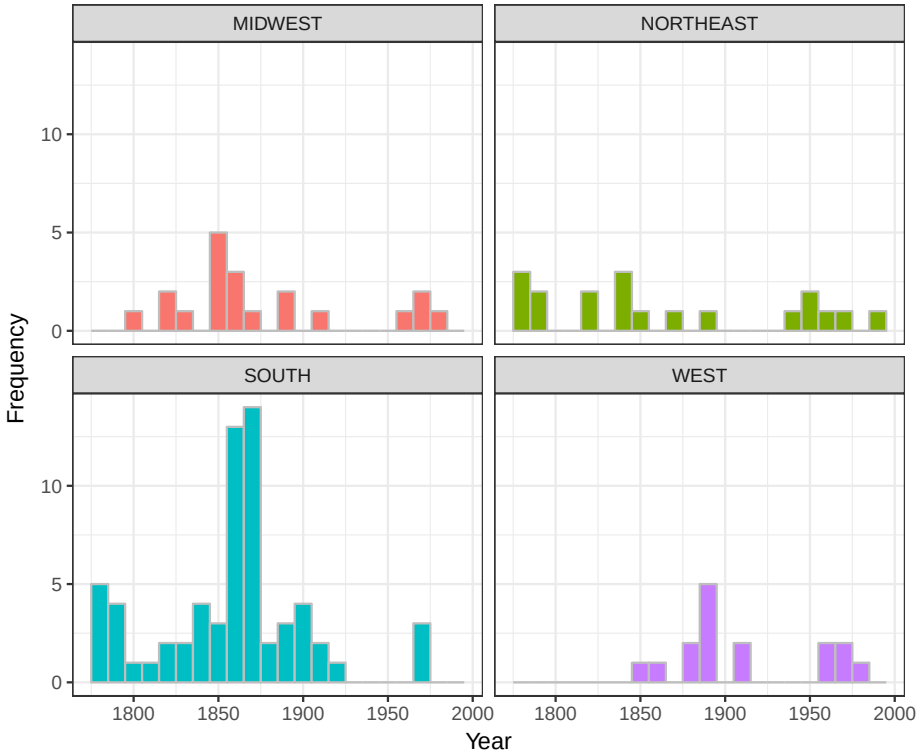

Methods

Several parameters must be extracted from our data to estimate the de jure power government actors have relative to one another. First, I identify entities described in each constitution, actors such as the governor, the Senate, or the House of Representatives. Second, I predict relationships between those entities. Third, I generate networks for each consti-

tution. Finally, I derive several inferential statistics from those networks. In this section, I describe all of the methods I use to generate a measure of de jure power for each actor within each state constitution.

Entity Identification

As described in the theory section, constitutions contain sets of actors, which I refer to as entities. These entities can be organizations, such as the House of Representatives and the Senate, or political offices, such as the Governor and Secretary of State. I only consider actors with independent responsibilities or powers to be entities.

To identify entities in the data, I have utilized a combination of two approaches: named entity recognition (NER) and the dictionary approach. A NER algorithm is designed to identify and label entities (individuals, organizations, etc.) from text and classify them into preexisting categories (Florian et al., 2004, 2003; Zhou and Su, 2001; Ritter et al., 2011). The algorithm I use was developed at Stanford by Finkel et al. (2005); their algorithm operates in the following way: it first tokenizes raw text into sentences; second, it breaks those sentences into individual words; third, it tags those words with parts of speech labels (ex. object, subject, verb); fourth, it detects entities using a trained naive Bayes machine learning model using the context and labels of the words; finally, it assigns a label to each entity which identifies whether the entity is a person, organization, or a location.

Stanford NER had some clear advantages over alternatives, such as MIT's NER algorithm and the python package nltk. First, it had a far lower error rate than the other algorithms. For instance, MIT's NER algorithm was guilty of pairing entities with unrelated words that resulted in them being misclassified as new entities. As an illustration, "the general assembly" would occasionally be misidentified as "the general assembly brings." Each

misclassification requires hand cleaning of entities which can be very time-consuming, as each algorithm identifies thousands of entities in the state constitution corpus of documents. Second, the Stanford NER algorithm is better documented than the alternatives and has been applied in many other academic settings (Benton et al., 2011; Yosef et al., 2011; Lingad et al., 2013).

However, all NER algorithms I examined featured a common flaw: they failed to identify entities related to the electorate, such as "voter," "district," or "citizen." Consequently, I found it necessary to supplement the entities identified by Stanford NER with those identified with a more reliable method: human coding.

One advantage of using US state constitutions relative to those at the national level is that there is relatively little variation in actors' names and their respective positions in state governments. For instance, all state constitutions refer to the executive as the "governor" and the legislature as the "General Assembly." Thus, even a small sample of constitutions should capture many of the most important entities that occur in other constitutions.

Accordingly, to identify entities that the Stanford NER algorithm may have missed, I first selected a random sample of 3 constitutions - New Hampshire's 1784 constitution, North Carolina's 1776 constitution, and Connecticut's 1965 constitution - and split them up into sets of sentences. Second, I manually identified entities missed by the Stanford NER algorithm. Using these entities hand-coded entities and those identified by Stanford NER, I created a dictionary of entities that may occur in other non-training constitutions.

Once I had gathered entities from the three hand-coded constitutions and with the Stanford NER algorithm, I assigned a common set of labels based on their function to allow for comparability between constitutions. For instance, while Virginia's constitution refers

to its lower house as "The House of Delegates," they are functionally identical to the lower houses in other states, which were typically referred to as "The House of Representatives." I assigned labels based on their underlying function in each constitution to these entities with common roles across states, such as "lower house" or "upper house."

One label of particular importance was that of the "public." The public was defined as the citizenry, specifically citizens who possessed a hierarchical relationship with other government entities. Thus, the label included references to districts – which are geographically segregated sets of voters – voters and other such entities.

Relationship Identification

Now that entities have been identified, I identify hierarchical relations between them to measure their de jure power. Identifying hierarchical relationships in the text of state constitutions was a three-stage process. First, I hand-coded each sentence of the three constitutions used to identify entities - New Hampshire's 1784 constitution, North Carolina's 1776 constitution, and Connecticut's 1965 constitution - to create a training set that could be used to identify relationships in the remaining 122 constitutions.[3] I identified a hierarchical relationship to be one in which a constitution grants one entity the power to induce a change in another's behavior or composition.

In the second stage, I evaluate and train a support vector machine (SVM) model using those sentences and use it to label sentences in other constitutions that are likely to contain hierarchical relationships. I chose a SVM model for two reasons: first, it proved to be more accurate in cross validation classification tests than virtually any other model or combination

[3]For each sentence, I was looking for any relationship in which one actor had the authority to influence another actor's behavior or incentives. I coded each sentence of North Carolina, New Hampshire, and Connecticut's constitutions as one if it contained at least one such relationship and 0 otherwise.

of models, with an average accuracy of 83%; second, it is far faster to compute than other candidate models, such as random forests. One drawback of this approach is that I could only identify sentences that were likely to contain relationships, but not the type or direction of the relationship itself. Thus, unlike the hand-coded data, none of the hierarchical relationships identified by prediction had any known direction. This has clear implications for my measurement strategy, which I discuss in the next section.

The final stage in identifying relationships was detecting the co-occurence of entities in those sentences that were predicted to contain relationships by the SVM model. Clearly, for a relationship to exist between two entities in a sentence, there must be at least two entities that co-occur in a sentence. Thus, even if a sentence was predicted to have a relationship by the SVM model, I considered it to only have a relationship it contained at least two entities. I then identified each entity in the sentence and made the assumption that each pair of entities in that sentence had a hierarchical relationship with one another. While this is a strong assumption, I don't expect it to bias the overall power of any particular entity and thus it should not be a threat to my measurement strategy.[4]

Network Statistics

Eigenvector centrality has some useful properties that make it an attractive measure of de jure power. First, there is substantial precedent in its use as a measure of power or influence in the context of social networks (Yan and Ding, 2009; Mizruchi and Potts, 1998). Second, Costenbader and Valente (2003) finds that eigenvector measures of centrality are robust to the removal of nodes relative to other measures of centrality. Subsequent studies on the robustness of centrality metrics have found eigenvector measures to be at least as robust

[4]Ideally, one would use a methodology that would use the context of a sentence to identify both the existence and direction of a relationship between pairs of actors. However, existing relationship identification algorithms have been trained on problems that to not translate well to identifying links between entities in state constitutions. In fact, when applied to this data set, MIT's relationship detection algorithm failed to identify a single relationship between entities. Retraining their algorithm requires a minimum of several thousand coded training observations, which is beyond the bounds of this study.

as other measures (Borgatti et al., 2006; Frantz et al., 2009). This is useful, as the SVM model's hierarchy labels almost certainly included false positives and negatives.

Using the identified entities and hierarchical links between them, I generate a network for each constitution from each set of relationships and entities I identified in the previous stages. These networks are represented as adjacency matrices. Each row and column has a corresponding entity, and each cell takes the value of the number of non-directed links between entities corresponding to the entities of row i and column j.

These networks are the basis upon which I estimate eigenvector centrality statistics which are used to represent the relative influence of each actor. The specific statistic used is the undirected eigenvector centrality algorithm as implemented in R's igraph library. To make each actor's centrality comparable to one another and to actors in other constitutions, I normalize the eigenvector centralities of each actor to 1. This means that the entity with the highest level of centrality in a given network will have a score of 1. The other entities will be assigned scores representing their centrality relative to the most central actor.

Using a non-directional measure of eigenvector centrality as the primary measure of power between entities has a cost: if eigenvector centrality is non-directional, then actors highly connected to key actors will receive high centrality scores, even if they are themselves agents of the other key actor. For example, in a state with a powerful governor, unimportant actors who frequently appear with the governor – such as the lieutenant governor – will be assigned a large eigenvector centrality, despite not having much formal power of their own. Consequently, I limit my analysis to those actors I expect to be at the top of the government hierarchy - the public, the governor, and the legislature, as I expect other entities' centralities will be biased depending on which large entity they are proximate to. This limitation also highlights the importance of identifying a method to determine the direction of hierarchical

relationships between entities, a problem I leave for future research.

Validation

This section presents the results of three validation exercises used to determine whether the eigenvector centrality statistics generated for each actor in each US state constitution are valuable and consistent with the academic understanding of de jure power. First, I compare actors' centralities from the three hand-coded constitutions with those from out of sample. Second, I compare predictions made by experts on state constitutions with the general trends I observe in the data. Third, I examine whether the de jure power estimates assigned to a few 'extreme cases' is consistent with a qualitative examination of the power described by the text of those constitutions.

Comparison with Hand-coded Data

Perhaps the most important benchmark of any metric generated by an unsupervised or semi-supervised computational process is how it compares to human coded data. In this case, I compare the human coded training constitutions' centralities with the centralities obtained through the process described above. The baseline power estimates are calculated using the directed eigenvector centrality using the hand-coded relationships and entities of the three hand-coded constitutions. I compare them with the undirected predicted centrality scores using OLS regression.

The purpose of this exercise is to determine how well predicted power scores correlate with the baseline measure of the power of various state actors. I expect that the predicted coefficient will be positive and significant, indicating that the predicted centrality scores ex-

plain some of the variation of the baseline centrality score. Additionally, I calculate the true undirected eigenvector centrality of the three hand-coded constitutions and compare them with the estimated eigenvector scores. In this case, the comparison is fairer as both the actual and predicted scores are using the same measure of centrality. Accordingly, I expect the relationship between these two sets of scores to be more strongly correlated than the former case.

Overall, I find that the predicted eigenvector centrality scores are effective at predicting both the undirected and directed baseline eigenvector scores for New Hampshire, North Carolina, and Connecticut. In Table 1, I regressed the actual directed and undirected scores on the predicted scores and found that the predicted scores were statistically significantly correlated with the "true" scores at more than a 99% confidence level. I find that a one-unit increase in predicted scores is associated with a 0.59 unit increase in directed centrality scores and a 0.65 unit increase in undirected centrality scores. As anticipated, the predicted scores were more effective at predicting the undirected scores, though the difference in the coefficients - 0.06 points - is not particularly large.

Table 1: OLS Comparison of Predicted vs Actual Eigenvector Centrality Scores

	True:	
	Directed	Undirected
	(1)	(2)
Predicted	0.586***	0.645***
	(0.066)	(0.061)
Constant	0.004	0.038
	(0.026)	(0.023)
Observations	87	87
R^2	0.477	0.571
Note:	*p<0.1; **p<0.05; ***p<0.01	

To better see which entities and states the predictive model was explaining best and most poorly, I visualized the relationship between predicted and actual scores for both the undirected and directed actual scores (see Figure 2). One clear difference between the two plots is that the directed measure of centrality assigned a centrality score of zero to several entities that received a non-zero score from the predictive model. The reason for this pattern is fairly intuitive: when estimating directed eigenvector centrality, entities' centralities increase in the number of hierarchical relationships that are directed towards them; in other words, they are only seen as central if they have authority to influence other entities. Since there are many entities in this data set that are subordinate to actors but have no subordinates of their own, those actors received a directed centrality score of zero. However, undirected eigenvector centrality makes no distinction between incoming and outgoing links; thus, actors who have no subordinates receive non-zero undirected eigenvector scores. Consequently, in the following sections, I avoid making inferences about actors with lower eigenvector centrality scores, as their scores are likely inflated due to their network position.

Figure 2: The Predicted vs Actual Power of State Entities

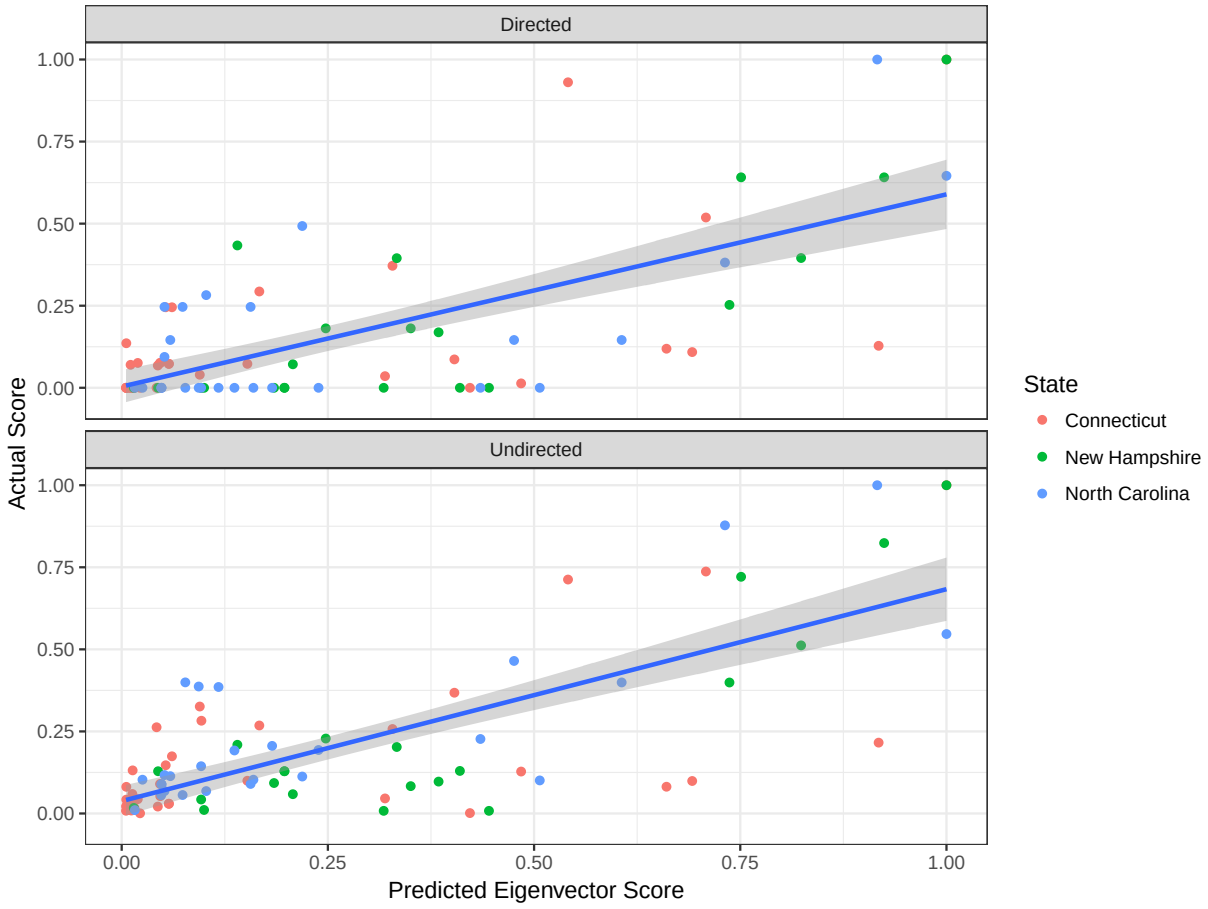

Overall, the comparison between the human and automatically scored constitutions is encouraging. There is a strongly significant and positive association between the predicted scores and the actual scores. While predicted scores are not as effective at predicting the directed eigenvector centrality scores, they still explain nearly 50% of the variation of the true directed eigenvector scores.

However, some of this metric's limitations are also made clearer with this validation exercise. First, there would likely be an improvement in the accuracy of predicted scores if I could predict the direction of hierarchical relationships. Second, there remains a large percentage of the variation in the true scores that is not well explained by the predicted scores. One

method of reducing this error would be to explore additional machine learning approaches that may improve out-of-sample prediction accuracy. An additional step I could take would be to increase the training data's size. While it is time-consuming to hand-code constitutional text, even a training set of only 750 observations could produce reasonably accurate predictions about the true eigenvector scores of constitutions. I imagine that adding a few thousand additional training observations would lead to meaningful improvements in model accuracy. Finally, these results imply that there remains a great deal of value in human coding of constitutional texts; were there sufficient funding, a crowd-sourced approach to entity and relationship identification would likely lead to dramatic increases in the accuracy of predicted centrality scores (Benoit et al., 2016).

Executive Power

If the method of measuring de jure political power described in this paper is valid, one would expect that its findings would be consistent with the findings of experts on state constitutions. One pattern that several scholars on state constitutions have identified is that the power of the governorship has increased in several states, beginning in the mid 20th century (Graves, 1966; Adrian, 1968). One of the most clear predictions about changes in the political power of various intra state political actors over time was made by Elazar (1982). He predicted that since the middle of the 20th century, state constitutions have adopted the Managerial style of governance, which placed more power in the hands of elected governors while taking power away from state legislatures. Moreover, he anticipated that Hawaii and Alaska are states that are representative of this broader set of changes to state constitutions. Elazar's predictions represent an opportunity to assess the validity of this measure of de jure power: if we observe an increase in the de jure power of the governor following the 1940-1950s, then that would suggest that this is a valid measure of de jure power.

Overall, the data is strongly consistent with Elazar's predictions. Figure 3 displays the trend lines for the eigenvector centralities of the public, the general assembly, and the governors of all 50 states from 1776 to 2003 along with 95% confidence intervals generated by a generalized additive model. The vertical yellow line represents the beginning of the period the Managerial style of state governance became popularlized, according to Elazar. Consistent with his prediction, there is a dramatic increase in the eigenvector centrality, or de jure power, of governors beginning in the late 1940s or early 1950s. Concurrently, we also see a decline in the relative influence of the general assembly and, surprisingly, the public as well.

Figure 3: De Jure Power of the Public, Governor, and GA over Time

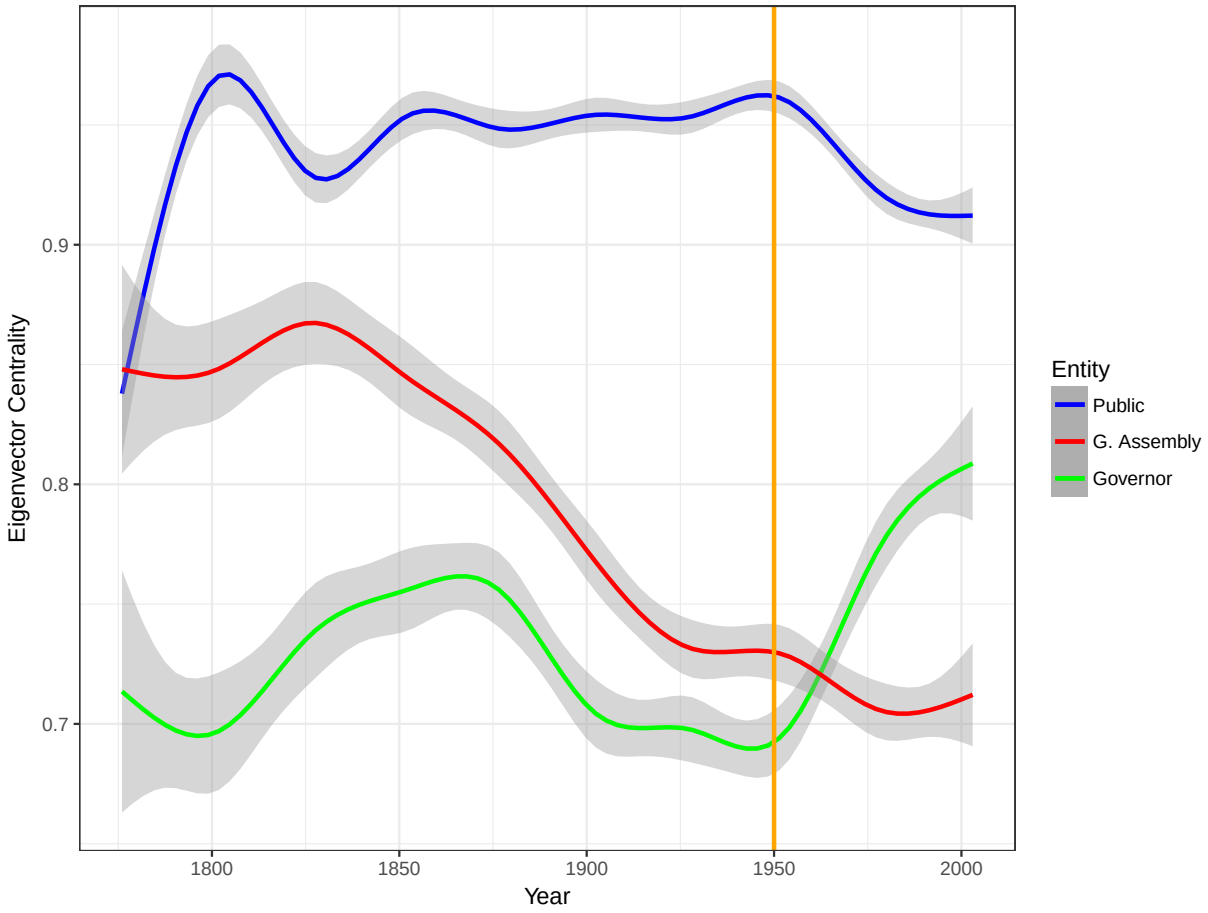

The general increase in power of the governor is not limited to a single region, as shown

in Figure 4. The Midwest, the Northeast, the South, and the West all see the power of the governorship generally increase after 1950, though the trends for the other actors are not as consistent. The de jure power of the public changes little in the South and the general assembly sees little change in its de jure power in the West. That said, the general trend in de jure power of the governor is remarkably consistent with Elazar's prediction.

Figure 4: De Jure Power of the Public, Governor, and GA over Time by Region

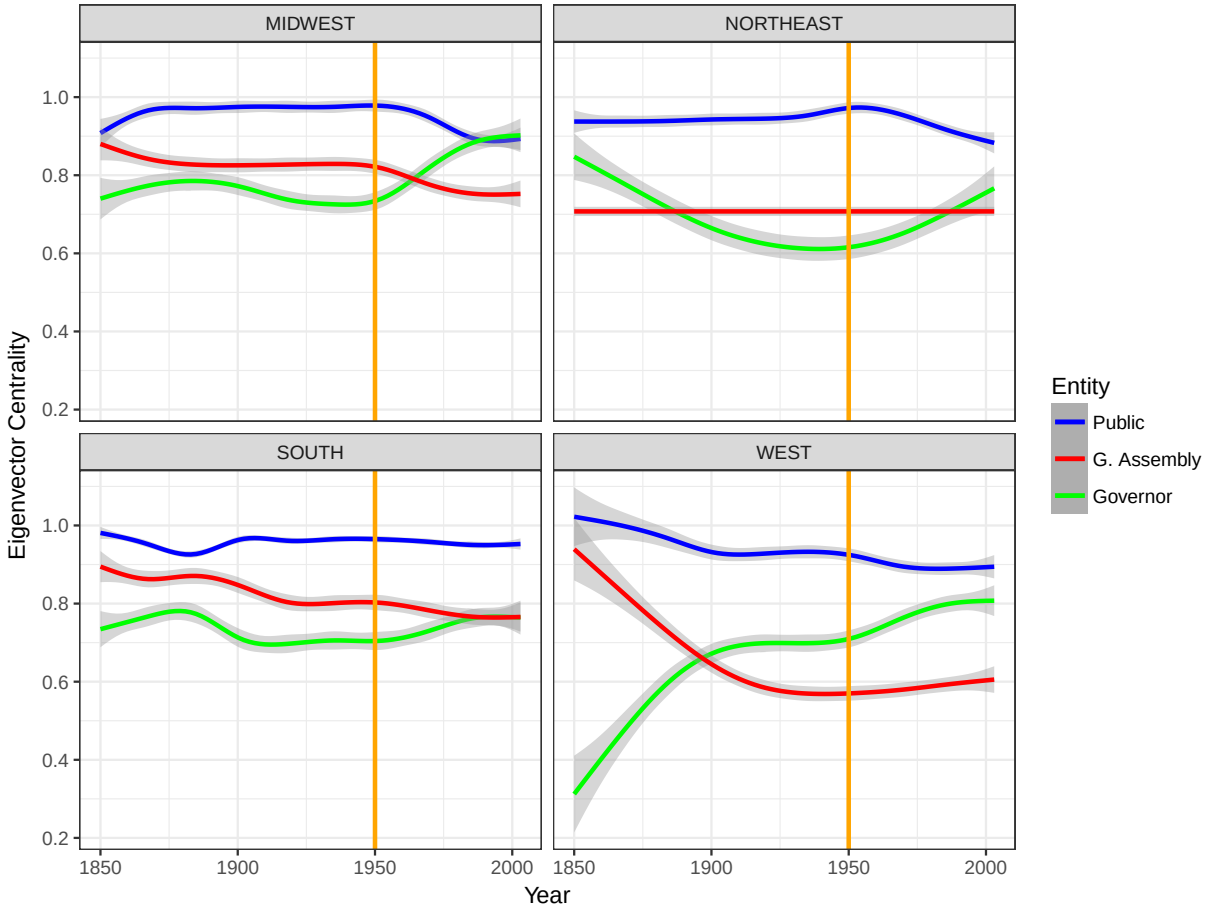

To confirm that the increasing de jure power is statistically significant, I also perform several OLS regressions using panel data in the unit of analysis is the state-year, which covers 1776 to 2003. Elazar's hypothesis indicates that the power of governors should increase beginning in 1950 and that this trend should be independent of the broader time trend. To

test his hypothesis, I predict the power of the governor of each state in each period (Y) with an interaction between a linear time trend with a binary indicator equal to 1 if the year is greater or equal to 1950 and is equal to zero otherwise. Equation 3 represents the anticipated data generating process where the parameter of interest is β_3, the marginal effect of time on the power of governors after the year 1950. I expect that β_3 is both statistically significant and greater than zero. To confirm that regional patterns are not driving the trend, I include dummy variables for each major geographic region in the US: the West, the East, the South, and the Midwest.

$$\mathrm{Y}_{it} = \beta_0 + \beta_1 \mathrm{Time}_t + \beta_2 \mathrm{Y1950}_t + \beta_3 \mathrm{Y1950} \times \mathrm{Time}_t + \beta_k X_k + \epsilon_{it} \tag{1}$$

The OLS results tell the same story as Figures 3 and 4: the de jure power of governors has increased since the 1950s. The positive and significant coefficients on the variable Y1950 x Year is consistent with Elazar's original prediction. However, the de jure power of the general assembly does not appear to have begun a decreasing trend since 1950. On the contrary, while the de jure power of the general assembly has seen a long term decline, reflected in loss of -0.001 normalized eigenvector units per year in Models 3 and 4, the decline has stabilized since 1950.

These results are robust to regional fixed effects. After controlling for them in Model 4, the general assembly's centrality is still decreasing, though at an overall rate of 0.0006 points per year. Overall, the OLS results for governors are consistent with Elazar's predictions; however, it turns out that there has been a general decline in the general assembly's power long before the 1950s and the decline has moderated or even disappeared since then.

Table 2: OLS Regression of Actor Type on Time Trends

	Dependent variable:			
	Governor		General Assembly	
	(1)	(2)	(3)	(4)
Year	-0.0003^{***}	-0.0004^{***}	-0.001^{***}	-0.001^{***}
	(0.0001)	(0.0001)	(0.0001)	(0.0001)
Y1950	-5.629^{***}	-5.803^{***}	-1.159^{**}	-0.825
	(0.597)	(0.592)	(0.553)	(0.503)
Y1950 x Year	0.003^{***}	0.003^{***}	0.001^{**}	0.0004^{*}
	(0.0003)	(0.0003)	(0.0003)	(0.0003)
Northeast		-0.081^{***}		-0.126^{***}
		(0.008)		(0.007)
South		-0.065^{***}		0.003
		(0.007)		(0.006)
West		-0.070^{***}		-0.200^{***}
		(0.008)		(0.007)
Constant	1.259^{***}	1.457^{***}	2.793^{***}	2.343^{***}
	(0.134)	(0.138)	(0.124)	(0.117)
Observations	6,634	6,634	6,634	6,634
R^2	0.022	0.041	0.070	0.236
Note:				*p<0.1; **p<0.05; ***p<0.01

Finally, Elazar described Alaska and Hawaii as models of the Managerial constitutional style; with the data generated in this project, it is possible to see if their scores are consistent with his prediction. In Table 3, I present the centrality scores for governors, the public, and the general assembly associated with each of constitution produced by Alaska and Hawaii since their formation in 1959. Again, we see results that are consistent with Elazar's predictions. Every version of both states' constitutions see governor eigenvector centrality scores of 1 – the highest possible score.

Table 3: Centrality of the Governor, Public, and GA in Hawaii and Alaska

	State	Year	Governor	Public	General Assembly
1	Alaska	1959	1.00	0.82	0.66
2	Hawaii	1959	1.00	0.63	0.56
3	Hawaii	1968	1.00	0.55	0.33
4	Hawaii	1978	1.00	0.82	0.59

Overall, I find substantial support for Elazar's prediction that the governor has become a more powerful actor in state politics since 1950. This trend is apparent both visually and in coefficient estimates derived from OLS regression. However, his prediction that this increase in power has come at the expense of the legislature's power is not entirely borne out by the evidence. The legislature of US states have seen a general decline in influence since the 1830s. This trend has been particularly pronounced in the West. Additionally, the fact that these estimates of de jure power appear to show trends in relative power that correspond well to predictions made in the literature supports the notion that my methodology effectively measures the empirical phenomenon of interest.

Qualitative Inspection of Extreme Cases

In this section, I examine the text of two constitutions to determine whether the centrality scores assigned to their respective governors have a basis in the text itself. Namely, I qualitatively analyze the text to identify what powers are associated with each governor and determine whether it appears as though one is more powerful than one another. To make the comparison as clear as possible, I examine two extreme cases: first, a constitution with an exceptionally weak governor; second, a constitution with a powerful governor. Michigan's 1909 constitution describes the governor with the lowest de jure power of all constitutions analyzed in this study – with a eigenvector centrality score of just 0.18. Coincidentally, Michigan's subsequent 1964 constitution contains one of the most powerful observed governors – with a score of 1, indicating that the governor is the most powerful actor in that

constitution (see Figure 5). Given this information, I expect to see that the text of the 1964 constitution will describe more hierarchical links between the governors and other actors than the 1909 constitution.

Figure 5: Power of Michigan's Governor over Time

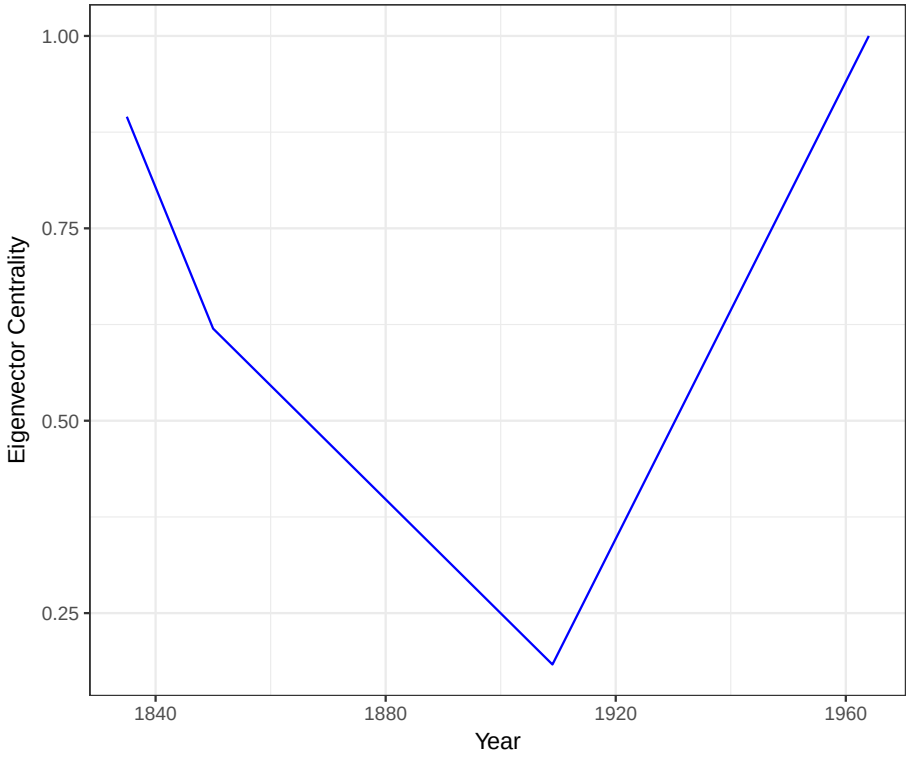

After examining Michigan's 1909 and 1964 constitutions, I find evidence that the overall de jure power of the governor has increased over time. First, I examined the hierarchical relationships of which the governor was a party. In the 1909 constitution, the governor was prohibited from appointing agents to fill vacancies in the state legislature; in contrast, the governor was explicitly given the power to do just that in the 1964 constitution. Since the state legislature is a very important actor in state politics, this change would certainly be seen as increasing the overall power of the governor. An additional change introduced by the 1964 constituion, was the bestowal of the authority to reorganize all government agencies as the governor saw fit; the governor in the 1909 constitution had no comparable power.

Another change seen in the 1964 constitution was that the governor was given the power to appoint all members of a committee that had the authority to determine the salaries of the governor, the liutenant governor, and the state legislators. This power was probably seen as a step too far by the Michigan state legislature, as a later amendment was adopted that gave the legislature the ability to reject salary proposals made by this commission. Finally, in 1964 the governor was given the new ability to remove judges from office, so long as they did so with the consent of a two thirds majority vote of both houses of the state legislature.

After examining Michigan's 1909 and 1964 constitutions, I find evidence that the overall de jure power of the governor has increased over time. First, I examined the hierarchical relationships of which the governor was a party. In the 1909 constitution, the governor was prohibited from appointing agents to fill vacancies in the state legislature; in contrast, the governor was explicitly given the power to do just that in the 1964 constitution. Since the state legislature is a very important actor in state politics, this change would certainly be seen as increasing the governor's overall power. An additional change introduced by the 1964 constituion, was the bestowal of the authority to reorganize all government agencies as the governor saw fit; the governor in the 1909 constitution had no comparable power. Another change seen in the 1964 constitution was that the governor was given the power to appoint all members of a committee that had the authority to determine the governor's salaries, the liutenant governor, and the state legislators. This power was probably seen as a step too far by the Michigan state legislature. A later amendment was adopted that gave the legislature the ability to reject salary proposals made by this commission. Finally, in 1964 the governor was given the new ability to remove judges from office, so long as they did so with the consent of a two thirds majority vote of both houses of the state legislature.

There were also some noteworthy changes to the governor's roles that were unrelated to their relationships with other government actors. In the 1909 constitution the governor was

limited to serving two year terms; the terms' length was increased in the 1964 constitution to four years with a two term limit. Additionally, the 1964 constitution introduced the power of pardon, commutation, and reprieves to the governorship. The governor was also given more flexibility in their ability to approve legislation: while in the 1909 constitution the governor was required to sign or veto legislation within five business days of receiving it from the legislature, the grace period was extended to 14 days in the 1964 constitution.

These changes to the governor's power in Michigan were noticed by Adrian (1968), an expert on state constitutions. He noted that George Romney, the governor of Massachusetts in 1964 went into the constitutional convention of 1964 with the explicit goal of increasing the power of the governor and that Romney was largely successful in doing so.

In summary, the transition from the 1909 constitution to the 1964 constitution saw the governor gain many new methods of influencing the legislature, government, and the judiciary while losing none; this pattern is entirely consistent with an overall increase in de jure power. While there were some new constraints placed on the governor, such as adding a term limit and the requirement that they provide annual reports to the state legislature, the net effect was clearly to make the governor a more powerful figure in Michigan politics. Thus, it appears that the content of the 1909 and 1964 constitutions in Michigan is consistent with the sharp increase in de jure power that was observed using eigenvector centrality scores.

Conclusion

This study aimed to create and validate a new measure for de jure power that has a basis in the extant literature on executive and de jure power. I did so using a combination of text as data methods to identify entities and links using the text of constitutions, while also

depending on human input for tasks such as classifying entities that would allow them to be more easily compared between constitutions. I then used those entities and links to generate eigenvector centrality statistics for each entity. Finally, I validated this methodology by comparing the estimated scores to those generated using hand-coded data, predictions made in the literature, and qualitatively examining two constitutions. Overall, all validation methods utilized in this paper support the notion that this measure of de jure power is capturing the intended qualities of state political actors.

Still, there remains a great deal of room for improvement on the methodology presented here. First, the constitutions described here do not account for changes in power that were introduced through amendments; were the scores estimated for every version of each constitution, I expect to see more variation in the power of actors within each constitution over time, particularly for constitutions which were never replaced. Second, while the Stanford NER algorithm was able to identify entities across a large number of constitutions very quickly, it systematically under reported some entities of interest, namely entities related to the public – such as "voter" or "district" – which thus required the identification of entities by hand. Ideally, I would generate a large database of constitutions in which all identities have been identified by hand and then train a new NER algorithm on that coded data so that future human hand-coding would not be required for documents of this type. Third, the co-occurrence method of identifying relationships is less then ideal, as it generates many false positives and is unable to identify the direction of hierarchical relationships. A better approach would involve a specifically trained model to identify directed relationships in constitutional text; as was the case for NER, this would require a large training data set.

With a valid measure of de jure power, many old and new hypotheses may be subjected to empirical evaluation. For instance, it would be possible to determine whether changes in the de jure power of the public was associated with lower levels of corruption. Additionally,

it may be possible to test whether increases in executive power were associated with higher economic growth levels, as anticipated by the Efficiency and Economy Movement reformers (Adrian, 1968; Swindler, 1971). I conclude with the hope that this paper contributes to an effort to quantify de jure power and apply text as data and network methods to that end.

References

Acemoglu, D. and J. A. Robinson (2008). Persistence of power, elites, and institutions. American Economic Review 98(1), 267–293.

Adrian, C. R. (1968). Trends in State Constitutions. Harvard Journal on Legislation 5(311), 243–258.

Ager, P. (2012). The Persistence of de Facto Power: Elites and Economic Development in the US South, 1840-1960. (November).

Benoit, K., D. Conway, B. E. Lauderdale, M. Laver, and S. Mikhaylov (2016). Crowd-sourced Text Analysis: Reproducible and Agile Production of Political Data. American Political Science Review 110(02), 278–295.

Benton, A., L. Ungar, S. Hill, S. Hennessy, J. Mao, A. Chung, C. E. Leonard, and J. H. Holmes (2011). Identifying potential adverse effects using the web: A new approach to medical hypothesis generation. Journal of Biomedical Informatics 44(6), 989–996.

Bertocchi, G. and A. Dimico (2017). De jure and de facto determinants of power: evidence from Mississippi. Constitutional Political Economy, 1–25.

Bois, W. E. B. D. (1910). Reconstruction and its Benefits. The American Historical Review 15(4), 781–799.

Bonacich, P. (2007). Some unique properties of eigenvector centrality. Social Networks 29(4), 555–564.

Borgatti, S. S. P., K. K. M. Carley, and D. Krackhardt (2006). On the robustness of centrality measures under conditions of imperfect data. Social Networks 28(2), 124–136.

Carpenter, D. R., E. Whittington, and B. Clinton (2013). Executive Power in American Power Institutional Executive Development. Perspectives on Politics 1(3), 495–513.

Costenbader, E. and T. Valente (2003). The stability of centrality measures when networks are sampled. Social networks.

Debs, A. (2010). Living by the Sword and Dying by the Sword? Leadership Transitions in and out of Dictatorships. Manuscript, Yale University.

Elazar, D. J. (1982). The Principles and Traditions Underlying State Constitutions. Publius: An Oxford Journal 12(1), 11–25.

Finkel, J., T. Grenager, and C. Manning (2005). Incorporating Non-local Information into Information Extraction Systems by Gibbs Sampling. roceedings of the 43nd Annual Meeting of the Association for Computational Linguistics (ACL 2005) (June), pp. 363–370.

Florian, R., H. Hassan, A. Ittycheriah, H. Jing, N. Kambhatla, X. Luo, N. Nicolov, S. Roukos, and T. Zhang (2004). A Statistical Model for Multilingual Entity Detection and Tracking. Proceedings of the Human Language Technologies Conference 2004 (HLT-NAACL'04), 1–8.

Florian, R., A. Ittycheriah, H. Jing, and T. Zhang (2003). Named entity recognition through classifier combination. Proceedings of the seventh conference on Natural language learning at HLTNAACL 2003 4, 168–171.

Frantz, T., M. Cataldo, and K. Carley (2009). Robustness of centrality measures under uncertainty: Examining the role of network topology. Computational and Mathematical.

Frye, T. (1997). A Politics of Institutional Choice: Post-Communist Presidencies. Comparative Political Studies 30(5), 523–552.

Graves, W. B. (1966). State Constitutional Law: A Twenty-Five Year Summary. William and Mary Law Review 8(1), 1–42.

Hossain, L. (2009). Effect of organisational position and network centrality on project coordination. International Journal of Project Management 27(7), 680–689.

Lingad, J., S. Karimi, and J. Yin (2013). Location extraction from disaster-related microblogs. Proceedings of the 22nd International Conference on World Wide Web - WWW '13 Companion, 1017–1020.

Metcalf, L. K. (2000). Measuring Presidential Power. Comparative Political Studies 33(5), 660–685.

Mizruchi, M. S. (1982). The American Corporate Network, 1904-1974.

Mizruchi, M. S. and B. B. Potts (1998). Centrality and power revisited: actor success in group decision making. Social Networks 20(4), 353–387.

Ritter, A., S. Clark, and O. Etzioni (2011). Named entity recognition in tweets: an experimental study. … of the Conference on Empirical Methods …, 1524–1534.

Samuels, D. J. and M. S. Shugart (2010). Presidents, Parties, and Prime Ministers: How the Separation of Powers Shapes Party Organization and Behavior.

Shugart, M. S. (2005). Comparative Executive–Legislative Relations: Hierarchies vs. Transactions in Constitutional Design Texts.

Shugart, M. S. and J. M. Carey (1992). residents and assemblies: Constitutional design and electoral dynamics (2nd ed.). Cambridge University Press.

Swindler, W. F. (1971). State Constitutions for the 20th Century. Nebraska Law Review 577, 577–599.

Wallis, J. J. (2005). Constitutions, Corporations, and Corruption: American States and Constitutional Change. The Journal of Economic History 65(1), 211–256.

Yan, E. and Y. Ding (2009). Applying centrality measures to impact analysis: a coauthorship network analysis. Journal of the American Society for Information Science and Technology 60(10), 2107–2118.

Yosef, M. A., J. Hoffart, I. Bordino, M. Spaniol, and G. Weikum (2011). AIDA: An Online Tool for Accurate Disambiguation of Named Entities in Text and Tables. Very Large Data Bases, 1450–1453.

Zhou, G. and J. Su (2001). Named entity recognition using an HMM-based chunk tagger. Proceedings of the 40th Annual Meeting on Association for Computational Linguistics - ACL '02 (July), 473.