Dear CPW,

Thank you in advance for taking the time to read my paper. This is an ongoing project of mine to use latent text analysis in order to understand constitutions. I thought this would be a good time to check in with the collective wisdom of CPW, especially before the summer break. In the conclusions of the paper, I mention a few extensions that I envision and will probably undertake over the next few weeks (some of these extensions are already running on the UM Statistics Server).

As always, I welcome your comments.

Sincerely,

Dominic J. Nardi, Jr.

It's not what you say, it's how much you say it: Comparing Authoritarian and Democratic Constitutions Using Latent Text Analysis

Dominic J. Nardi, Jr.
Political Science Department
University of Michigan

March 25, 2013

Abstract

While constitutions regulate political authority within a state, there is much we do not know about the interaction between political conditions and constitutional texts. Much of the literature relies upon hand-coded databases of constitutional texts. However, this approach raises both methodological and theoretical concerns. Advances in natural language processing, particularly latent text analysis, have the potential to provide a more reliable and refined means for coding constitutions. In this paper, I use a Latent Dirichlet Allocation topic model in order to compare constitutional topics in authoritarian and democratic regimes. Unlike previous articles, I find significant differences in the proportion of authoritarian and democratic constitution allocated to certain topics, particularly judicial and legislative institutions. Democratic constitutions are more likely to discuss some negative rights, while authoritarian constitutions are more likely to discuss positive rights. Democratic constitutions allocate proportionality greater text to elections, particularly district size. Finally, authoritarian constitutions are also more likely to allocate text to military affairs.

While constitutions are central to the regulation political authority and the division of power within a state (Brown, 2001), there is much we do not know about the interaction between political conditions and constitutional texts. Much of the literature over the past five years has relied upon hand-coded databases of constitutional texts, such as the Comparative Constitutions Project (CCP) (Elkins et al., 2009), in order to compare the presence or absence of constitutional features across countries. However, this approach raises both methodological and theoretical concerns. First, hand-coding is subject to bias and requires coders to make important decisions as to which topics are sufficiently important to merit coding. Second, it is unclear how much the mere presence or absence of a topic actually reveals about drafting decisions, especially because drafters often copy or mimic the constitutions of other countries (Go, 2003).

Advances in natural language processing have the potential to provide a more reliable and refined means of coding legal documents. Latent topic models derive topics latent within a corpus of text by finding the probability distribution of a cluster of related words over the entire vocabulary. This reduces coder discretion in that the substance of the topics is determined by the probability distribution, not by the coder. While other political scientists have applied topic models to Senate press releases (Grimmer, 2010) and judicial decisions (Rice, 2012), as of this writing the method has not been used on constitutions. Topic models also have the potential to provide much richer data on constitutional texts by revealing not just the presence or absence of a topic, but also the proportion of a constitution allocated to each topic.

One of the most fundamental questions in the comparative constitutions literature is whether or not the content of constitutions varies systematically by regime type. While a handful papers have explored this question using hand-coded data (Melton et al., 2011), hand-coded data seem particularly ill suited to answering this question. For various reasons, we might reasonably expect authoritarian and democratic regimes to converge with respect to the topics they cover, but to allocate varying levels of attention to those topics. Indirectly, this might help determine if and how authoritarian constitutions play any substantive political role.

In this paper, I use a Latent Dirichlet Allocation topic model in order to compare 184 constitutional texts in authoritarian and democratic regimes in 2006. I begin by reviewing the literature about constitutions and regime type, with a focus on potential hypotheses. Next, I discuss the methodological challenges that arise in hand-coding constitutional documents. I then explain the Latent Dirichlet Allocation methodology, as well as the process for converting the constitutional texts into data. After running the topic model, I present the resultant 67 topics and discuss potential substantive interpretations. Interestingly, the model produces some topics similar to those in the CCP database, such socioeconomic rights, but also reveals new topics, such as "government functions."

Finally, I use the results to test if democratic and authoritarian regimes differ significantly with respect to how much text they allocate to political institutions and rights.

I compare the proportion of constitutional text allocated to latent topics related to the branches of government, constitutional rights, and the military. Ultimately, I find that I find that democratic constitutions allocate significantly more text to judicial and legislative institutions, as well as to elections and negative rights. By contrast, there is no significant difference between the two with regards to rights broadly. Authoritarian constitutions are more likely to allocate constitutional text to socioeconomic rights the military.

1 Literature Review

For decades, it was commonly assumed that democracy and constitutionalism were closely correlated (Tate, 1997). However, Brown (2001) points out that, when read as a whole, the constitutions of authoritarian regimes are usually internally consistent, if illiberal. One portion of the constitution might list political rights, but another portion might clarify that those rights are not enforceable against the state. In such cases, authoritarian regimes can engage in illiberal practices without violating their constitutions. As such, scholars are increasingly taking authoritarian constitutions "seriously."

With the increased availability of hand-coded databases of constitutional texts, such as the Comparative Constitutions Project (CCP), scholars have used cross–national samples in order to test the impact of constitutional texts. Law and Versteeg (2012) find considerable variation in the extent to which political and economic conditions in a country reflect the rights "promised" in a constitution. They also find that rights conditions do not significantly improve with the presence of judicial review. By contrast, Blasi and Cingranelli (1996); Sandholtz (2012) find that a constitutionally guaranteed independent judiciary is a significant predictor of better human rights outcomes.

There is relatively less research about process of constitutional text generation, or why constitutions differ in the first place. While Elkins et al. (2009) uses the length of constitutions and number of topics covered as independent variables, we know less about the process that generates these variables in the first place. Does a longer constitution imply less agreement amongst the drafters or more agreement? Drafters might not feel compelled to discuss a particular topic at length if there exists an implied understanding (Ginsburg, 2010). Moreover, there is some evidence that legal origins impacts constitutional texts; common law countries tend to write longer, more detailed constitutions, while French civil law and socialist constitutions tend to be shorter (Glendon et al., 2008; Jackson and Tushnet, 2006; Helmke and Rosenbluth, 2009).

For the purposes of this paper, it is only necessary to establish the theoretical reasons why authoritarian and democratic constitutions might systematically differ from each other. The mere presence of a constitution reveals little – as of 2013, every single country aside from the United Kingdom has adopted a constitutional document.[1] Moreover, there has been

[1] In some cases, including national codes.

significance convergence in the topics that constitutions cover, especially since the end of the Cold War (Go, 2003; Law and Versteeg, 2011). Melton et al. (2011) conclude that, when controlling for the year in which the constitution was enacted, "the substantive differences between authoritarian and democratic constitutions are minimal" (Melton et al., 2011, 8).

Part of the problem in detecting constitutional differences across regime types is that, for constitutional drafters, simply listing a topic can be a relatively cheap signal, especially when the audience costs of its omission would be high. For example, both democratic and authoritarian constitutions frequently contain provisions guaranteeing judicial independence because not doing so would imply that the judicial system lacks credibility. Indeed, according to the CCP, 77.35% of constitutions in the year 2006 proclaim judicial independence. Moreover, several of the 22.65% of constitutions that do not explicitly guarantee judicial independence do so implicitly, such the United States, meaning that the signal is even noisier. In short, hand-coded topic labels for constitutions risk capturing such superficial similarities.

By contrast, the *amount* of text allocated to a particular topic should serve as a more reliable signal of the importance placed on that topic by the constitutional drafters. I assume, as do Elkins et al. (2009), that longer constitutions are costlier than shorter ones. Allocating text to a topic is costly – or at least costlier – because of the resources required to write more text and the binding effect of that specificity (Ehrlich and Posner, 1974). More text suggests that the constitution provides more explicit detail about how institutions will function and the exact content of rights, meaning that political actors have less discretion than they would have had with ambiguous text. Moreover, constitutional drafters are more likely to dedicate their time and resources on sections of the constitution of greatest importance or controversy. We should expect to see less text when there is an implied understanding about a particular topic (Ginsburg, 2010).

Perhaps the most common justification for constitutions is that they precommit governments to certain rights and government structures. North and Weingast (1989) note that institutional change is costly and therefore credibly commits the government to its stated policies. However, the mere mention of a topic in a constitution does not always provide a credible commitment. First, if omission of a topic would send a signal that the government does not intend to commit itself, the omission could be costly. Second, as mentioned above, in recent years constitutional drafters have mimicked their counterparts in other countries (Go, 2003; Law and Versteeg, 2011), meaning that there is likely to be convergence in the topics they cover. By contrast, the more text a constitution allocates to a topic, the more likely it is to signal a genuine commitment, both because of cost and specificity.

Political elites are more likely to strengthen constitutional courts in order to guarantee political and human rights commitments when they are at risk of losing power themselves (Ginsburg, 2003; Hirschl, 2004; Finkel, 2008). Doing so entails not just stating that judges are independent, but also explicating term lengths, control over judicial budgets, and jurisdiction. As such, we should expect democratic constitutions to allocate relatively more text to topics related to the judiciary than would authoritarian constitutions (*Hypothesis 1*).

While the logic of precommitment also extends to human rights (Erdos, 2009), it is more difficult to make generalized predictions about the relationship between the allocation of constitutional text to rights and the commitment to rights. Constitutions sometimes limit the applicability of rights through limitations clauses, which can require significant additional text in order to delineate exceptions to a right. For example, §354 of Myanmar's Constitution allows the legislature to limit freedom speech for "Union security, prevalence of law and order, community peace and tranquility or public order and morality...," which takes far more text than the U.S. First Amendment's blanket statement that "Congress shall make no law... abridging the freedom of speech." Therefore, the relationship between text about rights generally and regime type is ambiguous, although we should expect to see authoritarian regimes employ more text related to limitations clauses (*Hypothesis 2*).

Amongst the specific negative rights that are included in the topic model, we should expect democratic constitutions to provide more and more rigorous criminal rights procedures than authoritarian constitutions (*Hypothesis 3*). Unlike many other rights, criminal rights are in fact a package of related rights that must often be listed separately, such as rights to an attorney, to a fair trial, or against self-incrimination. By contrast, since the end of the Cold War, even socialist countries have sought private capital and thus not including a constitutional right to property can become a costly signal. Moreover, property protection can usually be handled with fairly minimal text by promising due compensation . As such, we should not expect to see such a difference with regard to property rights (*Hypothesis 4*).

Governments can also use constitutions in order to make broader policy or ideological statements. Because the costs of constitution-drafting and amendments are likely lower – but still extant – for authoritarian regimes, we should expect authoritarian constitutions to allocate a greater proportion of text to policy statements. Authoritarian rulers face less competition from internal competing groups, allowing them to promulgate a single ideology or policy within a constitutional document. Some authoritarian governments have even replaced their constitutions within a handful of years in order to announce new ideological directions (Elkins et al., 2009). For example, in 1980 Vietnam replaced its 1959 constitution in order to represent the reunification of North and South, which in turn was replaced in 1992 to represent the consolidation of economic reforms. Thus, we would expect authoritarian constitutions to allocate relatively more text to constitutional topics related to socioeconomic rights or ideological goals (*Hypothesis 5*).

Constitutions also serve as a coordination devices between the branches of government. Constitutions represent a bargain between elite groups that helps clarify the respective roles of government agents. The constitution serves as a focal point around which parties can rally against any unilateral defections (Weingast, 1997). On the one hand, the need for coordination is greater when the government contains multiple veto players (Tsebelis, 2002). On the other hand, even authoritarian regimes require a coordination mechanism, especially when divisions exist within the elite (Barros, 2002). However, for authoritarian regimes, the primary need for coordination is likely to be within the executive rather than the legislature, as authoritarian constitutions delegate greater authority to the executive (Brown, 2001). As such, we should expect that democratic constitutions allocate a greater proportion of text to legislative topics (*Hypothesis 6*), while the reverse is true for topics related to the executive (*Hypothesis 7*).

Finally, authoritarian and democratic constitutions will also likely differ with regard to a few specific topics. Given the centrality of elections to democracy, democratic constitutions should allocate more text to topics related to elections or the management of voting (*Hypothesis 8*). Authoritarian rulers, especially juntas, rely heavily upon the military for political support. Even if senior government leaders or constitutional drafters do not hail from the military themselves, authoritarian rulers will likely have greater need to address the military's role in the state explicitly in their constitutions (*Hypothesis 9*). Some constitutions, such as Myanmar's, dedicate an entire chapter to the military's rights and duties.

2 Methodological Problems with Hand-Coding

In addition to the question of coverage versus allocation, there are methodological concerns about using hand-coded data in order to measure constitutional topics. In hand-coding, the researcher must select the topics for which to code. While great care is usually taken in selecting potential topics, at the end of the day the coder's experience and expectations risk influencing the pool of topics. The priorities of the coder are not necessarily the same as those of constitutional drafters (Rice, 2012; McGuire and Vanberg, 2005). This is particularly true for cross–national comparisons because coders are not – and cannot become – familiar with the legal intricacies of every single country in their sample.

This risk leads to two types of coding error. First, potentially important constitutional topics might not considered or might be rejected by the coder. However, most professionally compiled databases receive input from multiple experts and tend to err on the side of overinclusivitiy. Second, this might lead researchers to include topics that constitutional drafters would not consider as substantively distinct. Constitutions often contain "laundry lists" of subjects, so for the purposes of cross–national comparison should the list or individual items be considered the relevant unit of analysis? For example, the CCP includes "artists or the arts" as a unique topic (Elkins et al., 2009, Appendix 2), something that constitutional drafters might not even consider separate from freedom of speech.

In either case, the exact number of topics has a significant effect on a proportion of topics covered in a given constitution. For example, the CCP uses 92 constitutional topics in order to generate a *scope* variable that assesses the proportion of potential topics each constitution covers. If the denominator is incorrect, any inferences using *scope* will be biased. If a coding scheme is systematically over- or underinclusive, then any bias should be minimized. However, if, as is more likely, inclusion or omission of potential topics is correlated with regime type, inferences will be biased in a manner that affects substantive interpretations. For example, if liberal democracies are more likely to mention artists, then they will also tend to score higher in *scope*.

In addition, there is a risk that perceptual bias will influence human coders. Harvey and Woodruff (2011) find evidence of confirmation bias in the Supreme Court Database, a collection of hand-coded U.S. judicial decisions (Spaeth et al., 2012). According to the study, issue areas generally align to the coder's expectations of how a case was decided along a liberal–conservative political spectrum. This risk is also present when hand-coding constitutions, especially as there are strong normative expectations with regard to democratic and authoritarian constitutions.[2] This becomes especially problematic when constitutional texts are ambiguous or subject to multiple interpretations.

3 Data

3.1 Constitutions

The data for this paper consists of 184 constitutions extant during the year 2006. The constitutions are the same as those code for the third phase of the Comparative Constitutions Project,[3] allowing for comparisons between the latent topic model and hand-coded results in future research. I use official English-language translations so I can compare words across constitutions, bearing in mind that English is a verbose language and there might be multiple correct translations of the same word.[4]

Before running a latent text model, all texts must be pre-processed in order to extract information about word frequency counts across documents. As recommended by Grimmer (2010), I stripped the texts of any numbers, punctuation, and extra whitespace. I also converted all words to lowercase so the same word in a different case would be recognized as such. I also removed a standard list of "stop words" (e.g., "and," "the," "will"). Finally, I employed a Porter stemming algorithm in order to stem words to their roots (e.g., "courts" becomes "court"). Because the Porter stemmer sometimes stems key words, such as "president" and "prime minister," into the same stems as other words (e.g., both "president" and "preside" are stemmed to "presid"), I created tags to mark these words as distinct.

Using these "tokens," I then process the texts to remove words that are less likely to contain relevant information about constitutional topics. I remove the top 20 terms from the

[2] At least for the CCP, this sort of confirmation bias appears minimal. The CCP guidelines were careful to instruct coders to focus on *de jure* constitutional text and not consider *de facto* practice. For example, the U.S. constitution is coded as not covering judicial independence, while China's is coded as covering judicial independence. Moreover, as noted above, Melton et al. (2011) finds little substantive difference between authoritarian and democratic constitutions; systematic confirmation bias would have biased against that finding.

[3] Available at http://www.comparativeconstitutionsproject.org/data.htm

[4] Along these lines, I converted a list of frequent British English words into American English.

entire corpus, including "constitut" (all forms of "constitution"), except for "president" and "prime minister" – words which are informative but also appear frequently. I also remove exceedingly rare terms. However, because it is important not to unintentionally remove rare but informative words, I adopt a relatively low threshold of words that appear in at least 10 documents (or 0.0033% of the total).[5] Finally, I remove any words that contain fewer than three letters, which are less likely to be informative (which also removes common articles and prepositions).

I created an additional stoplist with the names of each country. First, some constitutions repeat the name of the country throughout, stating "the parliament of the Republic of Slovakia" rather than simply "the parliament". The extent of this type of repetition varies across countries, so the topic model is more likely to incorporate frequently repeated names. Second, some constitutions, particularly those of Commonwealth countries, list the names of countries with which they are affiliated. Again, these "laundry lists" are not informative, but, unlike lists of subnational government units, they will not be removed when removing rare words.

An important question is the unit of analysis for each document. Latent topic models have typically been used to classify relatively short documents that contain a handful of topics (Grimmer, 2010; Rice, 2012). By contrast, constitutions are essentially long texts that can contain dozens – proverbial "laundry lists" – of distinct topics. Moreover, with only 184 constitutions, treating each constitution as an individual document might not provide the latent text model with enough information to obtain the distribution of topics across documents. As such, I run the model using the largest subunit of the document with substantive text. In practice, this means that each constitution is decomposed into its constituent "articles" or "sections." Doing so produces 30,977 unique documents.

As a robustness check, I rerun the model using each of the 184 constitutions and chunks of 500 words. I use the original training parameters for the topic model on the new disaggregations of the data. In both cases, I find the topics and distribution of topics hold [results will be posted on an online appendix].

In order to classify each country by regime type, I rely primarily on Freedom House political freedom scores. According to the Freedom House scale, countries that fall between 1–2 are considered "free," those between 3–4 are "partly free," while those between 5–7 are "not free" (House, 2011). The score is intended to measure the extent to which the government guarantees political rights and freedoms. Freedom House scores are provided by expert ratings and thus have been criticized. However, the main alternative, Polity (Marshall and Jaggers, 2010), suffers from severe missingness, providing scores for only 153 out of 184 constitutions in the sample. Moreover, these data are likely not missing at random.

[5]In practice, even removing words that only appear in at least 100 documents does not significantly affect the topic model.

4 Methodology

In this section, I provide the statistical basis for the Latent Dirichlet Allocation (LDA) model that I use. We start with $\omega = (\omega_1, ..., \omega_N)$ documents of corpus D containing N words from a vocabulary of V different words, i.e. $\omega_i \in \{1, ..., V\}$ for all $i = 1, ..., N$ (Grün and Hornik, 2011). Latent Dirichlet Allocation estimates the proportion of topic distributions θ for document ω as a latent variable.

4.1 Selecting the Number of Topics

The number of topics k must be set a priori in topic models. Selecting the number of topics is often the most difficult part of the process as there is currently no method that completely eliminates human discretion. Following Blei et al. (2003); Rosen-Zvi et al. (2004); Rice (2012), I use perplexity scores to guide my choice. Perplexity equals the geometric mean per-word likelihood, as follows:

$$Perplexity(\omega) = E^{-\frac{log(p(\omega))}{\sum_{d-1}^{D}\sum_{j-1}^{V} n^{(jd)}}}$$

, where $n^{(jd)}$ indicates how often the j^{th} term occurs in d^{th} document (Grün and Hornik, 2011).[6]

Perplexity scores indicate the ability of the model to read part of a document, learn the topic distribution, and then predict the remaining words in the rest of the document. The better the prediction, the less "perplexed" the model. However, perplexity is monotonically decreasing in the number of topics (Blei et al., 2003), so researchers cannot simply select a global or local minimum. Rather, best practice is to select based on the change in the rate of decrease (Blei et al., 2003; Rice, 2012). If perplexity decreases drastically at $k = k^*$ and then only improves gradually afterwards, then k^* provides the best improvement in perplexity.

I used the the Stanford Natural Language Processing (SNLP) Topic Modeling Toolbox (TMT),[7] to calculate perplexity scores for my corpus of texts ranging from 1 to 100 topics. First, 80% of texts are assigned to a training set while the remaining 20% are used for testing. I then used the first half of each document to estimate the topic distribution within that document based on the parameters derived from the training data. Finally, I estimated the number of equally probably word choices for the second half of the document. The perplexity score is simply the average perplexity across all testing documents.

[6]Unlike LDA, another type of latent text analysis, the Correlated-Topic Model (CTM), allows topics to be correlated. As such, Blei and Lafferty (2007) claim that CTM should encounter less perplexity than LDA because it bases its estimate not only on topics from the first half of the document, but also correlated topics. However, Chang et al. (2009) notes that CTM performs worse when evaluated by human interpreters. I am currently running my data in a CTM model on the University of Michigan Statistics server.

[7]Available at http://nlp.stanford.edu/software/tmt/tmt-0.4/

Perplexity provides a "best fit" for the model, but does not always produce substantively interpretable results. Chang et al. (2009) argue that topic models selected only according to perplexity scores suffer word intrusion (e.g., words that do not appear related to the topic) and topic intrusion (e.g., topics that do not appear related to the document). As such, they recommend visual inspection of the top topic terms in order to check that each topic has a substantive meaning (Grimmer, 2010).

I devised several rules for visual inspection. First, any k that produces topics with more than significant word intrusion – a non-informative personal, institutional, or geographic name – in the top 10 terms suggests the model was forced to create too many unique topics without enough informative terms. Second, the top 5 terms should suffer minimal topic intrusion. If there are too few topics such that $k < k^*$, then the topic model will clump distinct topics. Usually, clumping indicates that the model groups two terms together based upon their relationship to a third term, even if each relationship is of a different nature. For example, a model with fewer topics might group civil rights and institutional authority topics together because both are closely linked to the word "constitution".

4.2 Latent Dirichlet Allocation Model

Topic models utilize a Bayesian approach in order to determine complex posteriors for the probability that a particular document falls within a particular topic cluster. For Latent Dirichlet Allocation (LDA), the generative process relies upon a Dirichlet distribution for the topic distribution θ and the term distribution β for each topic as follows:

$$\theta \sim Dirichlet(\alpha)$$

$$\beta \sim Dirichlet(\delta)$$

, where α is a parameter that governs the Dirichlet distribution. The level of α helps determine pooling of topics. For this paper, I set the starting value of α at $\alpha = \frac{50}{k}$ and allow the variational expectation-maximization (VEM) to estimate α, as recommended by Griths and Steyvers (2004). The Dirichlet distribution assumes independence between the topics (Grimmer, 2010, footnote 6), which affects the number of topics it can support when compared to a Correlated-Topic Model (Blei and Lafferty, 2007).

Next, for each N words ω_i:

1. choose topic $z_i \sim Multinomial(\theta)$
2. choose word ω_i from multinomial probability distribution conditioned on topic z_i : $p(\omega_i|z_i, \beta)$

, where β is the term distribution of topics and contains the probability of a word occurring in a given topic k_i.

In order to estimate the model, maximum likelihood estimation is used to maximize the log-likelihood of the data with respect to the model parameters. For LDA, the variational expectation–maximization of the log-likelihood for one document $\omega \in D$ is:

$$l(\alpha, \beta) = log(p(\omega|\alpha, \beta))$$

$$= log \int \sum_z \prod_{i=1}^{N} p(\omega_i|z_i, \beta)p(z_i|\theta)p(\theta|\alpha)d\theta$$

, maximized for parameters α and β.

The VEM algorithm is an iterative method for determining ML estimates in a missing data framework. It iterates between an Expectation (E)-step, where given the data and current parameter it estimates the complete likelihood, and a Maximization (M)-step, where the expected complete likelihood is maximized in order to find new parameter estimates new parameter estimates. For topic models, the missing data are the latent variables θ, z for LDA (Grün and Hornik, 2011).

Finally, because each document is a disaggregated section of the constitution, I reestimate the model by country (called "splicing") in order to obtain posterior distributions for each topic in each constitution. The posterior reports the effective number of documents per constitution related to that particular topic. Because it is a probability, the posteriors are not integers.

5 Results

For my data, I selected $k = 67$ topics. As seen in Figure 1, the graph of perplexity scores decreases rapidly from 30 topics up to around 60 topics, and then varies afterwards. I then narrowed the intervals in order to better approximate the best number of topics. As seen in Figure 2, perplexity is lowest at $k = 67$, and then varies afterwards. Moreover, each of the 67 topics has a substantively interpretable meaning, whereas a greater number of topics $k > k^*$ produced topics that yielded little additional information.

Tables 2–4 list the top five terms derived from each of the 67 latent topics in the corpus. The topic model does not assign names or labels to the topics. However, visual inspect of the topics suggests that each topic has a plausibly substantive interpretation. There are no obvious examples of word or topic intrusion. For ease of interpretation, I assign labels to each of the topic (the word or phrase immediately after the topic number). The label is intended for reference, not to impose a substantive interpretation of value judgment onto the topic. Moreover, the label might aid non-lawyers in understanding how the words under each topic are related.

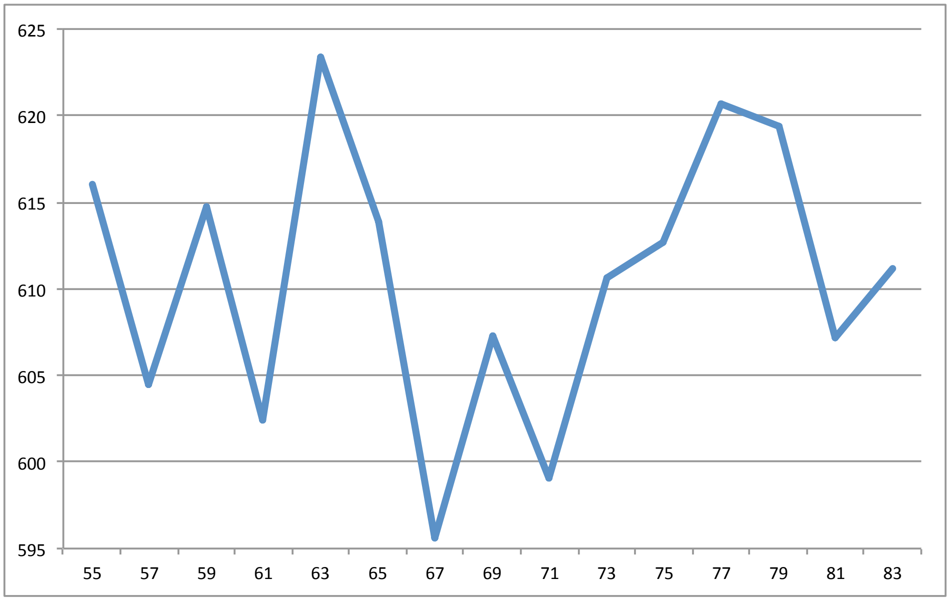

Figure 1: Perplexity scores with LDA for 55-83 topics

While 67 topics is fewer than the number of topics coded by the CCP (Elkins et al., 2009), as discussed above the CCP coding scheme probably overestimates the number of substantive topics by using coder–created tags for certain "low-frequency" or "niche" topics. By contrast, latent text analysis depends upon probability distributions over the entire corpus of documents. Many constitutions enumerate a large number of topics within a relatively small area of text. Words such as "speech" or "impeachment" might only appear once in a document. Moreover, there is no universally agreed upon language with which to express these rights and the latent text model cannot recognize, for example, that "speech" and "opinion" imply the same right. As such, the CCP and latent text data are not necessarily comparable.

6 Analysis

6.1 Case Studies

One of the greatest advantage of latent topic models is that researchers are not limited to selecting binomial or ordinal variables as a means of measuring the presence of topics in constitutional texts. LDA assigns probabilities across topics for each of the 30,977 documents. Combining the documents related to each of 184 constitutions yields a profile for the distribution of topics within each constitution. Thus, the results show not just whether a constitution covers a topic, but also the proportion of text allocated to the topic.

Figure 3 shows the text profiles of four constitutions: the United States, France, China, and Libya. A few differences are immediately clear. Consistent with our expectations above, the U.S. and French constitutions allocate greater text to topics related to judicial appointments and appeals – 2% and 2.23%, respectively, compared to less than 1% for both China and Libya (H1). The difference decreases when the "courts" topic is included, although that is perhaps because that topic deals more with the presence of a supreme court than the procedures for the court. For judicial independence, arguably the procedural protections are far more important.

The results are far more varied with respect to rights. As might be expected given the conventional wisdom, the U.S. allocates almost 6% of its constitution to sentencing, criminal trials, and criminal arrests, compared to negligible proportions for the authoritarian constitutions (France, at 2.2%, is in between). Moreover, the U.S. seems to earn its reputation for protecting private property – it allocates over 1% to property, more than double next highest (China). However, the two authoritarian countries also score slightly better with regard to citizens rights and rights protection, although the difference is less significant. Overall, this ambiguity accords with the expectation that the relationship between constitutional rights protections and text would be ambiguous, at best (H2).

On the other hand, both China and Libya allocate 9% and 7%, respectively, to socioeconomic rights, compared to less than 1% for the U.S. and 5% for France. The gulf is similarly vast for education & health, with China and Libya at 5.6% and 3.2%, respectively, compared to less than 1% combined for both the U.S. and France. As expected, this seems to suggest that these two countries are more likely to use their constitutions in order to proclaim broad policy goals (H3).

In order to compare how each constitution handles legislative and executive functions, I consider several related topics. First, for the legislature, I consider "legislative membership," "majority voting," "congress," and "parliament." For the executive, I consider "head of government," "cabinet responsibility," presidency," and "executive." Arguably, other topics also deal with legislative and executive matters, but the ones selected seem to provide the clearest linkage with those concepts without introducing potential noise. For example, without further context, "committees" could refer to legislative or executive committees. As seen in Table 1, the U.S. and France consistently allocate greater proportions of their texts to these four legislative topics than do China and the U.S., aside from the U.S. in the "parliament" topic (not surprising given that the U.S. is a presidential system). By contrast, the differences are much smaller for executive topics.

In order to compare how each constitution handles elections (H5), I combined the topics for "district size" and "elections." Doing so shows that the U.S. and France allocate a greater percentage of their constitutions to electoral topics – 3.69% and 2.34%, respectively – than do China or Libya – 1.15% and 0.05%, respectively. Interestingly, the difference is even greater when comparing the countries along district size, where the proportion of text

	Legislative membership	Legislative voting	Congress	Parliament
United States	0.0092	0.0185	0.0523	0.0002
France	0.0043	0.0152	0.0090	0.0237
China	0.0011	0.0031	0.0003	0.0002
Libya	0.0002	0.0003	0.0009	0.0007
	Cabinet	President	Head of government	Head of state
United States	0.0031	0.0388	0.0015	0.1119
France	0.0051	0.0159	0.0363	0.0183
China	0.0161	0.0123	0.0001	0.0296
Libya	0.0071	0.0003	0.0209	0.0214

Table 1: Legislative and Executive Topics for U.S., Libya, France, and China **(cont.)**

in the U.S. and French constitutions for district size are quadruple and double that of China, respectively. This might suggest a relationship between the extent to which a constitution encodes elections procedures and the procedural reliability of those elections.

Moreover, Libya, essentially a military dictatorship under Qaddafi, allocated almost 6% of its constitution to military affairs – almost four times the similar proportion in the U.S., French, and Chinese constitutions. As expected, this suggests that some types of authoritarian regimes feel compelled to enshrine the role of the military constitutionally (H6).

6.2 Cross–national Comparisons

While the results from the case studies seem to confirm all of the hypotheses, the results from the topic model also allow us to conduct cross–national tests. For Figures 3–10, I plot the average proportion of each topic listed against Freedom House political rights scores. With a few exceptions, my hypotheses are also confirmed at the cross–national level.

In Figure 3, more democratic countries tend to allocate greater portions of their constitutions to judicial appeals than are more authoritarian countries (H1). For the judicial appointments topic, democratic countries do tend to allocate more text, but the difference is small and does not fall outside the standard errors of each estimate. There is no significant difference by regime type with respect to the general "courts" topic, suggesting that most constitutions contain at least some provision for a supreme or constitutional court. However, it is also possible that older constitutions, whose drafters often made implicit assumptions about the appointment procedures, are driving the scores for these latter two topics downwards.

In Figure 4, it becomes clear that the U.S. was not such an outlier amongst democratic countries with regards to protecting criminal rights (H2). On average, countries with a Freedom House score between 1–2 allocate twice as much constitutional text to sentencing when compared to those that score 5–7. While the differences are not quite so dramatic for the criminal trials and arrests topics, there remains a significant difference between countries that score a 1 and those between 6–7. The effect is weaker with respect to criminal arrests, although even countries that score a 3 allocate a greater proportion of their constitutions to criminal arrests than do countries that score a 5. Interestingly, there is no discernible difference for the criminal penalties topic, although this topic of course does not say anything about the substance or harshness of those penalties.

The results for other rights in Figures 5 & 6 are much less clear (H2). The only right for which democratic countries outperform authoritarian countries is the right to property. At the extremes (Freedom House scores of 6 or 7), authoritarian constitutions generally include more text related to protecting rights, equal protection, socioeconomic rights, and education and health. As mentioned above, this might simply suggest that it is simply less costly for authoritarian regimes to promise positive rights – especially if those rights are not legally enforceable (H3). It might also suggest that authoritarian constitutions use more text for these rights as they need to explain limitations on those rights, as suggested with the example from Myanmar above.

There is a clear distinction between how democratic and authoritarian constitutions handle the legislative and executive branches (Figures 7 & 8). For all of the legislative topics except for "congress," democratic countries scoring between 1–2 outscore authoritarian countries scoring between 6–7, and even for "congress" there is a significant difference for countries with a Freedom House score of 1. The results are more ambiguous for those countries scoring between 3–4, which might reflect the fact that many of those countries tend to be quasi-authoritarian and often possess "democratic" institutions. By contrast, we see not significant difference for the four topics related to executive topics. These results suggest that democratic constitutions allocate relatively more text to legislative topics, although the results do not distinguish between the coordination and commitment theories.

Interestingly, the differences between democratic and authoritarian constitutions with respect electoral topics is not quite as clear as expected (Figure 9). On the one hand, democratic and even quasi-democratic constitutions (Freedom House scores 1–2) do outperform authoritarian constitutions with regard to "district size." On the other hand, there is no significant difference between regime types with respect to the broader "elections" topic. One possibility is that the type of election topic matters. While most countries do hold elections, the freeness and fairness of those elections vary. We might expect that constitutions will allocate more text to specific electoral procedures, such as district size, when they intend to commit themselves to those elections (H5).

Finally, Figure 10 shows a dramatic difference between how democratic with a Freedom House score of 1 and authoritarian constitutions address the armed forces. Countries with a Freedom House score of 7 tend to allocate proportionately 50% as much text to military affairs as are countries with a score of 1. Moreover, the difference remains significant when comparing countries with a Freedom House score of 1 versus a score of 5. The differences only become more ambiguous for the countries that score 3–4. By contrast, we see no significant difference with broader war powers or emergency powers. Unfortunately, the latent text model cannot unpack the content of the "armed forces" topic, so it is unclear whether it includes substantive regulations of the armed forces or merely mentions the military as an important political force.

7 Conclusion

For years, one of the most significant obstacles to research in comparative constitutional politics has been the lack of reliable data. Coding constitutions entails a significant investment of time and resources. More importantly, hand-coding might not capture the textual information most relevant to cross–national comparisons. Latent topic models provide a new way to take advantage of our most abundant source of data: the words in the constitutions themselves. While Latent Dirichlet Allocation and other topic models are not holy grails, they do help mitigate the risk of bias and coding error inherent in hand-coded data.

In this paper, I have attempted to contribute both methodological and theoretical advances. First, I propose a new use for latent topic models by analyzing the distribution of topics within each country's constitution. To provide the model with sufficient information over a small number of constitutions (184), I disaggregated each constitution into the largest constituent section with substantive text.

I then used the results to test several hypotheses about the relationship between regime type and constitutional texts. First, I found that democratic constitutions are more likely to allocate constitutional text to institutions that enforce constitutions, particularly courts. Democratic constitutions also allocate a greater proportion of text to negative rights, such as property and criminal procedure. By contrast, authoritarian constitutions spend more text on socioeconomic rights, perhaps because it is less costly for authoritarian regimes to provide rights that are not legally enforceable.

While democratic and authoritarian constitutions did not significant differ with respect to their treatment of the executive, democratic constitutions do allocate a proportionately greater amount of text to the legislative branch. This could signal that democratic regimes use the legislature as a commitment device, or simply reflect the greater need for those constitutions to assign responsibilities to the legislature. Democratic constitutions are also more likely to discuss some electoral topics, particularly district size. Finally, as expected, authoritarian regimes are significantly more likely to include topics related to the military in the constitution.

Future iterations of this paper will check the robustness of these results, both by utilizing different types of topic models and different political measures. The Correlated-Topic Model possesses some advantages over LDA in that it allows topics to be correlated (e.g., "criminal trials" and "criminal penalties"), allowing for more and more refined topics (Blei and Lafferty, 2007). I also plan to rerun the model after having broken the constitutions into smaller subunits (such as sentence-level rather than section- or article-level documents) to make sure that the unit of analysis does not bias results. Finally, I will substitute Polity scores for Freedom House scores to check if the measure of regime type influences the results (Marshall and Jaggers, 2010).

It will also be important to ensure that the results are not biased by political or economic variables correlated with regime type. For example, it is possible that authoritarian constitutions are more likely to protect socioeconomic rights because they also tend to be more recent and therefore reflect changing international norms. I plan to use data on the age of each constitution, as well as each country's legal origins and GDP per capita.

References

Robert Barros. *Constitutionalism and dictatorship: Pinochet, the Junta, and the 1980 constitution*. Cambridge Studies in the Theory of Democracy. Cambridge University Press, July 2002.

Gerald L. Blasi and David Louis Cingranelli. Do constitutions and institutions help protect human rights? In Stuart S. Nagel and David Louis Cingranelli, editors, *Human Rights and Developing Countries*. JAI Press, Inc., 1996.

David M. Blei and John D. Lafferty. A correlated topic model of science. *Annals of Applied Statistics*, 1(1):17–35, 2007.

David M. Blei, Andrew Y. Ng, and Michael I. Jordan. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022, 2003.

Nathan J. Brown. *Constitutions in a Nonconstitutional World: Arab Basic Laws and the Prospects for Accountable Government*. State University of New York Press, 2001.

Jonathan Chang, Jordan Boyd-Graber, Sean Gerrish, Chong Wang, and David M. Blei. Reading tea leaves: How humans interpret topic models. *Advances in Neural Information Processing Systems*, pages 288–296, 2009.

Isaac Ehrlich and Richard A. Posner. An economic analysis of legal rulemaking. *The Journal of Legal Studies*, 3(1):257–286, January 1974.

Zachary Elkins, Tom Ginsburg, and James Melton. *The Endurance of National Constitutions*. Cambridge University Press, 2009.

David Erdos. Postmaterialist social constituencies and political triggers: Ex- plaining bill of rights genesis in internally stable, advanced democracies. *Political Research Quarterly*, 62 (4):798–810, 2009.

Jodi S. Finkel. *Judicial Reform as Political Insurance: Argentina, Peru, and Mexico in the 1990s*. ND Kellogg Inst Int'l Studies. University of Notre Dame Press, May 2008.

Tom Ginsburg. *Judicial Review in New Democracies: Constitutional Courts in Asian Cases*. Cambridge University Press, July 2003.

Tom Ginsburg. Constitutional specificity, unwritten understandings and constitutional agreement. In Andras Sajo and Renaita Uitz, editors, *Constitutional Topography: Values and Constitutions*. Eleven International Publishing, 2010.

Mary Ann Glendon, Paolo G. Carozza, and Colin B. Picker. *Comparative Legal Traditions in a Nutshell (In a Nutshell)*. West Publishing, 2008.

Julian Go. A globalizing constitutionalism? *International Sociology*, 18(1):71–95, March 2003.

Justin Grimmer. A bayesian hierarchical topic model for political texts: Measuring expressed agendas in senate press releases. *Political Analysis*, 18:1–35, 2010.

T.L. Griths and M. Steyvers. Finding scientic topics. *Proceedings of the National Academy of Sciences of the United States of America*, 101:5228–5235, 2004.

Bettina Grün and Kurt Hornik. topicmodels: An r package for fitting topic models. *Journal of Statistical Software*, 40(13):1–30, 2011.

Anna Harvey and Michael J. Woodruff. Confirmation bias in the united states supreme court judicial database. *Journal of Law, Economics, and Organization*, 2011.

Gretchen Helmke and Frances Rosenbluth. Regimes and the rule of law: Judicial independence in comparative perspective. *Annual Review of Political Science*, 12:345–366, 2009.

Ran Hirschl. *Towards Juristocracy: The Origins and Consequences of the New Constitutionalism*. Harvard University Press, 2004.

Freedom House. Freedom in the world. Technical report, Freedom House, 2011.

Vicki C. Jackson and Mark V. Tushnet, editors. *Comparative Constitutional Law*. Foundation Press, 2006.

David S. Law and Mila Versteeg. Evolution and ideology of global constitutionalism. *California Law Review*, 99(1163):1219–20, 2011.

David S. Law and Mila Versteeg. Sham constitutions. *University of Virginia School of Law Public Law and Legal Theory Research Paper Series*, 2012-55, 2012.

Monty Marshall and Keith Jaggers. Polity iv project: Political regime characteristics and transitions, 1800-2004. Technical report, Polity IV Project, 2010.

Kevin McGuire and Georg Vanberg. Mapping the policies of the u.s. supreme court: Data, opinions, and constitutional law. In *Annual Meeting of the American Political Science Association*, September 2005.

James Melton, Zachary Elkins, and Tom Ginsburg. The content of authoritarian constitutions. In *Paper for Conference on Constitutions in Authoritarian Regimes*. University of Chicago, 2011.

Douglas C. North and Barry R. Weingast. Constitutions and commitment: The evolution of institutions governing public choice in seventeenth-century england. *The Journal of Economic History*, 49(4):803–832, 1989.

Douglas Rice. Measuring the issue content of supreme court opinions through probabilistic topic models. Prepared for presentation at the 2012 Midwest Political Science Association Conference, Chicago, Illinois, April 12-15, 2012, April 2012.

Michal Rosen-Zvi, Thomas Griths, Mark Steyvers, and Padhraic Smyth. The author-topic model for authors and documents. *In 20th Conference on Uncertainty in Artificial Intelligence. Banff, Canada*, 2004.

Wayne Sandholtz. Treaties, constitutions, courts, and human rights. *Journal of Human Rights*, 11(1):17–32, 2012.

Harold H. Spaeth, Lee Epstein, Ted Ruger, Keith Whittington, Jeffrey A. Segal, and Andrew D. Martin. *Supreme Court Database Code Book*, version 201201 edition, September 2012.

C. Neal Tate. The global expansion of judicial power: The judicialization of politics. In C. Neal Tate and Torbjorn Vallinder, editors, *The Global Expansion of Judicial Power*. NYU Press, 1997.

George Tsebelis. *Veto Players: How Political Institutions Work*. Princeton University Press, October 2002.

1. Sentencing	2. Land	3. Publication	4. In force	5. Parties
sentenc	land	decre	befor	parti
forc	union	forc	forc	polit
requir	amend	issu	commenc	associ
imprison	agricultur	offici	order	form
labor	reform	publish	exist	group
6. Agencies	7. Schedules	8. Federalism	9. Criminal trials	10. Councils
organ	claus	feder	offens	council
establish	schedul	para	court	member
administr	order	compet	crimin	presid
function	make	land	trial	appoint
determin	respect	matter	proceed	consult
11. Territory	12. Head of government	13. Commission members	14. Judicial appointment	15. Public funds
territori	minist	commiss	court	fund
degre	prime	servic	judg	appropri
state	appoint	appoint	justic	expenditur
integr	presid	member	chief	monei
line	member	commission	appoint	consolid
16. Criminal penalties	17. Oaths (religious)	18. Complaints	19. War powers	20. Legislative membership
crime	duti	action	declar	member
punish	oath	investig	emerg	speaker
penalti	swear	decis	resolut	elect
offens	defend	ombudsman	state	seat
commit	islam	violat	war	becom
21. Principles	22. Majority	23. Military	24. Expression	25. Labor
peopl	vote	forc	commun	work
nation	member	defens	freedom	employ
principl	major	secur	religi	worker
democrat	present	arm	religion	posit
state	two-third	militari	express	employe

Table 2: Top 5 terms with LDA when the number of topics set at k=67

Barry R. Weingast. The political foundations of democracy and the rule of law. *The American Political Science Review*, 91(2):245–263, June 1997.

26. Judicial appeals	27. Relational	28. Property rights	29. Taxes	30. Citizen rights
court	thei	properti	tax	nation
appeal	case	purpos	properti	citizen
decis	establish	right	servic	resid
high	correspond	reason	asset	right
jurisdict	concern	restrict	establish	foreign
31. Legislative chambers	**32. Policy implementation**	**33. Provincial government**	**34. Cabinet responsibility**	**35. Auditing**
deputi	polici	provinc	respons	report
chamber	implement	governor	cabinet	account
senat	ensur	provinci	secretari	audit
motion	plan	legislatur	minist	auditor-gener
majli	nation	capit	depart	submit
36. Appointments	**37. Delegation**	**38. District size**	**39. Initiating legislation**	**40. Education & health**
appoint	legisl	elector	propos	educ
remov	execut	number	submit	state
governor-gener	congress	constitu	approv	health
accord	nation	elect	referendum	famili
advic	deleg	member	adopt	right
41. Elections	**42. Protecting rights**	**43. Public salaries**	**44. Internal reference**	**45. Equal protection**
elect	right	grant	mean	equal
candid	freedom	allow	refer	discrimin
vote	guarante	pension	includ	subject
nomin	protect	benefit	relat	origin
ballot	citizen	servic	chapter	sex
46. Passing laws	**47. Oaths (secular)**	**48. Budgets**	**49. Criminal arrest**	**50. Congress**
pass	oath	budget	arrest	hous
amend	hold	revenu	detent	senat
assent	befor	financ	order	repres
alter	prescrib	financi	detain	member
motion	set	year	reason	joint

Table 3: Top 5 terms with LDA when the number of topics set at k=67
(cont.)

51. International affairs	52. Legislative sessions	53. Prosecution	54. Socioeconomic rights	55. Presidency
intern	session	proceed	econom	presid
treati	meet	prosecut	state	vice-presid
agreement	conven	director	social	case
languag	presid	attorney-gener	develop	resign
nation	extraordinari	crimin	cultur	absenc
56. Terms	**57. High court**	**58. Procedures**	**59. Parliament**	**60. Naturalize**
term	court	rule	parliament	citizen
elect	suprem	procedur	member	citizenship
year	judici	regul	subject	born
new	justic	matter	parliamentari	becom
mandat	tribun	provid	make	regist
61. Committees	**62. Time**	**63. Head of state**	**64. Local government**	**65. Functions**
committe	dai	state	local	function
people'	period	head	municip	perform
stand	time	presid	region	exercis
matter	date	repres	district	duti
order	month	foreign	bodi	discharg
66. Royalty	**67. Assemblies**			
king	assembl			
yang	nation			
majesti	member			
di-pertuan	presid			
sultan	befor			

Table 4: Top 5 terms with LDA when the number of topics set at k=67
(cont.)

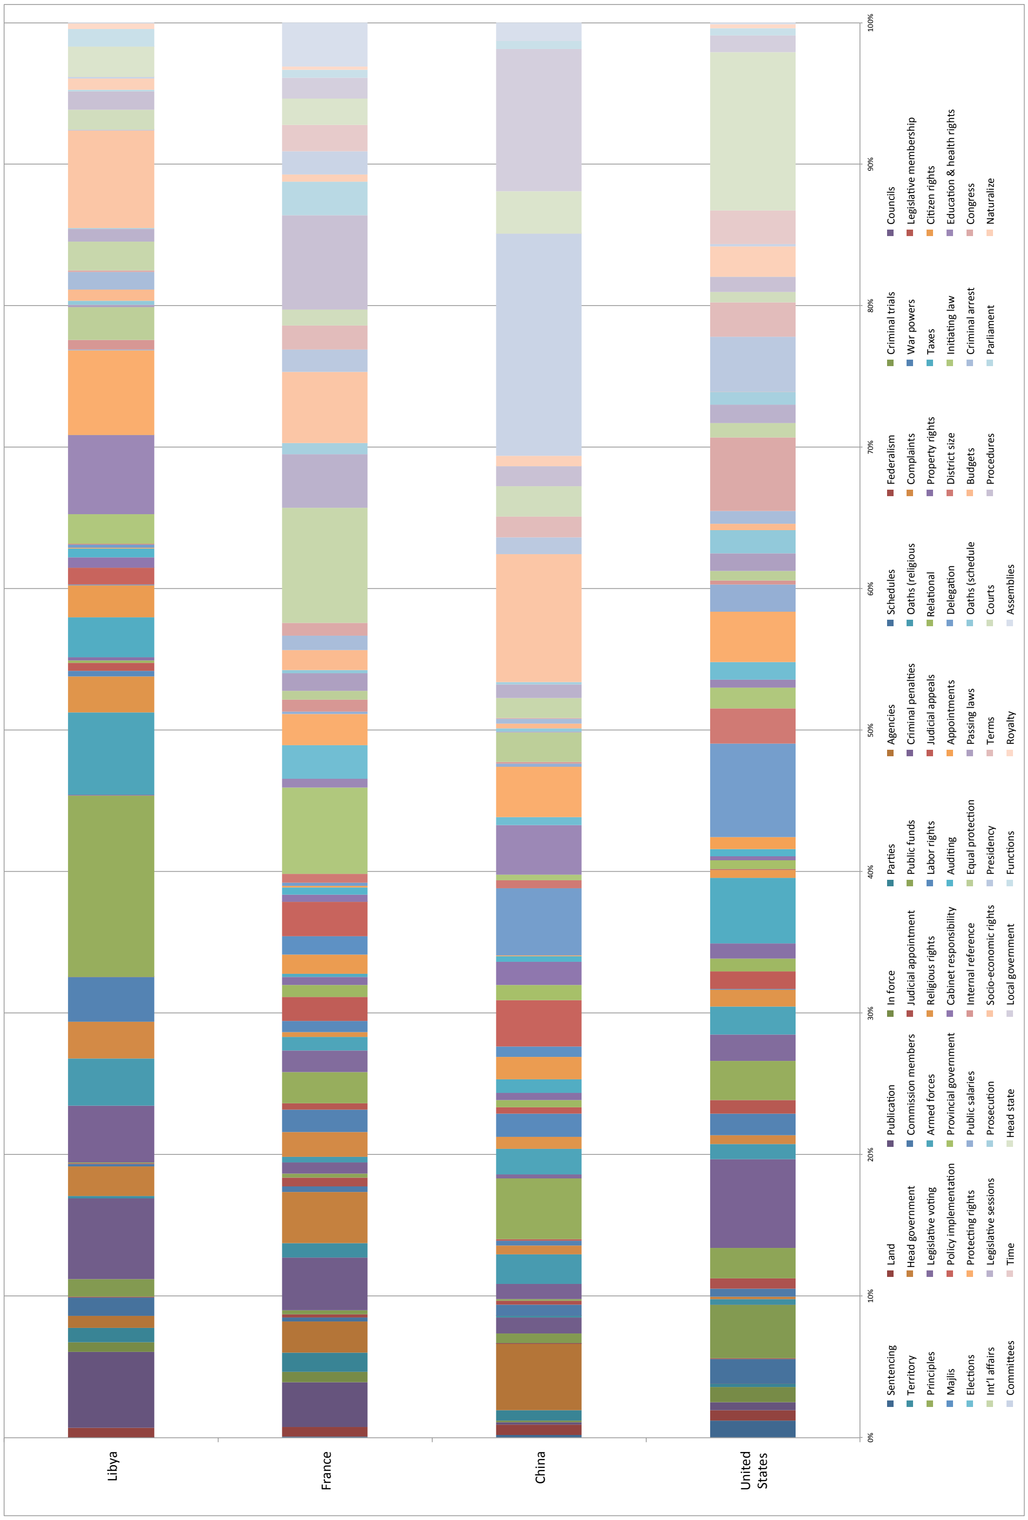

Figure 2: Distribution of topics for the United States, France, China, and Libya

Each color band represents the posterior probability assigned to that particular topic as a proportion of the total document.

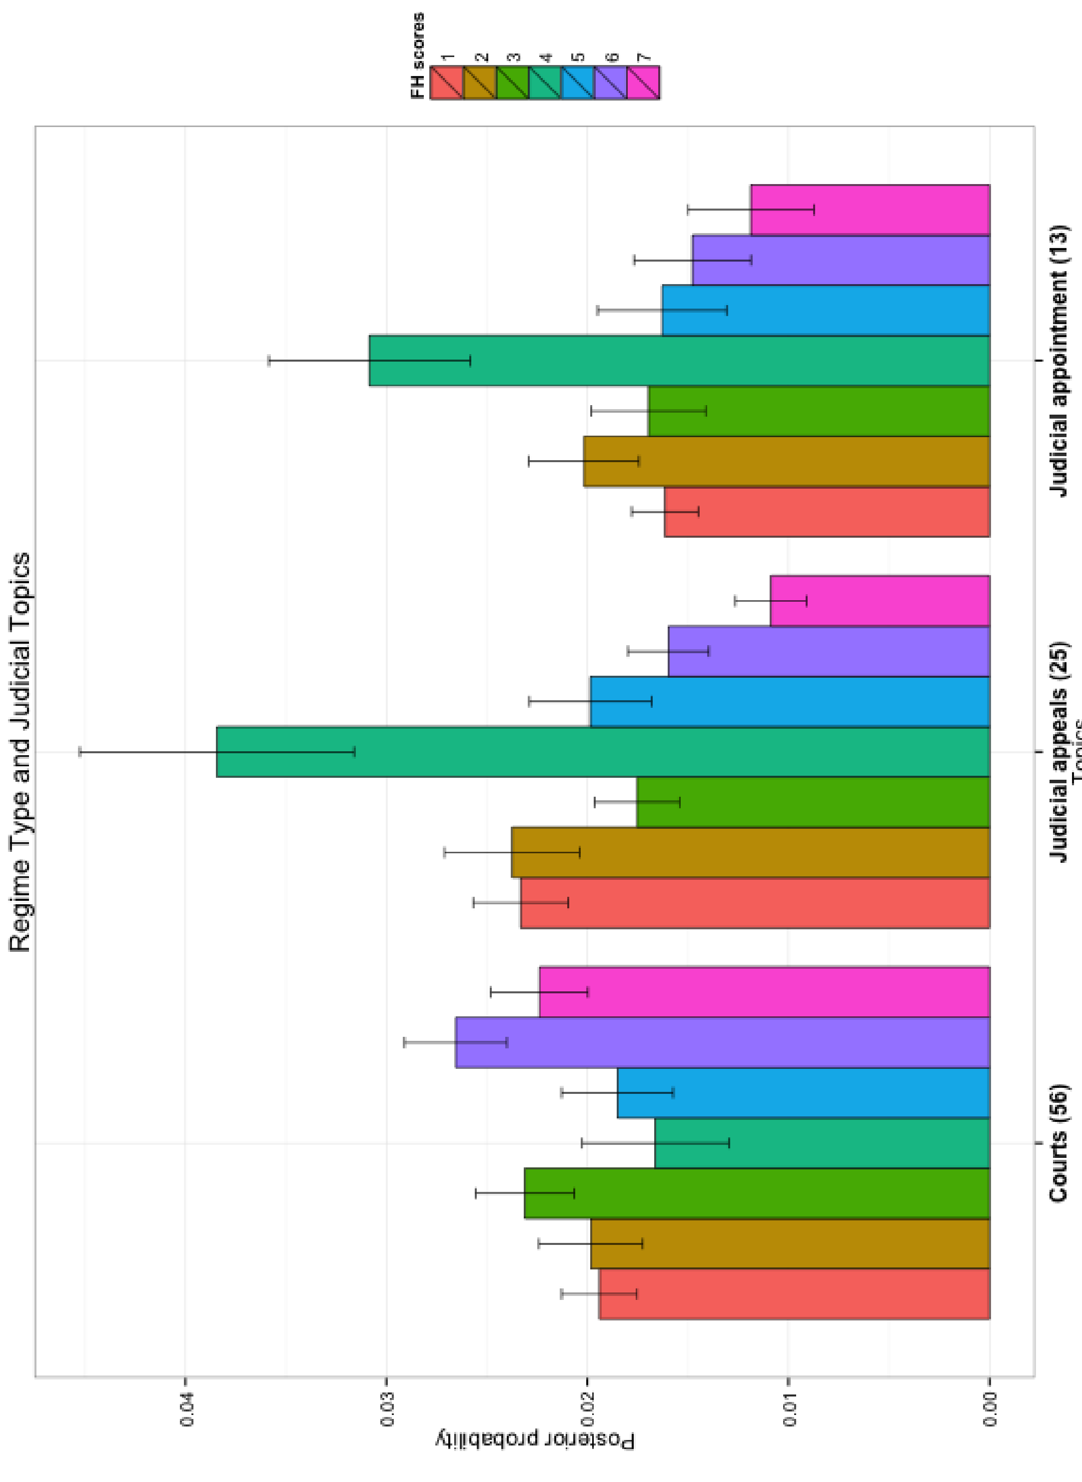

Figure 3: Freedom House Scores and Judicial Topics

Freedom House scores range from 1–7, with 1 indicating a more democratic/free political system. Complex posteriors are the probability estimate of the document allocated to a particular topic. Bars represent standard errors.

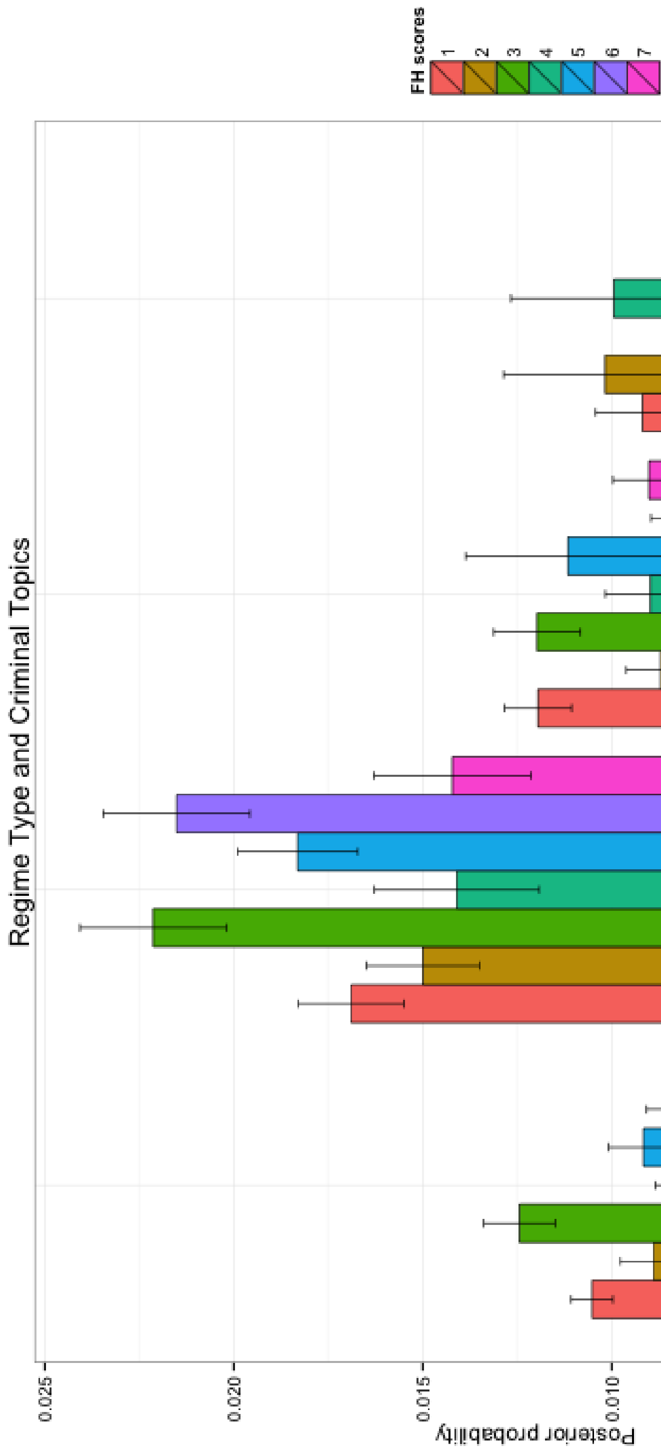

Figure 4: Freedom House Scores and Criminal Topics

Freedom House scores range from 1–7, with 1 indicating a more democratic/free political system. Complex posteriors are the probability estimate of the proportion of the document allocated to a particular topic. Bars represent standard errors.

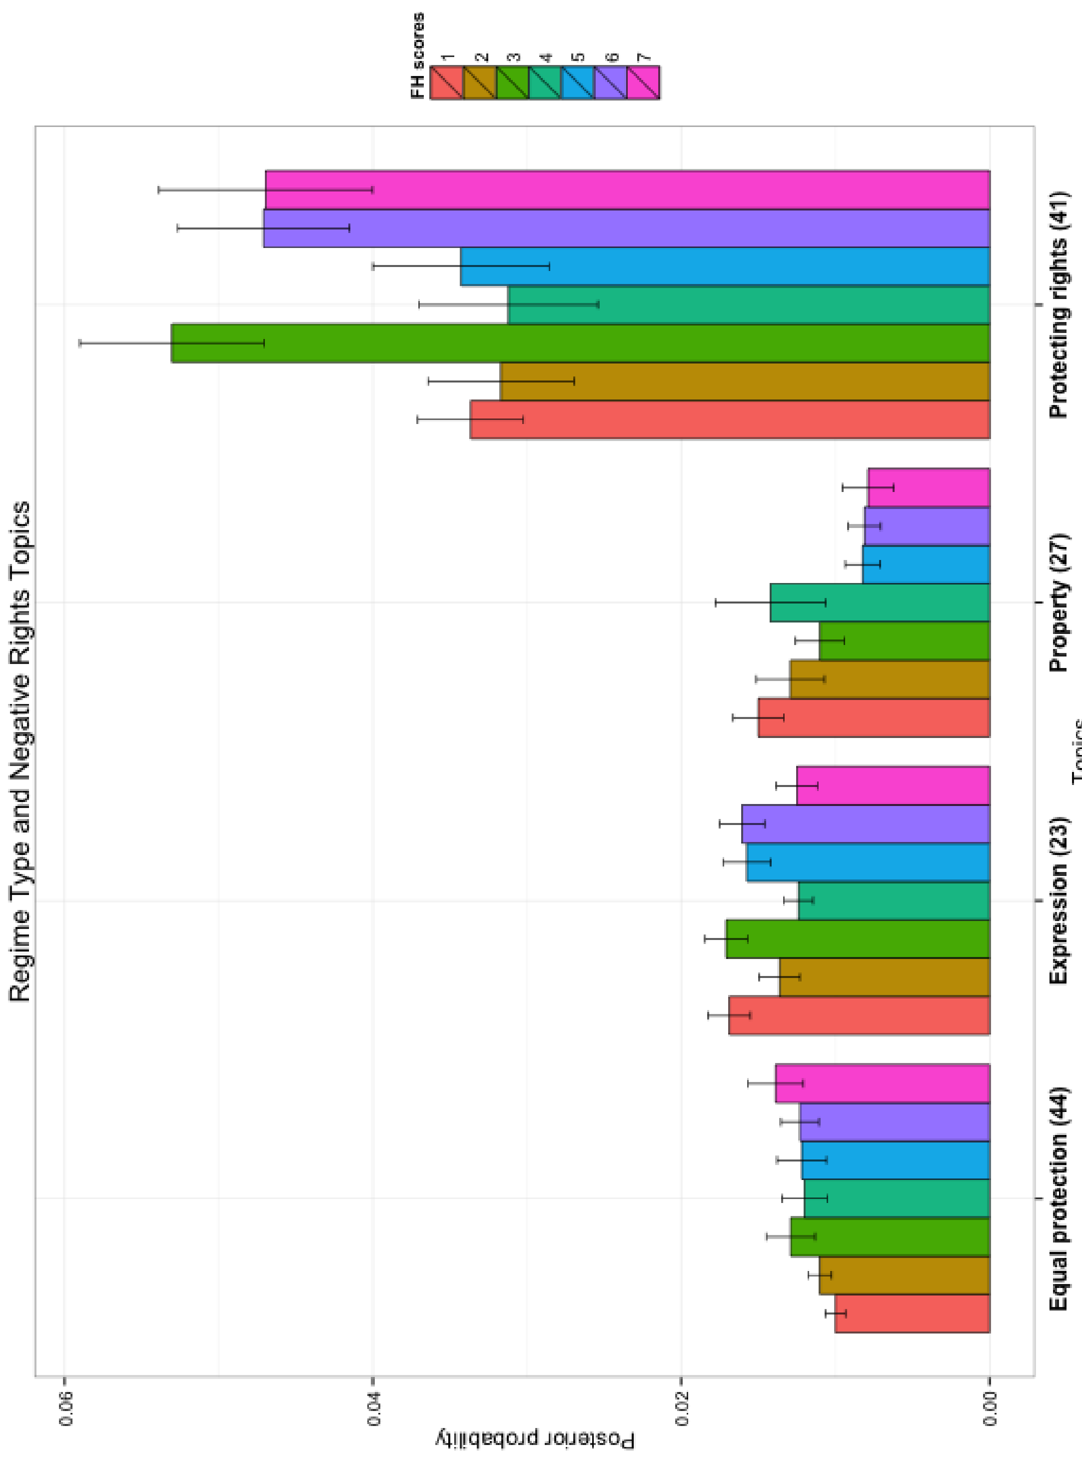

Figure 5: Freedom House Scores and Positive Rights Topics
Freedom House scores range from 1–7, with 1 indicating a more democratic/free political system. Complex posteriors are the probability estimate of the proportion of the document allocated to a particular topic. Bars represent standard errors.

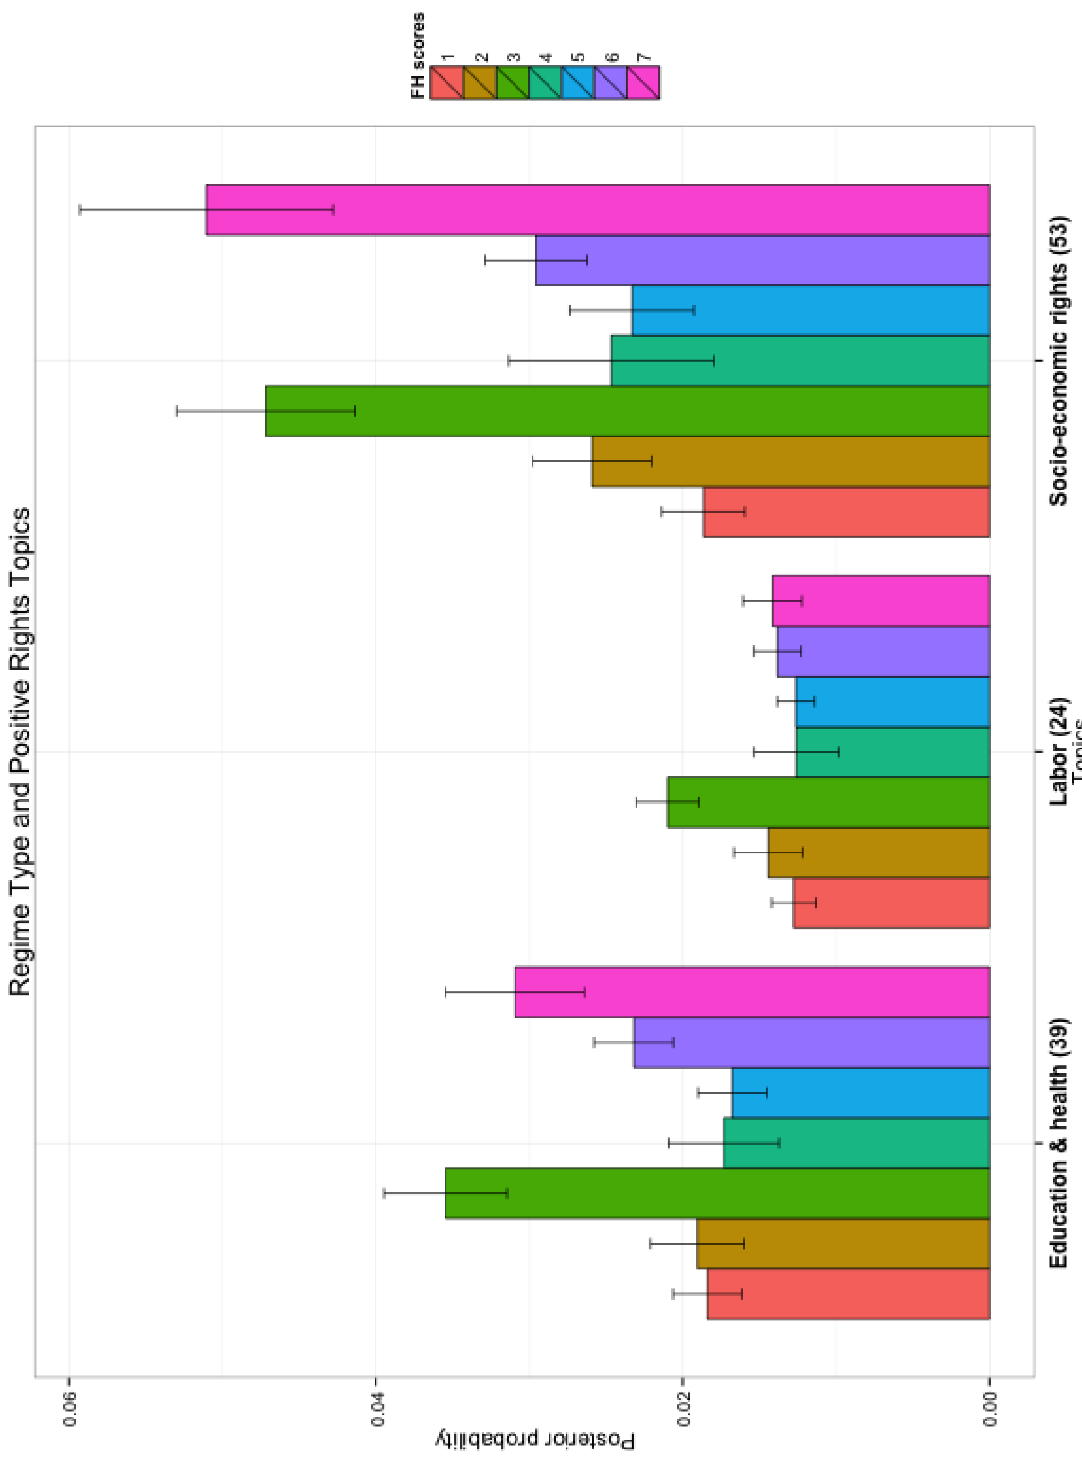

Figure 6: Freedom House Scores and Positive Rights Topics
Freedom House scores range from 1–7, with 1 indicating a more democratic/free political system. Complex posteriors are the probability estimate of the proportion of the document allocated to a particular topic. Bars represent standard errors.

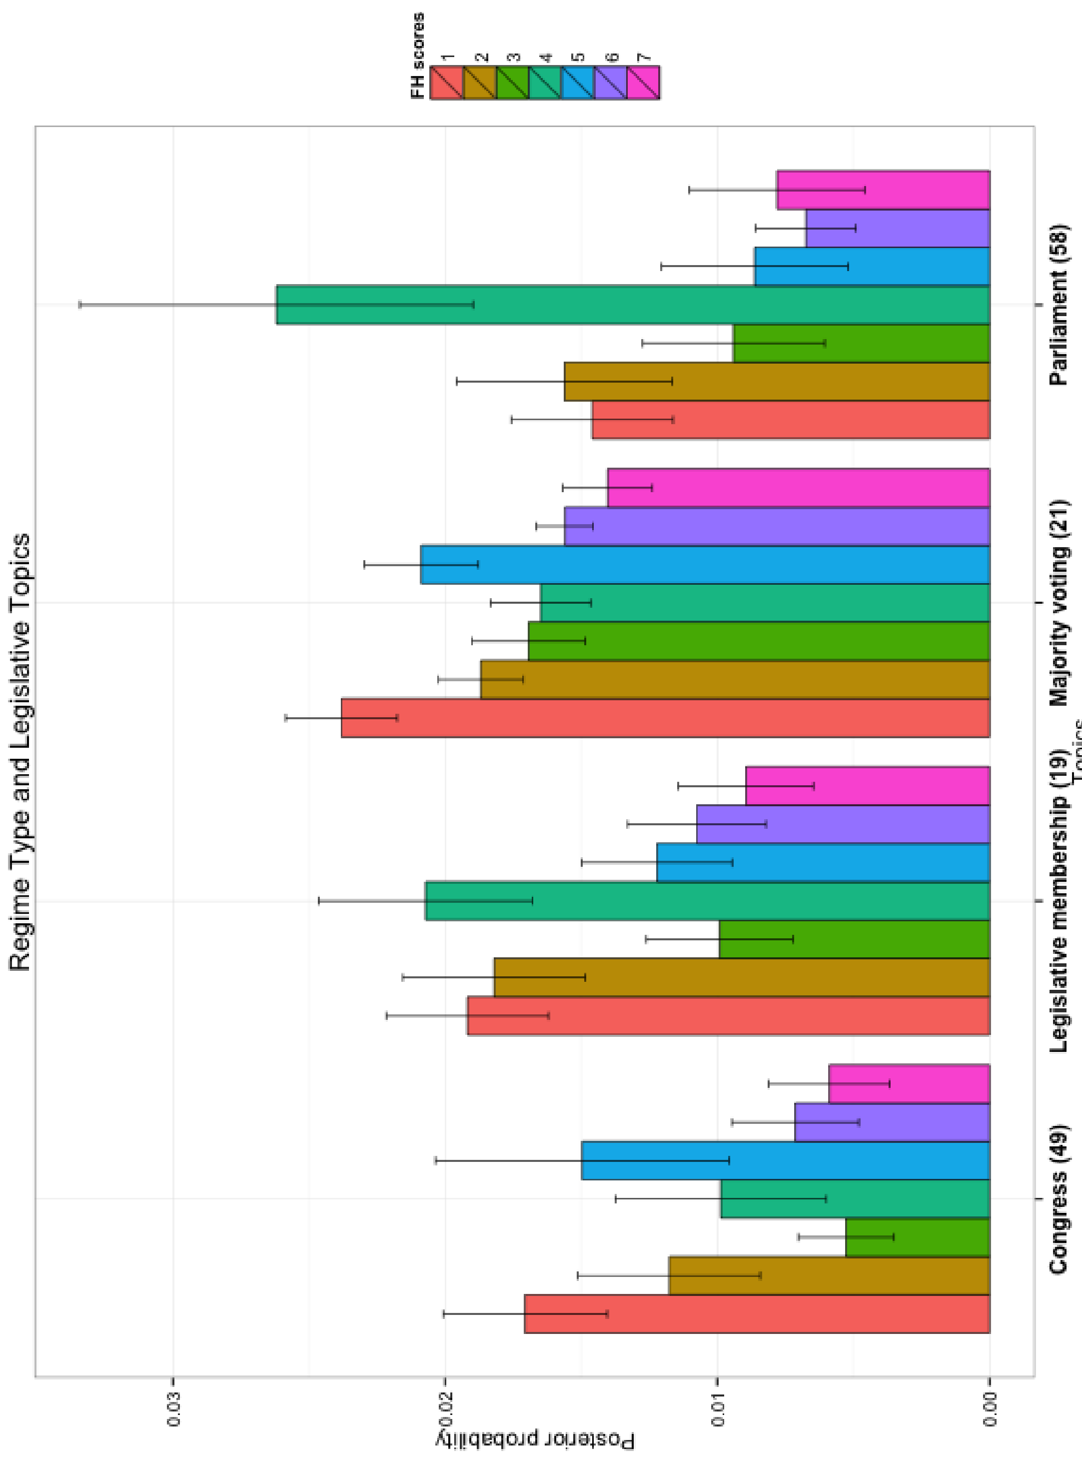

Figure 7: Freedom House Scores and Legislative Topics
Freedom House scores range from 1–7, with 1 indicating a more democratic/free political system. Complex posteriors are the probability estimate of the proportion of the document allocated to a particular topic. Bars represent standard errors.

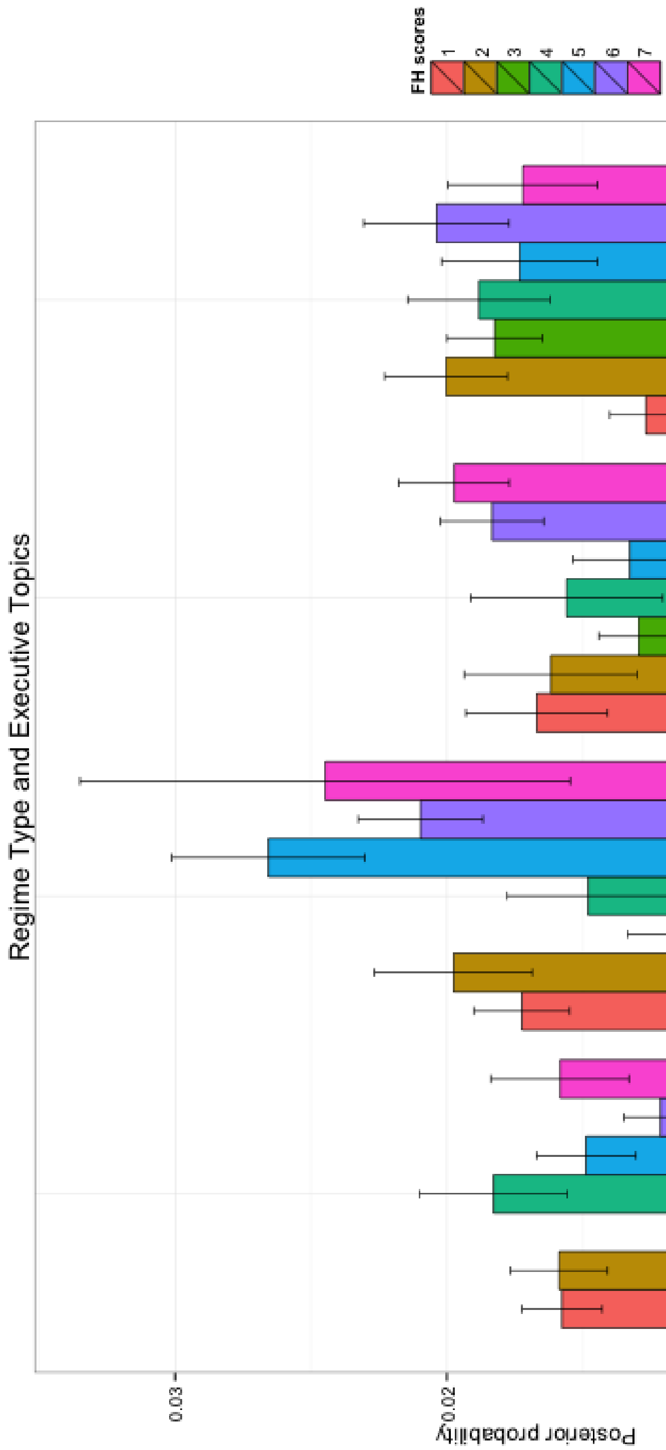

Figure 8: Freedom House Scores and Executive Topics

Freedom House scores range from 1–7, with 1 indicating a more democratic/free political system. Complex posteriors are the probability estimate of the proportion of the document allocated to a particular topic. Bars represent standard errors.

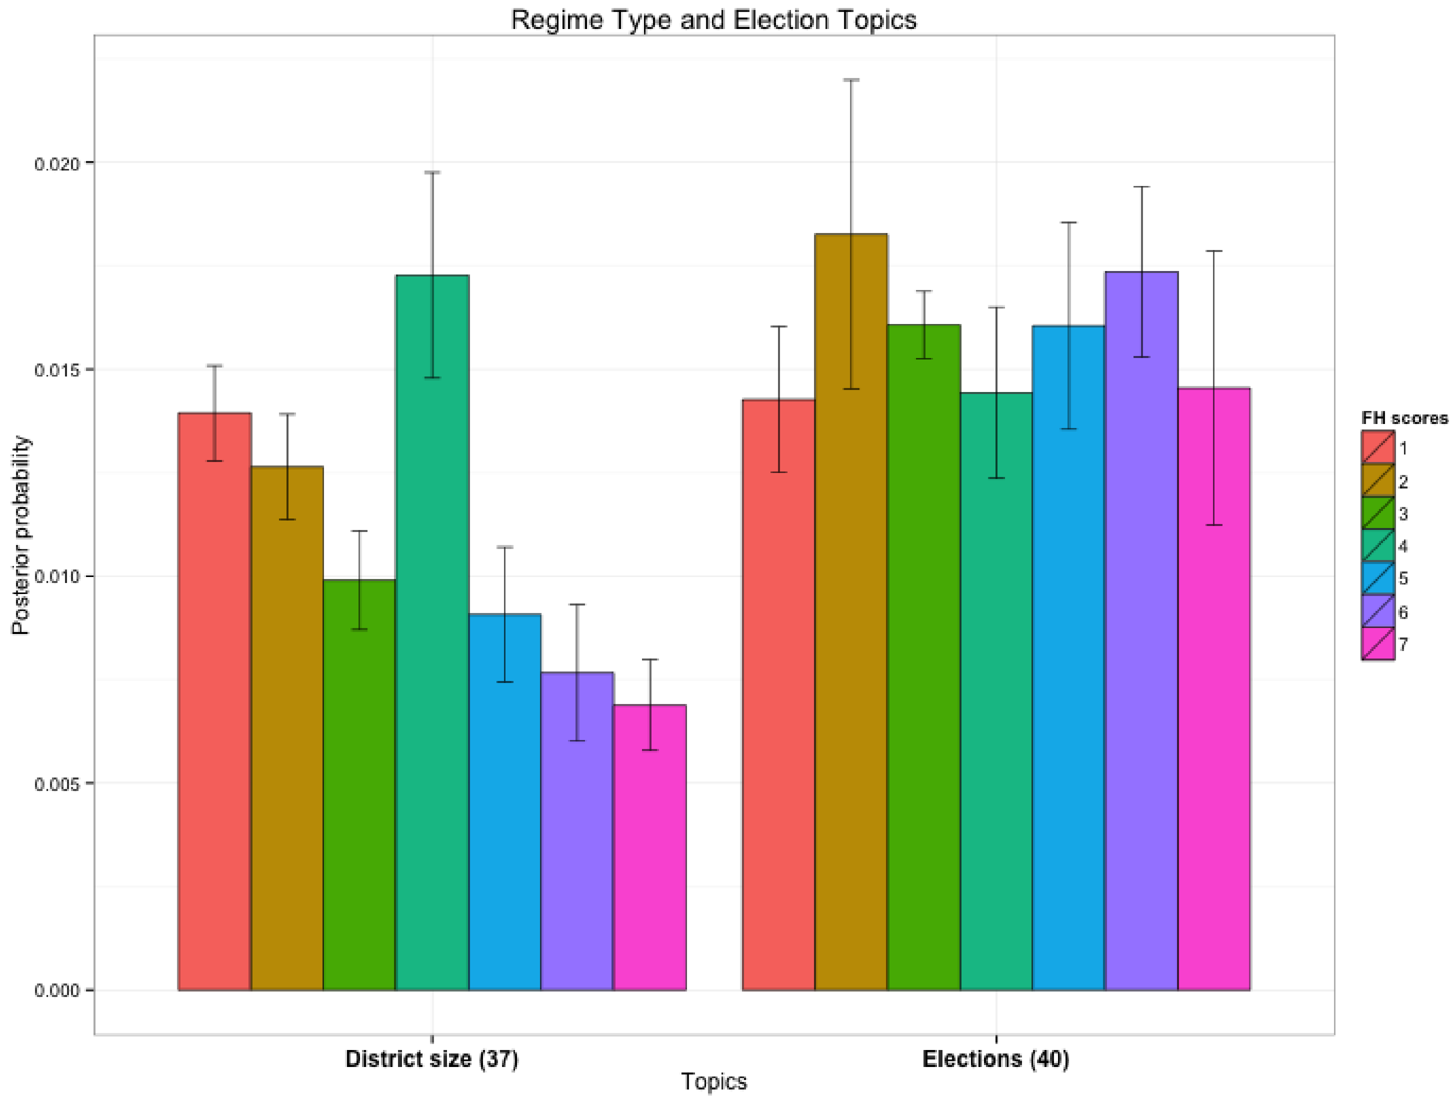

Figure 9: Freedom House Scores and Elections Topics
Freedom House scores range from 1–7, with 1 indicating a more democratic/free political system. Complex posteriors are the probability estimate of the proportion of the document allocated to a particular topic. Bars represent standard errors.

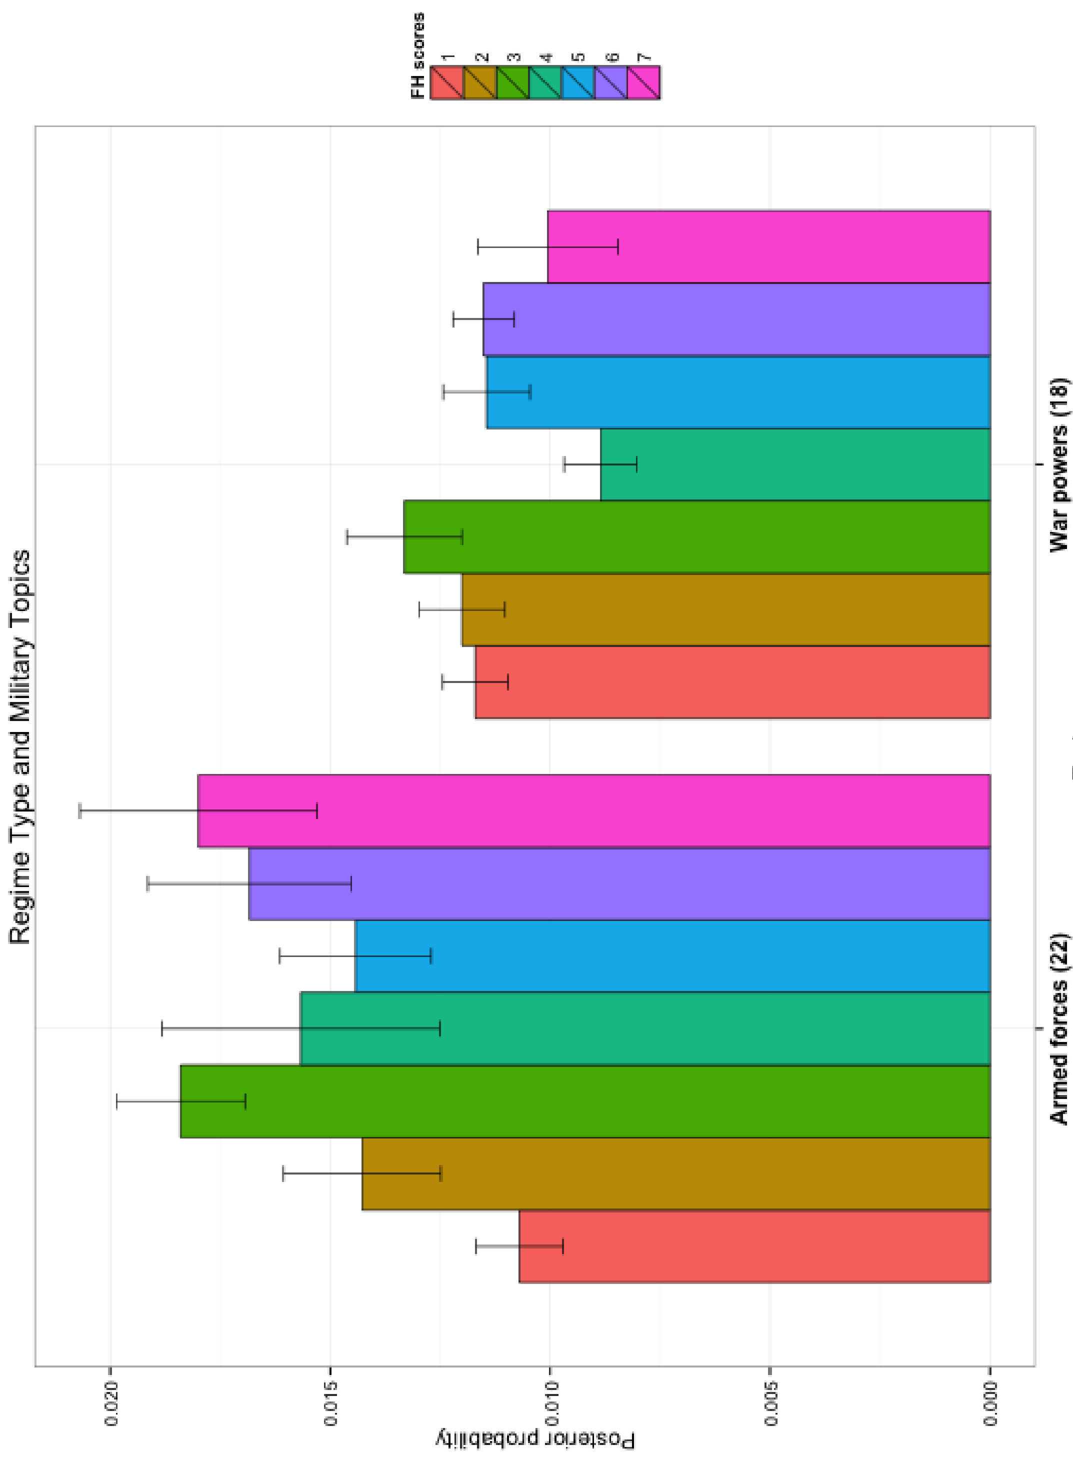

Figure 10: Freedom House Scores and Armed Forces Topics

Freedom House scores range from 1–7, with 1 indicating a more democratic/free political system. Complex posteriors are the probability estimate of the proportion of the document allocated to a particular topic. Bars represent standard errors.