Measuring a Contract's Breadth: A Text Analysis

Joshua Hall, Bryan McCannon, and Yang Zhou

MERCATUS WORKING PAPER

All studies in the Mercatus Working Paper series have followed a rigorous process of academic evaluation, including (except where otherwise noted) at least one double-blind peer review. Working Papers present an author's provisional findings, which, upon further consideration and revision, are likely to be republished in an academic journal. The opinions expressed in Mercatus Working Papers are the authors' and do not represent official positions of the Mercatus Center or George Mason University.

MERCATUS CENTER
George Mason University
3434 Washington Blvd., 4th Floor, Arlington, Virginia 22201
www.mercatus.org

Joshua Hall, Bryan McCannon, and Yang Zhou, "Measuring a Contract's Breadth: A Text Analysis," Mercatus Working Paper, Mercatus Center at George Mason University, Arlington, VA, March 2021.

Abstract

We use a computational linguistic algorithm to measure the topics covered in the text of school teacher contracts in Ohio. We use the topic modeling metrics in a calculation of the concentration of topics covered. This allows us to assess how expansive each contract is. As a proof of concept, we evaluate the relationship between our topic diversity measurement and the prevalence of support staff. This test is done on a subsample of the contracts in the state. If more specialized services are provided, then contracts must presumably be broader as they cover more employment relationships. We confirm a strong, statistically significant relationship between our measurement and the prevalence of these support staff. Thus, we have a valid measurement of contract breadth.

JEL codes: D83; L15; C81

Keywords: contracts; Latent Dirichlet Allocation; schools; teachers; text analysis; topic modeling; union

Author Affiliation and Contact Information

Joshua Hall
John Chambers College of Business and Economics, West Virginia University
joshua.hall@mail.wvu.edu

Bryan McCannon
John Chambers College of Business and Economic, West Virginia University
bryan.c.mccannon@gmail.com

Yang Zhou
Department of Economics, University of North Texas
yang.zhou2@unt.edu

© 2021 by Joshua Hall, Bryan McCannon, Yang Zhou, and the Mercatus Center at George Mason University

This paper can be accessed at https://www.mercatus.org/publications/regulation/measuring-contracts-breadth-text-analysis.

Measuring a Contract's Breadth: A Text Analysis

Joshua Hall, Bryan McCannon, and Yang Zhou

1. Introduction

The economic world occurs in words. Contracts are written, lawmakers make speeches, regulations are codified, and formal ownership of land is recorded in a title. The future of economic analysis is in the ability to quantify text.

The focus of this paper is on contracts. A typical contract is multidimensional and complex. Employment contracts, in particular, not only lay out the wage to be paid but also describe the monetary and nonmonetary benefits, working conditions, dispute resolution mechanisms, and more. Labor contracts must cover many dimensions. Our objective is to propose a method to measure the expansiveness of a contract.

We collect a corpus of school teacher contracts negotiated between the teacher's unions and school administrations in Ohio. These contracts cover compensation, benefits, leave of absence policies, school year and school day structure, and grievance procedures. They are lengthy and detailed.

We propose to use a topic modeling approach developed in computer science to analyze these texts. Specifically, we apply Latent Dirichlet Allocation (LDA). Details of LDA are provided in section 3. In short, though, LDA is an untrained, unsupervised algorithm. A set of documents is organized into T topics. Similar texts are grouped together to make up these topics. Each document, then, can be scored on the probability it falls into each of these. Using this approach, we classify the text within the contracts.

Our approach is to organize texts into a large number of topics. LDA defines a topic as a specific probability distribution over words in the dictionary. Thus, a topic is made up of words that are likely to occur in a text. Expansive contracts will cover many of the topics. Narrow contracts will cover only a few. To illustrate, a narrow contract may cover sick leave and personal days in the leave policy. An expansive contract will cover both, as well as policies for family health problems, military service, jury duty, and more. In a contract that covers missed work because of jury duty, to continue with the example, the words likely to arise (e.g., “jury,” “court,” “judge,” etc.) make up a topic. The narrow contract will put zero probability weight on this topic, whereas the broad contract will put a positive weight on it. Our proposal is to measure the concentration of topics covered in a contract, similar to the way scholars in industrial organization study market concentration. Specifically, as our baseline, we will consider 50 topics covered in contracts and measure the Hirschman-Herfindahl Index. A contract that is narrowly focused will have a topic concentration close to one. A broad contract, covering most of the possible topics, will have a lower value to our metric.

The novelty of our estimation strategy is that we are able to take long, expansive texts of contracts and condense them to a single, quantified measurement. Our outcome variable does not tell us what is specifically included in each contract, nor does it detail how generous the terms are. It is a measurement of contract broadness.

This important contribution to the economics of contracts addresses a central question in contract theory. That is, why do parties to a contract not create a fully comprehensive agreement that lays out the actions to be taken for every contingency that arises? According to the theory of incomplete contracts, transaction costs define the limits of a contract’s expansiveness. If external institutions exist that reliably solve disputes, or if social norms and conventions regulate individuals’ behaviors well, then the additional costs needed to fully articulate contingency in the contract are not needed. To date, though, a formal method to measure contract expansiveness has not been developed. We believe our contribution can provide a method needed in this research agenda.

The first step is to provide a proof of concept. We randomly select a sample of 60 school teacher contracts, which represents approximately 10 percent of our full sample. We use LDA to estimate the topic concentration for each contract. We note that school districts cover populations of various sizes and that the services provided at schools differ. For example, in this sample, only 40 percent of the schools hire at least one school social worker. Presumably, if a school does not hire workers to provide such support services, then the contract does not need to be as expansive as social workers' employment conditions do not need to be considered. In our subsample, to check the validity of our measurement, we ask whether our measured contract expansiveness is correlated with the existence of these support staff.

Looking across 10 distinct services, we find, even in our small sample, a strong correlation between our topic concentration measurement and the prevalence of these support services. We take this as strong, preliminary evidence that our proposed measurement is valid.

A few have used topic modeling in the social sciences. Outside of economics, LDA is somewhat popular. As one measure, Blei, Ng, and Jordan's (2003) paper that introduced the algorithm has more than 25,000 citations. It has been used effectively in related fields such as political science. Grimmer (2010) uses LDA to analyze press releases from US senators. Quinn et al. (2010) use it to evaluate speeches in the US Senate. It has been used recently in marketing to evaluate online discussions of products (Tirunillai and Tellis 2014) and in accounting to identify trends in 10-K disclosures (Dyer, Lang, and Stice-Lawrence 2017). The use of LDA in economics is new and represents an important contribution to the field. McCannon (2020) uses the method to classify descriptions of wine. He uses these as explanatory variables in a hedonic price equation to show that wine's price is determined in part by the described characteristics.

The work closest to ours is that of Hansen, McMahon, and Prat (2018). They evaluate Federal Open Market Committee (FOMC) transcripts. They use the variation in topics discussed, assessed using LDA, as the dependent variable to appreciate how experience on the FOMC and transparency interact. Thus, we employ the same estimation strategy in that the concentration of topics covered is the outcome variable of interest. They evaluate whether experience on the committee relates to the breadth of topics discussed by a committee member. We use it as a measurement of expansiveness of a contract.

2. Latent Dirichlet Allocation

Latent Dirichlet Allocation (LDA) is a computational linguistic algorithm (Blei, Ng, and Jordan 2003). It allows for the creation of communication measures based on topic models, which is a class of machine learning algorithms for natural language processing.

LDA allows for automatic clustering of any kind of text documents into a user-chosen number of clusters, known as *topics*. It uses a probabilistic model of text data. The method's logic is that when authors write about a particular theme (i.e., topic), they tend to use the same words. Hence, in texts about the same topic, similar words tend to co-occur. LDA describes each topic as a probability distribution over words and each document as a probability distribution over topics.

To explain, consider the following hypothetical example. Suppose an author is interested in writing a paper about law and economics. The concept behind a probabilistic model of writing is that the author first chooses which topics to include and assigns a probability distribution to these topics. A law and economics paper can be made up of (a) economics, (b) law, and (c) econometrics. Each of these three topics has words associated with it. For example, the topic of economics assigns a high probability to words such as *demand*, *supply*, and *equilibrium*, to name three examples. The topic of law will put a high likelihood on the words *crime*, *precedence*, and *antitrust* being selected.

The topic of econometrics puts large probability weight on words such as *coefficient*, *regression*, and *statistic*. The author chooses a probability distribution over topics, and each topic is defined as a probability distribution over words in the dictionary. By randomly selecting words according to these distributions (and filling in stop words such as *the* and *and*), the paper is written. LDA seeks to recover these topics by identifying the co-occurrence of words in the documents within the corpus of texts.

Here, we briefly describe the LDA algorithm. An interested reader is encouraged to consult Blei, Ng, and Jordan (2003) or Schwarz (2018) for further details. Each document d in a set of documents D is described as a probabilistic mixture of topics T. A document topic vector θ_d describes the document, and a document is determined by a probability distribution over topics. Each topic t in the set of topics T is described by a probability distribution over the vocabulary of words V present in all documents. Additionally, within each topic, there exists a probability distribution of words in the dictionary. For each document, the topic proportions are drawn from a Dirichlet distribution with parameter α, and for every topic the word probability distribution is drawn from a Dirichlet distribution with parameter β. Thus, the researcher must only select the number of topics to organize the documents into and the two hyperparameters α and β.

Gibbs sampling is used to estimate the conditional probabilities that best explain the corpus of documents.[1] We follow the conventional norms in the literature by setting, as the baseline, $\alpha = 0.25$ and $\beta = 0.1$ (Schwarz 2018). Furthermore, the primary specifications to be considered will organize the documents into 50 topics. We choose 50 to match Hansen, McMahon, and Prat (2018). In our smaller subsample, we consider 20 topics.

[1] Following Schwarz (2018), we use the ldagibbs command in Stata 15.

From this process, LDA calculates for each document the probability it falls within topic t, $\rho_d(t)$. With 20 topics, $t = 1, 2, \ldots, 20$, it follows that $\Sigma\rho_d(t) = 1$. These 20 variables, then, classify the topics discussed in the document.

With these 20 measurements, we calculate the concentration of topics for each document. Borrowing from the industrial organization literature, we use the Herfindahl- Hirschman Index (Hirschman 1945; Herfindahl 1950), which is defined as

$$\Sigma\rho_d(t)^2. \tag{1}$$

If a hypothetic contract only covers one of the fifty topics, then $\rho_d(t') = 1$ for topic t', and $\rho_d(t) = 0$ for all other topics. Hence, for this contract, the topic concentration measurement will be equal to 1. Instead, if the document is equally distributed across the 20 topics, then $\rho_d(t) = \frac{1}{20}$ for all t. For this other extreme, the measurement is 0.05. Therefore, our metric captures the diversity of topics covered in each contract. This becomes the outcome variable of interest in our analysis.

LDA is not the only way to quantify text. Previous efforts have relied on dictionary methods. With these the researcher must first select a set of words that are believed to be important. For each document in the corpus, then, the existence of a word or a number of words from that list is counted. See Gentzkow and Shapiro (2010) for a prominent example. This approach has a long history in wine reviews, as discussed previously, but has relied on hand coding of words researchers have deemed important. LDA, however, does not involve a researcher's prior knowledge or discretion. It provides a way of uncovering hidden themes in text without having to link themes to particular word lists prior to estimation. King, Lam, and Roberts (2017) point out that the human brain does not excel at recalling all keywords needed to adequately describe a topic. Instead, humans are good at making associations. Formal dictionary building techniques, such as LDA, do not rely on a researcher's ability to fully construct a keyword list. Thus, LDA is valuable when a researcher does not know a priori which words are the important ones to track. This allows researchers to avoid subjective judgments and to account for context. The work presented in this current paper is, to the best of our knowledge, the first to use LDA to study the breadth of contracts.

3. Ohio School Teacher Contracts and Ohio Data

We use Ohio as the setting for our analysis. We do so for two reasons. First, Ohio has a large number (600+) of local school districts, each of which has its own local teacher's union. Whereas many of these bargaining units are affiliates of the Ohio Federation of Teachers (OFT) or the Ohio Education Association (OEA), each bargaining unit is a separate entity comprising district teachers (and sometimes other employees) who bargain over district-level working conditions and salaries.[2] Second, Ohio's State Employee Relations Board (SERB) maintains a current archive of all active collective bargaining agreements in the state (Ohio State Employee Relations Board 2020).[3] These two conditions provide a great setting to test whether the LDA can be employed to measure contract scope.

Ohio has been the setting of a number of studies related to teacher collective bargaining, in part because of the two reasons cited in the previous paragraph. In the paper closest to ours, Hall, Lacombe, and Pruitt (2017) use the raw number of pages in each contract as a measure of the scope of collective bargaining. Using an education production function approach, they find that more pages in a contract is associated with lower scores on districtwide math tests, but that the effects are zero once spatial spillovers are taken into account. Whereas an improvement on the previous literature that compares unionized with nonunionized districts, the use of page lengths is

---

[2] The OFT and OEA are, in turn, affiliates of the American Federation of Teachers and National Education Association.

[3] SERB provides these documents online as part of its statutory requirement to be a clearinghouse of information on employment practices in the state. Geraci and DelRosso (2017) provide a list of states that have collections of collective bargaining agreements and those that do not.

problematic in measuring the scope of collective bargaining because agreements have no set format. That is, they differ in terms of font, page size, the inclusion of related information, and so on. For that reason, a data-driven approach on the scope of collective bargaining is needed. Ingle, Willis, and Fritz (2015) use Ohio teacher collective bargaining agreements negotiated after Ohio changed its teacher evaluation system in 2012 to see how agreements responded to the new system. The authors read and categorize contract language and coded provisions. The most interesting finding for our paper is that they find, contrary to their priors, that the law did not make contracts more standardized. Ingle, Willis, and Herd (2017) analyze reduction-in-force provisions from more than 500 teacher collective bargaining agreements in Ohio. Willis and Ingle (2018) use a similar approach to look at merit pay provisions in Ohio school districts, again finding wide variations across districts. Finally, Cook, Lavertu, and Miller (2020) use Ohio teacher agreements as an input to understanding how negotiations affect the use of inputs in school districts. They find evidence that districts with additional tax revenue—but not currently in contract negotiations—hire more teachers, whereas those who receive new revenue in the middle of negotiations increase salaries and benefits.

Collective bargaining agreements are available on the SERB website as PDF files. In November and December 2020, we downloaded all then-current agreements for Ohio school districts and extracted the text from each of the documents to create the raw data for the LDA modeling.

In addition, we obtained data from the Ohio Department of Education (2021) on district staffing levels for nonclassroom teachers. By focusing on school personnel who interact with students but who are not part of the regular classroom process, we aim to show that our measure of contract scope has meaningful content. For example, because many school districts do not employ audiologists, we assert that those that do employ audiologists will have agreements with a broader scope. We obtain data on the full-time equivalent number of employees in 19 categories of school employees. All of these 19 categories contain employees who engage with students and are typically covered by teacher bargaining agreements, not agreements with other school personnel (bus drivers, custodians, etc., who are covered by separate bargaining units and contracts in nearly all cases). This includes categories like general education teachers, which all school districts have as this is the common categorization of classroom teachers. It also includes teacher aides, which many—but not all—districts have. Our primary test of our concept comes from categories such as psychologists, interpreters, library and media specialists, audiologists, physical and occupational therapists, and so on, where the modal district employs zero in the category.

4. Results

As stated, we select randomly a subset of 60 teacher contracts in Ohio. From the text of each contract, we conduct an LDA estimation. For each contract, a probability distribution over 20 topics is derived. From this, we calculate the topic concentration metric, as described previously. Figure 1 depicts the distribution of topic concentration in our sample. Specifically, figure 1 depicts the kernel density function.

Figure 1. Distribution of Topic Concentrations

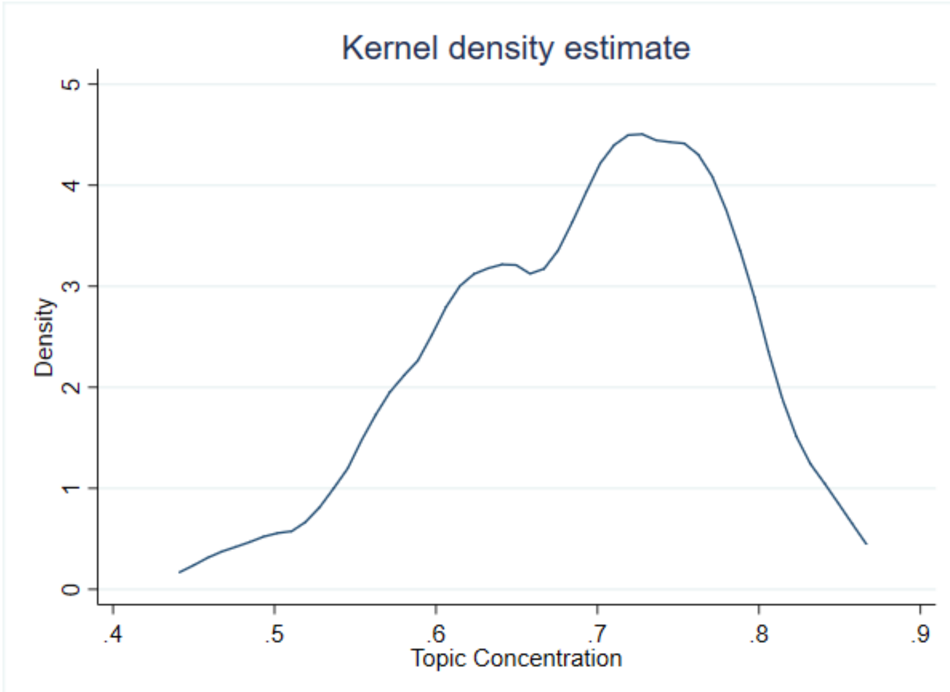

Source: Authors' calculations based on data from Ohio State Employee Relations Board (2020) and Ohio Department of Education (2021).

There is quite a bit of variation in our sample. Although the mean value is 0.691, the standard deviation 0.083. The 90th decile is approximately 36 percent larger than the 10th decile. It is this variation in contract breadth that we wish to evaluate.

Turning to information on school districts in Ohio, as a proof of concept test we consider staffing levels at each school. The dataset provides numbers of staff in 19 categories of job type. Table 1 provides descriptive information.

Table 1. Summary Statistics

Staff Type	Obs.	Mean	St. Dev.	Min	Max	% > 0
School Counselors	60	11.94	19.54	1	137.6	100%
School Nurses	60	7.22	16.25	0	107.2	86.7%
School Psychologists	60	7.42	16.02	0	80	85.0%
Interpreters	60	1.68	5.58	0	33	18.3%
Library/Media Specialists	60	6.83	13.02	0	92.4	30.3%
Audiologists	60	0.20	0.60	0	3	13.3%
Physical and Occupational Therapists	60	3.75	10.76	0	64	41.7%

Social Workers	60	2.41	6.63	0	44.1	40.0%
General Education Teachers	60	234.98	383.92	46.3	1979.1	100%
CTE Teachers	60	7.80	14.47	0	66	68.3%
Special Education Teachers	60	76.49	126.90	8	624.9	100%
Teacher Aides	60	18.85	35.34	0	220.4	76.7%
Gifted Intervention Specialists	60	3.725	7.89	0	55	78.3%
Fine Arts Teachers	60	13.78	23.95	2.8	132	100%
Music Teachers	60	13.75	21.07	2	135.5	100%
Physical Education Teachers	60	13.27	18.28	3	108.5	100%
TESOL	60	4.35	16.63	0	121	40.0%
Adaptive PE Teachers	60	0.61	2.31	0	17	21.7%
Speech Language Pathologists	60	7.82	15.55	0	88.2	95.0%

Source: Authors' calculations based on data from Ohio State Employee Relations Board (2020) and Ohio Department of Education (2021).
Note: TESOL = Teachers of English to Speakers of Other Languages.

As is to be expected, general education teachers are the most prevalent. Additionally, every school in our sample has counselors, special education teachers, and instructors in fine arts, music, and physical education. Providers of other services—namely nurses, speech pathologists, and psychologists—are nearly universal. In the analysis that follows, we focus on optional staff that not all school districts have.[4] We ask whether the existence and number of employees in each particular job type correlate with the expansiveness of the contracts.

The first step in our test is to consider the pairwise correlations between our topic concentration measurement and the number of staff for each job type. Table 2 provides these correlation coefficients and associated *p* values from a test of whether the correlation is zero.

Table 2. Pairwise Correlations

Staff Type	ρ	p value
Interpreters	**−0.246**	**0.058**
Library/Media Specialists	**−0.222**	**0.088**
Audiologists	−0.213	0.103
Physical and Occupational Therapists	−0.138	0.294
Social Workers	**−0.253**	**0.051**

[4] Hence, we drop staff types with coverage of 85 percent or more and teacher aides.

CTE Teachers	**−0.255**	**0.049**
Gifted Intervention Specialists	**−0.220**	**0.091**
TESOL	−0.198	0.130
Adaptive PE Teachers	−0.157	0.230
Speech Language Pathologists	−0.134	0.306

Source: Authors' calculations based on data from Ohio State Employee Relations Board (2020) and Ohio Department of Education (2021).
Note: Correlations significant at the 5 percent level (two-tailed test) are in bold typeface. CTE = career and technical education; PE = physical education; TESOL = Teachers of English to Speakers of Other Languages.

Five of the ten jobs have correlation coefficients that are statistically significant at the 10 percent level (bolded). Another two reached this level in one-tailed tests. Hence, across the board, the contract breadth is strongly correlated with staffing in these optional, support staff categories.

Importantly, each correlation coefficient is negative. This means that greater staffing levels are associated with a lower concentration of topics in the contract. This is what one would expect. Hence, we have a valid measurement.

A second approach to testing the validity of our measurement is to compare our topic concentration measurement of those contracts for school districts that have a non-zero number of staff with the contracts in school districts that do not have staff in each particular job type. The topic concentration metric of the latter is subtracted from the former. If the difference is zero, then the mean topic concentration in contracts of school districts with that particular job type is not different from the mean topic concentration of those without an employee within that job category. A positive difference measures how much more concentrated contracts are when no one is employed in the supporting jobs. Figure 2 depicts this difference for each of the 10 job types.

Most of the job types see a positive difference between the topic concentrations. That is, for schools without anyone in a given job, the topic concentration of the contract tends to increase. The two exceptions are physical therapists and speech pathologists. These two are statistically insignificant even at the 10 percent level, and even in one-tailed tests. Four of the eight recording a positive difference have a one-tailed difference-in-means t-test, significant at the 10 percent level. This again serves as evidence that our topic concentration measurement is effectively capturing the breadth of school teacher contracts in Ohio.

Figure 2. Contract Concentration: Zero vs. Non-Zero Staffing Levels

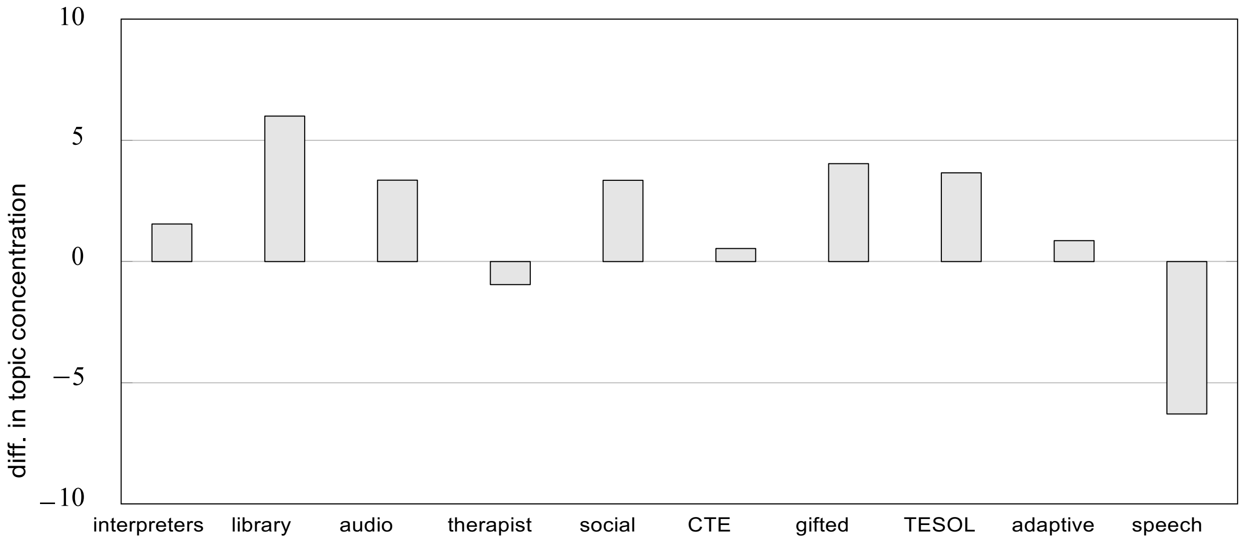

Source: Authors' calculations based on data from Ohio State Employee Relations Board (2020) and Ohio Department of Education (2021).
Note: CTE = career and technical education; TESOL = Teachers of English to Speakers of Other Languages. Each column represents the difference in the contract topic concentration between positive staff levels and observations with zero staffing (times 100). A positive difference means that the existence of staff of that job type decreases the concentration metric. The differences depicted for library staff, social workers, teachers of gifted students, and TESOL instructors have difference-in-means t-tests that are statistically significant (using a one-tailed t-test) at the 10 percent level.

Finally, the relationship between the topic concentration in the contracts and the staffing levels is not explained by overall staff size. To illustrate this, we estimate

$$Concentration_i = \beta_0 + \beta_1 Staff_i + \beta_2 Enrollment_i + {}_i, \quad (2)$$

where the dependent variable is our topic concentration measurement, *Staff* is the sum of the staffing numbers of the 10 jobs listed in table 2, and *Enrollment* is the number of students enrolled in the school. If the breadth of the contract is driven only by the size of the school, which happens to also require more staff, then $\beta_1 = 0$. If the diversity of staff has an independent effect, then $\beta_1 < 0$. In our sample, the estimated relationship is

$$Concentration_i = 0.692 - 0.0007 \times Staff_i + 0.000005 \times Enrollment_i. \quad (3)$$

The coefficient on *Staff* is statistically significant ($p = 0.083$), but the coefficient on *Enrollment* is not.

5. Conclusion

We use a computational linguistic algorithm named LDA to measure the topics covered in the text of school teacher contracts in Ohio. By using topic modeling metrics in a calculation of the concentration of topics covered, we assess how expansive a particular agreement is. To assess whether this concept is valid, we posit a negative relationship between our topic diversity measurement and the prevalence of nonclassroom teacher student support staff at the school. If more specialized services are provided, then contracts must presumably be broader as they cover more employment relationships. We confirm a strong, statistically significant relationship between our measurement and the prevalence of these support staff. We argue that this is evidence that LDA can be used to measure the breadth of collective bargaining. Future work involves extending the analysis to all contracts in the state of Ohio.

We view this as a methodological contribution and a first step into the rigorous analysis of contracts. Along with expanding these methods to the full dataset, there are a number of factors we are unable to account for in our work. For one, the length, breadth, and language used in a contract may likely be driven by the characteristics of the lawyers involved. Being able to sample numerous contracts within the same school district over time would allow us to separate the lawyer-specific influences on the text from any changes in staffing within the district. Further, we use the concept of topic here in the broad sense of the theme arising from the use of a specific set of words. By looking at a larger dataset of contracts and analyzing specific sections within the contracts, we expect to be able to home in on a clearer idea of what each topic covered is referencing.

References

Blei, David M., Andrew Y. Ng, and Michael L. Jordan. 2003. "Latent Dirichlet Allocation." *Journal of Machine Learning* 3 (January): 993–1022.

Cook, Jason, Stéphane Lavertu, and Corbin Miller. 2020. "Rent-Seeking through Collective Bargaining: Teachers Unions and Education Production." Working paper, University of Utah.

Dyer, Travis, Mark Lang, and Lorien Stice-Lawrence. 2017. "The Evolution of 10-K Textual Disclosure: Evidence from Latent Dirichlet Allocation." *Journal of Accounting and Economics* 64, no. 2-3 (November): 221–45.

Gentzkow, Matthew, and Jesse M. Shapiro. 2010. "What Drives Media Slant? Evidence from US Daily Newspapers." *Econometrica* 78, no. 1 (January): 35–71.

Geraci, Aliqae, and Jim DelRosso. 2017. "To Collect and Preserve: The State of State-Level CBA Collections in the United States." *Labor Studies Journal* 42, no. 3 (May): 165–79.

Grimmer, Justin. 2010. "A Bayesian Hierarchical Topic Model for Political Texts: Measuring Expressed Agendas in Senate Press Releases." *Political Analysis* 18, no. 1 (January): 1–35.

Hall, Joshua C., Donald J. Lacombe, and Joylynn Pruitt. 2017. "Collective Bargaining and School District Test Scores: Evidence from Ohio Bargaining Agreements." *Applied Economics Letters* 24 (1): 35–38.

Hansen, Stephen, Michael McMahon, and Andrea Prat. 2018. "Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach." *Quarterly Journal of Economics* 133, no. 2 (May): 801–70.

Herfindahl, O. C. 1950. "Concentration in the US Steel Industry." Doctoral dissertation, Columbia University.

Hirschman, Albert O. 1945. *National Power and the Structure of Foreign Trade*. Berkeley: University of California Press.

Ingle, William Kyle, Chris Willis, and James Fritz. 2015. "Collective Bargaining Agreement Provisions in the Wake of Ohio Teacher Evaluation System Legislation." *Educational Policy* 29, no. 1 (January): 18–50.

Ingle, William Kyle, Chris Willis, and Ann Herd. 2017. "Defining 'Comparable': An Analysis of Reduction in Force Provisions in Ohio School Districts." *Journal of School Leadership* 27, no. 1 (January): 68–93.

King, Gary, Patrick Lam, and Margaret E. Roberts. 2017. "Computer-Assisted Keyword and Document Set Discovery from Unstructured Text." *American Journal of Political Science* 61, no. 4 (October): 971–88.

McCannon, Bryan C. 2020. "Wine Descriptions Provide Information: A Text Analysis." *Journal of Wine Economics* 15, no. 1 (February): 71–94.

Ohio Department of Education. 2021. Ohio School Report Cards (archive). Accessed January 12, 2021. https://reportcard.education.ohio.gov/download.

Ohio State Employee Relations Board. 2020. “Collective Bargaining Agreements” (archive). Accessed December 2020. https://serb.ohio.gov/wps/portal/gov/serb/documents-and-decisions/collective-bargaining-agreements.

Quinn, Kevin M., Burt L. Monroe, Michael Colaresi, Michael H. Crespin, and Dragomir R. Radev. 2010. “How to Analyze Political Attention with Minimal Assumptions and Costs.” *American Journal of Political Science* 54, no. 1 (January): 209–28.

Schwarz, Carlo. 2018. “ldagibbs: A Command for Topic Modeling in Stata Using Latent Dirichlet Allocation.” *Stata Journal* 18, no. 1 (March): 101–17.

Tirunillai, Seshadri, and Gerard J. Tellis. 2014. “Mining Marketing Meaning from Online Chatter: Strategic Brand Analysis of Big Data Using Latent Dirichlet Allocation.” *Journal of Marketing Research* 51, no. 4 (October): 463–79.

Willis, Chris, and W. Kyle Ingle. 2018. “Profiles of Merit Pay Provisions in Ohio School Districts.” *Leadership and Policy in Schools* 17 (1): 78–114.