Online Appendix

Lexical Ambiguity in Political Rhetoric:
Why Morality Doesn't Fit in a Bag of Words

Patrick W. Kraft*† Robert Klemmensen‡

February 14, 2023

Forthcoming at the *British Journal of Political Science*

Contents

---

*Corresponding author. Email: patrickwilli.kraft@uc3m.es. Replication data for this article can be found on Harvard Dataverse at https://doi.org/10.7910/DVN/GNIJVE.

†Ramón y Cajal Fellow, Juan March Institute of Social Sciences, University Carlos III of Madrid, Spain

‡Professor, Department of Political Science, Lund University, Sweden

A Validating the Sentence Embeddings

As a first step, we want to validate our sentence embedding approach by examining whether our aggregation of pre-trained word embeddings results in a meaningful approximation of the underlying discourse vector of a given sentence. To show that this is the case, we are going to test whether we can recover the moral foundation mentioned in each sentence based on its underlying embedding vector alone—despite the fact that we exclude the moral term when computing the sentence embedding. In short, we test whether our approximation of "what is being talked about" allows us to differentiate rhetoric belonging to each moral foundation. To the extent that we can replicate the same clustering as the dictionary, we can be confident that our sentence embeddings provide a meaningful approximation of its underlying semantic context.

We proceed as follows. First, we extract all sentences that contain terms belonging to a single moral foundation.[1] Next, we randomly select 75% of the sentences and train a simple k-nearest neighbors algorithm to classify membership to each moral foundation based on our sentence embeddings using 5-fold cross-validation within the training set. Table A.1 presents the accuracy to predict moral foundations for the remaining 25% of sentences that were used as the hold-out set. Based on the embedding information alone, we can predict which moral foundation was mentioned in a sentence with an out-of-sample accuracy of 60%, which is significantly better than a random guess based on marginal distributions of foundation frequencies.

Accuracy	0.632
95% CI	(0.623, 0.641)
No Information Rate	0.326
P-Value [Acc > NIR]	< 0.001

Table A.1: Out-of-sample accuracy predicting moral foundations based on sentence embeddings using k-nearest neighbors. Model was trained on 75% of available speeches using 5-fold cross-validation.

We can further differentiate performance between moral dimensions. Figure A.1 compares the out-of-sample precision, recall, and F1-scores separately for each of the five foundations. Specifically, the figure shows that recall is lower for sentences that contain words belonging to the fairness and sanctity foundation. This pattern can be explained by the fact that fairness and sanctity are the least common moral foundations mentioned across all speeches, which makes it more difficult to predict positive instances in the test set. Of course, the overall performance in predicting each moral foundation is far from perfect, but the point of this exercise was not to train the most accurate model to predict moral foundations. Instead, the goal was to show that we are able to achieve sufficient differentiation between moral foundations based on sentence embeddings alone using a simple classification algorithm.

[1] In other words, we omit sentences that reference multiple moral foundations in order to ensure that each sentence can be uniquely matched to a single foundation.

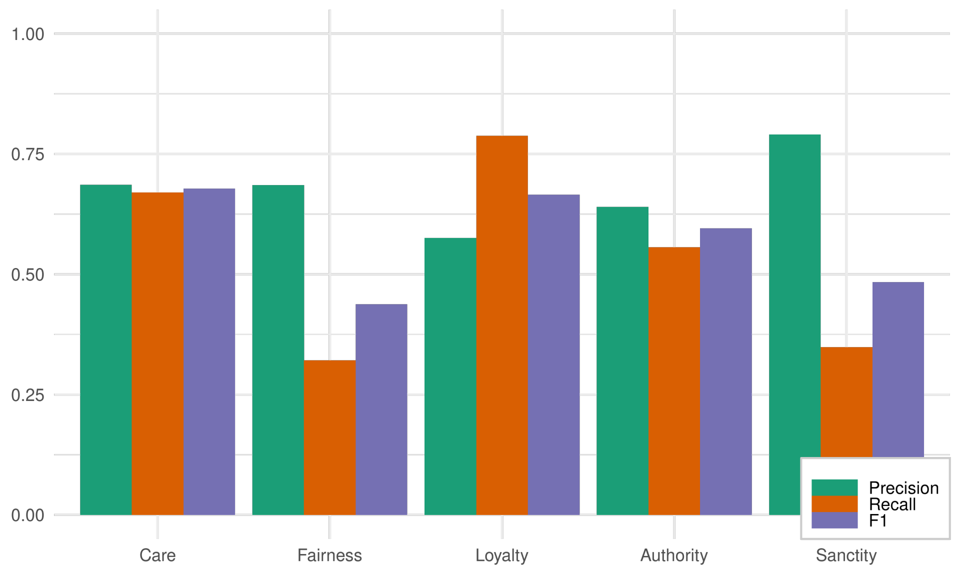

Figure A.1: Out-of-sample precison/recall/F1-score predicting each moral foundation based on sentence embeddings using k-nearest neighbors. Model was trained on 75% of available speeches using 5-fold cross-validation.

B Moral Language Over Time

Figure B.1 displays the average mention of moral dimensions over time. The patterns reveal no systematic temporal variation, for example due to gradual policy shifts or distinct historic events.

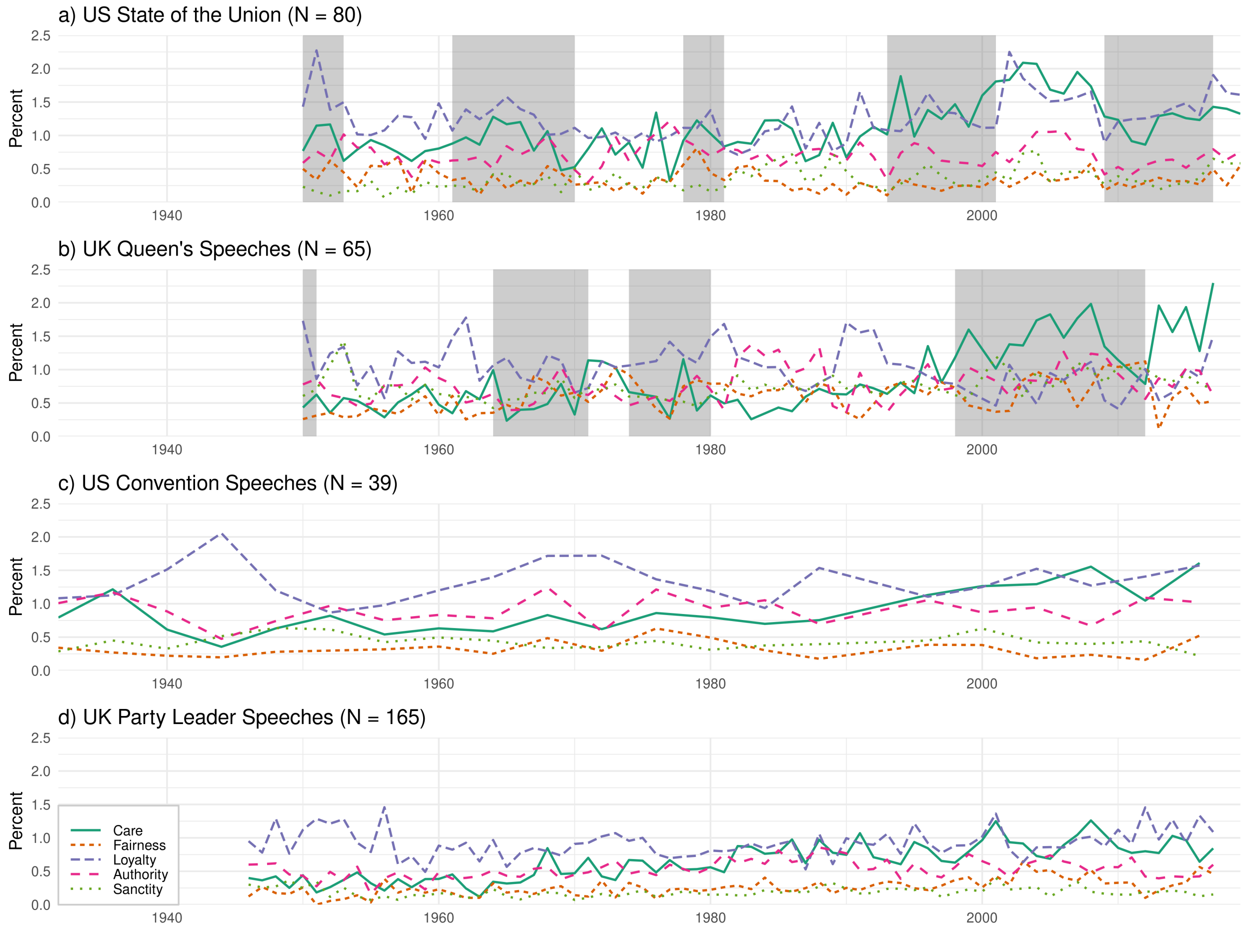

Figure B.1: Moral language over time. Shaded areas indicate Democratic/Labour government.

C Ideological Overlap in Dictionary Terms

Figure C.1 displays the most common dictionary terms across all moral foundations, along with the average number of mentions per speech. The similarity in relative term frequencies between speeches given by liberal and conservative politicians suggests that both sides not only emphasize the same moral foundations but furthermore use the same terms within each foundation.

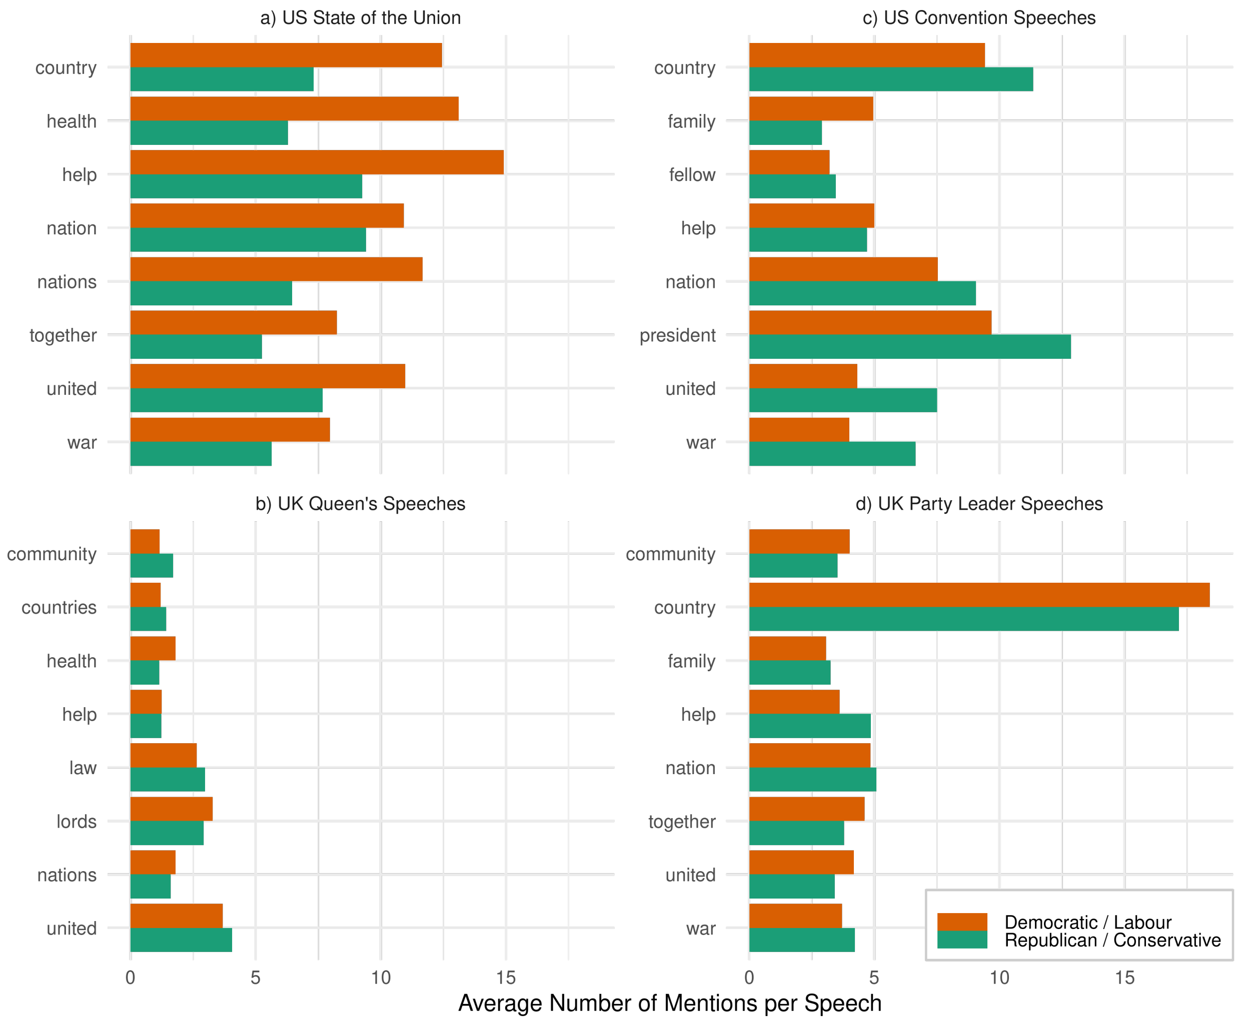

Figure C.1: Common shared MFT dictionary terms across speeches

D Replication Using conText Package

Figure D.1 displays the results of a direct replication of our analysis of semantic contexts using the embedding regression approach proposed by Rodriguez, Spirling, and Stewart (2023). Instead of computing individual sentence embeddings for each occurrence of a term included in the dictionary (as in our approach), they compute average embeddings of a target term itself across each set of documents. We implement their embedding regression framework for each moral foundation and compare the resulting embedding vectors between liberal conservative politicians. In contrast to our findings discussed in the main text, the replication using the method proposed in Rodriguez, Spirling, and Stewart (2023) suggests that *all* moral foundations across *all* types of documents are significantly different between liberal and conservative politicians. Crucially, however, their approach fails to account for the within-party variation in semantic context. Thus, Rodriguez, Spirling, and Stewart (2023) overestimate the extent to which context embeddings differ between liberal and conservative politicians relative to the varying semantic context within each group.

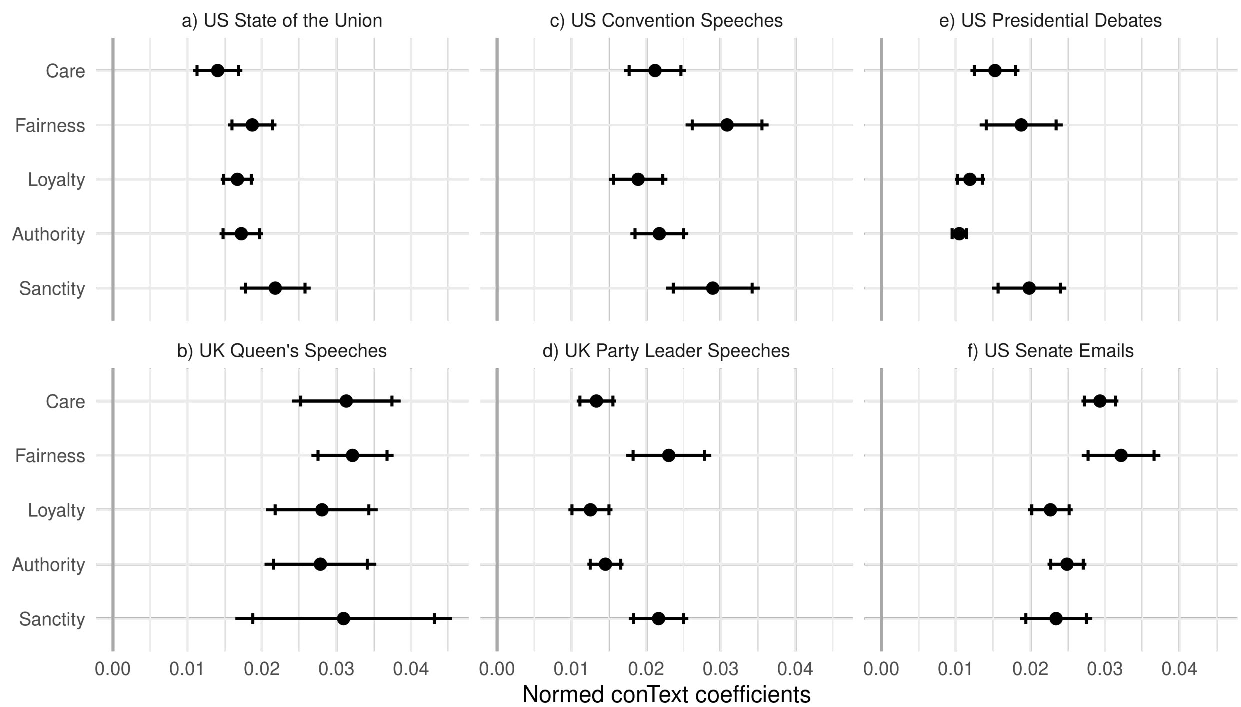

Figure D.1: Replication of main analysis using conText package. Positive coefficients indicate that embeddings of a given foundation are significantly different between parties. 95% and 90% confidence intervals based on a non-parametric bootstrap.

References

Rodriguez, Pedro L, Arthur Spirling, and Brandon M Stewart. 2023. "Embedding Regression: Models for Context-Specific Description and Inference." *American Political Science Review*: 1–20.