

A Moral Foundations Dictionary for the European Portuguese Language: The Case of Portuguese Parliamentary Debates

Mafalda Zúquete¹[0000-0003-4915-8784], Diana Orghian², and Flávio L. Pinheiro¹[0000-0002-0561-9641]

¹ NOVA IMS - Universidade Nova de Lisboa, Lisboa, Portugal

² FeedZai, Lisboa, Portugal

Abstract. Moral Foundations Theory (MFT) has shown that American liberals and conservatives rely on fundamentally different moral principles, offering a different perspective on the deepening political divide in US politics. However, results outside the US have been less clear, particularly in countries with a more diverse political landscape that does not fall into the traditional Liberal/Conservative dichotomy. Here, we expand the Moral Foundations Dictionary to European Portuguese, which we then use to analyze 10 years of transcripts of parliamentary sessions using standard Data Science and Text Mining techniques. Despite a larger number of represented parties, we show that no traditional parties fall into the Conservative or Liberal characterization and that the political landscape in Portugal is relatively homogeneous with the major difference observed concerning the dichotomy between Government and the parliament.

Keywords: Moral Foundations Theory · Moral Foundations Dictionary · Political Discourse · Text Mining.

1 Introduction

How can opposite beliefs co-exist, and which side is “morally right”? When trying to understand the co-existence of disparate moralities, Haidt and Joseph [18] proposed a structured way to study morality based on our intuitions and how they adapt to the culture we inhabit. This would give origin to what is now called *Moral Foundations Theory* (MFT) that assumes the existence of innate pillars which capture society’s moral variety [12].

According to MFT, human morality stands atop five different foundational dimensions: **Harm** (Care/Harm, it makes us sensitive to signs of suffering and need; it makes us despise cruelty and want to care for those who are suffering); **Fairness** (Fairness/Cheating, it makes us sensitive to indications that another person is likely to be a good (or bad) partner for collaboration and reciprocal altruism. It makes us want to shun or punish cheaters); **Ingroup** (Loyalty/Betrayal, it makes us sensitive to signs that another person is (or is not) a

team player. It makes us trust and reward such people, and it makes us want to hurt, ostracize, or even kill those who betray our group or us); **Authority** (Authority/Subversion, it makes us sensitive to signs of rank or status, and to signs that other people are (or are not) behaving properly, given their position); and **Purity** (Sanctity/Degradation, it includes the behavioral immune system, which can make us wary of a diverse array of symbolic objects and threats. It makes it possible for people to invest in objects with irrational and extreme positive and negative values, which are essential for binding groups together). Moreover, Graham et al. [12] considered the possibility of more than five moral foundations, namely Liberty/Oppression, but claimed that there was not enough evidence to support their foundationhood at this stage.

Graham et al. [13] developed a Moral Foundations Dictionary (MFD) to measure the moral load in political texts. By using the MFD in tandem with *Linguistic Inquiry and Word Count* (LIWC), Graham et al. [13] showed that liberals and conservatives relied on different sets of moral foundations: liberals tended to value the Harm and Fairness foundations primarily; where conservatives' moral discourse was more evenly distributed across foundations. When it comes to the creation of an MFD in languages other than (American) English, it is relevant to highlight the works of Matsuo et al. [26] (Japanese), Carvalho et al. [6] (Brazilian Portuguese), Alper et al. [1] (Turkish), Carvalho and Guedes [5] (Spanish), and Wan et al. [33] (Chinese). Here, we propose the creation of an MFD in European Portuguese, its validation, and an example analysis of the moral loading in Portuguese politics by studying the transcripts of the Diary of the Assembly of the Republic from 2011 to 2021. From this analysis, we noticed how the morality profile of the parties was quite similar among them, while being relatively different from the morality embedded in the Government discourse. In that sense, we also glimpsed a dimension in parliamentary dynamics that we found relevant: the parties morality profile reflected their proximity to governmental power.

2 Related Work

Across the literature, the relationship between moral foundations and American political ideology shows a consistent pattern: liberals primarily value the Harm and Fairness foundations, and conservatives' moral beliefs are based on all moral foundations [13, 16, 27].

We highlight the work of Graham et al. [13], which used methods from text mining to the study of morality under the paradigm of MFT, by creating the MFD to detect moral words in church sermons. Despite the replicability of these results when the study is precisely reproduced, extensions to include sermons from other churches, political transcripts, and texts from media outlets, meant that the conclusions did not hold [9]. This outcome highlights the shortcomings of this dictionary-based methodology [2]:

1. A limited amount of lemmas and stems of words;
2. "Radical" lemmas rarely used in everyday language;

3. An association with a moral binary scale, but lacking of a measure of “strength”.

New, more complex, algorithms have been developed to detect the underlying morality in the text, such as Garten et al.’s [10] *Distributed Dictionary Representations* (DDR) of the MFD, or Araque et al.’s [2]’s *MoralStrength*, which was obtained as an extension of the MFD, based on WordNet synsets. The creation of the *Moral Foundation Twitter Corpus* (MFTC), a corpus of tweets classified with regards to morality [21], is likely to promote the development of more MFD-related algorithms, following previous studies of morality in tweets [8, 24].

Indeed, the framework of MFT has been used to explore political chasms. The work of Koleva et al. [25] noted that the strongest unique predictor of a “culture war” opinion was often a sub-scale of the Moral Foundations Questionnaire (MFQ) (often Purity) instead of political ideology, interest in politics, religious attendance, or any demographic variable.

Despite the evidence for a binary view of American politics, a cluster analysis of answers to the MFQ resulted in four main clusters – secular liberals, libertarians, religious leftists, and social conservatives – which bound themselves to the two main political parties [17]. Furthermore, a deeper look at the differences between two groups of liberals – Clinton and Obama supporters during the 2008 Democratic primaries – showed how, despite both groups identifying themselves as equally liberal, the moral foundations that they valued were not the same: Clinton supporters showing more substantial support of the binding moral foundations, closer to the typical conservative voter; Obama supporters defaulting to the individualizing foundations [22]. Libertarians proved to be an interesting outlier to the bipartisan dynamic of American politics, as they appeared to reject all of the five moral foundations [17]. This led to the proposal of a sixth moral foundation, Liberty/Oppression, which would help characterize the moral profile of libertarians [23]. Still, Graham et al. [12] discarded this foundation, claiming that the evidence for foundationhood was still lacking.

While contemporary American politics shows a high appeal when studying how politics intersects with morality, there are works outside the American context. Parker et al. [30] focused on Australian politics, looking at how time and moral foundations were related, showing a more complex landscape to Australian politics than a myopic view of liberal/conservative would suggest. Moreover, outside the political sphere, relevant insights can be achieved by studying societal dynamics from the lenses of MFT. Examples include, but are not limited to, topics such as sports fandom [35], religion [11], education [14], health habits [32], social cognition [19], or media [31].

3 European Portuguese Moral Foundations Dictionary

3.1 Semi-supervised Translation of the English MFD

The English MFD created by Graham et al. [13] consisted of a set of full words and word stems related to each of the five moral foundations. The dictionary included both foundation-supporting words, or *Virtues*, and foundation-violating

words, *Vices*. Here, we adapted the methodology proposed by Matsuo et al. [26] to develop a European Portuguese Moral Foundations Dictionary. The proposed approach involves six steps, some automated and some manual (Figure 1).

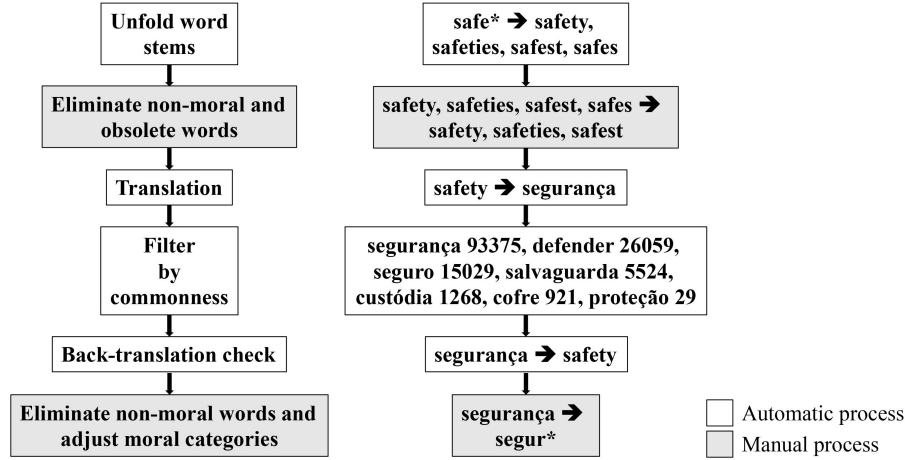


Fig. 1. Illustrative workflow of the translations process (left) using the stem “safe*” as an example (right).

Starting from the MFD developed by Graham [13], all the word stems were unfolded with the aid of *OneLook* [7] to find all words that started with each stem and then, using the *Common Words* filter, we selected the most common words associated with each stem.

The *Merriam-Webster* [29] and the *wikionary* [34] dictionaries were used to find and filter out words flagged as *obsolete* or *archaic*. We then used the *Cambridge University’s English (US) - Portuguese* [3] to translate each word to Portuguese. Translations that were expressions rather than words were discarded. In the case of words that translated to reflexive verbs, the verb in the infinitive was manually added to the list of translations. Since the MFD was meant to represent current speech, we then filtered out uncommon words using the *Corpus de Referência do Português Contemporâneo* (CRPC) [28]. For words that were radicals on the original dictionary, the 10 most common translations were kept, and the 5 most common translations were kept in the case of plain words. To ensure that translations kept the moral load of the original words, a backtranslation check was performed using the *Cambridge University’s Portuguese - English (US)* [4] dictionary. All the words for which the backtranslation differed from the original were removed. Finally, the words were checked to ensure that they were all moral, and the moral categories were adjusted to reflect the Portuguese language’s current use better.

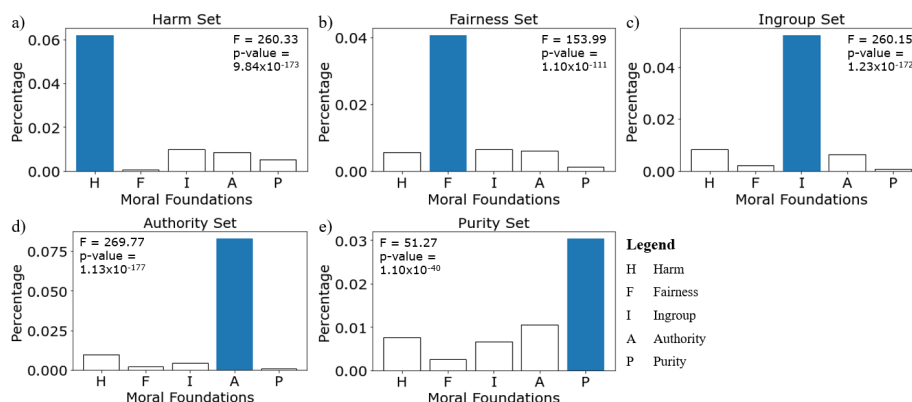


Fig. 2. Mean percentage of dictionary words in examples for each Moral Foundation after corrections to the MFD.

The outcome of the above steps is a list of words, unlike the original MFD, which was a list of words and word stems. As such, words that started the same way and were associated with the same moral foundations were manually gathered into lemmas, checking the CRPC to ensure that using the lemma on a word would not yield unwanted results. Also, because of the Portuguese Language Orthographic Agreement of 1990, double orthography was considered for some words and word stems. The European Portuguese MFD consists of a list of words and word stems, each associated with one or more moral foundations and the corresponding Vice/Virtue loading (except for the words associated with Morality General).

3.2 Experimental Validation

To test the validity of the resulting dictionary, we ran a survey-based study on Prolific that counted with the participation of 324 individuals who provided valid responses. The pool of participants were all native Portuguese speakers, including 36% females and 64% males, who were, on average, 24 years old.

During the survey, each participant was initially briefed on the definition of each moral foundation (which consisted of the translation of the definitions presented by Haidt [15]) and was asked to describe scenarios where the foundations were validated or violated. Participants who submitted short answers (under 20 words) or failed to answer for any moral foundations were considered to have invalid answers and, thus, were ignored in the subsequent analysis.

From the survey answers, we then determined the validity of the dictionary by conducting a one-way analysis of variance (ANOVA) on the the main effect of the moral foundations. That is, we checked if there was a significantly higher frequency of words from the moral foundation in question than the words from the remaining moral foundations.

Table 1. The five words with the strongest association to each of the Moral Foundations, discriminated by Virtue and Vice.

Moral Foundation	Words
Harm Virtue	abrigo, compaixão, defesa, empatia, segurança
Harm Vice	agressor, bruto, cruel, dano, indefeso
Fairness Virtue	balanço, equilíbrio, equivalente, razoável, reciprocidade
Fairness Vice	discriminação, desigual, desonesto, intolerância, preconceito
Ingroup Virtue	coletivamente, colónia, cooperativa, nacional, patriotismo
Ingroup Vice	abandonar, discordância, imigração, inimigo, terrorismo
Authority Virtue	cargo, comando, lei, obediência, ordem
Authority Vice	desobediência, infrator, insubordinação, rebelde, transgressão
Purity Virtue	igreja, puro, sagrado, santo, virgem
Purity Vice	mancha, obsceno, pecado, profano, sujo

Noteworthy to mention that after a pre-test of the survey materials we found an unexpected overabundance of Authority-related words in the Ingroup foundation. Under closer inspection, we found some ambiguity in the moral categorization of words starting with “traí*” and “lea*”. As such, we revised the MFD so that these words would only be Ingroup, rather than being both Authority and Ingroup. We also changed the category of “deslea*” in the same way, as this word was the direct antonym of “lea*” and it would not make sense to have them be classified differently. Furthermore, we noticed how “cumpr*” and “viol*” were consistently misclassified across all Moral Foundations, which was related to the presence of such words in the survey questions that primed participants to use them by default. As such, we decided to exclude those words from the MFD altogether.

Results show that, for most moral foundations, the most commonly identified MFD words were associated with the correct foundation, see Figure 2. Moreover, the MFD also passed the ANOVA, which supports the main effect of moral foundations (see Figure 2), and that these words effectively discriminate between different dimensions.

We can also measure which words show a stronger association with each moral dimension, based on which had the lowest entropy (seed words). Table 1 shows for each moral dimension the five words with the strongest association.

The European Portuguese Moral Foundations Dictionary is available for download [36].

4 Morality in the Portuguese Assembly of the Republic, a case study

4.1 Parliamentary Transcripts

As an example a possible application of the European Portuguese MFD, we studied the moral load embedded in the political discourse of the Portuguese




Fig. 3. A brief timeline of the parties in power in Portuguese politics.

Assembly of the Republic. In particular, how do the parties currently in parliament relate to each other with regards to morality?

To that end, we used the publicly available transcripts [20] of parliamentary sessions and debates between 20/06/2011 and 22/07/2021 from the *Diário da Assembleia da República* (DAR). This period covers the harsh austerity mandated by *Troika*, the economic growth that followed, the first *superavit* in 40 years, the slow but steady growth and normalization of the far right, and the COVID-19 pandemic. Figure 3 illustrates which parties were in power during the time period under study. This includes both presidential and governmental power.

Since the transcripts were in raw text format, some pre-processing steps were necessary to obtain a suitable structured working dataset. These included eliminating blank lines, replacing em-dashes with hyphens, removing page headers, adding name and party information to government members, identifying unidentified speakers, and checking for spelling errors.

Moreover, members of the government were often only identified by their function. Given that different people could occupy the same position over time, name and party information were manually added in those cases. If a line did not have an identified speaker, it was determined it was the same speaker as in the previous line, unless the line was referring to: applause, protests, laugh, pauses, votes, or changes in the presidency of the assembly. This way, lines that were not interventions by members of parliament were not wrongly attributed to a speaker. Finally, transcripts were manually checked for orthographic errors with the help of *Microsoft Word*'s spell check, as the text files contained some errors.

The edited transcripts were then processed line by line. Each line was classified as one of 6 categories if it did not have a specified speaker:

- **Summary:** the debate's summary;
- **Note:** notes at the end of the transcripts;
- **Proposals:** the transcripts featured legislative proposals that could be attributed to political parties;

- **Vote declaration:** members of parliament could submit a justification to their vote and these were included after the debate;
- **Ratification:** if a specific party wanted to correct a mistake that appeared in a previous diary;
- **Voices:** interventions during the debate that were not attributed to a single member of parliament. Usually interruptions and interjections;

If the line had a specified speaker, we saved their name and function (member of parliament (MP) or their role within the government) and their intervention. The final dataset included all parliamentary sessions from 20/06/2011 to 22/07/2021, spanning three legislatures, four governments, and ten years' worth of political interventions on the parliament floor.

4.2 Measuring Moral Load

To study the underlying morality in parliamentary speech we looked at the deviations in the use of moral words by MPs, compared to the expected number of moral words given the number of words spoken. This approach deviated from past works [13, 30]. Frimer [9] argued that American liberals and conservatives were not so different, as his replication and expansion of previous experiments [13] found no evidence of differences between the moral languages of liberals and conservatives. This informed our choice of how to encode the moral speech of Portuguese political parties, leading us to focus on how political speech would differ from expected.

In order to encode the moral load in the transcripts, we took the following steps:

1. Group the speech of all MPs for each legislative session and sum the total number of words and the number of words in each moral foundation;
2. Create log-log regressions for all moral foundations, where the x-axis was the number of words in a moral foundation and the y-axis was the total number of words used;
3. Calculate the residuals for each moral foundation, for each MP, for each legislative session, using the previously defined regressions;
4. Group the residuals by party, by year, and calculate the average residuals for each moral foundation, to get the morality profile of each party during each legislative session.

Only MPs who spoke more than 100 words in the entire legislative session were considered in step 2, and that threshold was raised to 1000 words when the party data was grouped in step 4. This was due to the fact that there were many MPs who had minimal interventions throughout the legislative session, and thus their contributions were not considered significant with regards to the morality of the party.

These regressions (see Figure 4) had R^2 scores that could be considered quite high, lending confidence to the results from these regressions.

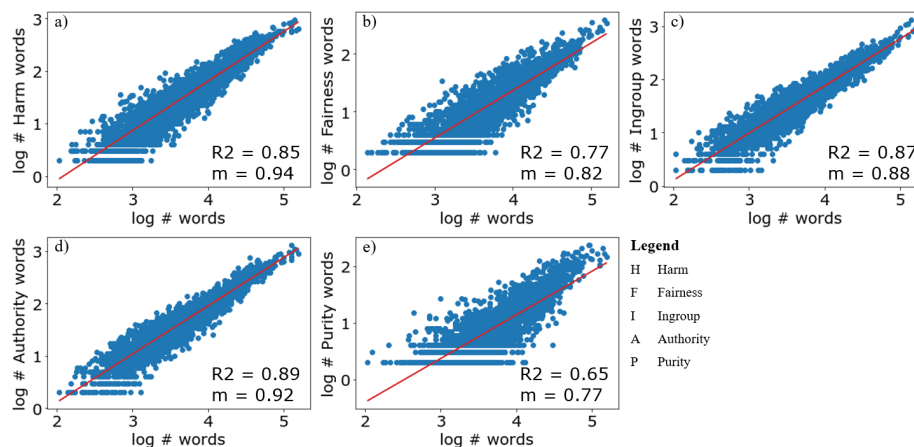


Fig. 4. log-log regressions of the total number of words by the total number of foundation words.

4.3 Political Analysis

Figure 5a shows the moral load of each foundation per party in the 2020/21 legislative session. While we were comparing parties, the sample sizes for each are quite varied. Furthermore, some of the parties in this analysis only had one MP and others had plenty. As such, the samples of some parties, like Chega! (CH) or Iniciativa Liberal (IL), do not fully represent that party but only one, or very few MPs.

We could see that most political forces had relatively similar moral profiles, with the government and Partido Ecologista “Os Verdes” (PEV) demonstrating some of the most extreme values. It appears that the parties do not differ much with regards to Harm, Ingroup, or Authority, and show the most disparate values in the Fairness and Purity foundations. More importantly we see that the government shows a substantially different profile from the remaining parties in the Parliament.

To better understand the dynamic between parties we ran a Factor Analysis (FA) on the party profiles for this year. Since the Bartlett sphericity test had a p-value of 0.00017 (under 0.05) and the Kaiser-Meyer-Olkin score was 0.76 (over 0.60) we could consider we could perform an FA on our dataset.

We chose to run our analysis with two factors as determined by the Kaiser criterion. The eigenvalues of the first two factors were 3.72 and 0.91. Though a strict application of the Kaiser criterion would dictate that only one factor was chosen, as the eigenvector of the second factor is lower than 1, since 0.91 is very close to 1 we decided to keep two factors. To confirm our choice of two factors, we analyzed the cumulative variance explained by one factor versus two. The cumulative variance went from 0.72 to 0.86, which we considered relevant enough to keep the two factors in our analysis. These values, as well as the factor loadings, are in Table 2.

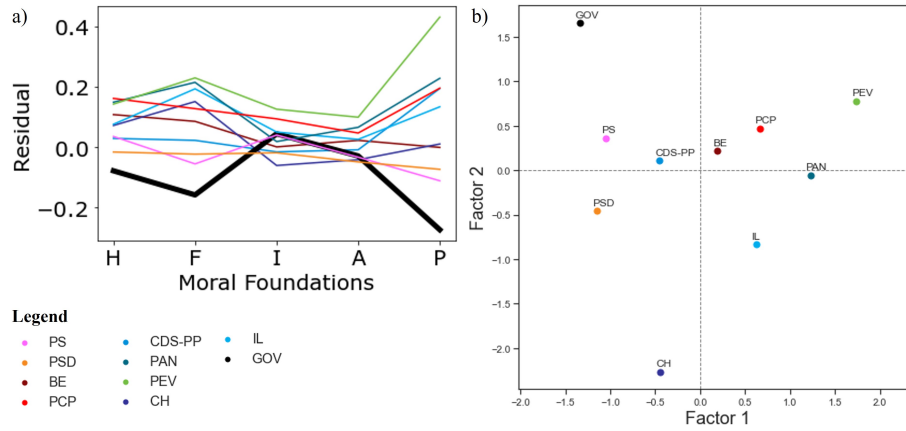


Fig. 5. Morality profile of every party currently in parliament (a) and a visualization of its factor analysis (b).

Though these factors are not easy to interpret, they allow us to better understand the dynamics of the parties. Parties in the first quadrant of Figure 5b exhibit high values for all moral foundations (PEV and Partido Comunista Português (PCP)), whereas parties in the second quadrant have lower values for most values but comparatively high Ingroup values (government and Partido Socialista (PS)).

Though these results do not reflect clearly the right/left-wing dynamic as its commonly understood, the visualization of the factors shows that the parties closer to the government are the ones which have been in power up to this point (PS, Partido Social-Democrata (PSD), and CDS-Partido Popular (CDS-PP)), and the ones farthest away are more recent parties, which elected MPs for the first time in 2019 (CH and IL).

The above conclusion is further reinforced when we look at the evolution of parliamentary dynamics over time. We excluded the most recent parties, IL and CH, from our analysis, since they have only been in Parliament for two legislative sessions. As such, to study the evolution of parliamentary speech we plotted the

Table 2. Factor loadings and corresponding variance.

	Factor 1	Factor 2
Harm	0.88	-0.18
Fairness	0.90	-0.43
Ingroup	0.55	0.61
Authority	0.96	0.28
Purity	0.89	-0.08
Variance	0.72	0.14
Cumulative variance	0.72	0.86




Fig. 6. Evolution of the distance between parties and the government.

euclidean distance between each party, in each year, and the government, over time (Figure 6).

In Figure 6 the dynamic we saw before is quite clear, with PS and PSD being quite close to the government throughout the entire period we considered. However, after the 2015 elections, when the PSD/CDS-PP coalition lost its power and PS became the governing party (see Figure 3), CDS-PP started drifting away from the government.

In opposition to the behavior of CDS-PP, we remark on how the moral profile of Bloco de Esquerda (BE) and PCP have evolved over time. These parties started getting closer to the government following the 2015 election, when their parliamentary support was instrumental for the PS minority government to be able to stand, but after the 2019 election, as PS won reelection but no longer had a formal parliamentary agreement with these parties, they started distancing themselves again.

Farthest away from the government we find the two environmental parties, PEV and Pessoas-Animais-Natureza (PAN).

5 Conclusions

We presented the steps employed in the development of a Moral Foundations Dictionary for European Portuguese. These dictionaries are key to support text-based analysis based on LIWC approaches.

As an example, we then studied the case of Political Discourse in the Portuguese Parliament transcripts between 2011 and 2021. Here, not only did the parties not fit into the Liberal and Conservative archetypes determined previously, but there was not a clear distinction between right and left wing parties

either. Instead, after running an FA we found that there is a dynamic with regards to proximity to the government, as the parties which have been traditionally in power are closer to the government, and less mainstream parties are farther away. This insight is consolidated through a temporal analysis of the distance between parties and the government over the years. Indeed, the parties which have shared governmental power since the establishment of the Third Republic have a moral profile which is very similar to that of the government.

Our work is not without limitations. For instance, we did not consider the different sample sizes between parties. This created limitations with regard to the representativeness of some samples, as not all parties and MPs speak with the same frequency or for the same length of time. Moreover, the methodology used here is relatively simple and lacks the ability to extract information from the context in which moral words are present.

Still, some of the most advanced algorithms used in the research of MFT, which are currently not available in European Portuguese, were built on top of the original MFD. This means that better and more accurate algorithms to detect moral foundations in the text could be created by expanding on the work we presented. Furthermore, the proposed MFD was a translation of an MFD created specifically for the American context. As such, this MFD likely failed to capture some nuances and peculiarities of the Portuguese language and Portugal’s moral context, which could be better grasped by an MFD created from scratch, by linguists or specialists in MFT. Nonetheless, our work provides a baseline upon which future work could be developed and improvements could be made.

While this project looked at 10 years’ worth of parliamentary data, our work could still be expanded by taking an even broader look to the Portuguese parliament, or even expanding the type of text we are taking into consideration. Close analysis of morality in speech during elections or referendums would undoubtedly create interesting insights. Another exciting avenue of research would be to consider texts extracted from the Internet, namely tweets. Studying this type of text would mean the expansion of the MFD via supervised learning, as well as a corpus of annotated tweets in Portuguese to study the accuracy of these methodologies.

Acknowledgements MZ and FLP acknowledge the financial support provided by FCT Portugal under the project UIDB/04152/2020 – Centro de Investigação em Gestão de Informação (MagIC). FLP acknowledges the financial support of the Portuguese Foundation for Science and Technology (“Fundação para a Ciência e a Tecnologia”) through grant DSAIPA/DS/0116/2019. The authors are thankful to César Hidalgo and Cristian Candia for the useful discussions and insightful comments.

References

1. Alper, S., Bayrak, F., Us, E.Ö., Yilmaz, O.: Do changes in threat salience predict the moral content of sermons? the case of friday khutbas in turkey. *European*

- Journal of Social Psychology **50**(3), 662–672 (2020)
2. Araque, O., Gatti, L., Kalimeri, K.: Moralstrength: Exploiting a moral lexicon and embedding similarity for moral foundations prediction. *Knowledge-based systems* **191**, 105184 (2020)
 3. Cambridge University Press: Cambridge Dictionary: Translations from English to Portuguese, <https://dictionary.cambridge.org/us/dictionary/english-portuguese/>
 4. Cambridge University Press: Cambridge Portuguese-English Dictionary: Translations from Portuguese to English, <https://dictionary.cambridge.org/us/dictionary/portuguese-english/>
 5. Carvalho, F., Guedes, G.: Dicionário de fundamentos morais em espanhol
 6. Carvalho, F., Okuno, H.Y., Baroni, L., Guedes, G.: A brazilian portuguese moral foundations dictionary for fake news classification. In: 2020 39th International Conference of the Chilean Computer Science Society (SCCC). pp. 1–5. IEEE (2020)
 7. Datamuse: OneLook Dictionary Search, <https://www.onelook.com/>
 8. Dehghani, M., Johnson, K., Hoover, J., Sagi, E., Garten, J., Parmar, N.J., Vaisey, S., Iliev, R., Graham, J.: Purity homophily in social networks. *Journal of Experimental Psychology: General* **145**(3), 366 (2016)
 9. Frimer, J.A.: Do liberals and conservatives use different moral languages? two replications and six extensions of graham, haidt, and nosek’s (2009) moral text analysis. *Journal of Research in Personality* **84**, 103906 (2020)
 10. Garten, J., Hoover, J., Johnson, K.M., Boghrati, R., Iskiwitch, C., Dehghani, M.: Dictionaries and distributions: Combining expert knowledge and large scale textual data content analysis. *Behavior research methods* **50**(1), 344–361 (2018)
 11. Graham, J., Haidt, J.: Beyond beliefs: Religions bind individuals into moral communities. *Personality and social psychology review* **14**(1), 140–150 (2010)
 12. Graham, J., Haidt, J., Koleva, S., Motyl, M., Iyer, R., Wojcik, S.P., Ditto, P.H.: Moral foundations theory: The pragmatic validity of moral pluralism. In: *Advances in experimental social psychology*, vol. 47, pp. 55–130. Elsevier (2013)
 13. Graham, J., Haidt, J., Nosek, B.A.: Liberals and conservatives rely on different sets of moral foundations. *Journal of personality and social psychology* **96**(5), 1029 (2009)
 14. Graham, J., Haidt, J., Rimm-Kaufman, S.E.: Ideology and intuition in moral education. *International Journal of Developmental Science* **2**(3), 269–286 (2008)
 15. Haidt, J.: *The righteous mind: Why good people are divided by politics and religion*. Vintage (2012)
 16. Haidt, J., Graham, J.: When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social Justice Research* **20**(1), 98–116 (2007)
 17. Haidt, J., Graham, J., Joseph, C.: Above and below left–right: Ideological narratives and moral foundations. *Psychological Inquiry* **20**(2-3), 110–119 (2009)
 18. Haidt, J., Joseph, C.: Intuitive ethics: How innately prepared intuitions generate culturally variable virtues. *Daedalus* **133**(4), 55–66 (2004)
 19. Hirsh, J.B., DeYoung, C.G., Xu, X., Peterson, J.B.: Compassionate liberals and polite conservatives: Associations of agreeableness with political ideology and moral values. *Personality and Social Psychology Bulletin* **36**(5), 655–664 (2010)
 20. Histórico-Parlamentar, A.: Debates Parlamentares, <https://debates.parlamento.pt>
 21. Hoover, J., Portillo-Wightman, G., Yeh, L., Havaladar, S., Davani, A.M., Lin, Y., Kennedy, B., Atari, M., Kamel, Z., Mendlen, M., et al.: Moral foundations twitter corpus: A collection of 35k tweets annotated for moral sentiment. *Social Psychological and Personality Science* **11**(8), 1057–1071 (2020)

22. Iyer, R., Graham, J., Koleva, S., Ditto, P., Haidt, J.: Beyond identity politics: Moral psychology and the 2008 democratic primary. *Analyses of Social Issues and Public Policy* **10**(1), 293–306 (2010)
23. Iyer, R., Koleva, S., Graham, J., Ditto, P., Haidt, J.: Understanding libertarian morality: The psychological dispositions of self-identified libertarians (2012)
24. Kaur, R., Sasahara, K.: Quantifying moral foundations from various topics on twitter conversations. In: 2016 IEEE International Conference on Big Data (Big Data). pp. 2505–2512. IEEE (2016)
25. Koleva, S.P., Graham, J., Iyer, R., Ditto, P.H., Haidt, J.: Tracing the threads: How five moral concerns (especially purity) help explain culture war attitudes. *Journal of research in personality* **46**(2), 184–194 (2012)
26. Matsuo, A., Sasahara, K., Taguchi, Y., Karasawa, M.: Development and validation of the japanese moral foundations dictionary. *PloS one* **14**(3), e0213343 (2019)
27. McAdams, D.P., Albaugh, M., Farber, E., Daniels, J., Logan, R.L., Olson, B.: Family metaphors and moral intuitions: How conservatives and liberals narrate their lives. *Journal of personality and social psychology* **95**(4), 978 (2008)
28. Mendes, A., Bacelar do Nascimento, F.: Corpus de Referência do Português Contemporâneo, <http://alfclul.clul.ul.pt/CQPweb/>
29. Merriam-Webster: Dictionary by Merriam-Webster: America’s most-trusted online dictionary, <https://www.merriam-webster.com/>
30. Parker, P., Sahdra, B.K., Ondaatje, M.: Individual differences are a better predictor of moral language use in australian prime ministers’ speeches than is their political party (2019)
31. Tamborini, R.: Moral intuition and media entertainment. *Journal of Media Psychology* (2011)
32. Van Leeuwen, F., Park, J.H., Koenig, B.L., Graham, J.: Regional variation in pathogen prevalence predicts endorsement of group-focused moral concerns. *Evolution and Human Behavior* **33**(5), 429–437 (2012)
33. Wan, Y., Peng, Z., Wu, F., Li, M., Gao, J.: Understanding public ethical acceptance and its antecedents and moderators of ai surveillance technology: Analysis from social media data. Available at SSRN 4246846
34. Wikimedia Foundation: Wiktionary, <https://www.wiktionary.org/>
35. Winegard, B., Deaner, R.O.: The evolutionary significance of red sox nation: sport fandom as a by-product of coalitional psychology. *Evolutionary psychology: an international journal of evolutionary approaches to psychology and behavior* **8**(3), 432–446 (2010)
36. Zúquete, M., Pinheiro, F.L., Orghian, D.: Moral Foundations Dictionary in European Portuguese, <https://github.com/mafalda-zuquete/mfd-pt>