

Expressions of Anxiety in Political Texts

Ludovic Rheault

Department of Political Science, University of Toronto

Abstract

Anxiety has a special importance in politics since the emotion is tied to decision-making under uncertainty, a feature of democratic institutions. Yet, measuring specific emotions like anxiety in political settings remains a challenging task. The present study tackles this problem by making use of natural language processing (NLP) tools to detect anxiety in a corpus of digitized parliamentary debates from Canada. I rely upon a vector space model to rank parliamentary speeches based on the semantic similarity of their words and syntax with a set of common expressions of anxiety. After assessing the performance of this approach with annotated corpora, I use it to test an implementation of state-trait anxiety theory. The findings support the hypothesis that political issues with a lower degree of familiarity, such as foreign affairs and immigration, are more anxiogenic than average, a conclusion that appears robust to estimators accounting for unobserved individual traits.

1 Introduction

Of the variety of emotions experienced by humans, anxiety ranks arguably among the most relevant for politics. Put simply, anxiety means feeling worried or concerned in the face of uncertainty (LaBar, 2016). We inherited this emotion from evolutionary processes with the critical function of keeping us alert to external threats (Suomi and Harlow, 1976; LeDoux, 2012; LaBar, 2016). Scholars have tied the emotion not only to uncertainty, but also to one's lack of control over events (Mandler, 1972; Mineka

and Kelly, 1989). Of all social contexts, democratic institutions appear to combine the full set of ingredients required to rouse anxiety. Indeed, parliaments gather individuals with the specific purpose of making decisions under uncertainty, sometimes in situations of emergency. The rules of the game are collective by nature, meaning that individual concerns are at times ignored by the majority, which can exacerbate the sense of helplessness. As a result, it appears natural to associate the domain of politics with anxiety. The emotion has already gained the attention of scholars interested in politics, although empirical evidence has focused mostly on actors outside political institutions, such as voters (Marcus et al., 2000, e.g.). This paper represents a systematic attempt to automatically measure the level of anxiety in an extensive corpus of parliamentary debates.

Previous studies seeking to measure anxiety in text using methods of NLP have been sporadic, yet influential (Gilbert and Karahalios, 2010; Bollen et al., 2011). Since the task of detecting anxiety involves the recognition of rather subtle forms of expression in language, the levels of accuracy reached using machine learning classifiers are typically lower than for other tasks. More than a decade ago, Mishne (2005) released a valuable resource for the study of emotions in text, namely a corpus of blog posts from the LiveJournal website where users were invited to self-report their current mood by selecting options from a list, including *anxious*. In the original study, classifiers trained with Bag-of-Words (BoW) features reached an accuracy rate of 54.25% for the detection of the anxious mood, in balanced class problems. In a more recent study using an

unbalanced dataset, Gilbert and Karahalios (2010) reached a very high level of accuracy for the detection of non-anxious texts, but only around 30% of correct predictions for anxious ones. Papers focusing on the related concept of fear have reported F1-scores between 0.20 and 0.74 (Strapparava and Mihalcea, 2007; Roberts et al., 2012; Mohammad, 2012).¹ It should be noted, however, that trying to map specific emotions such as anxiety into binary categories may not be the optimal approach. Feelings can be of various intensities, and binary annotations amount to losing precious information. Several modern approaches have instead relied upon emotion-specific lexicons that can be used to create continuous indicators (Tumasjan et al., 2010; Mohammad et al., 2013; Zhao et al., 2014). Using such a methodology, for instance, Bollen et al. (2011) were able to show that anxiety in social media is related to movements in the stock market.

The objective of this study is twofold. First, I seek to develop a methodology for the automatic measurement of levels of anxiety in writings, in particular, one that is adaptive to domain-specific corpora. The third section of this report describes the method and evaluates its performance using two annotated datasets: a sample of sentences from the Canadian House of Commons debates (the Hansard), annotated for anxiety using a crowd-sourcing platform, and the above-mentioned LiveJournal mood corpus. Second, my goal is to explain the prevalence of anxiety in real-life political speeches using a large section of the Hansard corpus.² Based on state-trait anxiety theory, which I discuss next, I implement predictive models in which some of the topics debated are expected to induce anxiety among parliamentarians.

2 Theoretical Background

I derive my expectations about the presence of anxiety in politics from the seminal state-trait anxiety theory proposed by Spielberger (1966). The theory emphasizes the distinction between anxiety as

¹Although Freud (1920, S. 3, Ch. 25) originally made a distinction between anxiety and fear, scholars have since emphasized the close connection between the two concepts, some advocating their conflation (McReynolds, 1976; Gray and McNaughton, 2000; LaBar, 2016).

²The full corpus is released on the www.lipad.ca website.

a *trait*, a more permanent feature of one's personality, and anxiety as a *state*, a temporary reaction to stressful situations. In this model, state anxiety is caused by the combination of both external events and individual predispositions. Anxiety as a trait has been the object of a voluminous literature, and corresponds to one of the "Big Five" personality traits (neuroticism) (Eysenck, 1997). In fact, other popular theories in psychology emphasize the role of factors such as personal experiences and biases in explaining anxious responses to external stimuli, for instance appraisal theory (Smith and Lazarus, 1993). Transposed to the realm of politics, this means that individual members of parliament (MPs) should exhibit various levels of anxiety in response to a situational trigger, depending on their personalities. An advantage of the Hansard corpus is that it allows to observe the same individuals repeatedly, under various situations. Thus, it becomes possible to account for individual traits, even if they are unobserved.

Regarding external factors, a central idea guiding my empirical analysis is that anxiety is more likely to occur the less an issue is known to MPs. Such an expectation directly follows from the nature of anxiety, which entails uncertainty, and parallels other theories in political psychology arguing that unfamiliar situations make individuals more anxious (Marcus, 2000; Marcus et al., 2000). In the context of a parliament, I expect subject-matters that are external to the country, such as foreign affairs and immigration, to be less familiar, hence more anxietyogenic. One could view this as an argument from analogy: for individuals, strangers can be a source of anxiety since their intentions are unknown; for representatives of a state, other countries and their nationals can be a similar source of unpredictability. Of course, national issues such as natural disasters or economic downturns can also be worrisome to politicians. However, a reasonable working hypothesis is that the more foreign the object of debate, the more likely MPs are to experience anxiety.

Unlike ordinary social settings, however, members of democratic institutions have functions that may constrain the manner in which they respond to ongoing events. In particular, members face different incentives depending on the role that their political formation occupies in a legislature. Just like personality traits may predispose to anxiety, one's

partisan affiliation can affect the likelihood of expressing signs of that emotion. Parties forming the opposition are required by convention to challenge the government’s decisions; in fact, a period in the daily business of the Canadian House of Commons, the Oral Question Period, is dedicated to that duty. It is natural to expect MPs from opposition parties to communicate more forcefully their worries about the problems faced by their constituents, as opposed to members of the party in power, who are expected to bring reassurance. Figure 1 is a simplified depiction of the model guiding the present study. The figure is loosely based on the state-trait anxiety model in Spielberger (1966), amended to account for the peculiarities of parliamentary life. Basically, political speeches are expected to become anxious due to external triggers (unfamiliar topics under debate), individual-level factors (mostly unobservable and deeply entrenched personality traits) and party-level functions (the government/opposition division).

3 Automated Detection of Anxiety in Text

A key challenge with the detection of anxiety in the Hansard corpus is to account for the specific language of a parliament. In contrast to the casual English used in social media, politicians in parliament may be less explicit about their mood. To illustrate, one message from the LiveJournal corpus annotated with the *anxious* mood starts with the following words:

[I] have finals today :(I am very nervous..

In this case, the statement expresses an unfiltered feeling that can be easily detected from word usage by a computer. However, parliamentary debates are subject to a decorum and speeches are usually less personal. An efficient methodology needs to take into account the vocabulary used to communicate emotions in such a formal setting.

For this reason, I rely upon a methodology that makes use of a vector space representation—the transformation of the corpus under study into numerical vectors based on word-word co-occurrences—to identify the proximity of each new word with a small set of seed words closely related to anxiety. This approach has been used previously to perform related tasks where semantic targets can be organized as polar opposites (Turney and

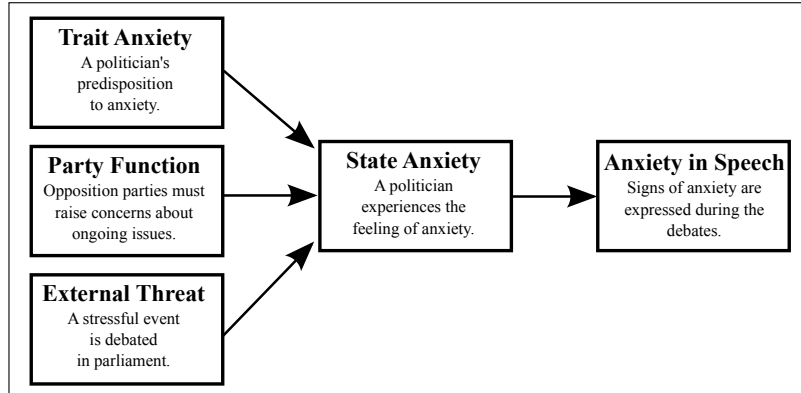
Littman, 2003). I have tested the approach with different lists of seed words. Although the selection of these initial seed words may be a sensitive task, accuracy results based on the annotated corpora were not severely affected by small changes in the seed list. For this analysis, I make use of 30 lemma/part-of-speech tuples associated with anxiety on the one hand, and with the absence thereof on the other hand. In the anxiety pole, seed lemmas include *threat, worry, fear, danger, anxious* and *tension*, for instance. These are general words selected recursively from a thesaurus for their association with the concept of anxiety. In the opposite pole, seed lemmas comprise *ease, protect, confidence, safe, trust* and *calm*. Considering pairs of lemmas and parts of speech (PoS) allows to take into account the distinction between the syntactic roles of each word, for instance when being used as a verb or as a noun.

Specifically, I start by preprocessing the full Hansard Corpus between 1980 and 2015. This stage includes word and sentence segmentation, the attribution of PoS tags and lemmatization, all performed using the Stanford CoreNLP library (Manning et al., 2014). This subset of the Hansard comprises above 245 million tokens, and the period was chosen to encompass that used in the rest of the empirical analysis. The vector space model is computed on a reconstructed corpus of combined lemma and PoS pairs using the Glove program (Pennington et al., 2014), which can be trained on the entire co-occurrence matrix. I create word vectors of 300 dimensions using a symmetric window size of 15 words.

Next, for each new lemma and PoS pair in the Hansard corpus, I compute the cosine similarity between their vectors and the vectors of each lemma and PoS pairing in the seed list. The sum of similarity scores with the no-anxiety seeds is then subtracted from the sum of similarity scores obtained for the anxiety pole, and scaled back into a $[-1, 1]$ range, where 1 means the most anxious lemma and -1 the least anxious. The created lexicon comprises 14,703 lemmas occurring 200 times or more in the corpus, after excluding digits and proper nouns.

Finally, I match the numerical scores of anxiety to the lemma/PoS pairs in the original corpus and compute average values for each individual speech. The result is a continuous indicator of the level anxiety, adapted to the specific register and tone of pol-

Figure 1: Simplified State-Trait Anxiety Model



itics, which can also be linked back to metadata on individual speakers, their party and the topic under debate.

To assess the performance of this approach, I retained the services of contributors from the CrowdFlower website. I randomly selected two series of 2,000 sentences from the Hansard, respectively from the Government Orders and Oral Question Period between 1994 and 2015.³ The decision to use sentences as a unit of annotation is justified by the large variation in the length of speeches during parliamentary debates. Preliminary assessments made with the help of a research assistant, based on entire speeches as units of analysis, revealed the task to be impractical for human coders, since full speeches may introduce several questions about how to properly code a document expressing different intensities of emotions. Shorter units of text can be more readily coded as anxious or non-anxious, and they reduce the risk of contradictory signals. I imposed the length of sentences to be above 8 tokens and removed procedural interventions (e.g. the introduction of motions). The 4,000 sentences were annotated by three different contributors via the crowd-sourcing platform, with an average pairwise agreement of 79.3%. Contributors were pre-screened with questions from a test set and their judgments considered reliable only if they reached a 70% agreement or above on these test questions. Although reliability may not reach the levels of more traditional methods, a re-

³Those two sections of the daily business of the House are the most important ones. The Government Orders comprise the debates on the bills introduced by the government. The Oral Question Period opens the floor to the opposition, who can question the government on its agenda and decisions.

cent study suggests that crowd-sourcing yields annotations comparable to those obtained from expert surveys (Benoit et al., 2016).

The proportion of sentences coded as anxious reaches 10.7% in the 4,000 Hansard sentences, larger than in the LiveJournal corpus. To avoid an unbalanced class problem, I have matched anxious sentences with the same number of randomly selected non-anxious ones. Although this is an admittedly small sample (854 sentences), the methodology reaches appreciable levels of accuracy using support vector machines (SVMs) with a linear kernel and a penalty parameter set at 0.5. Using the anxiety indicator as a single feature, the level of accuracy is around 63% (see Table 1), already above the levels of some previously cited attempts to detect the same emotion. Performance is slightly improved when adding a BoW including unigrams and bigrams.

I achieve important gains after including a simple categorical variable identifying the speaker, as well as title annotations from the original Hansard. With these metadata as features, the rate of accuracy reaches 83.6% (with a F1-score of 0.83, after selecting the 700 best features based on chi-square tests). This result is consistent with psychology theory, which emphasizes the importance of accounting for individual-specific traits, and illustrates the role of non-textual features in classification tasks. Accuracy scores are computed with randomly assigned training and testing sets; using 3K-fold cross-validation, the percentage of accuracy reaches 82% in the best model. Additionally, I have considered grammatical features such as verb tense, after

Hansard (Anxious v. Non-Anxious)

Features	Accuracy	F1-Score	Sample
Anxiety Score	62.9%	0.632	854
+BoW	69.0%	0.673	854
+Metadata	83.6%	0.830	854

LiveJournal (Anxious v. Non-Anxious Moods)

Features	Accuracy	F1-Score	Sample
Anxiety Score	57.2%	0.582	30000
+BoW	74.6%	0.745	30000

LiveJournal (Anxious v. All Others Moods)

Features	Accuracy	F1-Score	Sample
Anxiety Score	55.1%	0.565	30000
+BoW	72.9%	0.713	30000

Table 1: Accuracy Results for SVM Classifiers

identifying the verb phrase containing the root in the dependency grammar structure of each sentence. These features are significantly related to anxiety (the future tense being positively related to anxiety, and conversely for the past tense). However, they did not improve accuracy further once accounting for speaker identification.

Finally, I also assess the methodology against the LiveJournal corpus. The latter comprises over 700,000 entries associated with several moods. Since the proposed methodology is adaptive, I replicated the same steps with the LiveJournal data, creating anxiety scores using word vectors of the entire corpus and the same seed lists as before. Thus, the measure is now tailored to the specific genre of casual online discussions. To select anxious posts, I combined semantically equivalent reported moods (namely *anxious*, *stressed*, *worried* and *restless*). Following the approach in Mishne (2005), I matched a random sample of 15,000 anxious posts to the same number of non-anxious ones (merging *calm*, *relaxed*, *optimistic* and *relieved*) and to a sample of all other moods except anxiety.⁴ The two bottom panels of Table 1 report the results. Accuracy using the anxiety score alone is lower than before, around 57%, still using SVMs ($C = 80$). Yet, the fact that a univariate indicator is sufficient to achieve a performance similar to that reported in previous studies using BoW models brings some additional legitimacy to the methodology. When including unigrams and

⁴I have excluded irregular postings such as lists of tracks from musical albums and messages in foreign language.

bigrams, the predictive model reaches a level of accuracy exceeding 70%. Despite the possible limitations of the two corpora discussed in this section, these results suggest that the proposed score is significantly related to anxiety, as perceived by human coders or self-reported by Internet users.

4 Application to Parliamentary Debates

To examine the prevalence of anxiety in political speeches, I rely upon the previously introduced Hansard of the Canadian House of Commons. Using the anxiety scores generated with the method described in the previous section, speeches in that corpus are now associated with a level of anxiety. The database is constructed at the speech level, each observation corresponding to one intervention, and each being associated with metadata. I removed non-substantive speeches, for instance introductory remarks, questions of procedure, and interventions from the Speaker of the House.⁵ I have also removed speeches for which the MP was not identified and those comprising less than 8 tokens, to avoid a number of common brief utterances such as “Yes/No” responses.

A portion of the Canadian parliamentary debates had been manually annotated by other researchers using the Comparative Agendas Project (CAP) topic scheme, namely a set of over 40,000 questions asked during the Oral Question Period between 1983 and 2004 (Penner et al., 2006). Coders attributed a topic to each new question in a given exchange (an exchange may include a question, an answer, and sometimes follow-up questions), which means that each annotation can be matched to more than one speech in the Hansard corpus. Using variables common to both datasets, I was able to match most of those annotations back with their original text.⁶ This leaves a dataset of 119,623 speeches annotated by topic, which I focus on for the empirical analysis below. The topic scheme used by the CAP team is

⁵In contrast to Speakers in the US House of Representatives, who are high-ranking legislators, the role of Speaker in Westminster systems is procedural. Speakers in Canada do not play a partisan role in the debates.

⁶Manual inspection after running the matching algorithm revealed a few cases of mismatch due to discrepancies in the two databases, often unresolvable. However, the large majority of annotations appear correctly attributed in the Hansard corpus.

Speech	Score
<i>Mr. Speaker, the Prime Minister is joking about the serious and even catastrophic situation of Canada's public finances.</i>	0.29
<i>That this House condemn those policies of the government which, having bred a sense of insecurity, powerlessness and doubt among Canadians, are a cause of the recent increase in regional tensions and incidents of racism and bigotry in Canada.</i>	0.27
<i>In those capacities we have absolute confidence in the judgment of the Secretary of State for External Affairs, and we can offer that categoric assurance to the Leader of the New Democratic Party.</i>	-0.48
<i>Mr. Speaker, as always we will strive to protect the best interests of all Canadians.</i>	-0.50

Table 2: Examples of Anxious and Non-Anxious Speeches

rather detailed, comprising 25 main topics and more than 200 subtopics. For the sake of this study, I have simplified the scheme to 10 categories by relabeling the original CAP annotations. These 10 topics should sound fairly intuitive to both political scientists and scholars from other disciplines: constitutional affairs, economy, education, environment, foreign affairs, health care, immigration, law and order, natural resources, and a residual category containing all other national topics such as land management and intergovernmental affairs.

Illustrating the methodology when applied to political debates, Table 2 reports two examples of speeches above the 99th percentile on the anxiety score, and two under the 1st percentile. As can be seen, the measure correctly ranks the first speech as an expression of anxiety, despite the presence of a word like *joking*, which could have been misinterpreted as indicative of a light tone. The examples may also serve to illustrate how political functions influence the emotional valence of speeches: the two anxious speeches come from members of the opposition parties, whereas the bottom two speeches of Table 2 come from the government. The large majority of speeches fall under 0 on the anxiety scale, which is consistent with the fact that the feeling remains infrequent. The average value in the full sample is -0.13 . In what follows, I normalize the indicator by transforming it into a z -score with a mean of 0, which facilitates the interpretation of findings.


The main objective from now on is to test the hypotheses laid out earlier in the theoretical section.

I start by looking at the bivariate relationship between the topic under debate and the level of anxiety of MPs. Figure 2 shows the average level of anxiety by topic category (I use a weighted average with token counts as weights to account for discrepancies in speech lengths). As can be seen, immigration and foreign affairs are the topics for which Canadian MPs displayed the highest average level of anxiety between 1983 and 2004 during oral questions, which is consistent with the starting hypothesis. It shall be noted that anxiety is not bound to external topics, however. The question of natural resources—historically important for the Canadian economy, and including debates on sectors such as oil production and nuclear energy—ends up closely behind, before environment and economy. Law and order (which includes crime) and health care range somewhere in the middle, being neither the most anxious nor the least. Education and constitutional affairs rank as the least anxiogenic topics during this period.

It may be helpful to consider topic averages broken down by year to understand why immigration turns out first in the ranking during that time span. The highest levels of anxiety on the topic of immigration are observed between 1993 and 2001, period during which one of the major parties—the Conservatives—had ruptured into factions. The new parties (Reform and Bloc) began to voice a more forceful opposition on that particular topic. The period culminated with a major reform of immigration legislation in 2001, year during which the average level of anxiety reaches 0.31.

The corpus also allows for an examination of anxiety by party function, another component of the model introduced above. Table 3 reveals that members from parties in the opposition have levels of anxiety larger than the global average ($+0.26$), while the converse holds for parties forming the government (-0.20). The specific parties exhibiting the highest levels of anxiety in Table 3 are also those that remained in the opposition the whole time. I have also examined the link between anxiety and available biographical information on MPs. Women, for instance, appear slightly less anxious than men, still using the 1983–2004 sample of oral questions, based on a one-tailed mean comparison t -test ($p < 0.02$). Age also appears to be a significant predictor of anx-

Figure 2: Anxiety by Topic in the House of Commons (1983–2004)



Political Party	Anxiety
Bloc Québécois	0.16
Canadian Alliance	0.35
Conservatives	-0.08
Liberals	-0.13
NDP	0.27
Reform	0.33
Government	-0.20
Opposition	0.26

Table 3: Weighted Average of Anxiety (Normalized Scale) by Political Group during Oral Question Period (1983–2004)

ity in speeches: older MPs are less likely to express anxiety than younger ones.

Bivariate relationships, however, ignore the fact that some of those speeches are uttered by the same individuals. As pointed out earlier in the theoretical section, individual characteristics (traits) may affect whether someone is likely to feel nervous at any given point (state). Some MPs may have a tendency to experience anxiety and express worries in the face of adversity, whereas others may have a more stoic personality. Such individual traits can represent a confounding factor, and the unexplained variance of the anxiety score may be clustered within individual MPs. For these reasons, I rely upon statistical methods accounting for individual characteristics. I first consider a “within” estimator that subtracts individual averages from each variable in the model (which amounts to including a specific intercept parameter for each MP). This is also known as a fixed effects

estimator that can be expressed in the form

$$y_{ij} = a_j + \mathbf{x}'_{ij}\boldsymbol{\beta} + \varepsilon_{ij}, \quad (1)$$

where y_{ij} is the measured anxiety of a speech i uttered by MP j , the a_j are MP-specific intercepts, \mathbf{x} is a vector of observed metadata for each speech, and ε_{ij} is an error component. The within transformation precludes the inclusion of variables that do not vary within individuals, such as gender. On the other hand, this transformation implies that the estimator accounts for all unobserved MP characteristics, such as personality traits. Note that the government-opposition status varies for some MPs, and can be included among observed covariates. The second estimator is a hierarchical model with random intercepts at the party and individual levels. The model corresponds to:

$$y_{ij} = \alpha + \mathbf{x}'_{ij}\boldsymbol{\beta} + \lambda_j + \theta_k + \varepsilon_{ij} \quad (2)$$

where λ_j and θ_k are random intercepts for MP j and party k . This estimator allows for the inclusion of variables such as age and gender.

Table 4 reports estimates from both these models. Starting with the fixed effects model in the left column, estimated by OLS ($F = 40.04; p = 0.000$), the results are consistent with several of the findings stressed earlier about the relation between topics and anxiety. Since the topic variable is categorical, estimates are to be interpreted with respect to a base category, in this case the residual category “other topics”. As can be seen, the estimates suggest that foreign affairs and immigration are significantly more

Variable	Fixed Effects		Hierarchical	
	Est.	p	Est.	p
<i>Topic</i>				
Constitutional	-0.079	0.008	-0.079	0.000
Economy	0.075	0.000	0.076	0.000
Education	-0.041	0.057	-0.045	0.011
Environment	0.065	0.072	0.071	0.000
Foreign Affairs	0.125	0.000	0.129	0.000
Health Care	0.070	0.033	0.067	0.000
Immigration	0.145	0.000	0.150	0.000
Law and Order	0.077	0.001	0.081	0.000
Nat. Resources	0.145	0.000	0.143	0.000
<i>Party Status</i>				
Government	-0.360	0.000	-0.357	0.000
Obs.	119,623		117,704	

Table 4: Multivariate Models of Anxiety in Parliament. Both models include year, month and day of the week dummies as well as a control variable for the length of speeches in tokens. The hierarchical model is reported with approximate p -values, and also includes age and gender as controls.

anxiogenic than the base category, supporting the results presented so far. These estimates are statistically significant, with reported p -values computed using Arellano’s heteroskedasticity and autocorrelation (HAC) consistent standard errors (Arellano, 1987). However, once accounting for individual-specific factors, the topic of national resources appears to produce the same effect on anxiety than immigration, and a larger effect than foreign affairs. Education and constitutional affairs remain among the least anxious topics. Members of the government appear significantly less likely to express anxiety during their interventions, in line with the observations made previously. The last column of Table 4 reports results from the three-level hierarchical model estimated by restricted maximum likelihood. This model includes age and gender as covariates, which now turn out as non-significant predictors. A few data points on the birth dates of MPs are missing, reducing the number of observations. However, the results are very close to those obtained with the first estimator, and again support some of the key findings discussed so far.

5 Conclusion

Overall, this study illustrates the potential of NLP methods for the detection of specific emotions such as anxiety in politics. The finding that some topics are a source of anxiety may help to inform future research on legislative politics, for instance to examine

the role of emotions in the formation of preferences over issues. From a methodological standpoint, the study also illustrates the importance of theory and non-textual features in predictive tasks. In particular, the results suggest that machine learning models can be improved by accounting for author identification in corpora where documents are clustered by individuals. Although the findings bring support to the proposed model, a limitation of this study is that expressions of anxiety may not reflect the true emotional state of a speaker. For example, politicians may appeal to emotions strategically, in an attempt to persuade. Disentangling the motives behind the use of language would require additional research. Nonetheless, the framework proposed in these pages appears reliable enough to derive substantive results of interest.

Looking forward, the methodology could serve to answer other important questions about the role of anxiety in politics, ones that fell beyond the scope of this study. Previous research on anxiety has shown that the emotion influences how people make decisions. The experience of anxiety may ultimately make some choices less appealing, or refrain individuals from action altogether (Raghunathan and Pham, 1999; Gray and McNaughton, 2000). Marcus et al. (2000) stressed the hypothesis that anxiety leads individuals to search for more information before making decisions. Whether this has positive effects on the quality of decisions made by anxious individuals, or whether the feeling clouds their judgment, remains a debated question, albeit an important one (Brader, 2011; Ladd and Lenz, 2008; Marcus et al., 2011). Moreover, improving computational tools for the detection of specific emotions in texts can have applications useful beyond the study of politics. Examples include the detection of health conditions such as anxiety disorders, stock market forecasting or, more generally, contributions to the development of an artificial intelligence able to accurately identify specific emotions from language.

Acknowledgments

The author thanks Andreea Musulan for research assistance as well as the Connaught Fund for financial support.

References

- Manuel Arellano. 1987. Computing Robust Standard Errors for Within-Groups Estimators. *Oxford Bulletin of Economics and Statistics*, 49(4):431–434.
- Kenneth Benoit, Drew Conway, Benjamin E. Lauderdale, Michael Laver, and Slava Mikhaylov. 2016. Crowdsourced Text Analysis: Reproducible and Agile Production of Political Data. *American Political Science Review*, 110(2):278–295.
- Johan Bollen, Huina Mao, and Xiaojun Zeng. 2011. Twitter Mood Predicts the Stock Market. *Journal of Computational Science*, 2(1):1–8.
- Ted Brader. 2011. The Political Relevance of Emotions: ‘Reassessing’ Revisited. *Political Psychology*, 32(2):337–346.
- Michael W. Eysenck. 1997. *Anxiety and Cognition: A Unified Theory*. Psychology Press, Hove.
- Sigmund Freud. 1920. *A General Introduction to Psychoanalysis*. Boni and Liveright, New York.
- Eric Gilbert and Karrie Karahalios. 2010. Widespread Worry and the Stock Market. In *Proceedings of the International Conference on Weblogs and Social Media (ICWSM 10)*.
- Jeffrey A. Gray and Neil McNaughton. 2000. *The Neuropsychology of Anxiety: An Enquiry into the Functions of the Septo-Hippocampal System*. Oxford University Press, Oxford.
- Kevin S. LaBar. 2016. Fear and Anxiety. In Lisa Feldman Barrett, Michael Lewis, and Jeanette M. Haviland-Jones, editors, *Handbook of Emotions*. The Guilford Press, New York.
- Jonathan McDonald Ladd and Gabriel S. Lenz. 2008. Reassessing the Role of Anxiety in Vote Choice. *Political Psychology*, 29(2):275–296.
- Joseph LeDoux. 2012. Rethinking the Emotional Brain. *Neuron*, 73(4):653–676.
- George Mandler. 1972. Helplessness: Theory and Research in Anxiety. In Charles D. Spielberger, editor, *Anxiety: Current Trends in Theory and Research*, pages 359–378. Academic Press, New York.
- Christopher D. Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky. 2014. The Stanford CoreNLP Natural Language Processing Toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60.
- George E. Marcus, W. Russell Neuman, and Michael MacKuen. 2000. *Affective Intelligence and Political Judgment*. University of Chicago Press, Chicago.
- George E. Marcus, Michael MacKuen, and W. Russell Neuman. 2011. Parsimony and Complexity: Developing and Testing Theories of Affective Intelligence. *Political Psychology*, 32(2):323–336.
- George E. Marcus. 2000. Emotions in Politics. *Annual Review of Political Science*, 3:221–250.
- William T. McReynolds. 1976. Anxiety as Fear: A Behavioral Approach to One Emotion. In Marvin Zuckerman and Charles D. Spielberger, editors, *Emotions and Anxiety: New Concepts, Methods, and Applications*, pages 281–316. Lawrence Elbaum Associates, Hillsdale.
- Susan Mineka and Kelly A. Kelly. 1989. The Relationship Between Anxiety, Lack of Control and Loss of Control. In Andrew Steptoe and Ad Appels, editors, *Stress, Personal Control and Health*, pages 163–191. John Wiley & Sons, Chichester.
- Gilad Mishne. 2005. Experiments with Mood Classification in Blog Posts. In *Style2005: 1st Workshop on Stylistic Analysis Of Text For Information Access*.
- Saif M. Mohammad, Svetlana Kiritchenko, and Xiaodan Zhu. 2013. NRC-Canada: Building the State-of-the-Art in Sentiment Analysis of Tweets. In *Proceedings of the Seventh International Workshop on Semantic Evaluation Exercises (SemEval-2013)*, Atlanta, Georgia, USA, June.
- Saif Mohammad. 2012. Portable Features for Classifying Emotional Text. In *2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 587–591.
- Erin Penner, Kelly Blidook, and Stuart Soroka. 2006. Legislative Priorities and Public Opinion: Representation of Partisan Agendas in the Canadian House of Commons. *Journal of European Public Policy*, 13(7):959–974.
- Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global Vectors for Word Representation. In *Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*.
- Rajagopal Raghunathan and Michel Tuan Pham. 1999. All Negative Moods Are Not Equal: Motivational Influences of Anxiety and Sadness on Decision Making. *Organizational Behavior and Human Decision Processes*, 79(1):56–77.
- Kirk Roberts, Michael A. Roach, Joseph Johnson, Josh Guthrie, and Sanda M. Harabagiu. 2012. EmpaTweet: Annotating and Detecting Emotions on Twitter. In *Proceedings of the 9th Language Resources and Evaluation Conference*, pages 3806–3813.
- Craig A. Smith and Richard S. Lazarus. 1993. Appraisal Components, Core Relational Themes, and the Emotions. *Cognition and Emotion*, 7(3/4):233–269.
- Charles D. Spielberger. 1966. Theory and Research on Anxiety. In Charles D. Spielberger, editor, *Anxiety and Behavior*, pages 3–20. Academic Press, New York.
- Carlo Strapparava and Rada Mihalcea. 2007. SemEval-2007 Task 14: Affective Text. In *Proceedings of the*

- 4th International Workshop on Semantic Evaluations*.
Association for Computational Linguistics.
- Stephen J. Suomi and Harry F. Harlow. 1976. The Facts and Functions of Fear. In Marvin Zuckerman and Charles D. Spielberger, editors, *Emotions and Anxiety*, pages 3–34. Lawrence Erlbaum Associates, Hillsdale.
- Andranik Tumasjan, Timm O. Sprenger, Philipp G. Sandner, and Isabell M. Welp. 2010. Predicting Elections with Twitter: What 140 Characters Reveal about Political Sentiment. In *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media*.
- Peter D. Turney and Michael L. Littman. 2003. Measuring Praise and Criticism: Inference of Semantic Orientation from Association. *ACM Trans. Inf. Syst.*, 21(4):315–346.
- Jian Zhao, Liang Gou, Fei Wang, and Michelle Zhou. 2014. PEARL: An Interactive Visual Analytic Tool for Understanding Personal Emotion Style Derived from Social Media. In *Proceedings of the IEEE Symposium on Visual Analytics Science and Technology*, 203–212.