Detecting emotional polarity in Finnish parliamentary proceedings

Suvi Lehtosalo*
University of Tübingen
suvi.lehtosalo@student.uni-tuebingen.de

John Nerbonne
Groningen, Freiburg & Tübingen
j.nerbonne@rug.nl

Abstract

Few studies have focused on detecting emotion in parliamentary corpora, and none have done this for the Finnish parliament. In this paper, this gap is addressed by applying the polarity lexicon–based methodology of a study by Rheault et al. (2016) on speeches in the British Parliament to a Finnish corpus. The findings show an increase in positive sentiment over time. Additionally, the findings indicate that politicians' emotional states may be impacted by the state of the economy and other major events, such as the Covid-19 pandemic and the Russian invasion of Ukraine.

1 Introduction

1.1 Goal

The goal of this paper is to apply the methodology used by Rheault et al. (2016) on British parliamentary speeches to a Finnish dataset, in order to determine whether the findings – increase in emotional polarity over time and correlation with the state of the economy – can be replicated with a different corpus in another language. Additionally, comparison to other major events, namely the Covid-19 pandemic and the Russian invasion of Ukraine, is drawn in order to investigate which other topics may impact emotional polarity.

1.2 Background

This study is a contribution to the field of *affective computing*, that is, the detection of emotion in human language and expression. Recent studies in this field have examined various topics, such as the applications of emotion detection in medicine (Saffar et al., 2023), and using emotion analysis to detect fake news on social media (Hamed et al., 2023; Luvembe et al., 2023).

Rheault et al. (2016) performed a study detecting emotion in British parliamentary proceedings, and found an increase in polarity over time, as well as a correlation between the state of the economy and polarity. Since then, other works have addressed detecting anxiety (Rheault, 2016) and general emotionality (Gennaro and Ash, 2022) in parliamentary proceedings. A systematic analysis performed by Abercrombie and Batista-Navarro (2020) summarises efforts in the area of detecting sentiment or opinion in parliamentary debates. Detecting emotion is rarely done in isolation; most studies have focused on tasks such as detecting agreement/disagreement, vote prediction, and ideology detection. The use of a polarity lexicon is one of the common methods, the other most common method being the use of statistical machine learning.

A few studies have been performed involving automated analysis of Finnish parliamentary proceedings. These include the use of semantic tagging to detect discussions of Everyman's Rights (Kettunen and La Mela, 2022), and identifying topics in parliamentary speech using a Machine Learning approach (Ristilä and Elo, 2023). No previous studies could be found that focus on detecting emotion or use a polarity lexicon–based approach.

In addition to detecting emotion, this study aims to uncover some of the factors affecting emotionality in parliament. The impact of the economy on happiness is well-studied (Frey and Stutzer, 2000; Oswald, 1997), and a correlation between state of the economy and emotion in parliament was already found by Rheault et al. (2016). In recent years, the outbreak of war in Ukraine, as well as the Covid-19 pandemic have had a profound effect on peoples' lives, with the mental and emotional impact of the latter having been widely studied

*SL took part in a course "Digital Humanities from a Computational Linguistics perspective" taught by JN. SL presented the paper by Rheault et al. (2016) there and suggested trying a similar technique on the Finnish parliamentary protocols. SL found data for this, suitable prior work for detecting sentiment in Finnish and wrote the first version of the paper. JN supervised the work but finds that he added only a little.

Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics, pages 90–100
July 5–10, 2020.

(Boden et al., 2021; Talevi et al., 2020; Terry et al., 2020). In this study the correlation between these major events and emotion in parliament is studied, in addition to the correlation with the economy.

2 Method

2.1 Data and preprocessing

The parliamentary proceedings were acquired from the ParliamentSampo project, which has compiled all speeches in the Parliament of Finland from 1907 onwards (Hyvönen et al., 2024). As this study focuses solely on Finnish, all Swedish-language speeches were removed; after this, the dataset consisted of 168 million tokens. The corpus was lemmatised and Part of Speech–tagged using the *libvoikko* Python library (Voikko, n.d.). All proper nouns and digits were removed. As the word *arvoisa*, meaning *honourable*, is frequently used to address the Speaker or Members in expressions such as "Honourable Speaker", it was also removed from the corpus.

The economic data was obtained from Statistics Finland (Statistics Finland, 2024a,b,c,d). As their economic measures, Rheault et al. (2016) used GDP growth, unemployment rate, the misery index and a measure of labour disputes. The same measures were used here, but as the unemployment rate and misery index data were only available from 2010 onwards, two additional economic measures were added: growth in the Cost of Living Index (CLI) and the Consumer Price Index (CPI). The economic measures used are shown in Table 3 in Appendix A.

The Covid-19 incidence data was obtained from the Finnish National Infectious Diseases Register (Finnish Institute for Health and Welfare, 2024). The dataset includes country-wide data on Covid testing and the number of cases. A measure was calculated from these data by dividing the number of cases by the number of tests done, in order to approximate the proportion of positive test results.

2.2 Detecting emotion

Emotion detection was performed using a polarity lexicon–based approach. As existing Finnish polarity lexicons are not specifically tailored for political data, a new lexicon was constructed based on the parliamentary corpus. This involved generating word embeddings and selecting positive and negative seed lemmas, and then using these to calculate polarity scores to create the lexicon. This lexicon could then be used to calculate polarity scores for time periods of arbitrary length.

The word embeddings were generated from the corpus using the GloVe algorithm (Pennington et al., 2014), also used by Rheault et al. (2016). GloVe embeddings were employed to maintain consistency with the original work; additionally, they appear to perform well in comparison with other context-independent word embedding methods, as shown in works such as Toshevska et al. (2020) which compared multiple word embedding methods with human similarity judgements, and Jain et al. (2021), which compared multiple word embedding methods for hate speech detection.

Words with fewer than 50 occurrences were removed to eliminate typos and rare words. A list of seed lemmas with 100 positive and 100 negative lemmas was compiled, filtering out words relating to political topics such as war and disease. Further details on embedding generation and seed lemma selection are available in Appendix A.

Calculation of word polarities Word polarities were calculated based on cosine similarity, as done by Rheault et al. (2016):

$$s_i = \sum_{p=1}^{P} \cos(\mathbf{v}_i, \mathbf{v}_p) - \sum_{q=1}^{Q} \cos(\mathbf{v}_i, \mathbf{v}_q) \quad (1)$$

where s_i is the score for lemma i, positive seed lemmas are indexed by p and negative seed lemmas by q, and $\cos(\mathbf{v}_i, \mathbf{v}_p)$ represents the cosine of the angle between vectors $\mathbf{v}_i$ and $\mathbf{v}_p$.

After calculating the scores, the 2000 lemmas with the highest scores and the 2000 lemmas with the lowest scores were added to the polarity lexicon, making the size of the lexicon 4200, including the original seed lemmas.

Calculation of polarity in a time period For calculating the polarity in a time period, the equation by (Rheault et al., 2016) was used:

$$y_t = \frac{\sum_{i=1}^{n_i} \mathbf{1}\{w_{it} \in L\} s_i \theta_{it}}{\sum_{i=1}^{n_i} \mathbf{1}\{w_{it} \in L\}} \quad (2)$$

Here, $\mathbf{1}\{w_{it} \in L\}$ is a filter, returning 1 when the word w_{it} is included in the polarity lexicon, and 0 otherwise; s_i is the word score retrieved from the polarity lexicon. θ_{it} is another filter, which is intended to eliminate negated words. This is to avoid, for example, evaluating phrases like *not happy* as positive or *not angry* as negative. The filtering is

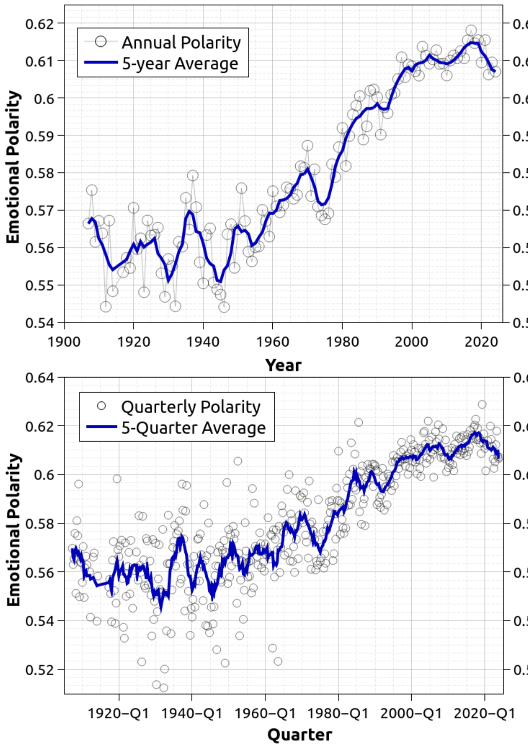

Figure 1: Emotional polarities calculated for each year and each quarter, showing both the original data points and 5-point rolling averages.

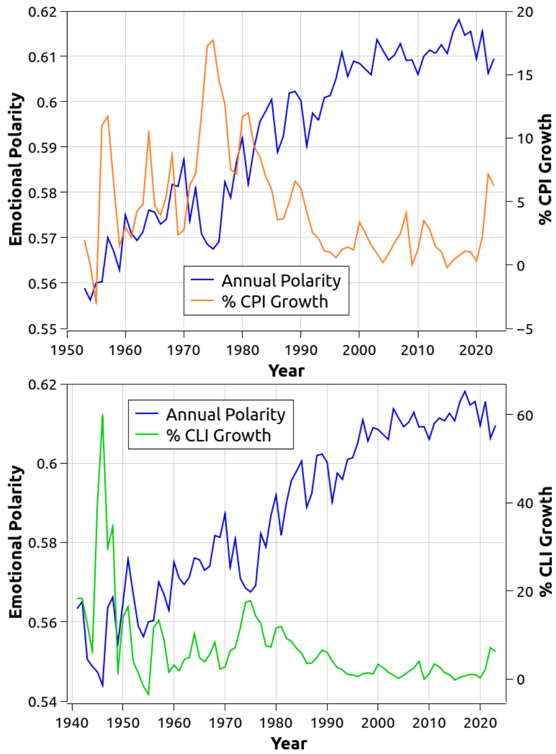

Figure 2: Change in Consumer Price Index and Cost of Living Index plotted against emotional polarity.

performed by checking whether the word w_{it} is located between a negation word and a punctuation mark.

3 Results

3.1 Emotion in the Finnish parliament

The annual and quarterly emotional polarities are shown in Figure 1. The sentiments become increasingly positive over time, mirroring the similar findings by Rheault et al. (2016). This clear rise can be seen from the late 1940s to the late 2010s. More variation in the polarities can be seen before the 1960s. Overall, the polarities fall in a range [0.54,0.62], with the possible range being [-1.0,1.0]. As the focus is change in polarity over time, the specific values do not matter, but this shows that words with negative polarities were much less common in the corpus than positive words.

3.2 Comparison with economic data

The different economic measures were plotted against emotional polarity to examine potential correlations. Figure 2 shows the plots for CPI and CLI growth, which showed the clearest visual correlation. Both measures show a peak in the 1970s, which coincides with more negative sentiments. CLI growth also shows a similar phenomenon in the 1940s. Both graphs show a decrease in growth from the 1980s to the late 2010s, which is reflected as an increase in positive sentiments. These graphs seem to strongly suggest a relationship between emotional polarity and measures of inflation. Plotting the misery index also showed potential correlation, though not as clearly (see Figure 3).

In order to further investigate the findings from plotting the data, pairwise Granger causality tests were performed on each of the economic measures and emotional polarity. Bonferroni correction (Dunn, 1961) was applied to the p-values to address the multiple comparisons problem. The resulting p-values can be seen in Table 4 (see Appendix A). Before correction, the p-values for CLI growth and emotional polarity and CPI growth and emotional polarity fall below the threshold of 0.05, supporting a possible relationship; after correction, only CLI growth is below this threshold at 0.034. The p-value of the Misery Index–Polarity relationship is well above this threshold, indicating no relationship.

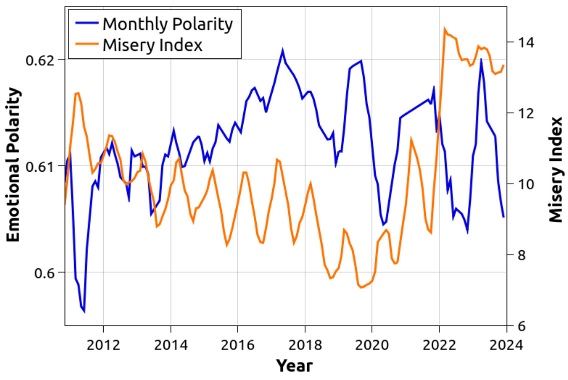

Figure 3: 4-month rolling averages of the misery index and emotional polarity.

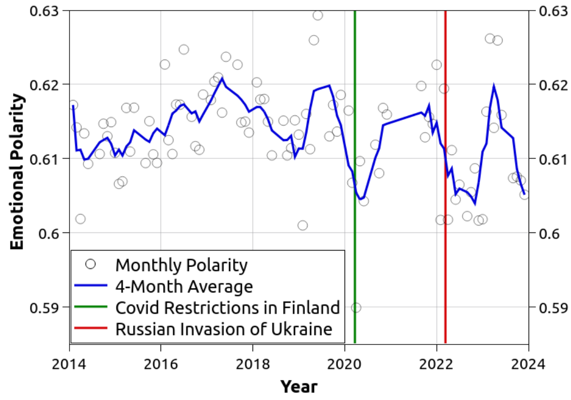

Figure 4: Emotional polarity and recent events.

3.3 Comparison with recent events

Another topic of investigation was the impact of recent major events on emotional polarity. The two events chosen for this were the Covid-19 pandemic and the Russian invasion of Ukraine.

Figure 4 shows monthly polarities in the last decade, as well as vertical lines indicating the significant dates. The first date is the first implementation of Covid-19 restrictions in Finland, the 17th of March 2020. This specific date was chosen (as opposed to other dates, such as the first incidence of Covid-19 in Finland), as it shows when Covid-19 was first seen as a national emergency in Finland. The second date shown is the start of the ongoing Russian invasion of Ukraine, the 24th of February 2022. Both dates coincide with periods of significant negativity in the parliament, suggesting a possible correlation.

As statistical data is easily available for Covid-19, the correlation between Covid-19 and emotional polarity was further investigated with statistical testing. Granger causality testing performed on the measure of positive test rates resulted in a Bonferroni-corrected p-value of 0.029 (see Table 4 in Appendix A), supporting the possibility of a correlation.

4 Discussion

Interpretation of results An overall rise in emotional polarity was seen from the mid-20th century to the 2010s. This mirrors the findings by Rheault et al. (2016), showing that this may be a universal phenomenon rather than specific to the British parliament.

A correlation was found between measures of inflation and emotional polarity. This also aligns with the findings by Rheault et al. (2016), although the specific measures used were different; a correlation could not be shown using labour disputes or the misery index as the economic measures, and after Bonferroni correction only one economic measure showed a statistically significant relationship to emotional polarity. As the findings are dependent on the specific measures used, it is difficult to assert a general relationship between economy and emotional polarity.

A possible correlation was seen between emotional polarity and the Covid-19 pandemic and the Russian invasion of Ukraine, although the latter was assessed only visually due to scarcity of statistical data. It is difficult to assess whether these correlations may simply be caused by the words associated with these events being labelled as negative. While the seed lemma selection attempted to avoid this effect by filtering out any related lemmas, it cannot be guaranteed that this was entirely effective.

Limitations Rheault et al. (2016) evaluated their approach using an IMDB film review dataset; as the approach here closely follows theirs, no new evaluation was performed. It is, however, possible that some minor differences in the approaches (or differences between the languages) could have caused a difference in performance, so a new evaluation could be useful. For example, the FinnSentiment social media corpus (Lindén et al., 2020) could be used for this, although the casual and contemporary nature of the language used on social media makes it less applicable to parliamentary speech. Additionally, as the evaluation performed by Rheault et al. (2016) also used a contemporary film review dataset, the approach of using a single polarity lexicon to measure emotion over many decades has not been validated. In practice, this kind of validation would be difficult to perform, due to a lack of labelled data spanning a long period of time.

A more careful approach to generating the polarity lexicon may be beneficial. While the initial list of seed lemmas was filtered to remove words related to politics or other words that could skew the results, the other 4000 lemmas in the polarity lexicon were selected solely based on the calculated similarity to the seed lemmas, with no further filtering. Searching the lexicon for words relating to topics such as disease or war and filtering these out could help avoid skew, though performing a thorough search would be quite labour-intensive.

5 Conclusion

The results broadly align with those shown by Rheault et al. (2016), with an increase in positive sentiment over time, and correlation between polarity and certain economic measures. However, this correlation could only be seen with certain measures, and the measures involved were slightly different between this study and Rheault et al. Links between labour disputes and polarity and the misery index and polarity could not be demonstrated. Further investigation into the topic is needed in order to conclusively determine whether economic measures impact the emotion shown in parliamentary speech. One future avenue of research that would likely prove fruitful is sentiment analysis through a Machine Learning approach, as this can achieve a higher accuracy compared to polarity lexicon–based methods, although at the expense of interpretability (Hartmann et al., 2023). Additionally, utilising contextual word embedding methods such as ELMo (Peters et al., 2018) or BERT (Devlin et al., 2019) would be worthwhile, as these may perform better compared to older word embedding methods such as GloVe (Jain et al., 2021).

Some correlation was also seen between the other studied events and emotional polarity. Further evaluation could determine whether this can be explained by words related to these events being evaluated as negative, or whether discussions unrelated to these topics also show a lowered polarity. Future study could classify speeches by topic, a task which has already been performed with Finnish parliamentary data (Ristilä and Elo, 2023), and eliminate speeches related to the events in question, in order to only focus on the unrelated speeches.

References

Gavin Abercrombie and Riza Batista-Navarro. 2020. Sentiment and position-taking analysis of parliamentary debates: a systematic literature review. *Journal of Computational Social Science*, 3(1):245–270.

Matt Boden, Lindsey Zimmerman, Kathryn J. Azevedo, Josef I. Ruzek, Sasha Gala, Hoda S. Abdel Magid, Nichole Cohen, Robyn Walser, Naina D. Mahtani, Katherine J. Hoggatt, and Carmen P. McLean. 2021. Addressing the mental health impact of COVID-19 through population health. *Clinical Psychology Review*, 85:102006.

CSC - IT Center for Science. 2004. Frequency Lexicon of the Finnish Newspaper Language.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *arXiv preprint*. ArXiv:1810.04805 [cs].

Olive Jean Dunn. 1961. Multiple comparisons among means. *Journal of the American statistical association*, 56(293):52–64.

Finnish Institute for Health and Welfare. 2024. Covid-19 cases in the infectious diseases registry. https://sampo.thl.fi/pivot/prod/en/epirapo/covid19case/summary_tshcddaily.

Bruno S Frey and Alois Stutzer. 2000. Happiness, economy and institutions. *The economic journal*, 110(466):918–938.

Gloria Gennaro and Elliott Ash. 2022. Emotion and Reason in Political Language. *The Economic Journal*, 132(643):1037–1059.

Suhaib Kh Hamed, Mohd Juzaiddin Ab Aziz, and Mohd Ridzwan Yaakub. 2023. Fake news detection model on social media by leveraging sentiment analysis of news content and emotion analysis of users' comments. *Sensors*, 23(4):1748.

Jochen Hartmann, Mark Heitmann, Christian Siebert, and Christina Schamp. 2023. More than a Feeling: Accuracy and Application of Sentiment Analysis. *International Journal of Research in Marketing*, 40(1):75–87.

Eero Hyvönen, Laura Sinikallio, Petri Leskinen, Senka Drobac, Rafael Leal, Matti La Mela, Jouni Tuominen, Henna Poikkimäki, and Heikki Rantala. 2024. Publishing and Using Parliamentary Linked Data on the Semantic Web: ParliamentSampo System for Parliament of Finland. In *Semantic Web*.

Minni Jain, Puneet Goel, Puneet Singla, and Rahul Tehlan. 2021. Comparison of various word embeddings for hate-speech detection. In *Data Analytics and Management: Proceedings of ICDAM*, pages 251–265. Springer.

Kimmo Kettunen and Matti La Mela. 2022. Semantic tagging and the Nordic tradition of everyman's rights. *Digital Scholarship in the Humanities*, 37(2):483–496.

Krister Lindén, Tommi Jauhiainen, and Sam Hardwick. 2020. FinnSentiment – a Finnish Social Media Corpus for Sentiment Polarity Annotation.

Alex Munyole Luvembe, Weimin Li, Shaohua Li, Fangfang Liu, and Guiqiong Xu. 2023. Dual emotion based fake news detection: A deep attention-weight update approach. *Information Processing Management*, 60(4):103354.

Andrew J Oswald. 1997. Happiness and economic performance. *The economic journal*, 107(445):1815–1831.

Jeffrey Pennington, Richard Socher, and Christopher D. Manning. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543.

Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations. *arXiv preprint*. ArXiv:1802.05365 [cs].

Ludovic Rheault. 2016. Expressions of Anxiety in Political Texts. In *Proceedings of the First Workshop on NLP and Computational Social Science*, pages 92–101, Austin, Texas. Association for Computational Linguistics.

Ludovic Rheault, Kaspar Beelen, Christopher Cochrane, and Graeme Hirst. 2016. Measuring Emotion in Parliamentary Debates with Automated Textual Analysis. *PLOS ONE*, 11(12):e0168843. Publisher: Public Library of Science.

Anna Ristilä and Kimmo Elo. 2023. Observing political and societal changes in Finnish parliamentary speech data, 1980–2010, with topic modelling. *Parliaments, Estates and Representation*, 43(2):149–176. Publisher: Routledge _eprint: https://doi.org/10.1080/02606755.2023.2213550.

Pedro L. Rodriguez and Arthur Spirling. 2022. Word Embeddings: What Works, What Doesn't, and How to Tell the Difference for Applied Research. *The Journal of Politics*, 84(1):101–115. Publisher: The University of Chicago Press.

Alieh Hajizadeh Saffar, Tiffany Katharine Mann, and Bahadorreza Ofoghi. 2023. Textual emotion detection in health: Advances and applications. *Journal of Biomedical Informatics*, 137:104258.

Statistics Finland. 2024a. Official Statistics of Finland (OSF): Annual national accounts [online publication]. https://www.stat.fi/en/statistics/vtp.

Statistics Finland. 2024b. Official Statistics of Finland (OSF): Consumer price index [online publication]. https://www.stat.fi/en/statistics/khi.

Statistics Finland. 2024c. Official Statistics of Finland (OSF): Labour force survey [online publication]. https://www.stat.fi/en/statistics/tyti.

Statistics Finland. 2024d. Official Statistics of Finland (OSF): Statistics on labour disputes [online publication]. https://www.stat.fi/en/statistics/tta.

Synonyymit.fi. n.d. Synonyymit.fi: Synonyymisanakirja netissä.

Dalila Talevi, Valentina Socci, Margherita Carai, Giulia Carnaghi, Serena Faleri, Edoardo Trebbi, Arianna di Bernardo, Francesco Capelli, and Francesca Pacitti. 2020. Mental health outcomes of the CoViD-19 pandemic. *Rivista di Psichiatria*, 55(3):137–144.

Peter C. Terry, Renée L. Parsons-Smith, and Victoria R. Terry. 2020. Mood Responses Associated With COVID-19 Restrictions. *Frontiers in Psychology*, 11. Publisher: Frontiers.

Martina Toshevska, Frosina Stojanovska, and Jovan Kalajdjieski. 2020. Comparative analysis of word embeddings for capturing word similarities. *arXiv preprint arXiv:2005.03812*.

Voikko. n.d. Voikko – Free linguistic software for Finnish. https://voikko.puimula.org/.

A Appendix

A.1 Word embedding parameters

The GloVe algorithm (Pennington et al., 2014) was used to generate word embeddings. Words with fewer than 50 occurrences were ignored to eliminate typos and rare words. Rheault et al. (2016) set this threshold at 200; due to the smaller size of the Finnish dataset, a lower threshold was selected here. The window size was set at 15, and the generated vectors are 300-dimensional. These parameters were chosen as they were used by Rheault et al. (2016); based on other research (Rodriguez and Spirling, 2022), the window size may be larger than necessary, but it did not pose a significant computational burden for this project.

A.2 Seed lemma selection

The method for selecting seed lemmas followed that used by (Rheault et al., 2016).

The initial positive words *hyvä*, *rakkaus*, and *onnellinen* (*good*, *love*, and *happy*), and negative words *paha*, *viha*, and *surullinen* (*bad*, *anger*, and *sad*) were chosen. Synonyms for these words were then recursively extracted from the online synonym dictionary Synonyymit.fi (n.d.). The resulting lists were filtered, first automatically to filter out any words that could not be lemmatised using *libvoikko*, and then manually. In the manual filtering, words were removed if they were not unambiguously either negative or positive, or if they were related to topics often discussed in parliament, such as “predator” (wolves are a recurring topic in the Finnish parliament) or “illness” (this was likely more common during the Covid-19 pandemic).

After filtering, the 100 most common words in each of the two lists were retained as seed lemmas. The word frequency data for this was obtained from the Frequency Lexicon of the Finnish Newspaper Language (CSC - IT Center for Science, 2004). As no general or political frequency lexicon was found, this lexicon was considered suitable for the task.

A.3 Seed lemmas

Lemma	PoS	Translation
hyvä	adj.	good
tärkeä	adj.	important
apu	noun	aid
voitto	noun	win
mahdollinen	adj.	possible
vahva	adj.	strong
selvä	adj.	clear
helppo	adj.	easy
merkittävä	adj.	significant
tyytyväinen	adj.	satisfied
palkinto	noun	reward
kaunis	adj.	beautiful
reilu	adj.	fair
hieno	adj.	fine
taito	noun	skill
myönteinen	adj.	positive
selkeä	adj.	clear
kiinnostunut	adj.	interested
arvostaa	verb	appreciate
tehokas	adj.	efficient
mukava	adj.	nice
mielenkiintoinen	adj.	interesting
hyöty	noun	utility
rakkaus	noun	love
ilo	noun	joy
turvallinen	adj.	safe
arvokas	adj.	valuable
kyky	noun	ability
tasainen	adj.	even
aktiivinen	adj.	active
luottamus	noun	trust
iloinen	adj.	joyful
suosikki	noun	favourite
täydellinen	adj.	perfect
aito	adj.	genuine
erinomainen	adj.	excellent
positiivinen	adj.	positive
edullinen	adj.	affordable, beneficial
rauhallinen	adj.	calm
upea	adj.	gorgeous
järkevä	adj.	rational
sankari	noun	hero
mahtava	adj.	mighty
rakastaa	verb	love
lahja	noun	gift
onnellinen	adj.	happy
into	noun	enthusiasm
voimavara	noun	asset
kiva	adj.	nice
viisas	adj.	wise

Lemma	PoS	Translation
tyydyttää	verb	satisfy
lämpö	noun	warmth
järki	noun	sense
yhtenäinen	adj.	united
vakaa	adj.	steady
sopu	noun	agreement
kirkas	adj.	clear
maltillinen	adj.	moderate
inhimillinen	adj.	humane
innokas	adj.	eager
vankka	adj.	robust
ihana	adj.	lovely
ainutlaatuinen	adj.	unique
ihailla	verb	admire
halukas	adj.	eager
huima	adj.	wild
palkkio	noun	reward
huikea	adj.	breathtaking
ylpeä	adj.	proud
suosittu	adj.	popular
luotettava	adj.	reliable
yhtyä	verb	unite
innostunut	adj.	enthusiastic
kultainen	adj.	golden
rakas	adj.	beloved
tykätä	verb	like
oleellinen	adj.	essential
laadukas	adj.	high quality
luonteva	adj.	natural
yhteisymmärrys	noun	understanding
riemu	noun	jubilation
reipas	adj.	brisk, cheerful
armo	noun	mercy
ahkera	adj.	diligent
lahjakas	adj.	talented
perusteellinen	adj.	comprehensive
rehellinen	adj.	honest
kunnollinen	adj.	proper
asiallinen	adj.	proper
plussa	noun	plus
tyylikäs	adj.	stylish
uskottava	adj.	credible
siisti	adj.	neat
uskollinen	adj.	loyal
viisaus	noun	wisdom
ystävyys	noun	friendship
ehjä	adj.	intact
hyödyllinen	adj.	useful
intohimo	noun	passion
mieluisa	adj.	pleasing

Table 1: Positive seed lemmas used.

Lemma	PoS	Translation
vaikea	adj.	difficult
paha	adj.	bad, evil
huono	adj.	bad
kriisi	noun	crisis
kaivata	verb	yearn
vaikeus	noun	difficulty
virhe	noun	error
heikko	adj.	weak
vastainen	adj.	against
vahinko	noun	damage
loukkaantua	verb	be offended
mahdoton	adj.	impossible
pelko	noun	fear
vaarallinen	adj.	dangerous
hankala	adj.	difficult
hyökkäys	noun	attack
raju	adj.	fierce
alhainen	adj.	low
pettymys	noun	disappointment
vaativa	adj.	demanding
väärin	seikkasana	wrong
vaiva	noun	inconvenience
kielteinen	adj.	negative
ankara	adj.	strict
kriittinen	adj.	critical
rankka	adj.	tough
synkkä	adj.	gloomy
syyllinen	adj.	guilty
ahdas	adj.	cramped
suru	noun	grief
kohtuuton	adj.	unreasonable
tuska	noun	pain
kiusata	verb	bully
yksinäinen	adj.	lonely
väkivaltainen	adj.	violent
väkivaltaisuus	noun	violence
ongelmallinen	adj.	problematic
tarpeeton	adj.	superfluous
harmi	noun	harm
surkea	adj.	poor
kehno	adj.	bad
kauhu	noun	dread
tuhoisa	adj.	disastrous
heikkous	noun	weakness
murhe	noun	grief
surullinen	adj.	sad
viha	noun	hate
negatiivinen	adj.	negative
kiusallinen	adj.	awkward
virheellinen	adj.	incorrect

Lemma	PoS	Translation
puutteellinen	adj.	inadequate
karu	adj.	barren
tylsä	adj.	boring
tyytymätön	adj.	unsatisfied
epäkohta	noun	fault
haukkua	verb	insult
haitallinen	adj.	damaging
katkera	adj.	bitter
verinen	adj.	bloody
likainen	adj.	dirty
ahdistus	noun	anxiety
onneton	adj.	unhappy
liiallinen	adj.	excessive
julma	adj.	barbarous
surra	verb	mourn
helvetti	noun	hell
tyly	adj.	rude, harsh
traaginen	adj.	tragic
tyytymättömyys	noun	dissatisfaction
moite	noun	reproach
vihainen	adj.	angry
huonokuntoinen	adj.	in poor condition
syrjäytyä	verb	become alienated
työläs	adj.	arduous
synti	noun	sin
karkea	adj.	rough
erehdys	noun	mistake
toivoton	adj.	hopeless
julmuus	noun	cruelty
aggressiivinen	adj.	aggressive
kurja	adj.	miserable
hävetä	verb	be ashamed
rasite	noun	encumbrance
riesa	noun	nuisance
ankea	adj.	bleak
myrkyllinen	adj.	toxic
levoton	adj.	restless
ilkivalta	noun	vandalism
kaipuu	noun	yearning
epätoivoinen	adj.	desperate
armoton	adj.	merciless
jäykkä	adj.	stiff
tappelu	noun	fight
sopimaton	adj.	improper
riittämätön	adj.	insufficient
mitätön	adj.	puny
kurjuus	noun	misery
typerä	adj.	stupid
vääryys	noun	injustice
turhautua	verb	get frustrated

Table 2: Negative seed lemmas used.

A.4 Additional tables and figures

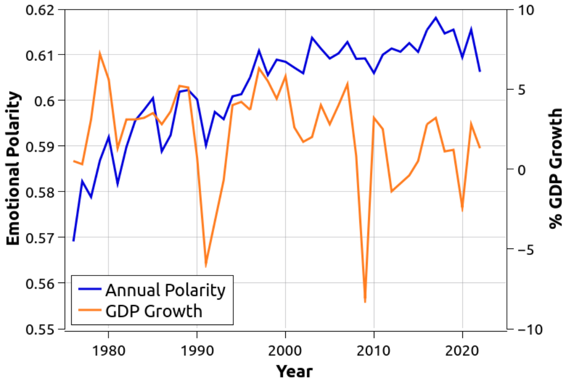

Figure 5: GDP growth and emotional polarity.

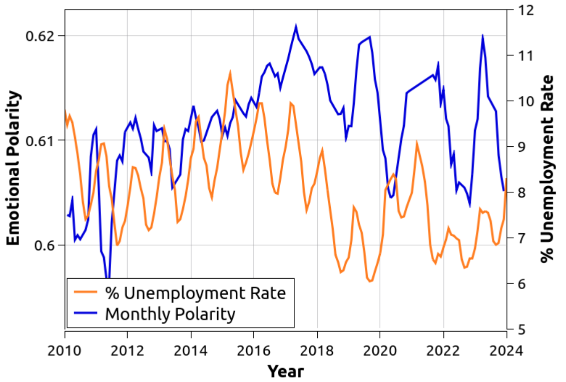

Figure 6: Unemployment rate and emotional polarity (4-month rolling average).

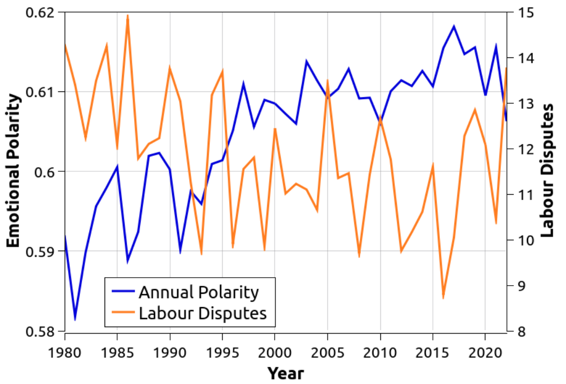

Figure 7: Labour disputes and emotional polarity.

Name of Measure	Description	Time Period	Database
CLI Growth	% change in Cost of Living Index, annual	1939-2023	2024b
CPI Growth	% change in Consumer Price Index, annual	1952-2023	2024b
GDP Growth	% change in Gross Domestic Product, annual	1976-2022	2024a
Labour Disputes	Natural logarithm of days lost to strikes, annual	1980-2022	2024d
Misery Index	Sum of CPI growth and unemployment, monthly	2010-2024	2024b; 2024c
Unemployment	Rate of unemployment, monthly	2010-2024	2024c

Table 3: Economic measures used.

Cause	Effect	P	Corrected P
Polarity	CLI Growth	0.930	1.000
Polarity	CPI Growth	0.537	1.000
Polarity	GDP Growth	0.442	1.000
Polarity	Labour Disputes	0.091	1.000
Polarity	Misery Index	0.061	0.732
Polarity	Pos. Covid Tests	0.350	1.000
CLI Growth	Polarity	0.003	0.034
CPI Growth	Polarity	0.040	0.484
GDP Growth	Polarity	0.215	1.000
Labour Disputes	Polarity	0.157	1.000
Misery Index	Polarity	0.245	1.000
Pos. Covid Tests	Polarity	0.002	0.029

Table 4: Results of the Granger causality tests. Rows with a corrected p-value below 0.05 are highlighted.