Online Appendix: Do emotions affect political trust judgments?

Table of contents

A1 Descriptive statistics and self-selection analysis **2**

A2 Sentiment and emotion distributions by demographics **5**

A3 Sentiment classification: Evaluation of automated approaches, alternative dictionary approaches, and linear model results **7**
A3.1 Evaluation of automated approaches: BERT vs. GPT 7
A3.2 Alternative dictionary approaches 8
A3.3 Linear model sentiment and trust: Regression coefficients 10
A3.4 (Alternative) Ordinal logistic regression: Political trust on sentiment (GPT) 11
A3.5 (Alternative) Linear regression: Political trust on sentiment (BERT) 12

A4 Emotion classification: Regression coefficients and alternative classification approaches **13**
A4.1 Linear model emotion and trust: Regression coefficients 13
A4.2 Alternative emotion classification: NRC word-emotion lexicon 13

A5 Prompting strategy **15**

A6 Question wording and answer categories **15**

A7 Robustness check: External events during the survey period **17**

A8 References **19**

A1 Descriptive statistics and self-selection analysis

In this section we provide both descriptive statistics for the overall sample as well as a self-selection analysis of providing open-ended responses.

Table A1: Descriptive statistics of key variables (full sample).

Variable	Mean	SD	N
Age	45.31	15.45	1,478
Education	4.48	1.22	1,476
Socioeconomic Status	5.88	1.75	1,478
Trust in Government	2.40	0.92	1,464
Trust in Parties	6.21	2.40	1,478
Trust in Parliament	5.98	2.73	1,478
Trust in Elections	4.39	2.20	1,478
Satisfaction with Democracy	5.79	2.58	1,478
Political Interest	1.94	0.82	1,478
Life Satisfaction	6.97	2.41	1,478
Financial Situation	5.43	2.88	1,478

Table A1 provides an overview of the key variables in the full sample (n=1,478). Of the total sample of n=1,478 respondents, n=491 (33%) provided an open-ended answer, while n=987 (67%) did not respond to the open-ended question or provided unusable responses. A potential concern with the voluntary open-ended voice question is that respondents who chose to answer might systematically differ from those who did not (self-selection bias). For instance, particularly emotional respondents may be more inclined to provide an answer, or conversely, less emotional respondents may skip it. To address this concern, we compare the two groups across key demographic variables, trust measures, and other relevant characteristics. Our self-selection analysis compared 491 respondents who provided open-ended answers with 987 who did not, across several key variables. Table A2 provides comparisons between respondents who did and did not provide open-ended answers. Generally, we do not find substantively meaningful differences between the two groups. Out of 11 variables, only socioeconomic status reaches statistical significance (p=0.033), with respondents who provided an answer reporting slightly lower socioeconomic status. However, the corresponding effect size is negligible (Cohen's d=-0.12). Most importantly, none of the trust measures – which are the central variables in our analysis – differ significantly between the two groups, alleviating the concern that self-selection into the open-ended question is related to the outcome of interest. While one statistically significant difference was detected (socioeconomic status), all effect sizes remain small or negligible (all $|d| \leq 0.12$). Crucially, the two groups do not differ on any of the political trust variables, which are the outcome variables in our main analyses. Furthermore, political interest – which could serve as a rough proxy for motivation to engage with a political survey question – also does not differ between the groups. Taken together, these results suggest that self-selection bias is minimal and unlikely to substantially affect the validity of our main findings. Figure A1 and Figure A2 show the

distribution of demographic and trust variables by response status visually highlighting that there do not seem to be relevant selection effects.

Table A2: Comparison of key variables between respondents who did and did not provide open-ended answers.

	No Resp. (n=987)		**Yes Resp.** (n=491)		**Statistical Test**			
Variable[1]	Mean	SD	Mean	SD	t	p-value	d	
Age	44.95	15.50	46.02	15.34	−1.26	0.207	0.07	
Education	4.49	1.22	4.46	1.22	0.49	0.622	−0.03	
Socioeconomic Status	5.95	1.76	5.74	1.73	2.13	0.033	−0.12	*
Trust in Government	2.41	0.93	2.38	0.89	0.56	0.575	−0.03	
Trust in Parties	6.18	2.41	6.29	2.37	−0.81	0.417	0.04	
Trust in Parliament	5.91	2.71	6.12	2.76	−1.39	0.166	0.08	
Trust in Elections	4.39	2.25	4.40	2.11	−0.05	0.960	0.00	
Satisfaction with Democracy	5.77	2.59	5.83	2.55	−0.44	0.657	0.02	
Political Interest	1.93	0.82	1.98	0.83	−1.18	0.240	0.07	
Life Satisfaction	6.91	2.43	7.09	2.38	−1.39	0.165	0.08	
Financial Situation	5.34	2.86	5.63	2.92	−1.80	0.072	0.10	

[1]Rows highlighted indicate statistically significant differences ($p < 0.05$); Effect sizes are Cohen's d values.

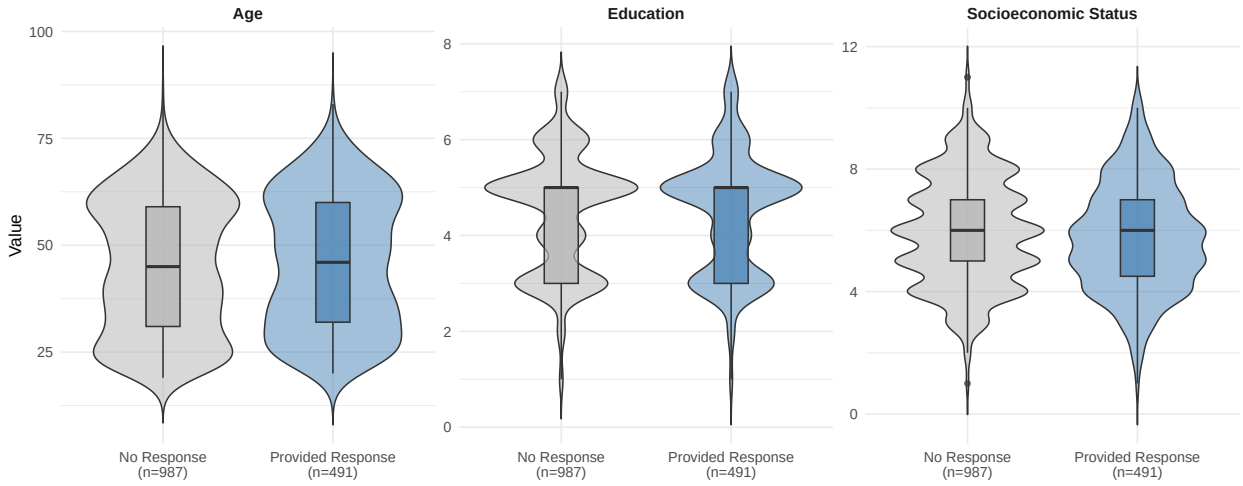


Figure A1: Distribution of demographic variables by response status. Violin plots show the full distribution, with boxplots overlaid to show median and quartiles.

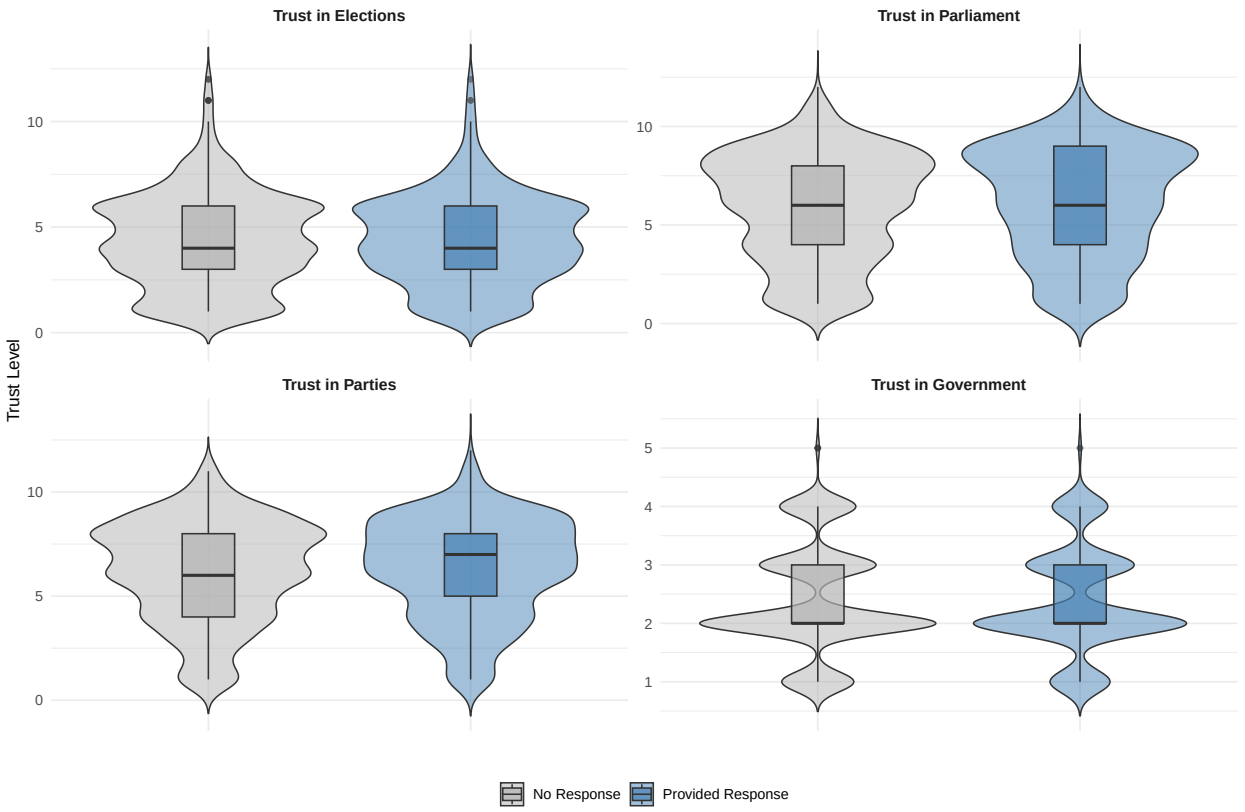


Figure A2: Distribution of political trust measures by response status. Higher values indicate greater trust.

A2 Sentiment and emotion distributions by demographics

To illustrate who expresses different types of affect, we present cross-tabulations of sentiment (GPT classification) and emotion (SpeechBrain classification) by party preference, age group, and education level.

Regarding sentiment, the most notable pattern is partisan (Table A3): Republican and Independent respondents express substantially more negative sentiment (69% and 68%, respectively) than Democrats (53%), while Democrats show a higher share of positive sentiment (10% vs. 6% for Republicans). This is consistent with the survey being conducted during the Biden presidency, when Republican respondents – as supporters of the opposition party – would be expected to express more negative evaluations of the federal government. Differences across age groups (Table A4) and education levels (Table A5) are comparatively modest.

For the emotion classification, the dominant pattern is the high share of neutral classifications across all groups (74–88%), reflecting the limitations of the four-category SpeechBrain taxonomy discussed in the main text. The most notable variation is by age (Table A7): anger increases markedly with age (1% among 18–30 year-olds vs. 20% among respondents aged 59+), while sadness is most prevalent among the youngest group (7%). Partisan differences in emotion are small (Table A6), with anger slightly more common among Republicans and Independents (11% and 10%) than Democrats (8%). Differences by education (Table A8) are similarly modest. These patterns should be interpreted cautiously given the small absolute numbers of non-neutral emotion classifications.

Table A3: Sentiment (GPT) by party preference (column percentages)

	Republican	Democrat	Independent
negative	68.6	53.4	68.4
neutral	25.7	36.6	23.5
positive	5.7	9.9	8.1

Table A4: Sentiment (GPT) by age group (column percentages)

	18-30	31-45	46-58	59+
negative	66.0	59.0	63.2	59.4
neutral	32.0	34.5	23.7	30.4
positive	2.0	6.5	13.2	10.1

Table A5: Sentiment (GPT) by education (column percentages)

	High school or less	Bachelor	Graduate
negative	61.5	63.6	53.9
neutral	31.5	27.2	39.5

positive	7.0	9.2	6.6

Table A6: Emotion (SpeechBrain) by party preference (column percentages)

	Republican	Democrat	Independent
Anger	10.5	8.2	10.3
Happiness	4.8	4.7	2.9
Neutral	81.0	84.1	81.6
Sadness	3.8	3.0	5.1

Table A7: Emotion (SpeechBrain) by age group (column percentages)

	18-30	31-45	46-58	59+
Anger	1.0	3.6	9.6	19.6
Happiness	4.0	5.8	4.4	2.9
Neutral	88.0	87.8	84.2	73.9
Sadness	7.0	2.9	1.8	3.6

Table A8: Emotion (SpeechBrain) by education (column percentages)

	High school or less	Bachelor	Graduate
Anger	8.4	9.6	7.9
Happiness	2.8	4.8	5.3
Neutral	85.3	80.9	86.8
Sadness	3.5	4.8	0.0

A3 Sentiment classification: Evaluation of automated approaches, alternative dictionary approaches, and linear model results

A3.1 Evaluation of automated approaches: BERT vs. GPT

In this subsection of the Appendix, we assess the performance of the two automated sentiment classification approaches introduced in the main paper: BERT and GPT. Table A9 presents the confusion matrix for GPT and BERT sentiment classification. Notably, when these classifiers diverge, their discrepancies are limited to the neutral category, except for a single instance where BERT identifies positive sentiment and GPT identifies negative sentiment. Both classifiers exhibit an overall agreement of 79%.

Table A9: Confusion Matrix: Sentiment Classification GPT vs. BERT

	negative	neutral	positive
negative	255	36	0
neutral	46	104	12
positive	1	9	28

To systematically evaluate the performance of the two sentiment classifications presented in our paper, we manually annotated a randomly drawn subset of the data (n=197). This dataset was generated by an independent human coder, given no specific instructions beyond categorizing survey responses into positive, negative, or neutral.
By using the human-annotated labels as a reference, we calculated the accuracy (number of correct predictions / total number of predictions) for the two automated approaches, GPT and BERT. Our BERT classification result in an agreement of 78%. In contrast, the GPT classification exhibited slightly better higher accuracy with an agreement of 81% in relation to our manually annotated dataset.
Table A10 illustrates the distributions over categories for the three classifiers. Notably, instances of disagreement between the human and the automated approach predominantly involve negative versus neutral assignments.

Table A10: Sentiment Classification: BERT vs. GPT vs. manual

Category	BERT	GPT	manual
negative	117	112	99
neutral	66	69	87
positive	14	16	11

A3.2 Alternative dictionary approaches

Sentiment Classification can be achieved through a variety of modeling architectures. In the main paper we pursued approaches with Deep Learning Architectures (BERT and GPT), however in this appendix we wanted to provide results from a more simple approach, namely dictionary approaches to this task. Historically, the analysis of sentiment in survey responses was achieved with theory-based dictionaries of affectively scored words (Mossholder et al. 1995) and still nowadays there are various, popular dictionaries available. For this appendix, we utilize two of the most popular ones: AFINN (Nielsen 2011) and VADER (Hutto and Gilbert 2014). Both AFINN and VADER are open-source lexicons, however VADER includes more capabilities as it is more comprehensive, context-aware, and granular than AFINN. For example, it is able to correctly handle negations as it incorporates knowledge about grammatical rules. Additionally, VADER is better suited for analyzing short texts (e.g., survey answers), since it is specifically attuned to sentiment expressed in social media.
Since both dictionaries operate on different outcome scales (i.e., -5,5 for AFINN and -1,1 for VADER) we create a categorical outcome variable utilizing rule-based categorization: Assigning '0' as neutral, categorizing the lower 50% of negative scores as highly negative and the upper 50% as slightly negative. Conversely, the upper 50% of positive scores were designated as highly positive, while the lower 50% were marked as slightly positive.

Figure A3 depicts the shares of sentiment classifications detected by modality (AFINN vs. VADER). The two classifiers show slight variances, but both convey the impression that the data consists of mainly positive sentiment (70% for AFINN, and 60% for VADER). Additionally, the correlation between the sentiment categories and political trust scores, as indicated by the red line in the plot, demonstrates a slight increase in political trust as response sentiment is more positive.

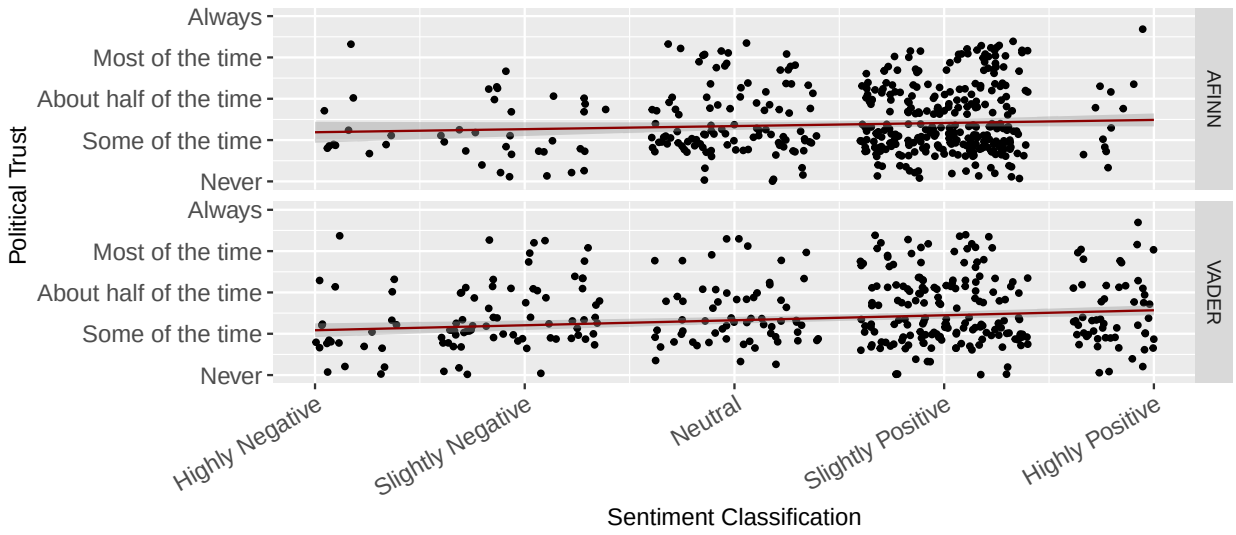


Figure A3: Sentiment Classification with five categories by classifier (AFINN vs. VADER). *Note*: Analysis of n=491 open-ended answers, "afinn" achieved using the valence-based AFINN dictionary, "Vader" achieved via Vader compound score, 5-category breaks were achieved with '0' as neutral, categorizing the lower 50% of negative scores as highly negative and the upper 50% as slightly negative, the upper 50% of positive scores as highly positive, while the lower 50% as slightly positive.

In summary, these findings deviate from the results outlined in the main paper, primarily attributed to the dictionary approaches displaying a higher prevalence of neutral and positive sentiment classifications. This discrepancy may stem from various factors. Notably, dictionary approaches struggle to identify sentiment when the response contains terms absent from the dictionary. For example, in survey responses like "Joe Biden is a pedophile. Hillary Clinton is a satanic," both dictionary methods failed to detect any sentiment, leading to their classification as neutral in our coding procedure. In contrast, BERT and GPT successfully identified the unequivocally negative sentiment in such responses.

Dictionary-based classifiers "work in cases where clearly defined sets of words indicate the presence of particular content but struggle with nuance and generalization" (Rytting et al. 2023, 2). One set of solutions to these challenges stem from the field of machine learning and in particular supervised machine learning such as naive bayes, random forests, and SVMs have been shown to work well for sentiment classification. The downside however is that they require large datasets of human-annotated training data. Most recent developments and the emergence of large language models that are based on deep learning circumvent even this necessity and the use of fine-tuned deep learning models (fine-tuning only requires a small set of fine-tinung data) or even zero-shot prompting "requires no training data and minimal programming to implement" (Burnham 2023, 2), however, "[u]nlike sentiment dictionaries, it produces results comparable to, and sometimes better than, supervised classification" (Burnham 2023, 2). Learning-based approaches (Bello, Ng, and Leung (2023); Chiorrini et al. (2021) use BERT models for sentiment classification of Twitter Data; Sailunaz and Alhajj (2019) use Naive Bayes), such as large deep learning models have shown better performance

in terms of sentiment classification compared to lexicon-based approaches (Bharti et al. 2022; Chiorrini et al. 2021). Their largest contribution might be that they move away from a simple bag-of-words approach but instead the context of the respective words. Eventually, similarly to the preceding subsection, we computed the degree of agreement between these dictionary-generated classifications and the manually annotated ones. This analysis serves as a demonstration that our more complex, deep-learning-based classifiers are significantly more effective in capturing nuanced sentiment.

Results from using AFINN compared to the reference of or manually annotated data, achieves an agreement of 21%. For VADER, the same comparison results in an agreement of 26%.

A3.3 Linear model sentiment and trust: Regression coefficients

Table A11 shows regression model estimates underlying Figure 3 in the main paper.

Table A11: Linear Model: Sentiment (GPT) and Political Trust

	Bivariate	Multivariate
(Intercept)	2.040***	2.036***
	(0.041)	(0.239)
sentiment_GPTneutral	0.625***	0.574***
	(0.072)	(0.071)
sentiment_GPTpositive	1.860***	1.773***
	(0.121)	(0.119)
age		0.003
		(0.002)
education		0.024
		(0.030)
socioeconomic_status		−0.067**
		(0.021)
party_preferenceDemocrat		0.310***
		(0.083)
party_preferenceIndependent		0.037
		(0.090)
party_preferenceOther		0.624*
		(0.295)
party_preferenceNo preference		0.236
		(0.230)
party_preferencePrefer not to answer		−0.436
		(0.499)
Num.Obs.	491	491
R2	0.356	0.410
R2 Adj.	0.353	0.398

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

A3.4 (Alternative) Ordinal logistic regression: Political trust on sentiment (GPT)

Figure 3 in the main paper showed findings from regression analysis of sentiment on trust scores. This appendix provides robustness measures for an alternative model specification, namely findings for an ordinal (or ordered) logistic regression. Ordinal logistic regression is a suitable model when outcome variables are ordinal in nature, meaning they possess a meaningful and ordered categorical structure, but the intervals between categories are not assumed to be equal – which might be the case for our 5-category political trust item.
Figure A4 shows predicted probabilities for each of the outcome categories (e.g., trust-never (1), trust-some of the time (2), etc) by predicting different intercepts and slopes for each category.

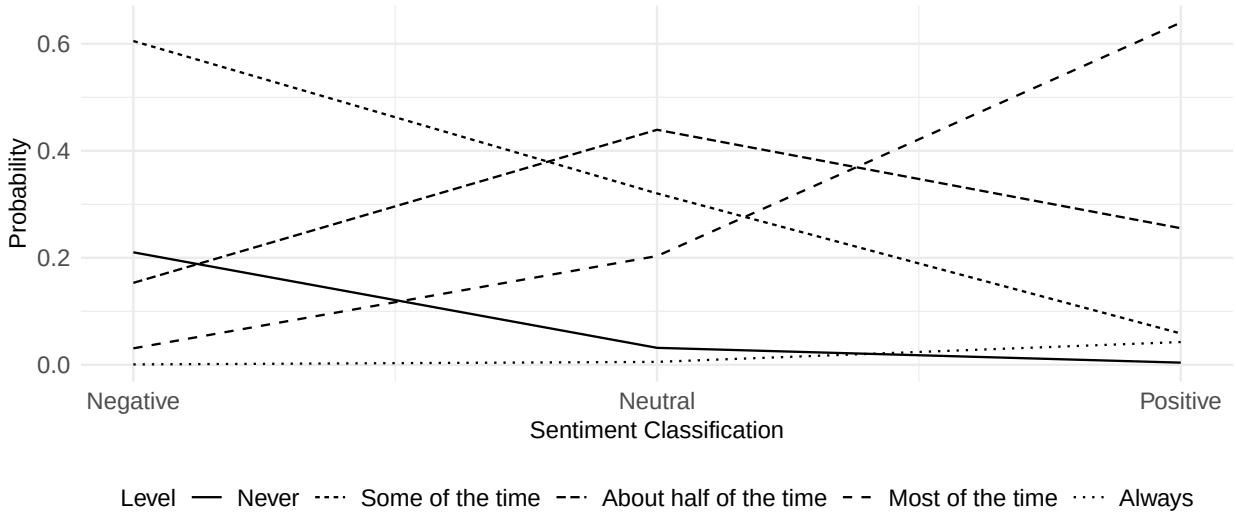


Figure A4: Ordered logit model of sentiment and a five-category trust score (bivariate).

In the main paper we intentionally chose to treat our outcome variable as continuous due to various reasons. Firstly, it enhances the overall readability and accessibility of our research findings as we choose the simplest model that is appropriate. Secondly, it. yields similar results as an ordered logistic model and thirdly, our decision aligns with previous research that describes the understanding that it "[…] has become common practice to assume that Likert-type categories constitute interval-level rather than ordinal-level measurement" (Blaikie 2003, 231). Monte Carlo Simulations have also shown that parametric tests, such as a F-Test in a linear regression, was strongly robust to the interval data assumption (as well as moderate skewing) when data was collected using a 5 to 7 point Likert response format with no resulting bias (Carifio and Perla 2007; Glass, Peckham, and Sanders 1972).

A3.5 (Alternative) Linear regression: Political trust on sentiment (BERT)

Figure A5 shows findings for a linear regression of political trust (5 categories) on our sentiment variable achieved with BERT. Figure A5 is the counterpart to Figure 3 in the main paper.

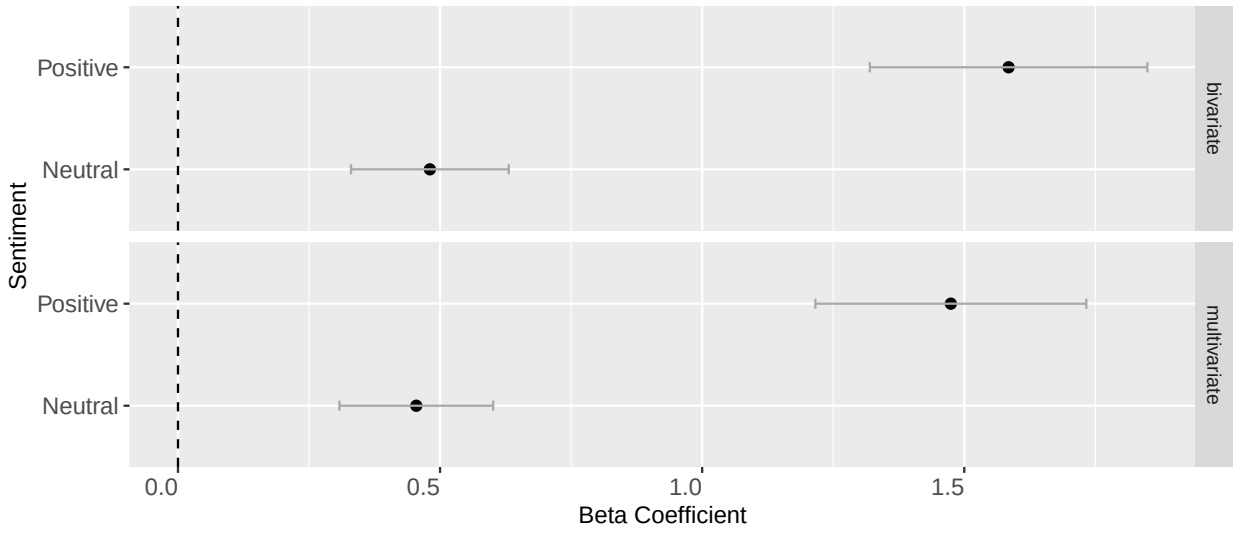


Figure A5: Linear model of sentiment (BERT) and a five-category trust score (bi- and multivariate). *Note*: The data for the sentiment classification displayed in this figure stems from the BERT-based classification. Respective findings for the GPT classification can be found in the main paper.

A4 Emotion classification: Regression coefficients and alternative classification approaches

A4.1 Linear model emotion and trust: Regression coefficients

Table A12 shows regression model estimates underlying Figure 5 in the main paper.

Table A12: Linear Model: Emotion and Political Trust

	Bivariate	Multivariate
(Intercept)	2.360***	2.290***
	(0.044)	(0.292)
speechbrain_classification['ang']	−0.019	−0.118
	(0.141)	(0.139)
speechbrain_classification['sad']	−0.027	0.053
	(0.214)	(0.206)
speechbrain_classification['hap']	0.544**	0.472*
	(0.199)	(0.191)
age		0.007**
		(0.003)
education		0.006
		(0.036)
socioeconomic_status		−0.090***
		(0.025)
party_preferenceDemocrat		0.456***
		(0.101)
party_preferenceIndependent		0.081
		(0.111)
party_preferenceOther		0.492
		(0.362)
party_preferenceNo preference		0.128
		(0.282)
party_preferencePrefer not to answer		−0.033
		(0.609)
Num.Obs.	491	491
R2	0.015	0.116
R2 Adj.	0.009	0.096

+ p < 0.1, * p < 0.05, ** p < 0.01, *** p < 0.001

A4.2 Alternative emotion classification: NRC word-emotion lexicon

For the task of emotion classification, in the main paper, we chose to analyze the original, "raw" audio files, as they might contain paralinguistic features, such as intonation, pitch, volume, pauses, but also laughter or breathing noises. However, one can also research text (i.e., our transcripts from the audio files) for emotions. Despite the lack of paralinguistic characteristics in text, they can still convey important information since the emotional

essence of an utterance can also be significantly conveyed through the choice of words. For example, words such as "lovely" and "awesome" carry strong emotions compared to more generic, non-emotional words, such as "person" and "day" (Yoon, Byun, and Jung 2018, 1) and a sentence like "This phone is a piece of junk" has a stronger valence than "I think this phone is fine" (Liu 2010, 632). Hence, we additionally classified emotions using a text-based emotion recognition classifier, namely the NRC word-emotion lexicon (Mohammad and Turney 2013). While there are multiple off-the-shelf methods available to achieve this task, for example pysentimiento's emotion classifier (Pérez et al. 2023), the NRC lexicon can be considered to be one of the most popular resources for the analysis of emotions in texts (Nandwani and Verma 2021). It contains associations of words with eight basic emotions: joy, surprise, sadness, anger, disgust, fear, contempt, and anticipation. By analyzing a given text, words are assigned to the corresponding word-emotion pairs in the dictionary. Since each word is assigned to a specific emotion (in some cases more than one), we decided to use a multi-class classification. In this multi-class classification each sentence can be assigned multiple emotions at the same time (only the neutral category is of course exclusive). Figure A6 displays the findings from our emotion classification results obtained through applying the NRC lexicon. For comparison, Figure 4 in the main paper displays the audio-based analysis using SpeechBrain.

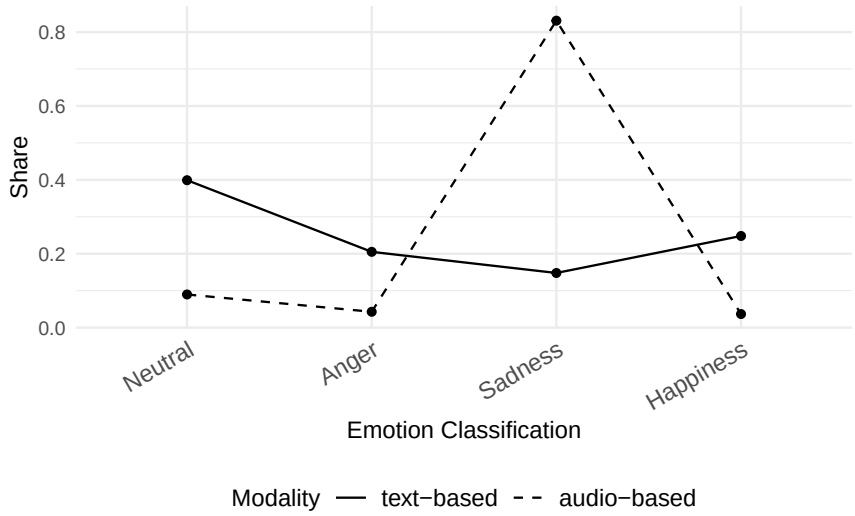


Figure A6: Emotion Classification obtained from EmoLex. *Note*: Text-based analysis achieved with the NRC Emotion Lexicon, Audio-based analysis achieved with SpeechBrain. Classification of n=491 open-ended answers, note that text-based NRC classifier allows for multi-class assignment.

Figure A6 demonstrates that irrespective of modeling approach the majority of open-ended responses do not contain explicit emotional language and are characterized by a neutral emotional tone. However and notably, the share of neutral statements is significantly higher in the audio-based condition compared to the text-based one (85% vs. 39%). In other words, the text-based analysis identifies a greater share of emotions within the survey answers.

A5 Prompting strategy

This appendix aims at describing some of the details we considered when drafting the prompt for our zero-shot classification of sentiment using GPT-3.5-turbo. As was described in the main paper, our prompt aimed at provide a very clear and straightforward instruction which returns a single token (i.e, positive, negative, neutral) to detect the nature of associations in its most standard form: positive, negative or neutral:. Our prompt read as follows: | "Classify the sentiment of the following open-ended survey answer into neutral, negative or positive. Text: i Sentiment:", where i is the survey response. Since sentiment is a comparatively easy task, we refrained from providing specific examples, which allowed us to use zero-shot (compared to few-shot) prompting. Prompt engineering plays a crucial role in influencing the final results, as slight modifications in the prompt, such as formatting changes, can alter the probability distribution over tokens and, consequently, impact the outcome. Nevertheless, in a series of experiments, Rytting et al. 2023 demonstrated that, on the whole, prompt engineering has limited influence on GPT-3's performance when coding social science datasets. Their findings emphasize the importance of selecting unique first tokens for each category, avoiding overly extreme descriptors like "very positive" and "very negative" in favor of more neutral terms like "positive" and "negative." Moreover, it is recommended to employ substantive alternatives in the prompt, replacing binary choices like "yes" and "no" with more meaningful options such as "positive" and "negative." Interestingly, the format of the prompt itself does not appear to significantly impact results. This includes practices like enclosing categories in quotes or using various delimiters such as slashes or pipes to separate the prompt task and input. Such findings underscore the relative stability of the information retrieval process in the face of diverse prompt formats.

A6 Question wording and answer categories

Table A13: Survey variables, question wording, and answer categories with numeric coding.

Variable Name	Question Wording	Answer Categories
`trust_in_government`	How often can you trust the federal government in Washington to do what is right?	Always = 5; Most of the time = 4; About half of the time = 3; Some of the time = 2; Never = 1; Don't know = NA
`political_trust_parties`	political parties	0 = No trust at all to 10 = Complete trust; Don't know = NA
`political_trust_parliament`	parliament	0 = No trust at all to 10 = Complete trust; Don't know = NA

Variable Name	Question Wording	Answer Categories
political_trust_elections	elections	0 = No trust at all to 10 = Complete trust; Don't know = NA
satisfaction_democracy	On the whole, how satisfied are you with the way democracy works in the United States? 0 means extremely dissatisfied and 10 means extremely satisfied.	0 = Extremely dissatisfied to 10 = Extremely satisfied; Don't know = NA
party_preference	In this final section, we kindly request your assistance in providing us with some additional information related to you. Generally speaking, do you usually think of yourself as a Republican, Democrat, Independent, or what?	Republican = 1; Democrat = 2; Independent = 3; Other = 4; No preference = 5; Don't know = NA; Prefer not to answer = NA
political_interest	How interested would you say you are in politics?	Very interested = 1; Somewhat interested = 2; Not very interested = 3; Not at all interested = 4; Don't know = NA
financial_situation	Overall, thinking of your assets, debts and savings, how satisfied are you with your current personal financial condition? 0 means extremely dissatisfied and 10 means extremely satisfied.	0 = Extremely dissatisfied to 10 = Extremely satisfied; Don't know = NA
life_satisfaction	All things considered, how satisfied are you with your life as a whole these days? 0 means extremely dissatisfied and 10 means extremely satisfied.	0 = Extremely dissatisfied to 10 = Extremely satisfied; Don't know = NA

Variable Name	Question Wording	Answer Categories
`education`	What is the highest level of education you have completed?	Less than a high school diploma = 1; GED = 2; High school diploma = 3; Associates degree = 4; Bachelor's degree = 5; Master's degree = 6; Doctorate = 7; Don't know = NA
`age`	What is your current age in years?	*[Numeric Input: 18-99]*
`socioeconomic_status`	Think of a ladder as representing where people stand in society… Where would you put yourself on the socioeconomic ladder?	10 = Top; 9; 8; 7; 6; 5; 4; 3; 2; 1 = Bottom; Prefer not to answer = NA

A7 Robustness check: External events during the survey period

Our survey was fielded from September 6 to October 27, 2023, a period that included two notable political events: the announcement of impeachment proceedings against President Biden (September 12) and the start of the Hamas–Israel war (October 7). To assess whether these events influenced our main findings, we re-estimate the sentiment and emotion regression models with two binary indicators: one for responses collected on or after September 12 (post-impeachment announcement) and one for responses collected on or after October 7 (post-Hamas–Israel war). In our analytical subsample, 390 responses were collected before September 12, 67 between September 12 and October 6, and 34 on or after October 7.

Table A14: Robustness check: Sentiment (GPT) and political trust, controlling for survey period

	Post Sep 12	Post Oct 7
(Intercept)	2.039***	2.047***
	(0.244)	(0.242)
sentiment_GPTneutral	0.574***	0.574***
	(0.071)	(0.071)
sentiment_GPTpositive	1.774***	1.774***
	(0.119)	(0.119)
age	0.003	0.003
	(0.002)	(0.002)
education	0.024	0.024
	(0.030)	(0.030)
socioeconomic_status	−0.067**	−0.067**
	(0.021)	(0.021)
party_preferenceDemocrat	0.309***	0.307***
	(0.083)	(0.083)
party_preferenceIndependent	0.036	0.034
	(0.091)	(0.091)
party_preferenceOther	0.623*	0.616*
	(0.296)	(0.296)
party_preferenceNo preference	0.236	0.232
	(0.230)	(0.231)
party_preferencePrefer not to answer	−0.438	−0.437
	(0.502)	(0.499)
post_sep12	0.005	
	(0.093)	
post_oct7		0.037
		(0.132)
Num.Obs.	491	491
R2	0.410	0.411
R2 Adj.	0.397	0.397

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A15: Robustness check: Emotion (SpeechBrain) and political trust, controlling for survey period

	Post Sep 12	Post Oct 7
(Intercept)	2.238***	2.293***
	(0.298)	(0.296)
speechbrain_classification['ang']	−0.118	−0.118
	(0.139)	(0.140)
speechbrain_classification['hap']	0.466*	0.471*
	(0.192)	(0.192)
speechbrain_classification['sad']	0.069	0.053
	(0.206)	(0.206)
age	0.008**	0.007*
	(0.003)	(0.003)
education	0.007	0.006
	(0.036)	(0.036)
socioeconomic_status	−0.090***	−0.090***
	(0.025)	(0.025)
party_preferenceDemocrat	0.464***	0.456***
	(0.101)	(0.101)
party_preferenceIndependent	0.087	0.081
	(0.111)	(0.112)
party_preferenceOther	0.509	0.490
	(0.363)	(0.364)
party_preferenceNo preference	0.137	0.127
	(0.283)	(0.283)
party_preferencePrefer not to answer	0.016	−0.033
	(0.611)	(0.609)
post_sep12	−0.099	
	(0.114)	
post_oct7		0.011
		(0.161)
Num.Obs.	491	491
R2	0.118	0.116
R2 Adj.	0.096	0.094

+ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table A14 and Table A15 show that the main findings are robust to the inclusion of both period indicators. The coefficients for sentiment and emotion categories remain substantively unchanged compared to the main models reported in the paper, and neither the post-September 12 nor the post-October 7 indicator is statistically significant, suggesting that these external events did not systematically alter the sentiment–trust or emotion–trust relationships in our data.

A8 References

Bello, Abayomi, Sin-Chun Ng, and Man-Fai Leung. 2023. "A BERT Framework to Sentiment Analysis of Tweets." *Sensors* 23 (1): 506.

Bharti, Santosh Kumar, S Varadhaganapathy, Rajeev Kumar Gupta, Prashant Kumar Shukla, Mohamed Bouye, Simon Karanja Hingaa, and Amena Mahmoud. 2022. "Text-Based Emotion Recognition Using Deep Learning Approach." *Comput. Intell. Neurosci.* 2022: 1–8.

Blaikie, Norman. 2003. *Analyzing Quantitative Data*. SAGE Publications Ltd.

Burnham, Michael. 2023. "Stance Detection with Supervised, Zero-Shot, and Few-Shot Applications." *arXiv*. https://doi.org/10.48550/arXiv.2305.01723, 30.

Carifio, James, and Rocco J Perla. 2007. "Ten Common Misunderstandings, Misconceptions, Persistent Myths and Urban Legends about Likert Scales and Likert Response Formats and Their Antidotes." *J. Soc. Sci.* 3 (3). Science Publications: 106–16.

Chiorrini, Andrea, C Diamantini, Alex Mircoli, and D Potena. 2021. "Emotion and Sentiment Analysis of Tweets Using BERT." *EDBT/ICDT Workshops* 2841: 1–14.

Glass, Gene V, Percy D Peckham, and James R Sanders. 1972. "Consequences of Failure to Meet Assumptions Underlying the Fixed Effects Analyses of Variance and Covariance." *Rev. Educ. Res.* 42 (3). American Educational Research Association: 237–88.

Hutto, Clayton, and Eric Gilbert. 2014. "VADER: A Parsimonious Rule-Based Model for Sentiment Analysis of Social Media Text." *ICWSM* 8 (1): 216–25.

Liu, Bing. 2010. "Sentiment Analysis and Subjectivity." *Handbook of Natural Language Processing* 2 (2010). Oxfordshire: 627–66.

Mohammad, Saif M, and Peter D Turney. 2013. "Crowdsourcing a Word–Emotion Association Lexicon." *Comput. Intell.* 29 (3): 436–65.

Mossholder, Kevin W, Randall P Settoon, Stanley G Harris, and Achilles A Armenakis. 1995. "Measuring Emotion in Open-Ended Survey Responses: An Application of Textual Data Analysis." *J. Manage.* 21 (2): 335–55.

Nandwani, Pansy, and Rupali Verma. 2021. "A Review on Sentiment Analysis and Emotion Detection from Text." *Soc Netw Anal Min* 11 (1): 81.

Nielsen, Finn. 2011. "A New ANEW: Evaluation of a Word List for Sentiment Analysis in Microblogs." *arXiv*. https://doi.org/10.48550/arXiv.1103.2903, 6.

Pérez, Juan Manuel, Mariela Rajngewerc, Juan Carlos Giudici, Damián A Furman, Franco Luque, Laura Alonso Alemany, and María Vanina Martínez. 2023. "Pysentimiento: A Python Toolkit for Opinion Mining and Social NLP Tasks." *arXiv*. https://doi.org/10.48550/arXiv.2106.09462.

Rytting, Christopher Michael, Taylor Sorensen, Lisa Argyle, Ethan Busby, Nancy Fulda, Joshua Gubler, and David Wingate. 2023. "Towards Coding Social Science Datasets with Language Models." *arXiv*. https://doi.org/10.48550/arXiv.2306.02177, 27.

Sailunaz, Kashfia, and Reda Alhajj. 2019. "Emotion and Sentiment Analysis from Twitter Text." *J. Comput. Sci.* 36: 1–18.

Yoon, Seunghyun, Seokhyun Byun, and Kyomin Jung. 2018. "Multimodal Speech Emotion Recognition Using Audio and Text." *arXiv*. https://doi.org/10.48550/arXiv.1810.04635, 7.