A Gladiatorial Arena: Incivility in the Canadian House of Commons

Online Appendix

R. Michael Alvarez and Jacob Morrier

Division of the Humanities and Social Sciences

California Institute of Technology

September 2023

Table of Contents

A Descriptive Statistics

Table A1 contains descriptive statistics of the distribution of the five emotional attributes measured by the Perspective API. Table A2 displays the count of QP interventions categorized by the party of the MP who pronounced them and the language in which they were delivered. Table A3 presents the correlation coefficient matrix summarizing the relationships between the emotional attributes. In Figure A1, we visually represent the distribution of estimated probabilities that interventions in our corpus exhibit each of the five emotional attributes measured by the Perspective API. Lastly, Figure A2 is a counterpart to Figure 1 in which both smoothed time series and individual data points are represented.

Table A1: Descriptive Statistics of the Emotional Attributes of QP Interventions

	IDENTITY ATTACK	INSULT	PROFANITY	THREAT	TOXICITY
N	122,084	122,084	122,084	122,084	122,084
Mean	0.010	0.033	0.014	0.008	0.054
Std. Deviation	0.023	0.054	0.014	0.012	0.062
Minimum	0.001	0.006	0.008	0.005	0.001
25^{th} Percentile	0.003	0.010	0.010	0.006	0.016
Median	0.005	0.019	0.012	0.007	0.034
75^{th} Percentile	0.009	0.031	0.014	0.007	0.066
Maximum	0.572	0.716	0.765	0.560	0.720

Table A2: Number of QP Interventions by Party and Language

		Party			*Total*
		CPC	LPC	NDP	
Language	English	45,195	29,133	13,963	88,291
	French	13,301	11,545	8,947	33,793
Total		58,496	40,678	22,910	122,084

Table A3: Correlation between the Emotional Attributes of QP Interventions

	IDENTITY ATTACK	INSULT	PROFANITY	THREAT	TOXICITY
IDENTITY ATTACK	1	0.413	0.259	0.339	0.556
INSULT		1	0.594	0.158	0.917
PROFANITY			1	0.147	0.595
THREAT				1	0.298
TOXICITY					1

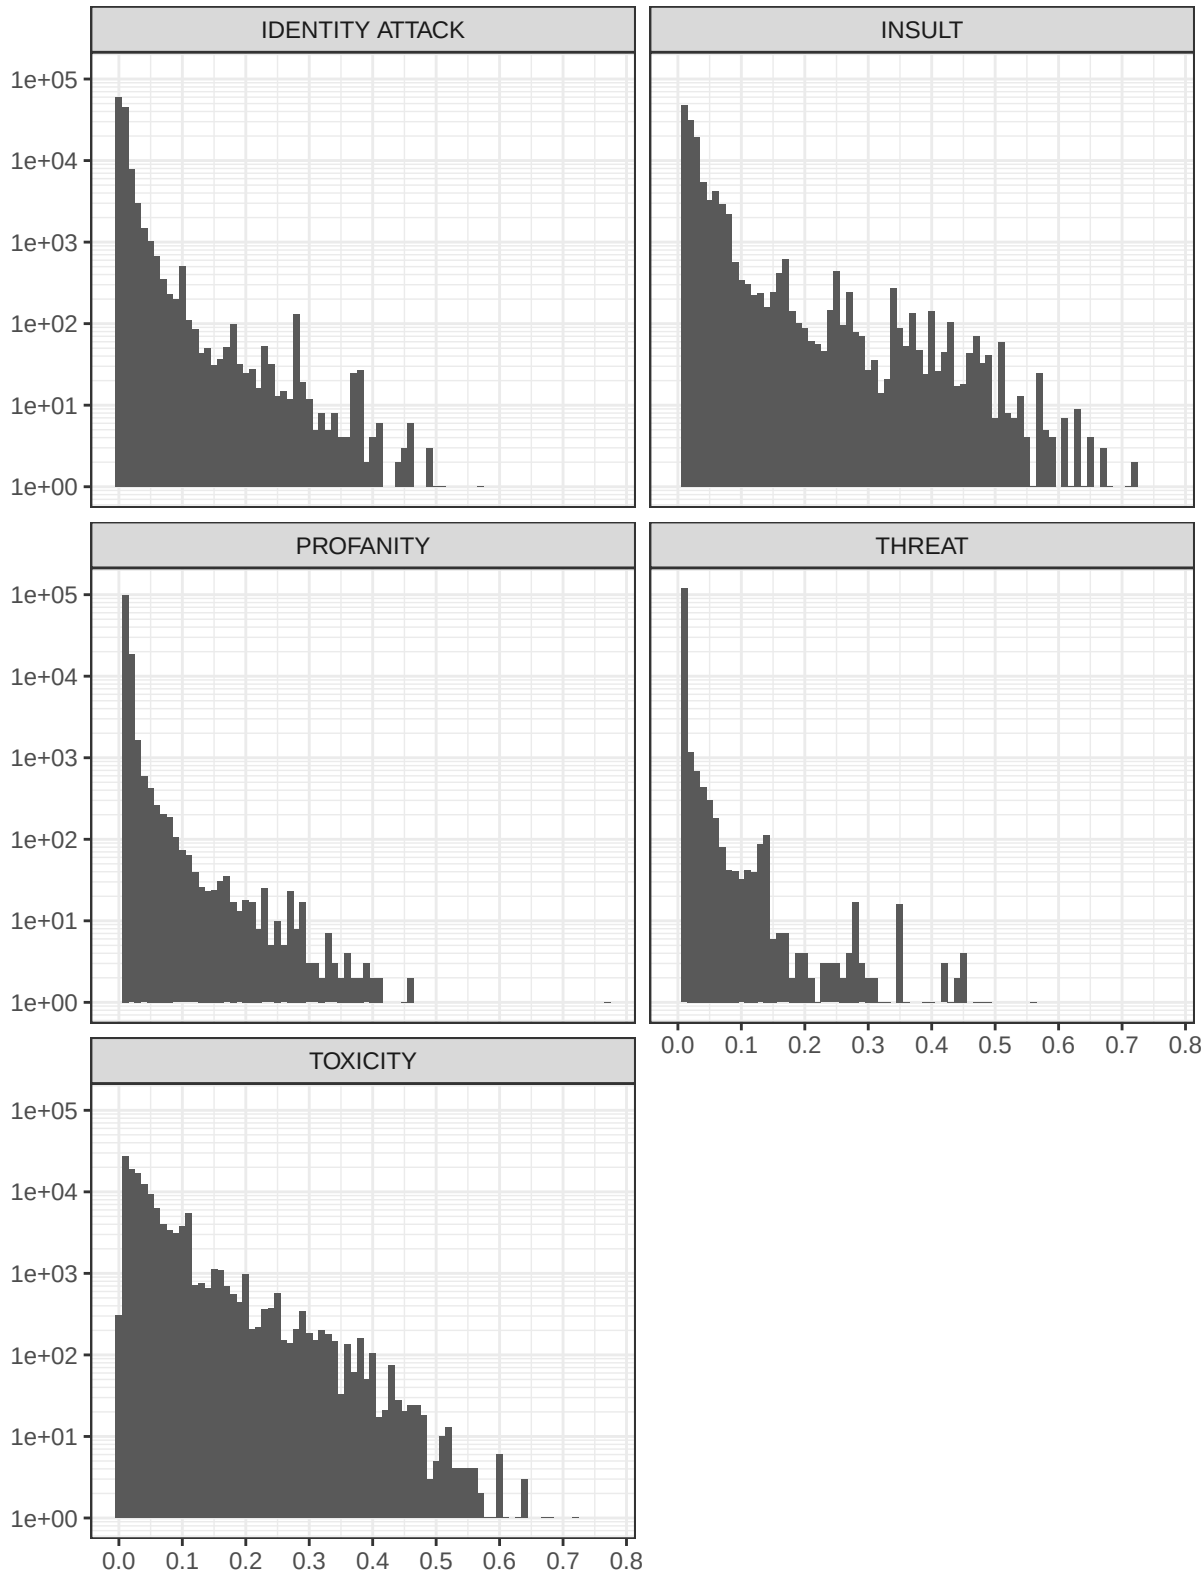

Figure A1: Distribution of Estimated Probabilities by Emotional Attribute

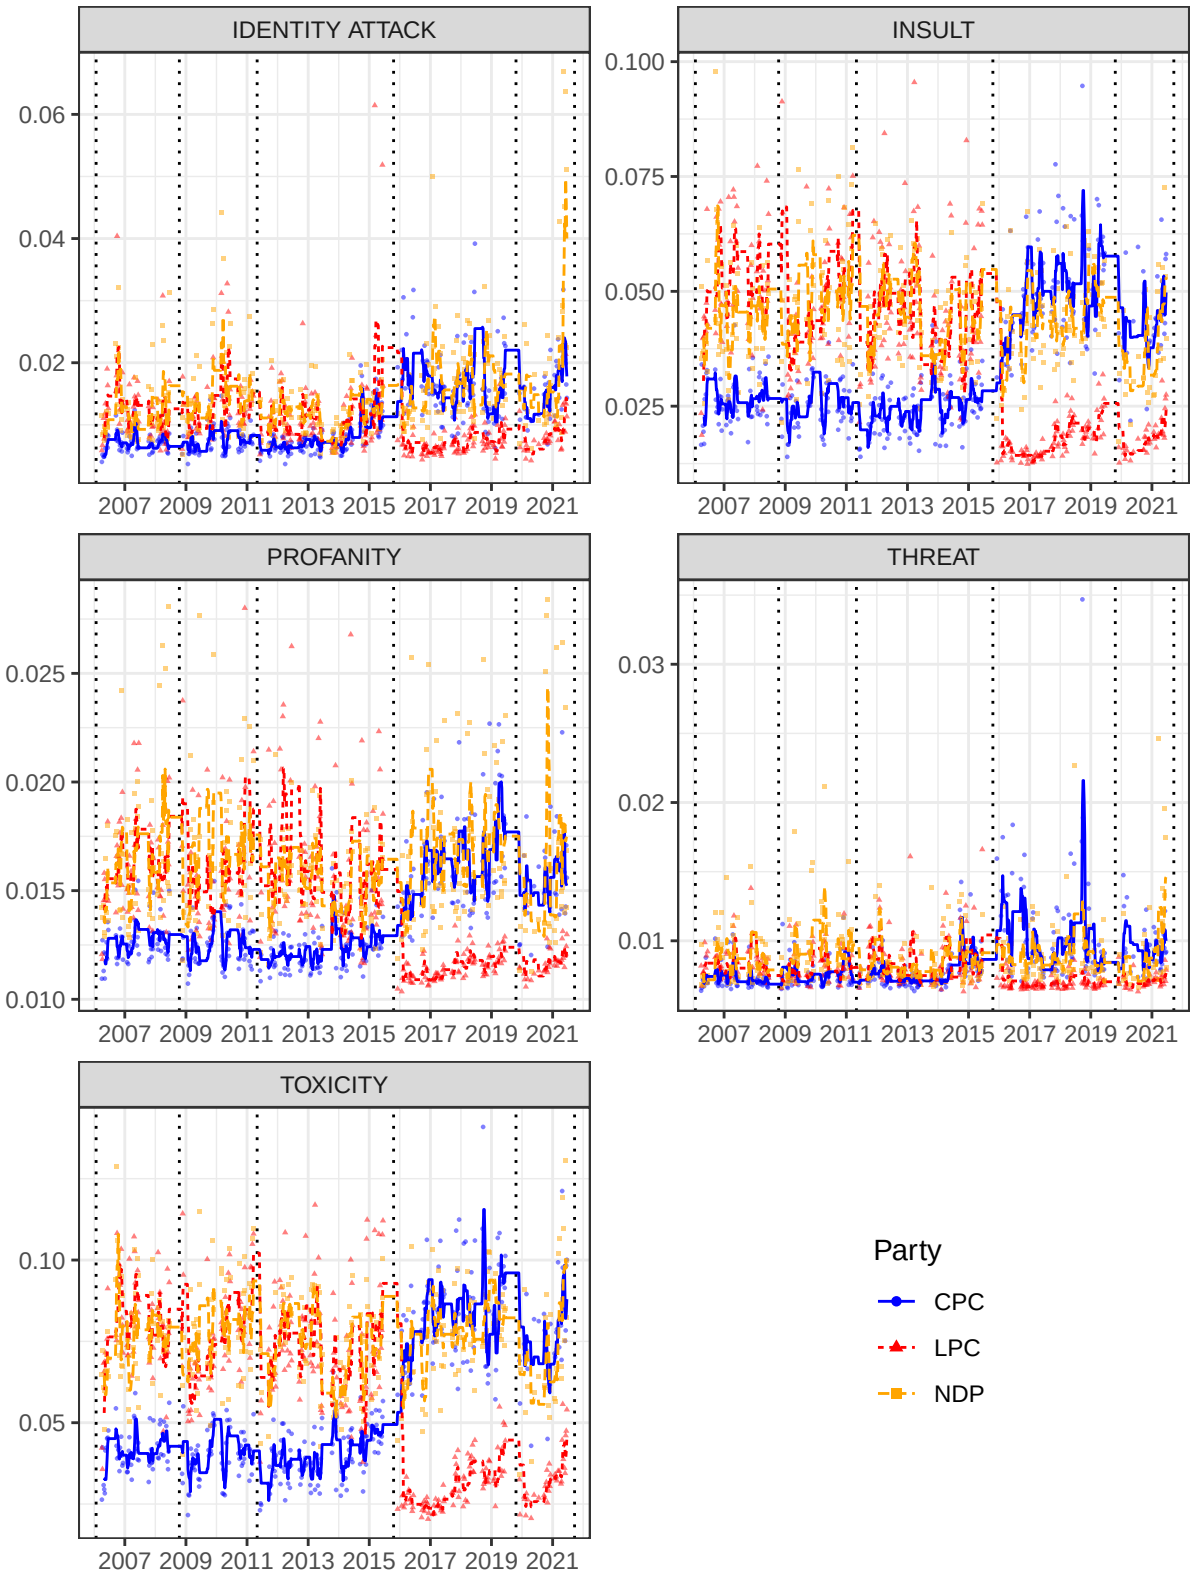

Figure A2: Weekly Evolution of the Emotional Attributes of QP Interventions by Party

B Validity of the Perspective API's Toxicity Measure

The Perspective API was intended to analyze online discussions rather than political speeches or debates in deliberative assemblies. Consequently, one may worry that the measures derived from its models are unsuitable for our analysis. To dissipate these concerns and assess the validity of the Perspective API's estimates within the scope of this study, we conducted a validity experiment.

The experiment began by randomly selecting 500 QP interventions from our corpus. Each co-author of this paper was tasked with independently reviewing them and indicating whether, in their professional judgment, they exhibit toxicity as per the definition found in Table 1. This procedure closely emulates the method used to generate the models' training labels.

For each coauthor, the correlation coefficient between the Perspective API's toxicity scores and our labels stands at 0.442 and 0.408, respectively. Additionally, the precision of our labels, representing the proportion of interventions classified as toxic with estimated toxicity scores above average, is 0.707 and 0.833, respectively. Finally, the recall of our labels, representing the share of interventions with estimated toxicity scores above average classified as toxic, is 0.653 and 0.242, respectively.

Figure B1 visually depicts the distribution of estimated toxicity scores conditional on whether each coauthor classified interventions as toxic. Each panel represents the labels provided by a coauthor. The main observation here is that the distribution of toxicity scores for the interventions we classified as toxic consistently exhibits first-order stochastic dominance over the distribution for interventions not classified as toxic. In simpler terms, this means that documents we labeled to be toxic systematically have higher estimated toxicity scores than those we did not identify as toxic.

An illustration of our consolidated labels' distribution is contained in Figure B2. This figure illustrates the distribution of documents' estimated toxicity scores conditional on the number of coauthors who classified them as toxic. A notable pattern emerges as the number of coauthors labeling a document as toxic increases: the distribution of estimated toxicity scores shifts towards higher values. Also, there is a correlation of 0.432 between the proportion of coauthors classifying documents as toxic and their estimated scores. This correlation entails that the interventions classified as toxic by a higher share of coauthors exhibit, in general, higher estimated toxicity scores.

In summary, our experiment provides convincing evidence that the Perspective API's toxicity scores align with the assessments made by subject-matter experts, including the coauthors of this paper.

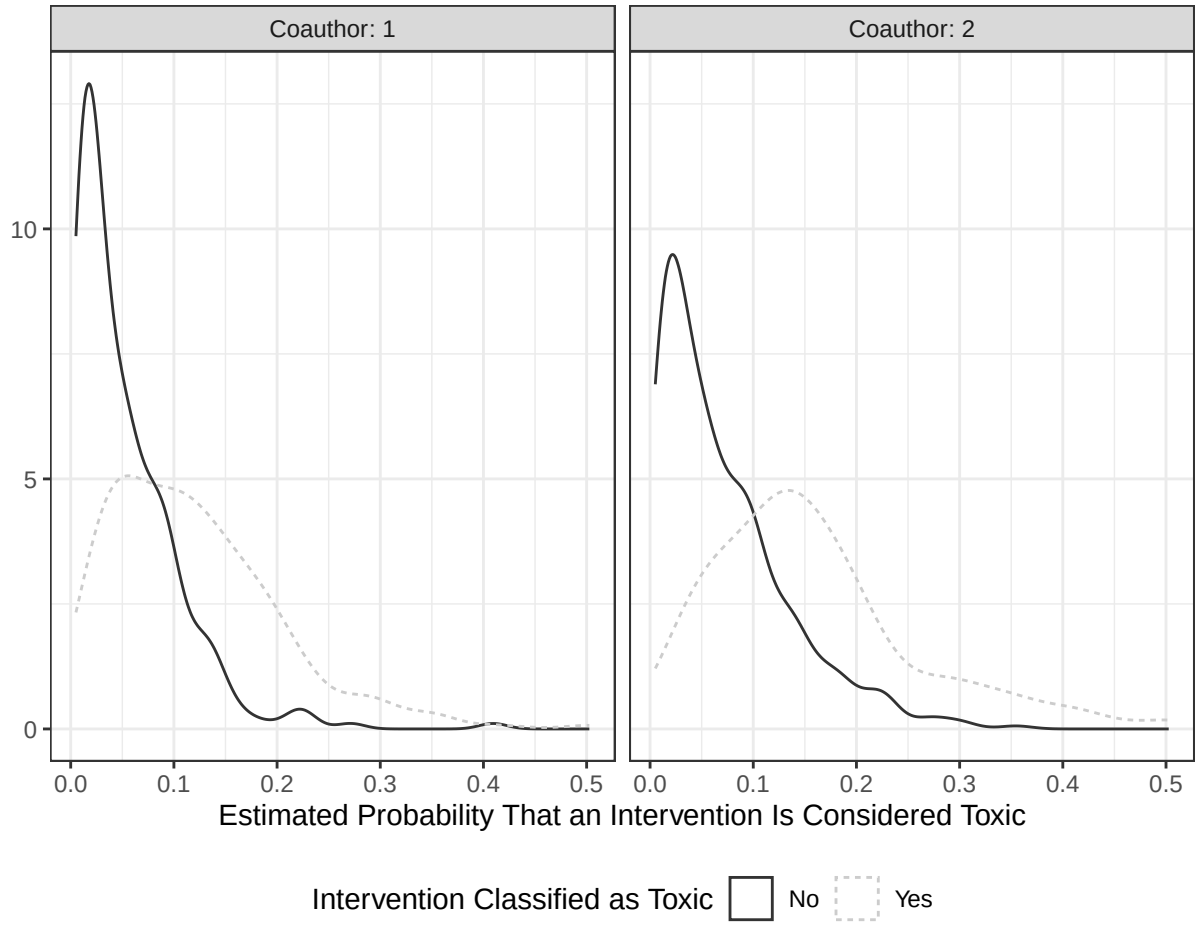

Figure B1: Distribution of the Estimated Toxicity Scores of QP Interventions Conditional on whether a Coauthor Classified Them as Toxic

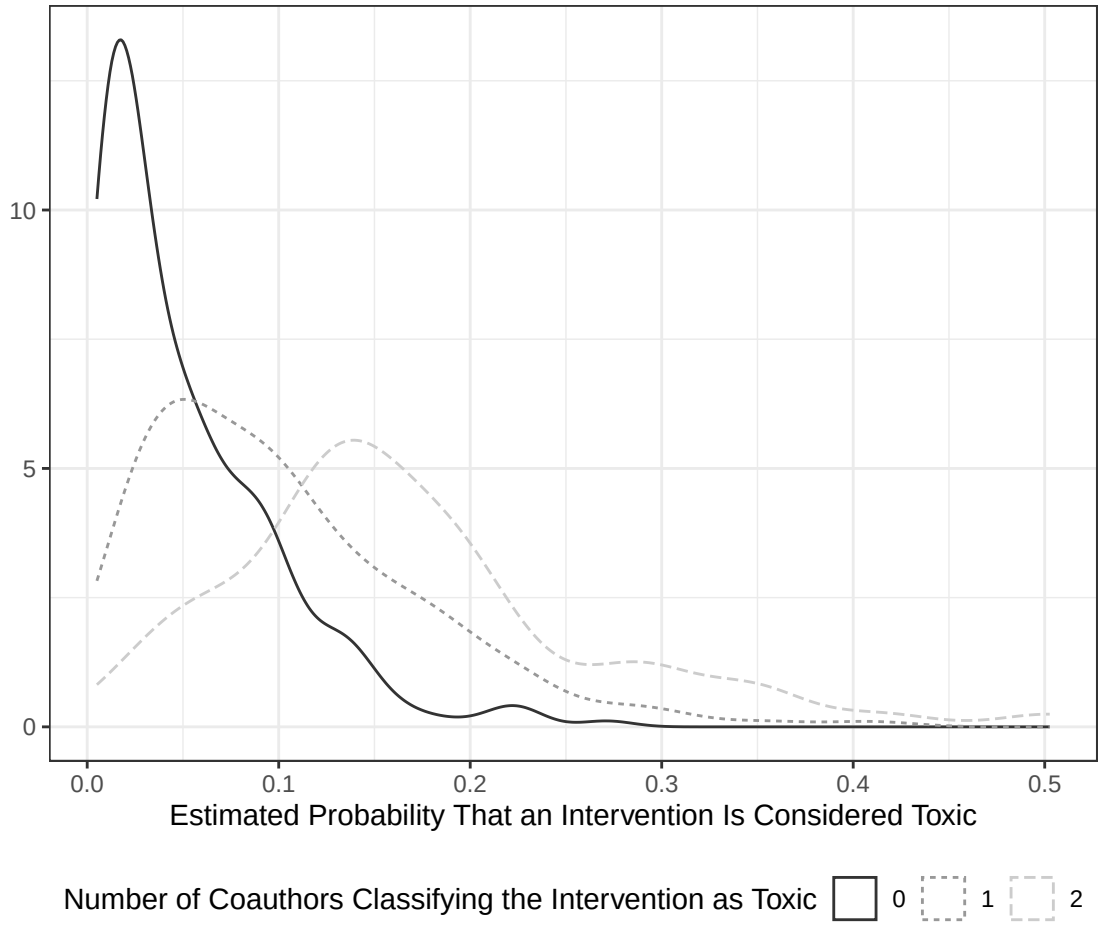

Figure B2: Distribution of the Estimated Toxicity Scores of QP Interventions by the Number of Coauthors Who Classified Them as Toxic

C Ordinary Least Squares Regression Analysis

Our regression models describe an inherently dynamic process. The absence of the dependent variable's lagged value on the right-hand side of our models makes them vulnerable to the serial correlation of their residuals. This is a critical issue because serial correlation in a model's error terms renders standard estimates inconsistent and invalidates the associated inference. To mitigate serial correlation in the residuals of our models, we implement the Cochrane–Orcutt estimation procedure. The results of this estimation approach are contained in Table 2.

For comparison, Table C1 presents estimates of our regression models obtained through the conventional ordinary least squares (OLS) estimation procedure. The Durbin–Watson statistic is a test statistic used to identify first-order autocorrelation in a regression model's residuals. The values of the Durbin–Watson statistic reveal statistically significant first-order autocorrelation in our models' error terms. This implies that the estimates derived from OLS estimation and the associated inferences are unreliable. In contrast, the values of the Durbin–Watson statistic found in Table 2 do not reveal a significant autocorrelation in our models' error terms. Therefore, the Cochrane–Orcutt estimation procedure efficaciously mitigates the serial correlation in our regression models' error terms.

Table C1: Regression Results

	(1) IDENTITY ATTACK	(2) INSULT	(3) PROFANITY	(4) THREAT	(5) TOXICITY
Time Until Next Election	−0.002***	−0.001***	−0.0003***	−0.0002*	−0.001***
	(0.0002)	(0.0002)	(0.0001)	(0.0001)	(0.0002)
Time Until Next Election × Minority	−0.002***	−0.001**	−0.001***	−0.0004	−0.001***
	(0.0005)	(0.0004)	(0.0002)	(0.0002)	(0.0003)
Government	−0.613***	−0.878***	−0.280***	−0.187***	−0.892***
	(0.025)	(0.021)	(0.011)	(0.012)	(0.018)
Minority	−0.032	0.007	0.015	−0.014	−0.002
	(0.041)	(0.035)	(0.018)	(0.020)	(0.029)
Language: French	−0.277***	−0.215***	−0.110***	−0.114***	−0.267***
	(0.019)	(0.017)	(0.008)	(0.009)	(0.014)
Party: LPC	−0.150***	−0.164***	−0.039***	−0.103***	−0.164***
	(0.025)	(0.021)	(0.011)	(0.012)	(0.018)
Party: NDP	−0.078***	−0.199***	−0.022*	−0.068***	−0.146***
	(0.028)	(0.024)	(0.012)	(0.014)	(0.020)
Constant	−3.982***	−2.757***	−4.046***	−4.612***	−2.178***
	(0.036)	(0.031)	(0.016)	(0.017)	(0.026)
Observations	2,130	2,130	2,130	2,130	2,130
R^2	0.357	0.522	0.358	0.190	0.639
Adjusted R^2	0.355	0.521	0.356	0.188	0.638
Durbin–Watson Statistic	1.412***	1.675***	1.915***	1.550***	1.589***

Note: *p<0.1; **p<0.05; ***p<0.01

D Language-Agnostic Regression Analysis

Table D1 contains the results of the estimation of regression models in which the language in which QP interventions are delivered is not included as a covariate. These regression models predict the average probability that interventions from a party's members each week exhibit one of the five emotional attributes measured by the Perspective API regardless of the language in which they were delivered. This supplementary analysis is meant to assess whether accounting for disparities in the incidence of uncivil behavior based on the language in which QP interventions impacts our findings. Overall, our substantive results appear robust to disregarding the language in which interventions were delivered.

Table D1: Regression Results

	(1) IDENTITY ATTACK	(2) INSULT	(3) PROFANITY	(4) THREAT	(5) TOXICITY
Time Until Next Election	-0.002^{***}	-0.001^{***}	-0.0004^{***}	-0.0002	-0.001^{***}
	(0.0003)	(0.0002)	(0.0001)	(0.0002)	(0.0002)
Time Until Next Election × Minority	-0.002^{**}	-0.0004	-0.0005^{**}	-0.001	-0.001
	(0.001)	(0.001)	(0.0002)	(0.0003)	(0.0005)
Government	-0.591^{***}	-0.826^{***}	-0.282^{***}	-0.216^{***}	-0.831^{***}
	(0.040)	(0.028)	(0.012)	(0.019)	(0.025)
Minority	0.051	0.018	0.030	0.013	0.025
	(0.064)	(0.046)	(0.020)	(0.031)	(0.041)
Party: LPC	-0.179^{***}	-0.189^{***}	-0.041^{***}	-0.124^{***}	-0.185^{***}
	(0.040)	(0.029)	(0.012)	(0.019)	(0.025)
Party: NDP	-0.053	-0.194^{***}	-0.016	-0.088^{***}	-0.139^{***}
	(0.046)	(0.033)	(0.014)	(0.021)	(0.029)
Constant	-4.077^{***}	-2.793^{***}	-4.073^{***}	-4.617^{***}	-2.246^{***}
	(0.056)	(0.040)	(0.017)	(0.026)	(0.035)
Observations	1,065	1,065	1,065	1,065	1,065
R^2	0.261	0.502	0.43	0.148	0.584
Adjusted R^2	0.258	0.500	0.428	0.145	0.582
Durbin–Watson Statistic	2.080	2.076	2.013	2.019	2.085

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

E Document-Level Regression Analysis

As revealed in Table A2, there is a significant variation in the number of QP interventions across parties. Approximately half of interventions are initiated by members of the government. This is because each question from an opposition member calls for a response from either a cabinet minister or a parliamentary secretary appointed to assist them. While backbench members of the ruling party can ask questions, their participation is considerably lower than opposition MPs. The allocation of questions to opposition parties is proportional to the weight of their representation in the House. Accordingly, we consistently observe that more interventions emanate from members of the official opposition than from members of third parties.

Our regression models neutralize disparities in the number of QP interventions emanating from each party by predicting the average likelihood that an intervention from a specific party's members during a particular week is considered to display one of the emotional attributes measured by the Perspective API. This transformation ensures that each party has an equal number of observations in the models' estimation.

To assess the robustness of our substantive findings to this modeling choice, we conduct additional regression analyses in which individual interventions serve as the unit of analysis. To address potential correlation in the residuals of observations from the same speaker or meeting, we report standard errors clustered by both speaker and meeting. Note that this does not account for the possible correlation in the residuals of observations from two meetings that took place closely in time. These models' estimates can be found in Table E1.

While estimates of some coefficients vary between the two modeling approaches, this supplemental analysis generally supports our substantive findings. We note two exceptions. Firstly, when all other factors are held constant, the presence of a minority government does not appear to be associated with any statistically significant differences in the incidence of incivility and its progression over time. Secondly, only the LPC consistently exhibits, all else equal, less uncivil behavior than other parties.

Table E1: Regression Results

	(1) IDENTITY ATTACK	(2) INSULT	(3) PROFANITY	(4) THREAT	(5) TOXICITY
Time Until Next Election	−0.001***	−0.001**	−0.0003**	−0.0002***	−0.001***
	(0.0002)	(0.0005)	(0.0001)	(0.0001)	(0.001)
Time Until Next Election × Minority	−0.0003	0.0001	−0.0001	−0.0001	−0.00004
	(0.001)	(0.001)	(0.0001)	(0.0002)	(0.001)
Government	−0.552***	−0.749***	−0.203***	−0.120***	−0.914***
	(0.042)	(0.039)	(0.013)	(0.011)	(0.045)
Minority	−0.067	−0.046	0.0002	−0.011	−0.052
	(0.093)	(0.054)	(0.019)	(0.014)	(0.055)
Language: French	−0.189***	−0.227***	−0.082***	−0.055***	−0.301***
	(0.024)	(0.029)	(0.013)	(0.005)	(0.032)
Party: LPC	−0.042	−0.107***	−0.030***	−0.038***	−0.124***
	(0.038)	(0.035)	(0.011)	(0.010)	(0.043)
Party: NDP	0.036	−0.046	−0.002	−0.023*	−0.025
	(0.050)	(0.059)	(0.022)	(0.014)	(0.065)
Constant	−4.631***	−3.204***	−4.191***	−4.802***	−2.521***
	(0.055)	(0.074)	(0.023)	(0.011)	(0.080)
Observations	122,084	122,084	122,084	122,084	122,084
R^2	0.097	0.178	0.086	0.033	0.206
Adjusted R^2	0.097	0.178	0.086	0.033	0.206

Note: *p<0.1; **p<0.05; ***p<0.01

F Regression Analysis with Interventions by Members of the Bloc Québécois

Our analysis disregards QP interventions by members of BQ. From January 2006 to May 2011 and from October 2019 onwards, BQ has held third-party status, giving its members the right to participate daily in QP. However, since it did not maintain official party status throughout our period of interest, our data on the interventions by members of BQ is scarce, if not virtually inexistent, from May 2011 to October 2019. Thus, we excluded this party from our analysis, as we did for the Green Party of Canada (GPC) or the People's Party of Canada (PPC), two other parties who were represented in the House of Commons during some part of our period of interest but never had official party status.

Here, we consider whether our findings are robust to this decision. To do so, we reproduce some elements of our analysis after having incorporated into our corpus interventions by members of BQ over the time this party held official party status, that is, during the 39th, 40th, and 43rd legislatures. Those are the only periods during which BQ has consistently taken part in QP, hence, during which we have meaningful data about the incidence of incivility in interventions from members of this party.

Table F1 presents descriptive statistics of the distribution of the five emotional attributes measured by the Perspective API. Table F2 displays the count of interventions conditional on the party of the MP they emanated from and the language in which they were delivered. Table F3 presents the correlation coefficient matrix for the five emotional attributes. In Figure F1, we visually represent the distribution of the estimated probabilities that interventions exhibit each of the five attributes. Figure F2 illustrates rolling averages of the likelihood that a QP intervention from a member of a given party in some week exhibits one of the emotional attributes measured by the Perspective API over the last four weeks with available data.

Table F4 contains estimation results of our main regression models on the corpus with interventions by members of BQ. Table F5 describes the results of the estimation of document-level regression models on the corpus with interventions from members of BQ. While the estimates of some coefficients change after adding QP interventions from members of BQ in our dataset, this supplementary analysis entirely supports our substantive findings. Furthermore, both sets of regression models indicate that all else equal, including the language in which interventions are delivered, those by members of BQ exhibit significantly less incivility than those of other parties.

Table F1: Descriptive Statistics of the Emotional Attributes of QP Interventions

	IDENTITY ATTACK	INSULT	PROFANITY	THREAT	TOXICITY
N	129,194	129,194	129,194	129,194	129,194
Mean	0.010	0.033	0.014	0.008	0.054
Std. Deviation	0.023	0.054	0.014	0.011	0.062
Minimum	0.001	0.006	0.008	0.005	0.001
25^{th} Percentile	0.003	0.010	0.010	0.006	0.017
Median	0.005	0.019	0.012	0.007	0.034
75^{th} Percentile	0.009	0.031	0.014	0.007	0.065
Maximum	0.572	0.813	0.765	0.560	0.836

Table F2: Number of QP Interventions by Party and Language

		Party				*Total*
		BQ	CPC	LPC	NDP	
Language	English	1	45,195	29,133	13,963	88,292
	French	7,109	13,301	11,545	8,947	40,902
Total		7,110	58,496	40,678	22,910	129,194

Table F3: Correlation between the Emotional Attributes of QP Interventions

	IDENTITY ATTACK	INSULT	PROFANITY	THREAT	TOXICITY
IDENTITY ATTACK	1	0.410	0.258	0.336	0.554
INSULT		1	0.594	0.156	0.918
PROFANITY			1	0.146	0.596
THREAT				1	0.295
TOXICITY					1

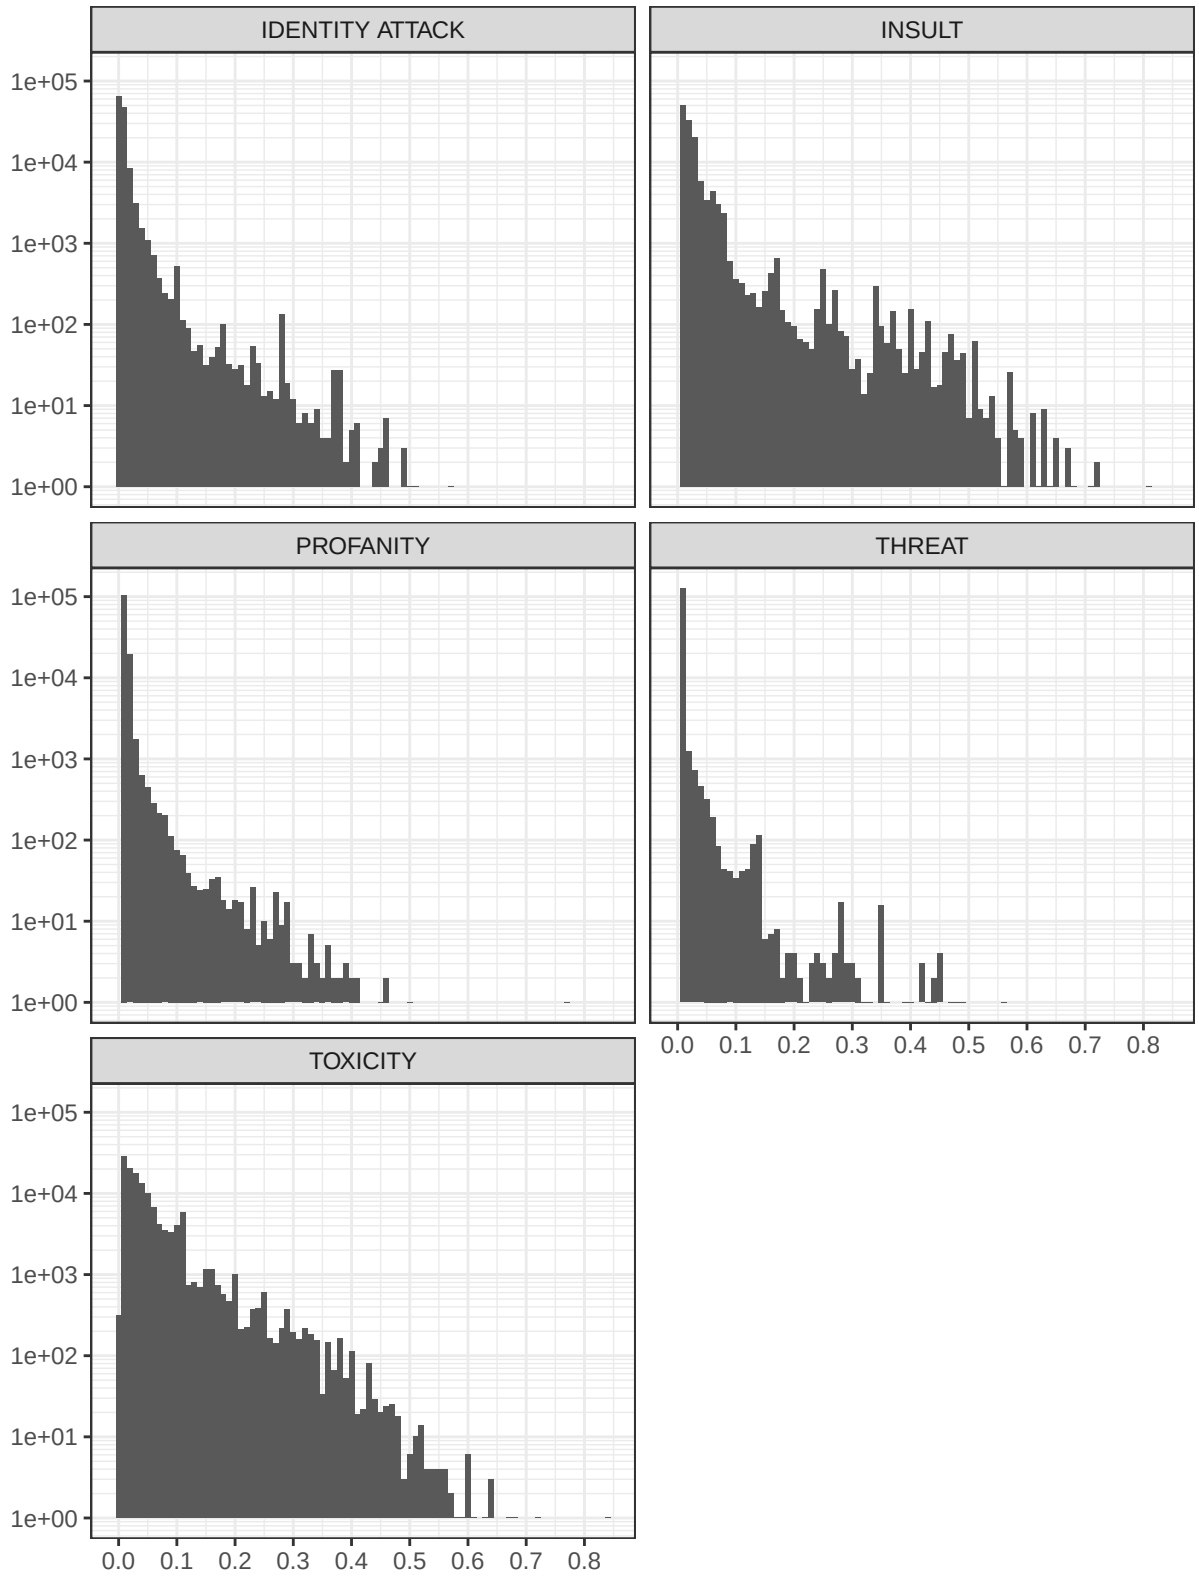

Figure F1: Distribution of Estimated Probabilities by Emotional Attribute

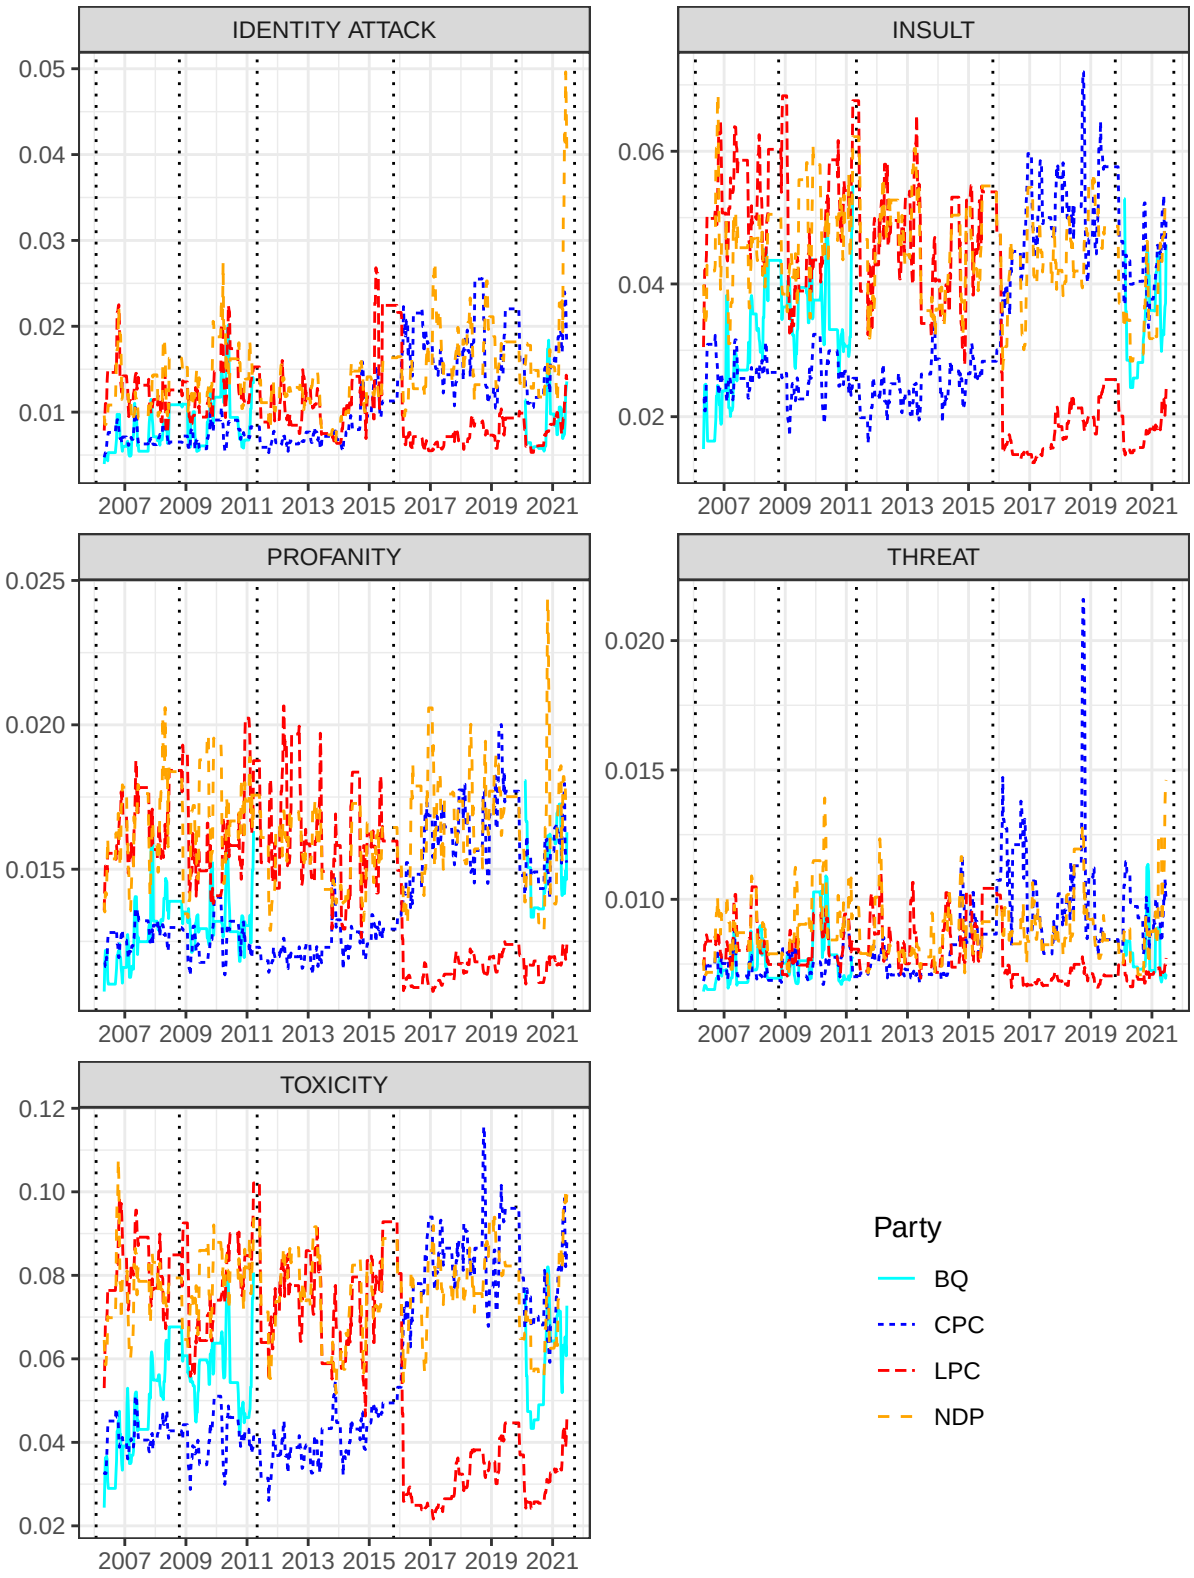

Figure F2: Weekly Evolution of the Emotional Attributes of QP Interventions by Party

Table F4: Regression Results

	(1) IDENTITY ATTACK	(2) INSULT	(3) PROFANITY	(4) THREAT	(5) TOXICITY
Time Until Next Election	−0.002***	−0.001***	−0.0003***	−0.0002	−0.001***
	(0.0003)	(0.0002)	(0.0001)	(0.0001)	(0.0002)
Time Until Next Election × Minority	−0.002***	−0.001***	−0.001***	−0.0004	−0.001***
	(0.001)	(0.0004)	(0.0002)	(0.0003)	(0.0004)
Government	−0.614***	−0.878***	−0.280***	−0.188***	−0.892***
	(0.033)	(0.025)	(0.011)	(0.015)	(0.022)
Minority	−0.004	0.027	0.023	−0.013	0.017
	(0.054)	(0.040)	(0.018)	(0.024)	(0.035)
Language: French	−0.276***	−0.214***	−0.109***	−0.113***	−0.267***
	(0.014)	(0.014)	(0.008)	(0.007)	(0.011)
Party: BQ	−0.331***	−0.326***	−0.098***	−0.106***	−0.298***
	(0.060)	(0.045)	(0.020)	(0.027)	(0.040)
Party: LPC	−0.150***	−0.163***	−0.039***	−0.103***	−0.164***
	(0.034)	(0.025)	(0.011)	(0.015)	(0.022)
Party: NDP	−0.079**	−0.199***	−0.022*	−0.069***	−0.146***
	(0.039)	(0.029)	(0.013)	(0.017)	(0.025)
Constant	−3.982***	−2.752***	−4.045***	−4.611***	−2.172***
	(0.048)	(0.036)	(0.016)	(0.021)	(0.032)
Observations	2,285	2,285	2,285	2,285	2,285
R^2	0.294	0.441	0.337	0.174	0.552
Adjusted R^2	0.293	0.440	0.336	0.172	0.551
Durbin–Watson Statistic	2.109	2.078	2.008	2.061	2.116

Note: *p<0.1; **p<0.05; ***p<0.01

Table F5: Regression Results

	(1) IDENTITY ATTACK	(2) INSULT	(3) PROFANITY	(4) THREAT	(5) TOXICITY
Time Until Next Election	−0.001***	−0.001**	−0.0003**	−0.0002***	−0.001***
	(0.0002)	(0.0005)	(0.0001)	(0.0001)	(0.001)
Time Until Next Election × Minority	−0.001	−0.0002	−0.0002	−0.0001	−0.0004
	(0.001)	(0.0005)	(0.0001)	(0.0002)	(0.001)
Government	−0.552***	−0.749***	−0.203***	−0.120***	−0.914***
	(0.042)	(0.039)	(0.013)	(0.011)	(0.045)
Minority	−0.051	−0.028	0.007	−0.010	−0.032
	(0.093)	(0.056)	(0.020)	(0.014)	(0.060)
Language: French	−0.189***	−0.227***	−0.082***	−0.055***	−0.301***
	(0.024)	(0.029)	(0.013)	(0.005)	(0.032)
Party: BQ	−0.298***	−0.265***	−0.085**	−0.064***	−0.295***
	(0.079)	(0.100)	(0.034)	(0.018)	(0.112)
Party: LPC	−0.043	−0.108***	−0.030***	−0.038***	−0.125***
	(0.037)	(0.034)	(0.010)	(0.010)	(0.042)
Party: NDP	0.036	−0.046	−0.002	−0.023*	−0.026
	(0.050)	(0.059)	(0.022)	(0.014)	(0.066)
Constant	−4.630***	−3.203***	−4.191***	−4.802***	−2.520***
	(0.055)	(0.074)	(0.023)	(0.011)	(0.080)
Observations	129,194	129,194	129,194	129,194	129,194
R^2	0.094	0.169	0.083	0.032	0.197
Adjusted R^2	0.094	0.169	0.083	0.032	0.197

Note: *p<0.1; **p<0.05; ***p<0.01

G Comparison between the Perspective API's English and French Models

Our analysis exploits English transcripts of QP interventions published by the Clerk of the House of Commons. These include professionally translated transcripts of interventions delivered in French. We also have access to French transcripts of all QP interventions, including professionally translated versions of those initially delivered in English. To quantify the degree of incivility in these QP interventions, we rely on the models hosted in the Perspective API. While our analysis primarily uses models trained for English documents, some models can analyze French documents. Accordingly, we replicate our study using interventions' French transcripts.

This exercise is of interest for two reasons. Firstly, it allows us to evaluate whether the professional translation of QP interventions distorts our results. For instance, translators may consciously or unconsciously edit certain parts of interventions, especially if they contain highly discourteous language. In such a scenario, the lower incidence of incivility in interventions delivered in French may reflect this censorship more than a genuine relationship between language and uncivil behavior. Also, if this correlation resulted from conscious or unconscious expurgation, we would expect, when considering the French transcripts of QP interventions, an opposite relationship in which English interventions exhibit a lower incidence of uncivil behavior.

Secondly, this supplementary analysis allows us to assess whether linguistic biases are embedded in the Perspective API's models. The models' training process begins with training multilingual BERT-based models. These models are then distilled into single-language Convolutional Neural Networks. Therefore, although models for different languages are built from the same basis, they do not identically mirror each other. We have no assurance that estimates of the incidence of uncivil behavior derived from the English and French transcripts of the same intervention will be equal.

Table G1 contains descriptive statistics of the distribution of the five emotional attributes derived from the French transcripts of QP interventions. Table G2 displays the count of QP interventions categorized by the party of the MP who pronounced them and the language in which they were delivered. Note that due to the condition of the data furnished by the House of Commons, the data collection process did not yield an equal number of transcripts in English and French. Table G3 presents the correlation coefficient matrix summarizing the relationships between emotional attributes. In Figure G1, we visually represent the distribution of estimated probabilities that interventions in our corpus exhibit each of the five emotional attributes.

Given the data's structure, it is infeasible to match the English and French transcripts of documents. For this reason, we cannot make individual comparisons between the estimates derived from English and French transcripts of each intervention. Instead, we compare estimates of the likelihood that QP interventions from a specific party each week display one of the emotional attributes derived from the English and French transcripts. This represents the best alternative under our technical constraints.

Figure G2 compares estimates of the likelihood that QP interventions from members of a specific party in a given week exhibit each emotional attribute derived from interventions' English and French transcripts. These estimates are positively correlated, with correlation coefficients above 0.9 for the three most prevalent attributes. Also, estimates of the incidence of identity attacks and threats are, on average, higher for English transcripts. In contrast, estimates of the incidence of insults and toxicity are systematically higher for French transcripts.

Figure G3 is analogous to Figure G3 except it highlights each party's estimates. The magnitude of the correlation between estimates from English and French transcripts varies with the parties' relative size and, by extension, the number of interventions by their MPs. Unsurprisingly, estimates for parties with more interventions by their members consistently exhibit a higher correlation when compared to those for parties with fewer interventions by their members. There do not seem to be biases in estimates based on the language in which interventions are pronounced.

For each emotional attribute, Figure G4 compares estimates of the likelihood that QP interventions from MPs of a specific party in a given week delivered in a particular language derived from the English and French transcripts of QP interventions. The correlation between these estimates is lower than that displayed in Figure G2. This is, at least partially, because these estimates are derived from fewer interventions. It appears that these estimates demonstrate similar biases as those previously described.

Figure G5 dissects the relationship illustrated in Figure G4 based on the language in which interventions were delivered. Estimates of the incidence of uncivil behavior for interventions delivered in English tend to exhibit a higher level of correlation than estimates for interventions pronounced in French. This discrepancy can be attributed to the higher number of interventions delivered in English, resulting in estimates of incivility in the latter being based on a higher number of observations.

To conclude, we consider whether our substantive findings hold when we estimate models to predict the incidence of incivility derived from the French transcripts of QP interventions. Table G4 presents the estimation results of our regression models. Table G5 contains estimation results for our intervention-level regression models. While specific coefficients may change, overall findings remain consistent when considering French transcripts. This is especially true for the most prevalent forms of incivility, namely, insults and toxicity. We note two meaningful exceptions. Firstly, a minority government appears not to be associated with any significant variation in the incidence of incivility. Secondly, per the intervention-level regression model, the evolution of the prevalence of insults in QP interventions is not significantly correlated with the time left until the next election.

Table G1: Descriptive Statistics of the Emotional Attributes of QP Interventions

	IDENTITY ATTACK	INSULT	PROFANITY	THREAT	TOXICITY
N	123,520	123,520	123,520	123,520	123,520
Mean	0.009	0.046	0.014	0.007	0.066
Std. Deviation	0.016	0.051	0.015	0.006	0.055
Minimum	0.00000	0.005	0.007	0.005	0.00003
25^{th} Percentile	0.003	0.020	0.011	0.006	0.029
Median	0.005	0.030	0.012	0.006	0.049
75^{th} Percentile	0.009	0.057	0.015	0.007	0.093
Maximum	0.451	0.568	0.698	0.348	0.498

Table G2: Number of QP Interventions by Party and Language

		Party			*Total*
		CPC	LPC	NDP	
Language	English	45,630	29,170	13,963	88,763
	French	14,105	11,703	8,949	34,757
Total		59,735	40,873	22,912	123,520

Table G3: Correlation between the Emotional Attributes of QP Interventions

	IDENTITY ATTACK	INSULT	PROFANITY	THREAT	TOXICITY
IDENTITY ATTACK	1	0.423	0.219	0.412	0.549
INSULT		1	0.479	0.178	0.930
PROFANITY			1	0.114	0.483
THREAT				1	0.272
TOXICITY					1

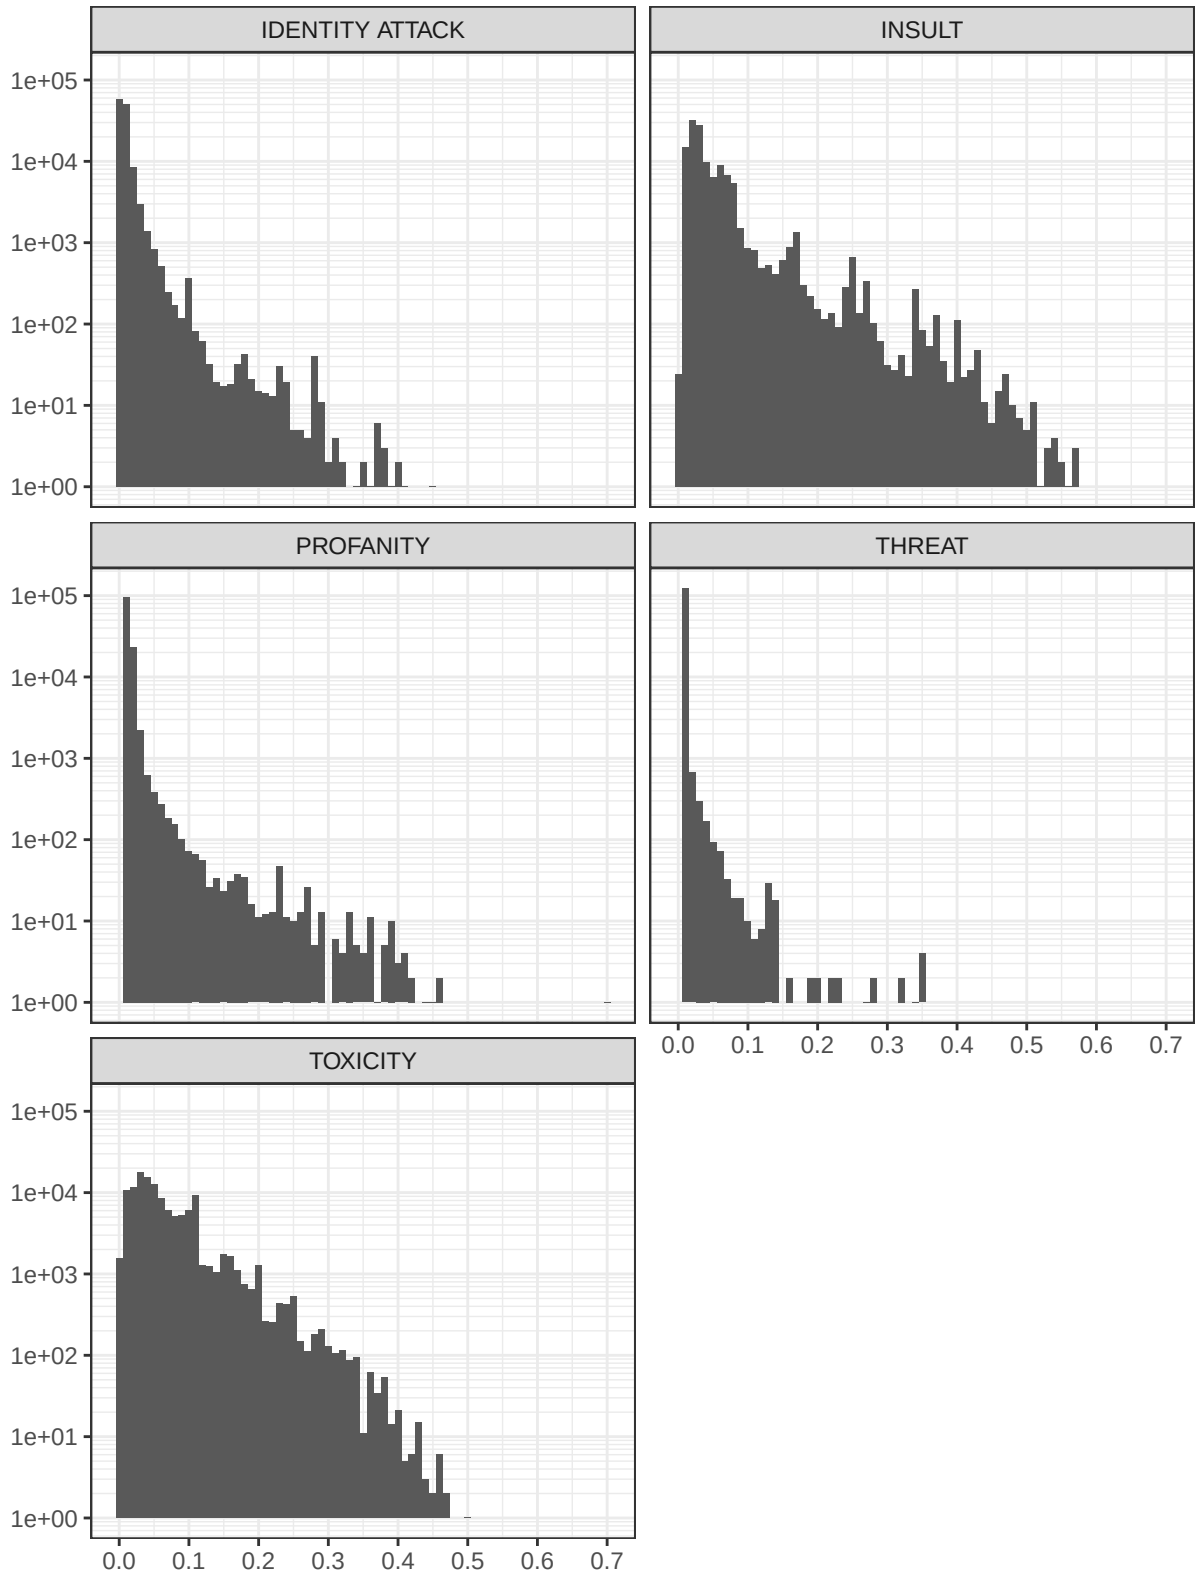

Figure G1: Distribution of Estimated Probabilities by Emotional Attribute

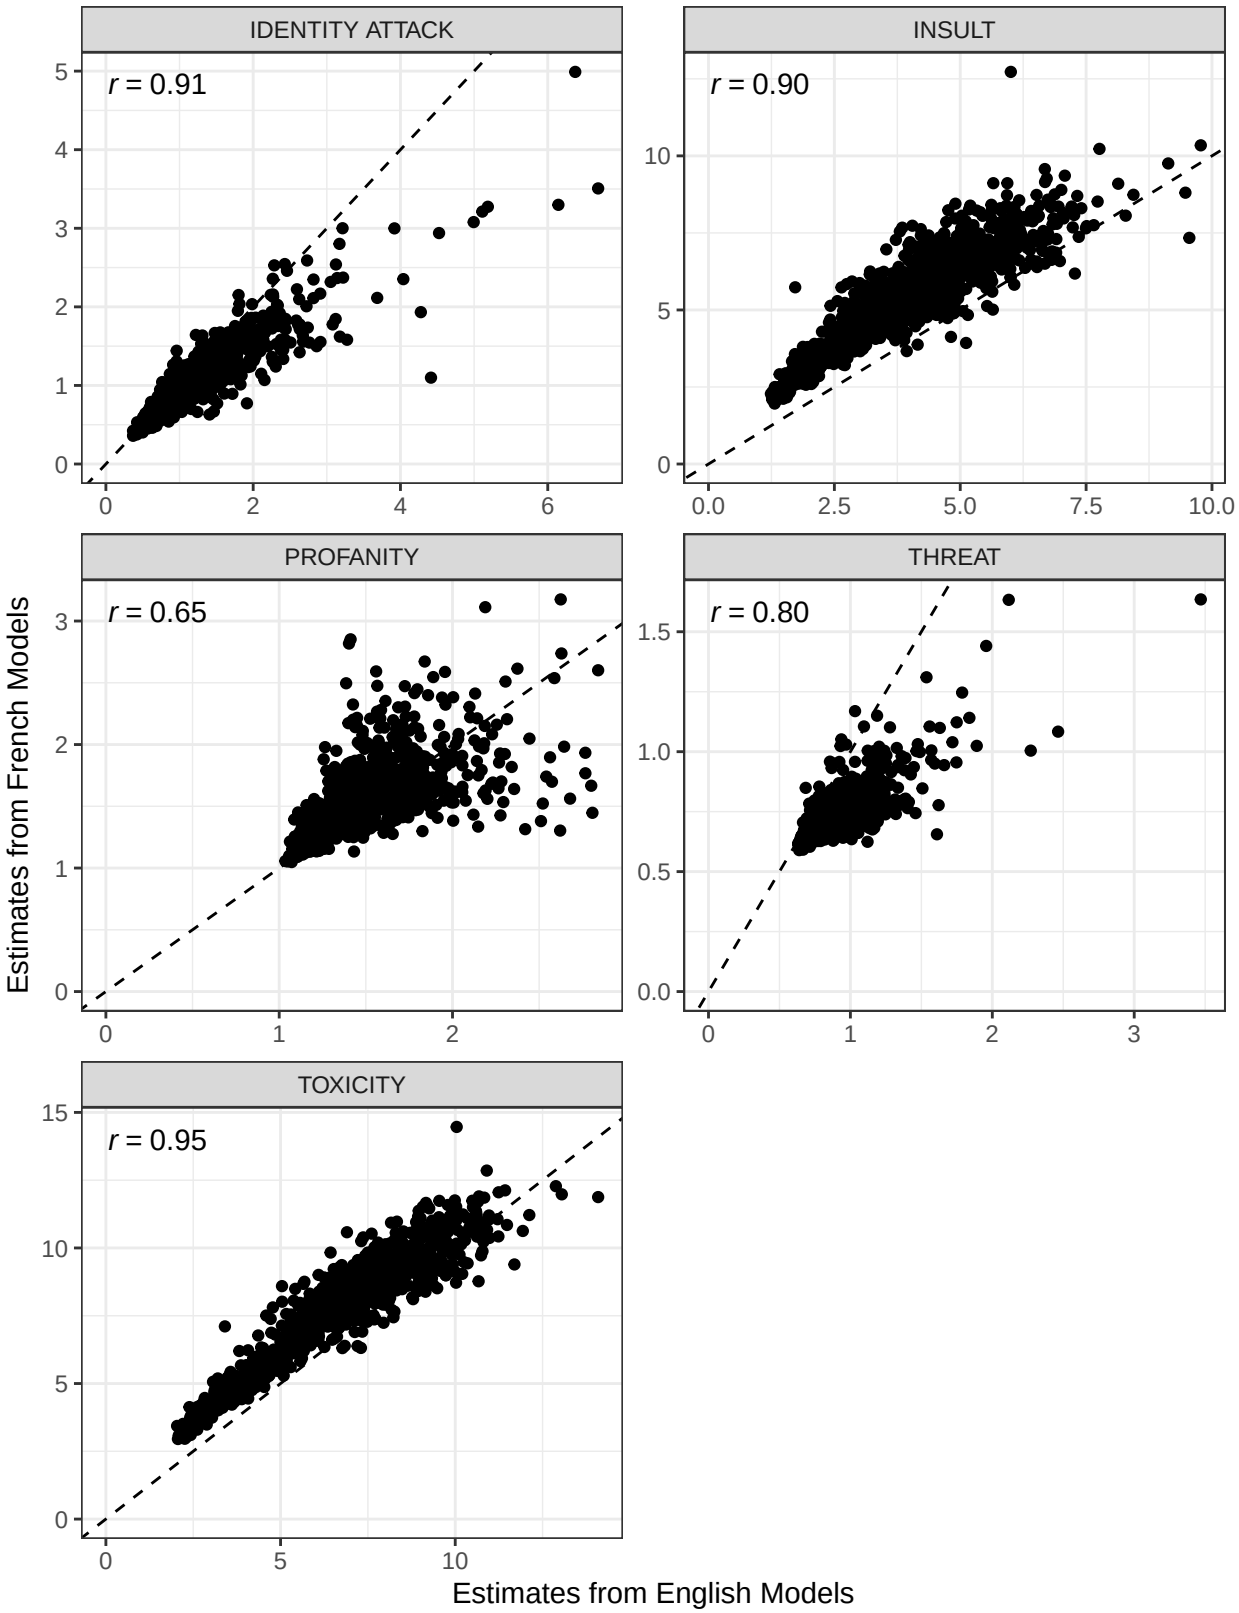

Figure G2: Relationship between Estimates from the English and French Models

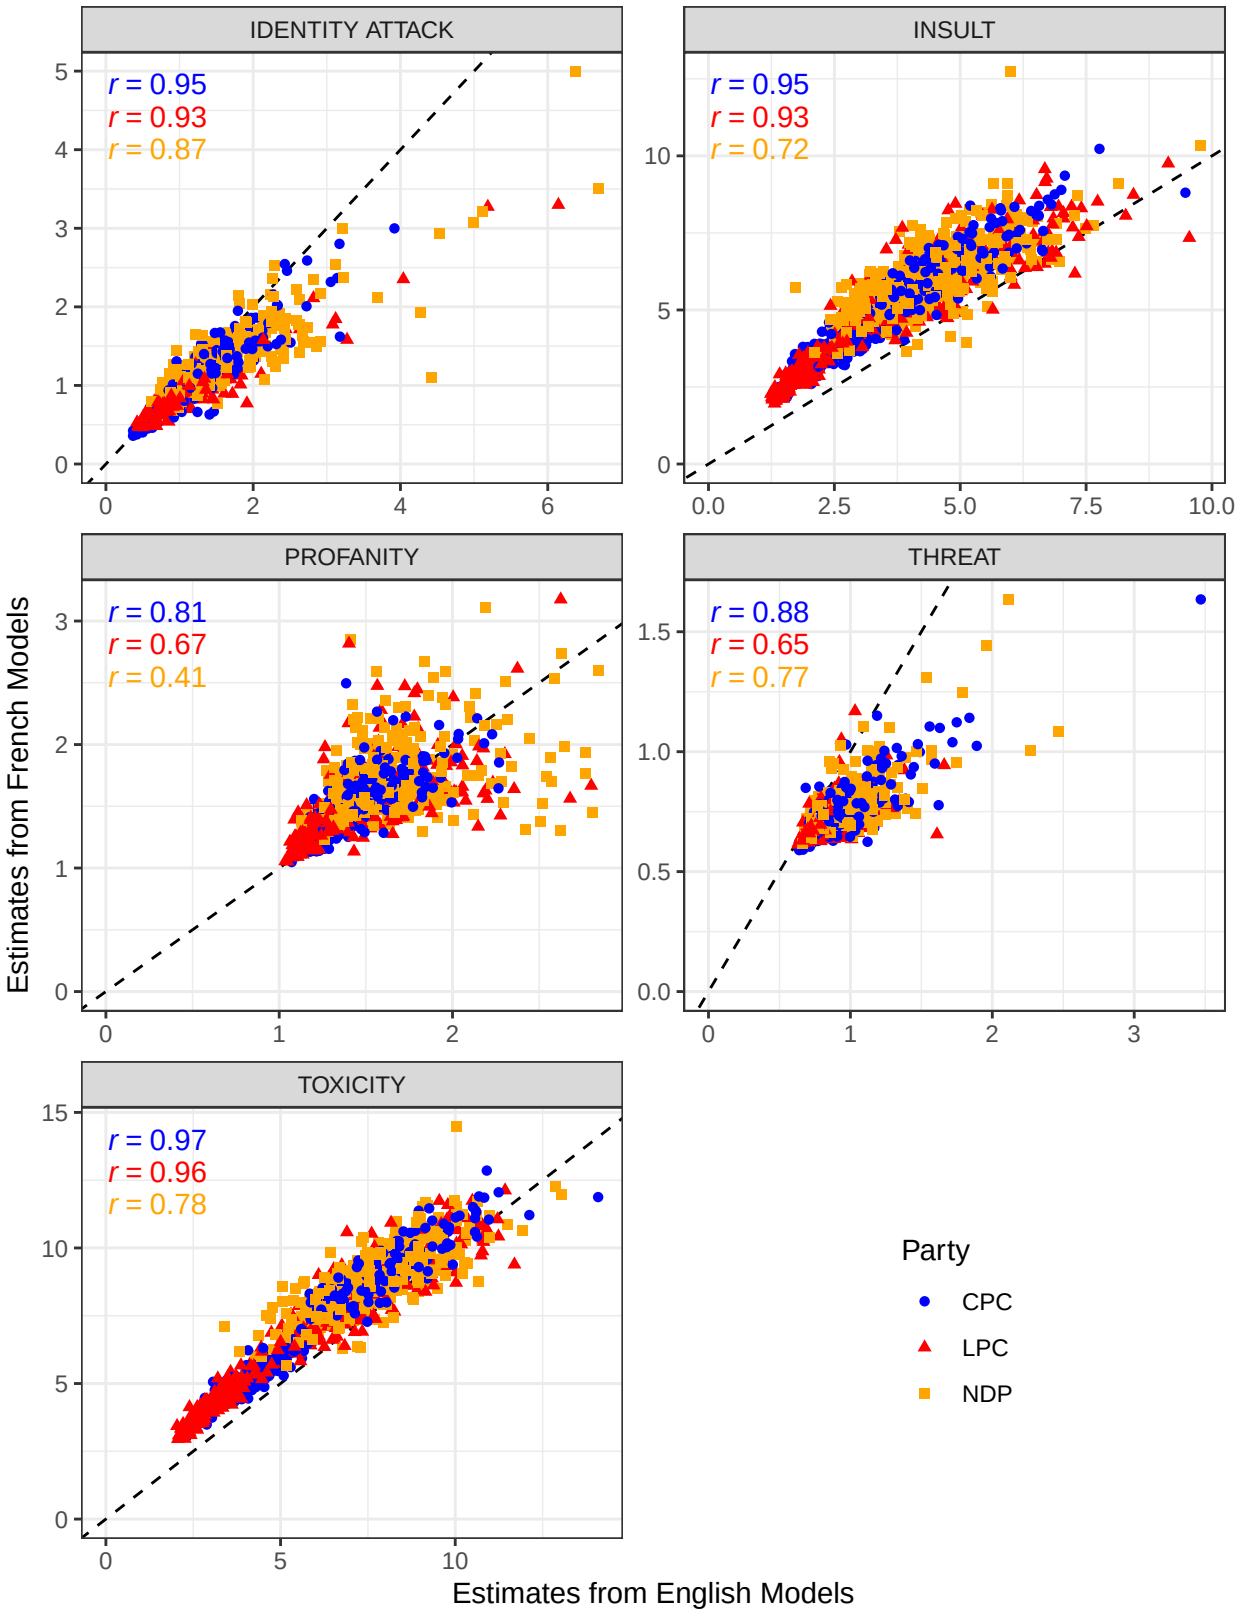

Figure G3: Relationship between Estimates from the English and French Models by Party

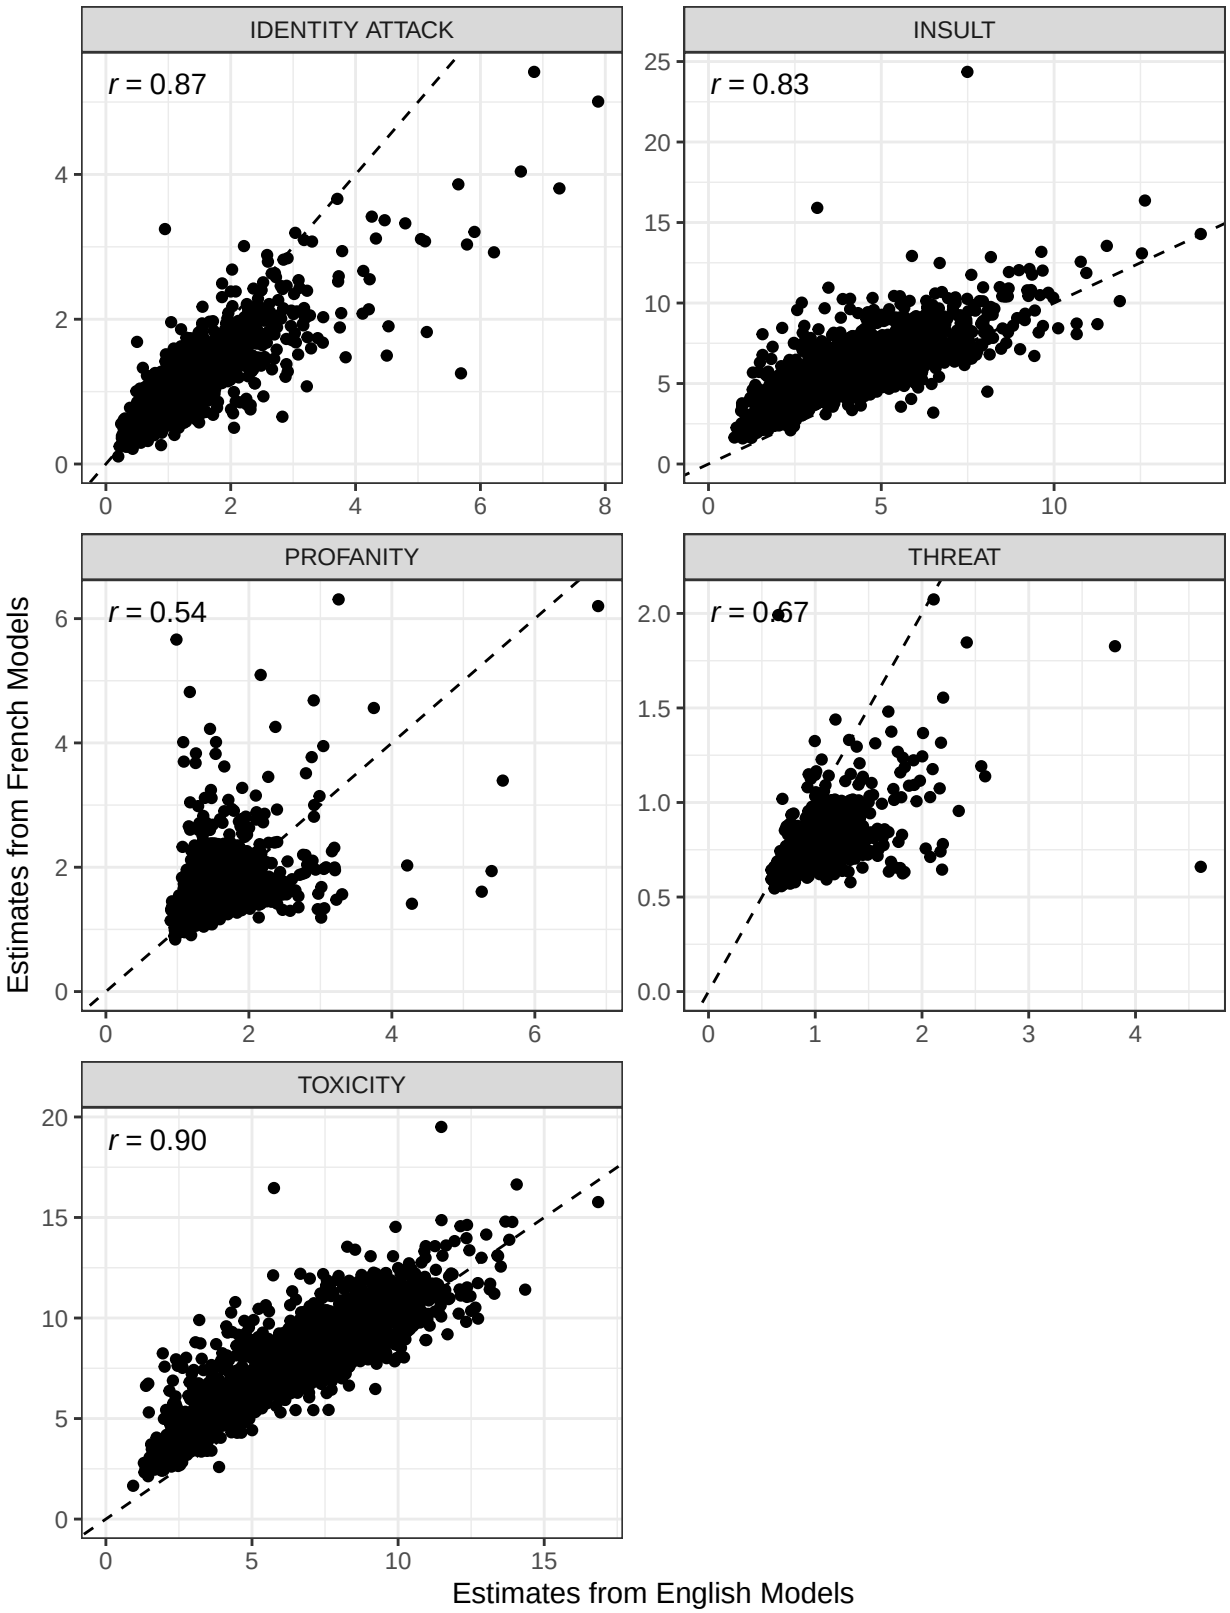

Figure G4: Relationship between Estimates from the English and French Models

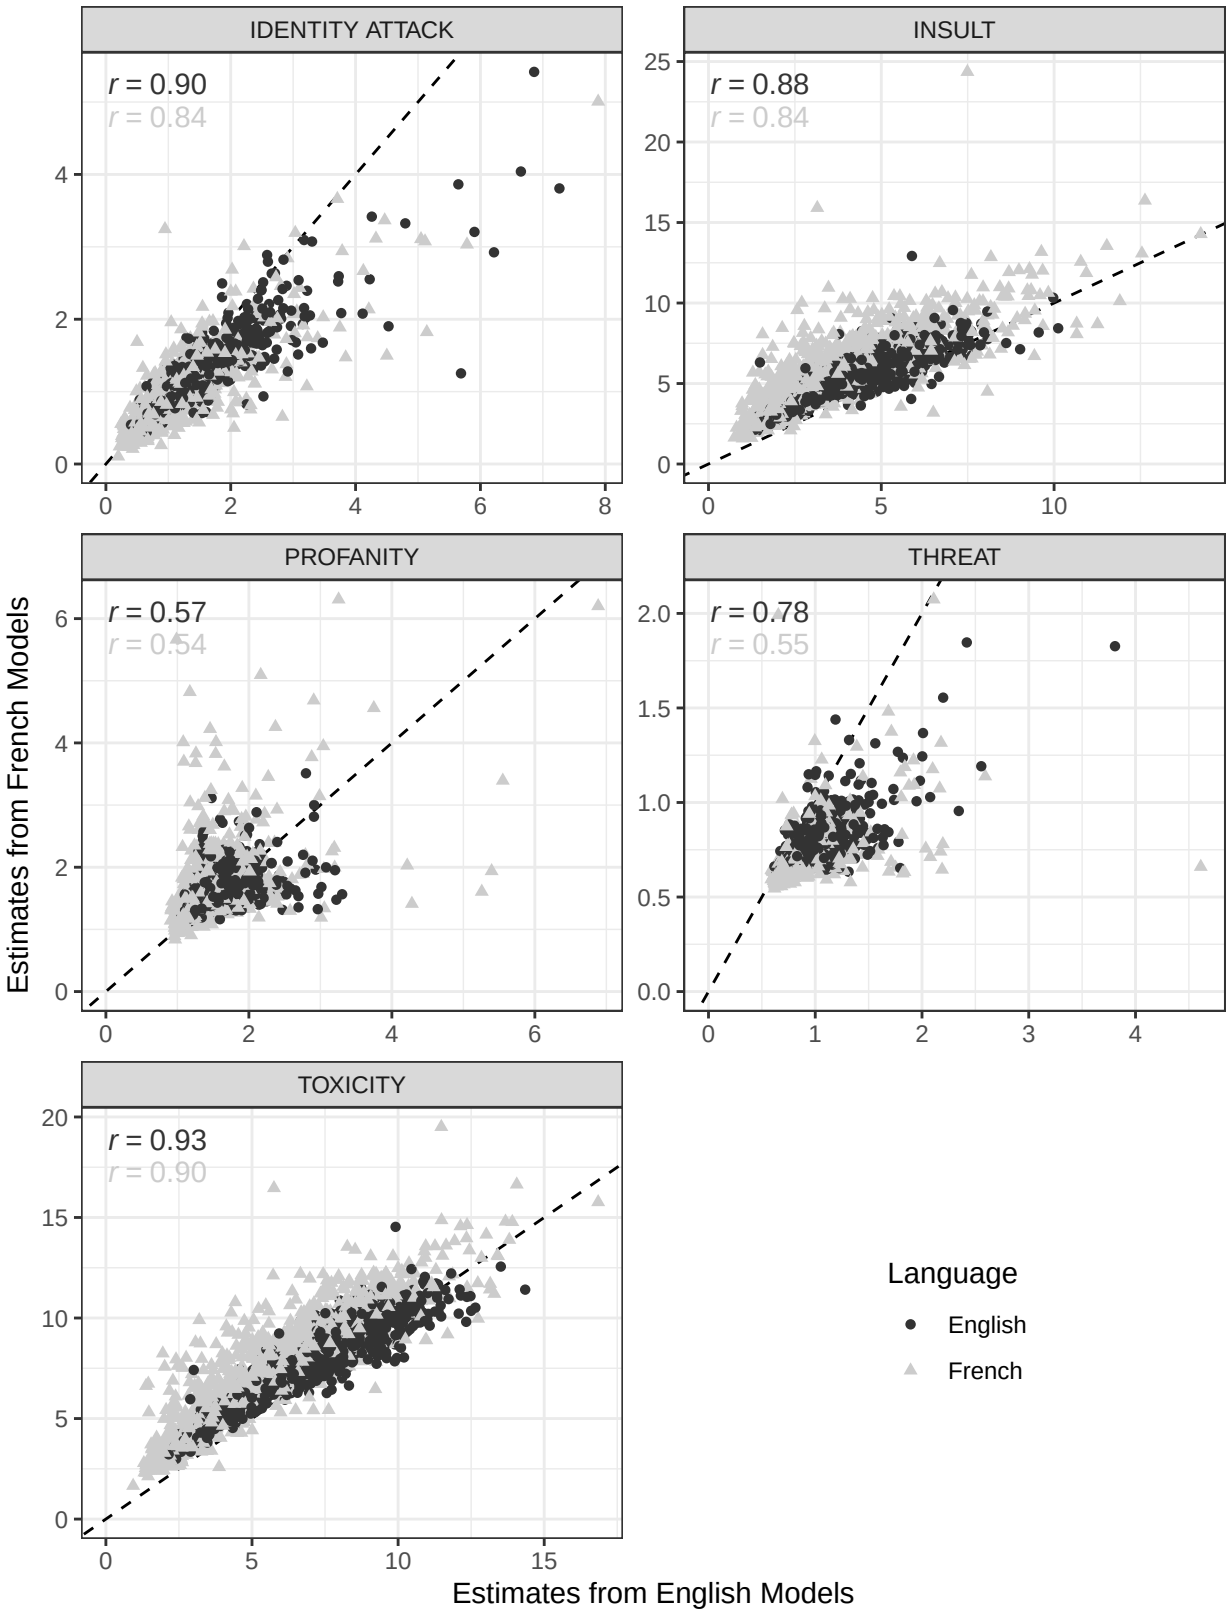

Figure G5: Relationship between Estimates from the English and French Models by Language

Table G4: Regression Results

	(1) IDENTITY ATTACK	(2) INSULT	(3) PROFANITY	(4) THREAT	(5) TOXICITY
Time Until Next Election	−0.002***	−0.001***	−0.0004***	−0.0002***	−0.001***
	(0.0002)	(0.0002)	(0.0001)	(0.0001)	(0.0001)
Time Until Next Election × Minority	−0.001**	−0.0004	−0.001***	−0.0001	−0.0003
	(0.001)	(0.0003)	(0.0002)	(0.0002)	(0.0003)
Government	−0.536***	−0.747***	−0.278***	−0.114***	−0.746***
	(0.028)	(0.019)	(0.012)	(0.008)	(0.016)
Minority	−0.089*	−0.004	0.019	−0.023*	−0.031
	(0.045)	(0.031)	(0.020)	(0.014)	(0.026)
Language: French	−0.248***	−0.001	−0.023***	−0.086***	−0.076***
	(0.012)	(0.010)	(0.008)	(0.004)	(0.008)
Party: LPC	−0.137***	−0.165***	−0.038***	−0.047***	−0.134***
	(0.028)	(0.019)	(0.012)	(0.008)	(0.016)
Party: NDP	−0.0001	−0.146***	0.004	−0.017*	−0.080***
	(0.032)	(0.021)	(0.014)	(0.010)	(0.018)
Constant	−4.161***	−2.525***	−4.035***	−4.828***	−2.126***
	(0.039)	(0.027)	(0.018)	(0.012)	(0.023)
Observations	2,130	2,130	2,130	2,130	2,130
R^2	0.350	0.493	0.288	0.245	0.600
Adjusted R^2	0.348	0.492	0.286	0.244	0.599
Durbin–Watson Statistic	2.154	2.121	2.028	2.062	2.137

Note: *p<0.1; **p<0.05; ***p<0.01

Table G5: Regression Results

	(1) IDENTITY ATTACK	(2) INSULT	(3) PROFANITY	(4) THREAT	(5) TOXICITY
Time Until Next Election	-0.001^{***}	-0.001^{*}	-0.0004^{***}	-0.0002^{***}	-0.001^{**}
	(0.0001)	(0.0004)	(0.0001)	(0.00004)	(0.0005)
Time Until Next Election × Minority	−0.0003	0.0002	−0.0001	−0.0001	0.0001
	(0.001)	(0.001)	(0.0001)	(0.0002)	(0.0005)
Government	-0.594^{***}	-0.700^{***}	-0.235^{***}	-0.088^{***}	-0.818^{***}
	(0.056)	(0.034)	(0.015)	(0.009)	(0.042)
Minority	−0.126	−0.047	−0.016	−0.016	−0.073
	(0.115)	(0.053)	(0.021)	(0.017)	(0.049)
Language: French	-0.245^{***}	-0.049^{*}	-0.047^{***}	-0.059^{***}	-0.156^{***}
	(0.038)	(0.025)	(0.017)	(0.006)	(0.036)
Party: LPC	0.015	-0.135^{***}	−0.019	−0.011	-0.105^{***}
	(0.050)	(0.029)	(0.013)	(0.008)	(0.040)
Party: NDP	0.108	−0.060	0.005	0.002	−0.013
	(0.068)	(0.053)	(0.022)	(0.011)	(0.058)
Constant	-4.684^{***}	-2.795^{***}	-4.140^{***}	-4.915^{***}	-2.301^{***}
	(0.075)	(0.064)	(0.026)	(0.008)	(0.076)
Observations	123,520	123,520	123,520	123,520	123,520
R^2	0.104	0.180	0.095	0.045	0.199
Adjusted R^2	0.103	0.180	0.095	0.045	0.199

Note: *p<0.1; **p<0.05; ***p<0.01