

# Observing the Southern US Culture of Honor Using Large-Scale Social Media Analysis

**Juho Kim**

Faculty of Applied Science and Engineering  
University of Toronto  
Toronto, Ontario, Canada  
juho.kim@mail.utoronto.ca

**Michael Guerzhoy**

Division of Engineering Science  
University of Toronto  
Toronto, Ontario, Canada  
guerzhoy@cs.toronto.edu

## Abstract

A *culture of honor* refers to a social system where individuals' status, reputation, and esteem play a central role in governing interpersonal relations. Past works have associated this concept with the United States (US) South and related with it various traits such as higher sensitivity to insult, a higher value on reputation, and a tendency to react violently to insults. In this paper, we hypothesize and confirm that internet users from the US South, where a *culture of honor* is more prevalent, are more likely to display a trait predicted by their belonging to a *culture of honor*. Specifically, we test the hypothesis that US Southerners are more likely to retaliate to personal attacks by personally attacking back. We leverage OpenAI's GPT-3.5 API to both geolocate internet users and to automatically detect whether users are insulting each other. We validate the use of GPT-3.5 by measuring its performance on manually-labeled subsets of the data. Our work demonstrates the potential of formulating a hypothesis based on a conceptual framework, operationalizing it in a way that is amenable to large-scale LLM-aided analysis, manually validating the use of the LLM, and drawing a conclusion.

## 1 Introduction

A *culture of honor* refers to a social system where individuals' status, reputation, and esteem play a central role in governing interpersonal relations. In such cultures, maintaining and defending one's honor and that of one's family or group is of paramount importance. A perceived slight, insult, or challenge to one's honor often necessitates a response, which could range from verbal defense to physical retaliation, to restore the lost esteem and reputation. These cultures are sometimes theorized as emerging in societies where centralized authority is weak or absent, and where individuals must rely on their reputation and the fear of retaliation to deter aggression or mistreatment by others.

We hypothesize and confirm that internet users who belong to a region where a *culture of honor* is more prevalent would be more likely to display traits predicted by their belonging to a *culture of honor*. Specifically, we test the hypothesis that US Southerners are more likely to retaliate to personal attacks by attacking their attacker back.

To analyze data at scale, we leverage OpenAI's API for GPT-3.5 (Brown et al., 2020) to both geolocate internet users and to automatically detect whether users are insulting each other.

The concept of honor has been explored in social psychology, philosophy, and literature, with authors often developing a nuanced and intricate conceptual framework. The concept has been empirically validated, notably by Nisbett and Cohen (1996), on small samples.

However, large-scale empirical analysis of the concept of honor has been very expensive until very recently. We show that LLMs, and specifically GPT-3.5, can be used to address this. In this work, we explore a specific prediction made by the conceptual framework of the *culture of honor* (and the claim that the US South is such a culture).

We emphasize connecting the conceptual framework of *culture of honor* to testable empirical hypotheses. We propose further testable hypotheses arising from philosophical and descriptive work on *culture of honor* that are also testable by leveraging LLMs to label data while validating the outputs of the LLMs.

## 2 Related Works

Nisbett and Cohen (1996) pioneered the concept of a *culture of honor* as it pertains to the US South, identifying traits such as a higher sensitivity to insult, a higher value on reputation, and a tendency to react violently to insults.

To further support their hypothesis, Nisbett and Cohen (1996) have conducted both observational

and empirical studies on groups of Northern and Southern college students. In a series of experiments, the subjects were put into a situation where one would feel “diminished.” Then, the researchers examined whether the subjects would “take [an] aggressive action to compensate for the diminishment they experience.” The results conclusively demonstrated that people in the Southern US tend to exhibit the traits associated with *cultures of honor* than their Northern counterparts. Taking inspiration from their previous work, our study extends Nisbett and Cohen (1996)’s work in the online domain.

The concept of honor has been explored since ancient times, e.g., in the Bible and by Sophocles in, e.g., *Antigone*<sup>1</sup>. More recently, Appiah (2011) and Sommers (2018) connected the ancient concept, Nisbett and Cohen (1996)’s work, and modern philosophy. We draw on more recent philosophical work to make concrete predictions about internet denizens who belong to *cultures of honor*.

## 2.1 How a Culture of Honor Might Manifest on the Internet

In a *culture of honor*, a perceived attack or slight calls for retaliation to defend one’s honor and reputation (Sommers, 2018; Appiah, 2011). On an internet forum where users have fixed usernames, and particularly for users whose username can be connected to other social media possibly to their real name, it is plausible that members of *cultures of honor* would be more likely to retaliate against verbal attacks. We hypothesize that more members of *cultures of honor* would be geolocated to the Southern US and that those geolocated to the Southern US would therefore be more likely to verbally retaliate in internet arguments.

There is prior work on detection aggression with both classical and Transformer-based methods (Warner and Hirschberg, 2012; Sadiq et al., 2021; Ghosh et al., 2023), with the best results reporting > 90% precision/recall figures. Ziems et al. (2024) predict future (as opposed to present) toxicity (as opposed to aggressiveness) on a different subset of the dataset we use (Zhang et al., 2018), with the best results obtained by fine-tuning RoBERTa (Liu et al., 2019), at 64% F1.

<sup>1</sup>We distinguish between exploring and conceptualizing honor from referring to the psychological and cultural construct. Other ancient sources that do not necessarily conceptualize and explore honor but do refer to it extensively include Homer and Confucius.

## 3 Methodology

We obtain data on internet user conversations and fine-tune GPT-3.5 models to assign users to the corresponding US regions (South, non-South, or N/A) as well as to label each part of the conversation as whether or not it constitutes a personal attack. We then analyze the data to compare the rate at which Southern and non-Southern users retaliate when personally attacked by personally attacking back.

The subreddit-changemyview dataset from ConvoKit (Chang et al., 2020) was considered for our analysis. While we also analyzed other datasets in ConvoKit – reddit-corporus-small, wikiconv (Hua et al., 2018) (from 2015 to 2018, inclusive), and wiki-corporus (Danescu-Niculescu-Mizil et al., 2012) – we were able to obtain data of interactions involving personal attacks with sufficient statistical power in the subreddit-changemyview dataset only.

We attempted to locate the users by fetching the self-declared locations from user profiles with the matching usernames on the social media platform X, formerly known as Twitter. Note that not all self-declared locations could be associated with a US region because some locations are not in the United States or do not specify a US region (e.g. “Moon,” “United States,” or “Barcelona, Spain”). For the subreddit-changemyview dataset, the number of speakers for which we were able to associate a location and the number of speakers for which we were able to associate a US region are tabulated in Table 1. For an expanded table showing the numbers for datasets other than subreddit-changemyview, please consult Appendix A.

In order to obtain the location-US region pairs for fine-tuning and validating our geolocation model, samples of the collected locations were manually labeled by a human annotator and divided into training and validation sets (both of size 100).

For the personal attack classifier, we used a pre-labeled dataset of personal attacks (Zhang et al., 2018). The dataset Zhang et al. was heavily unbalanced with 2094 positive and 27927 (3833 of which are section headers) negative labels. 2094 non-section header negative entries were randomly selected and merged with the positively labeled entries. Then, they were evenly divided into perfectly balanced training and validation sets.

The training sets for each task were then used to fine-tune (Bommasani et al., 2021) GPT-3.5 mod-

Table 1: The number of speakers with whom we were able to associate a location (through the matching X profile’s self-declared location) and the number of speakers with whom we were able to associate a US region (as the associated location contained enough information to discern their US region). In this table, we only show the `subreddit-changemyview` dataset since the others do not contain a sufficient number of attacks. An expanded version of this table, showing all other datasets is shown in Appendix A.

Dataset	# Speakers	With Location (%)	With US Region (%)
<code>subreddit-changemyview</code>	119889	7.6	2.9

els for the tasks of detecting personal attacks and identifying the US regions (for more details, see Appendix B). We then test the fine-tuned GPT-3.5 models by using the validation datasets. This step is essential: we empirically demonstrate that the fine-tuned models work as expected rather than blindly relying on GPT outputs.

Note that the high validation accuracies (96.0% for personal attack detector and 100.0% for US region classifier) show that the training set size used to fine-tune the two models is sufficient and that the fine-tuned models perform very accurately.

We are interested in testing whether US Southerners are more likely to retaliate with personal attacks when personally attacked, compared to their non-Southern counterparts. We measure this tendency by comparing their “retaliation” rates. Our main hypothesis is that the users from the US South have a higher “retaliation” rate than their non-Southern counterparts.

We also introduce two other metrics: “aggression” and “response” rates to see if the Southerners and non-Southerners show different tendencies in other regards. The three metrics we use are defined as follows:

1. Aggression: The rates of posting personal attacks.
2. Response: The rates of responding back if personally attacked.
3. Retaliation: The rates of personally attacking back when responding to a personal attack.

These three rates are computed from the aforementioned conversational datasets in two different counting schemes: per user and per interaction (for users in multiple interactions, the rate is the average of 0s and 1s, with 1s indicating a positive label).

We also review some conversations containing the postings we flag as retaliatory to see if they indeed are examples of online fights in order to ensure the correctness of our analysis.

## 4 Analysis




Figure 1: The bar graphs of all rates for the `subreddit-changemyview` dataset by metrics and US regions. The metric labels “AGG,” “RESP,” and “RET” represent aggression, response, and retaliation rates, respectively. Note that the retaliation rates corresponding to the US South are greater than or equal to those of the US non-South.

The computed rates for each scenario are summarized in Table 2. Only the dataset from `subreddit-changemyview` contains samples with enough statistical power and therefore is the only one considered for further analysis.

Note that it is not appropriate to test statistical hypotheses when a test is underpowered (Button et al., 2013), nor is it appropriate to draw conclusions if a statistical test cannot be run. We include a summary for datasets other than `subreddit-changemyview` for completeness and transparency in Appendix C.

For the `subreddit-changemyview` dataset, both the per-interaction and per-speaker rates for this dataset are graphed in Figure 1. While the differences in the per-speaker rates across the US regions are marginal for aggression (+0.3%) and non-existent for response (0.0%), a notable difference (+5.6%) in the retaliation rate is observed.

We fit a mixed-effect model for the per-speaker retaliation rates in `subreddit-changemyview`:

Table 2: The per-speaker and per-interaction rates for each metric and US region. The metric labels “AGG,” “RESP,” and “RET” represent aggression, response, and retaliation rates, respectively. For each rate, the sample size is equal to the number of people who posted, were personally attacked, and responded to a personal attack at least once, respectively. In this table, we only show the subreddit-changemyview dataset since the others do not contain a sufficient number of attacks. An expanded version of this table, showing all other datasets is shown in Appendix C.

Dataset	Per	Metric	Non-South		South	
			Rate (%)	# Samples	Rate (%)	# Samples
subreddit-changemyview	Speaker	AGG	9.1	4491	9.4	1674
		RESP	51.4	1318	51.4	453
		RET	20.6	849	25.3	292
	Interaction	AGG	8.8	8507/96146	9.4	4820/51365
		RESP	54.5	4113/7552	56.3	2658/4717
		RET	23.0	947/4113	28.6	761/2658

$$\text{retaliation}[i] \sim \text{Bernoulli}(\text{logit}^{-1}(\beta_0 + \alpha_{\text{regions}[\text{speakers}[i]]} + \gamma_{\text{regions}[i]})) \quad (1)$$

$$\alpha_j \sim \mathcal{N}(0, \sigma_\alpha) \quad (2)$$

In the above Equations 1 and 2, for interaction  $i$  between two users, the original poster “speakers[ $i$ ]” (who is also a potential retaliator) did retaliate if “retaliation[ $i$ ] = 1”. The index “regions[speakers[ $i$ ]]” is the region (South/non-South) to which the speaker was geolocated.  $\gamma_{\text{regions}[i]}$  is the fixed effect of regions “regions[ $i$ ]” (the main coefficient of interest in this analysis), and  $\alpha_j$  is the random effect of speaker  $j$ .

Priors over parameters that are not mentioned above are flat.

The mixed-effect model estimates the speaker random effects using partial pooling, which is appropriate in our situation, where the large majority of speakers have only a few (or a single) interactions.

We find that people from the US South are more likely to retaliate (odds ratio  $1.2 \pm 0.1$ , p-value  $< 0.002$ ) than people from outside the US South. The R code used to carry out this analysis is shown in Appendix D.

To further validate our analysis, 100 random samples of flagged retaliatory conversations for all datasets were reviewed to see if they were part of an ongoing online fight. For the subreddit-changemyview dataset which we used to prove our hypothesis, we find the actual fight percentages to be 86.0%. Similar numbers for datasets other than subreddit-changemyview are observed in Appendix E.

## 5 Discussion

The three landmark experiments performed by Nisbett and Cohen (1996) “did not produce any truly violent behavior in [their] subjects, so it is an extrapolation to say that [they] have shown the process by which an insult results in actual violence for members of a *culture of honor*.” Indeed, they measured more indirect attributes of *cultures of honor* such as ratings given by third-party observers and physiological changes to draw their conclusions. In contrast, our methodology demonstrated that the Southerners are more likely to signal verbal violence or capacity for verbal violence to protect their reputation in their online interactions.

## 6 Conclusion

In this study, we were able to confirm a theoretically motivated prediction: internet users from the US South are more likely to retaliate against verbal aggression on the theory that reputation must be aggressively defended in a *culture of honor*.

We show that analyses such as ours, on an extremely large scale, are possible by fine-tuning GPT-3.5 and manually verifying its outputs on samples of the inputs. By following similar processes laid out in this paper, analyses other than comparing the US South and the US non-South populations or testing hypotheses presented by social science frameworks aside from *culture of honor* are also possible.

## Limitations

Our analysis is limited to the US and to English-speaking users on Reddit. *Cultures of honor* as they manifest online may in principle be analyzed for other countries and languages.

The geolocation of the user may be incorrect, and, while the concentration of people who can be said to belong to a *culture of honor* to some extent in the US South is thought to be higher, it is neither the case that all US Southerners are members of a *culture of honor* or that there are no member of a *culture of honor* elsewhere. The geolocation we obtain is merely a proxy variable. Note, however, that one in general would expect it to be more difficult to detect a trend using a proxy variable than using a direct measurement.

We validate our LLM-based analysis by checking that the LLM output is consistent with our hand-labeled sample. However, it is possible that LLMs are poorly calibrated for, for example, detecting insults in certain underrepresented dialects. Although our analysis of the hand-labeled sample means that our classifiers are correct on average, it is possible that a manual analysis would drive researchers to discover particular patterns for subgroups of the users. For example, if insults in a particular small dialect are not detectable by an LLM, it is conceivable that for speakers of that dialect, our findings would be reversed.

Our statistical analysis does not account for attenuation and error-in-variables. That is accepted practice in many analyses, where the effect of those is very difficult to estimate (and one can argue that the arbitrary 95% significance threshold would be different if one were to upper-bound the effects of attenuation and error-in-variables).

## Ethics Statement

All research carried out in this paper was based on publicly available datasets. However, some users in the datasets may not want to be associated with *cultures of honor*. We only use the geolocation of users in our analysis.

Membership in a *culture of honor* can be viewed as a possibly harmful stereotype. It is important to approach broad cultural labels with nuance, humility, and respect, and to not apply broad stereotypes to individuals. However, we believe that research into *cultures of honor* is appropriate: honor, as a sociological and psychological phenomenon, has fascinated humanity for millennia, and understanding *cultures of honor* is important both philosophically and practically. Further, most members of *cultures of honor* are proud of their membership (Sommers, 2018) and do not view affinity to that kind of culture as a harmful stereotype.

Our research only uses publically-available data and as such is not human-subjects research and is exempt from ethics board review.

The relevant subset of the datasets we used with labels from the fine-tuned models is available on Zenodo (Kim and Guerzhoy, 2024).

## References

- K.A. Appiah. 2011. *The Honor Code: How Moral Revolutions Happen*. WW Norton.
- Rishi Bommasani, Drew A Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, et al. 2021. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language Models are Few-Shot Learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Katherine S Button, John PA Ioannidis, Claire Mokrysz, Brian A Nosek, Jonathan Flint, Emma SJ Robinson, and Marcus R Munafò. 2013. Power failure: why small sample size undermines the reliability of neuroscience. *Nature reviews neuroscience*, 14(5):365–376.
- Jonathan P. Chang, Caleb Chiam, Liye Fu, Andrew Wang, Justine Zhang, and Cristian Danescu-Niculescu-Mizil. 2020. [ConvoKit: A toolkit for the analysis of conversations](#). In *Proceedings of the 21th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 57–60, 1st virtual meeting. Association for Computational Linguistics.
- Cristian Danescu-Niculescu-Mizil, Lillian Lee, Bo Pang, and Jon Kleinberg. 2012. Echoes of power: Language effects and power differences in social interaction. In *Proceedings of WWW*, pages 699–708.
- Soumitra Ghosh, Amit Priyankar, Asif Ekbal, and Pushpak Bhattacharyya. 2023. [A transformer-based multi-task framework for joint detection of aggression and hate on social media data](#). *Natural Language Engineering*, 29(6):1495–1515.

Yiqing Hua, Cristian Danescu-Niculescu-Mizil, Dario Taraborelli, Nithum Thain, Jeffery Sorensen, and Lucas Dixon. 2018. [WikiConv: A corpus of the complete conversational history of a large online collaborative community](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2818–2823, Brussels, Belgium. Association for Computational Linguistics.

Juho Kim and Michael Guerzhoy. 2024. [Supplementary: Observing the Southern US Honor Culture Using Large-Scale Social Media Analysis](#).

Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.

Richard Eugene Nisbett and Dov Cohen. 1996. *Culture of Honor: The Psychology of Violence in the South*. Westview Press.

Saima Sadiq, Arif Mehmood, Saleem Ullah, Maqsood Ahmad, Gyu Sang Choi, and Byung-Won On. 2021. [Aggression detection through deep neural model on twitter](#). *Future Generation Computer Systems*, 114:120–129.

T. Sommers. 2018. *Why Honor Matters*. Basic Books.

William Warner and Julia Hirschberg. 2012. Detecting hate speech on the world wide web. In *Proceedings of the second workshop on language in social media*, pages 19–26.

Justine Zhang, Jonathan Chang, Cristian Danescu-Niculescu-Mizil, Lucas Dixon, Yiqing Hua, Dario Taraborelli, and Nithum Thain. 2018. [Conversations gone awry: Detecting early signs of conversational failure](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1350–1361, Melbourne, Australia. Association for Computational Linguistics.

Caleb Ziems, William Held, Omar Shaikh, Jiaao Chen, Zhehao Zhang, and Diyi Yang. 2024. [Can Large Language Models Transform Computational Social Science?](#) *Computational Linguistics*, 50(1):237–291.

## A Statistics on the Users with Identifiable US Region Information

We attempted to locate the users from both Reddit and Wikipedia by fetching the self-declared locations from user profiles with the matching usernames on the social media platform X, formerly known as Twitter. In Wikipedia, anonymous editors are assigned their IP addresses as their usernames from which the associated locations can be obtained (including the associated US region). The full statistics regarding user locations for each dataset considered are shown in Table 3.

## B Information on the Fine-Tuned GPT-3.5 Models

Tables 4 and 5 describe information relevant to the fine-tuning of the GPT-3.5 models for the tasks of US region classification and personal attack detection, respectively.

## C Aggression, Response, and Retaliation Rates

In Table 6, we give the per-speaker and per-interaction aggression, response, and retaliation rates for every considered dataset. Note that due to the insufficient statistical power in reddit-corpus-small, wikiconf, and wiki-corpus, their numbers should not be used to draw any conclusions about the hypotheses (Button et al., 2013).

## D R Code for Fitting a Mixed-Effect Model

The mixed-effect model for the per-speaker retaliation rates for the subreddit-changemyview dataset is fit using the lme4 package in R.

```
glmer(retaliation ~ speaker_2_us_region +  
      (1 | speaker_2_id), family = "binomial")
```

Here, the variables speaker\_2\_us\_region and speaker\_2\_id are the region and the id of a person who started an interaction, was personally attacked, and may potentially retaliate back.

## E Statistics on the Ongoing Fights Detected by our System

Up to 100 samples of user interactions in each dataset (including the ones that were not analyzed) flagged to involve personal attacks were analyzed by a human annotator to see if they truly constitute an ongoing online fight. As tabulated in Table 7, we found that, for the subreddit-changemyview dataset, the true ongoing fight rate was 86.0%, high enough to ensure our analysis is correct. Note that the ongoing fight rates vary from 71.0% (reddit-corpus-small, 75.0% (wiki-corpus), to 91.0% (wikiconv) for the other datasets.

Table 3: The number of speakers with an IP address as its username (relevant to Wikipedia only), the number of speakers with whom we were able to associate a location (either through the matching X profile’s self-declared location or their IP address), and the number of speakers with whom we were able to associate a US region (as the associated location contained enough information to discern their US region) for each dataset.

Dataset	# Speakers	With IP (%)	With Location (%)	With US Region (%)
subreddit-changemyview	<b>119889</b>	–	<b>7.6</b>	<b>2.9</b>
reddit-corpus-small	217100	–	7.5	2.8
wikiconv	621142	39.1	40.9	16.7
wiki-corpus	38462	7.8	11.4	4.8

Table 4: The system prompts, default classifications (in case the model produces an unexpected output), accuracies, and training/validation information of the fine-tuned GPT-3.5 model for the task of US region classification. Note that the system prompt was broken to fit it in the page. The actual prompt is single-line.

Task	US Region Classification
System Prompt	What US region is the following location in? Possible answers are “SOUTH”, “NON-SOUTH”, or “N/A”.
Expected input	User location information
Default Label	N/A
Accuracy (%)	100.0
# Training Samples	100
# Validation Samples	100
Data Source	Hand-crafted
Base Model	gpt-3.5-turbo-1106
Trained Tokens	13491
# Epochs	3
Final Training Loss	0.0000
# Steps	291
Training Time	00:13:18

Table 5: The system prompts, default classifications (in case the model produces an unexpected output), accuracies, and training/validation information of the fine-tuned GPT-3.5 model for the task of personal attack detection. Note that the system prompt was broken to fit it in the page. The actual prompt is single-line.

Task	Personal Attack Detection
System Prompt	Does the following text contain a personal attack? Possible answers are “YES” or “NO”.
Expected input	User post
Default Label	NO
Accuracy (%)	96.0
# Training Samples	2094
# Validation Samples	2094
Data Source	<a href="#">Zhang et al. (2018)</a>
Base Model	gpt-3.5-turbo-1106
Trained Tokens	816762
# Epochs	3
Final Training Loss	0.0000
# Steps	1501
Training Time	01:15:17

Table 6: The per-speaker and per-interaction rates for each metric, US region, and dataset. The metric labels “AGG,” “RESP,” and “RET” represent aggression, response, and retaliation rates, respectively. For each rate, the sample size is equal to the number of people who posted, were personally attacked, and responded to a personal attack at least once, respectively. In this paper, we only analyze the subreddit-changemyview dataset since the others do not contain a sufficient number of interactions involving personal attacks.

Dataset	Per	Metric	Non-South		South	
			Rate (%)	# Samples	Rate (%)	# Samples
subreddit-changemyview	Speaker	AGG	<b>9.1</b>	<b>4491</b>	<b>9.4</b>	<b>1674</b>
		RESP	<b>51.4</b>	<b>1318</b>	<b>51.4</b>	<b>453</b>
		RET	<b>20.6</b>	<b>849</b>	<b>25.3</b>	<b>292</b>
	Interaction	AGG	<b>8.8</b>	<b>8507/96146</b>	<b>9.4</b>	<b>4820/51365</b>
		RESP	<b>54.5</b>	<b>4113/7552</b>	<b>56.3</b>	<b>2658/4717</b>
		RET	<b>23.0</b>	<b>947/4113</b>	<b>28.6</b>	<b>761/2658</b>
reddit-corpus-small	Speaker	AGG	16.5	2522	16.1	943
		RESP	52.1	528	46.7	171
		RET	26.0	211	15.4	58
	Interaction	AGG	17.0	982/5780	14.7	277/1878
		RESP	34.4	342/993	31.8	90/283
		RET	29.2	100/342	15.6	14/90
wikiconv	Speaker	AGG	5.5	80965	6.1	22486
		RESP	46.3	1365	43.1	350
		RET	18.2	142	22.7	37
	Interaction	AGG	2.8	15627/556156	2.2	2793/125372
		RESP	14.2	271/1912	16.1	75/467
		RET	19.2	52/271	20.0	15/75
wiki-corpus	Speaker	AGG	5.4	1429	7.5	400
		RESP	35.8	114	35.7	29
		RET	19.1	29	16.7	6
	Interaction	AGG	3.0	287/9594	4.4	69/1578
		RESP	23.9	47/197	15.4	6/39
		RET	23.4	11/47	16.7	1/6

Table 7: The percentages of the sampled retaliatory interactions that a human annotator labeled as an ongoing fight. By showing that most samples we flagged as retaliatory to indeed be a part of an ongoing online fight, we validate the robustness of our analysis.

Dataset	Ongoing Fights	
	Rate (%)	# Samples
subreddit-changemyview	<b>86.0</b>	<b>86/100</b>
reddit-corpus-small	71.0	71/100
wikiconv	91.0	61/67
wiki-corpus	75.0	9/12