

Supplement to:

Zhang, Yinxian, and Di Zhou. 2025. "One Sentiment, Multiple Interpretations: Contrasting Official and Popular Anti-Americanism in China" *Sociological Science* 12: 511-536.

Online Supplement for

One Sentiment, Multiple Interpretations: Contrasting Official and Popular Anti-Americanism in China

1. Data Collection

We collected *Zhihu* discussions and *People's Daily* (*PD* hereafter) articles about the U.S. to investigate and contrast popular and official anti-Americanism.

1.1 *Zhihu*

Data collection

In December 2022, we retrospectively scraped all publicly accessible *Zhihu* questions and answers related to the U.S. that were posted between 2011, the year when *Zhihu* first launched, and 2022, the end of our data collection period.

Zhihu contents are mandatorily organized by topic tags. Under the first-level topic tag “USA”, we first ranked all subtopics by discussion volume, i.e., the number of answer posts each subtopic accumulates. Then, among the top 30 most popular subtopics, we only included topic tags that are relevant to social, political, and economic issues. Niche topics such as “Disney” and “NFL” and less relevant topics such as “U.S. Visa” and “U.S. Law” were excluded. Lastly, we limited to the subtopics that have been used since 2011 to ensure a full coverage of the study period. Consequently, three *Zhihu* topic tags – “U.S. politics”, “U.S. economy” and “U.S. society” were included.

To clarify, a *Zhihu* post can be tagged with multiple topics. For instance, posts about American law enforcement may be double tagged as “U.S. Law” and “U.S. politics”. Therefore, even if certain subtopics were excluded, posts under these topics could still be included in our data if they were relevant for American political or socioeconomic issues and were tagged as such.

Note that, as shown in Figure 1 in the main text, *Zhihu* had much fewer data points in earlier years. This is primarily because it was a smaller, invite-only community before 2013, but it was also due to the fact that earlier online content has been disappearing from the internet. A recent Pew Research Center reports that 38% of webpages that existed in 2013 are no longer accessible a decade later¹, potentially due to costs of

¹ See the report here: <https://www.pewresearch.org/data-labs/2024/05/17/when-online-content-disappears/> (accessed March 31, 2025)

storage and maintenance. Finally, while censorship could be another factor, Chinese censorship was ramped up in recent years (especially after 2019), not the other way around. See the main text for a detailed discussion of the censorship effect.

To partly address the issue of data loss in earlier years, we supplemented the data with a secondary *Zhihu* dataset obtained from Zhou (2022), which was retrospectively collected in August 2020. It helped us recover some but certainly not all of the earlier *Zhihu* posts. As a result, *Zhihu* trends in earlier years were based on fewer data points and thus had a larger error margin.

Segmentation and processing

In total, we collected 215,708 *Zhihu* answers associated with the three topic tags. We then cleaned the raw data and retrieved text segments substantively about the U.S. To do this, we first compiled a list of keywords denoting “the U.S. (*meiguo*)” and its colloquial variants, including America (*meiguo*), the imperial America (*meidi/mei diguo zhuyi*), United State of America (*meilijian*), American government (*meiguo zhengfu*), US-China (*meizhong*), American style (*meishi/meiguo shi*), and a few internet slangs referring to the U.S. (*miguo/laomei/meili guo/piaoliang guo/dengta guo*). These keywords were compiled based on the results of a word embedding model that identified semantically close terms to the term “the U.S. (*meiguo*)”, and our local knowledge of the Chinese online discussions.

Then, we extracted the nearest three sentences before and after each of the keywords from each *Zhihu* answer. The nearest 6 sentences (or less, if the answer had less than 3 sentences before/after the keyword) surrounding each of the focal keywords formed a text segment. In total, we retrieved 171,807 text segments from the raw data.

The rest of the analysis were conducted on the segment level, and some metrics were then aggregated back to the answer level, as reported in the main text.

1.2 People's Daily (PD)

We collected *PD* articles from the *People's Daily Database* (<https://www-oriprobe-com.proxy.library.nyu.edu/peoplesdaily.shtml>) in April and May 2024. We collected all articles published between 2011 and 2022 whose titles contain the keyword “the U.S. (*meiguo*).” Unlike social media content, the language in official media content is formal and standard, we thus did not include the colloquial variants or internet slangs. In total, we collected 3,548 articles.

Given the nature and structure of official media news articles, we assumed that *PD* articles with a title mentioning “the U.S.” would be focused on the U.S. and thus did not further segmentize the articles. We directly conducted sentiment analysis and qualitative thematic analysis on the article level.

2. Sentiment classification

To identify anti-American content, we leveraged GPT models to conduct sentiment analysis on both *Zhihu* and *PD* datasets.

2.1 *Zhihu*

Sentiment analysis on *Zhihu* was conducted on the segment level. We followed a two-step workflow. First, we trained a relevance classifier to identify text segments that were substantively related to the U.S.; Second, we trained a sentiment classifier to identify anti-American content from the *relevant* segments.

Relevance Classifier

We finetuned a GPT babbage-002 model for relevance classification after testing a range of GPT models including curie, davinci-002 and gpt-3.5-turbo-1106. The finetuning was based on our hand coding of a sample of text segments (N=529) proportionally sampled by year from the full set. Both authors hand coded the sampled segments independently following the coding scheme in Table S1. The intercoder reliability statistic, Krippendorff's alpha, is 0.96, indicating a near perfect agreement.

We then used the hand-coded data to finetune the babbage-002 model with default parameters. The prompt for the model was each text segment, and the completion was 1 if relevant and 0 otherwise. The classifier reached an accuracy of 0.83 and weighted F1 of 0.83, which is highly satisfactory. Eventually, among the 171,807 text segments, 150,996 were identified as relevant².

Table S1. Hand Coding Scheme for Relevance Classification, *Zhihu*

Code	Coding scheme
0 / irrelevant	The text segment is not about the U.S. or any typical incidents or figures that represent the U.S., such as Trump or Obama. However, mentioning "the U.S." or a person's name does not automatically pass the test unless the substantive content is related to the country or its figures.
1 / relevant	The text segment is substantively about the U.S., including any incidents or figures that represent the U.S.

Sentiment Classifier

For sentiment classification, we finetuned a gpt-3.5-turbo-1106 model after testing a range of GPT models including curie, davinci-002 and babbage-002. The finetuning was based on our hand coding of the same sample from the first step. After identifying the

² Twelve segments did not receive a valid classification from the relevance classifier and were recorded as "NA" in the dataset.

relevant text segments from the sample, both authors hand coded these segments independently following the coding scheme in Table S2. The intercoder reliability statistic, Krippendorff's alpha, was 0.92, indicating a near perfect agreement.

Table S2. Hand Coding Scheme for Sentiment Classification, *Zhihu*

Code	Coding Scheme
1 / Positive	The text segment expresses an approval of American political/economic/diplomatic/cultural/social values, institutions, and practices. It <i>explicitly</i> supports the U.S. and/or is optimistic about its future even though there may be problems and issues right now. Or it shows support/approval of certain events or figures with the assumption that they have made or will make the U.S. better.
0 / Neutral	The text segment is neutral and/or ambivalent about American political/economic/cultural/social values, institutions, and practices. It does <i>not</i> explicitly support or oppose the U.S. Or it recognizes that there are good and bad aspects about the U.S. and selectively approves some elements but not others.
-1 / Negative	The text segment expresses disapproval and/or doubt of American political/economic/diplomatic/cultural/social values, institutions, and practices. It <i>explicitly</i> condemns or opposes the U.S. Or it shows disapproval of certain incidents or figures representing the U.S.; Or it supports these incidents or figures with the assumption that they have undermined or will undermine the U.S.

We used the hand-coded data to finetune the GPT-3.5-turbo model. We set the temperature parameter to 0 and the max number of tokens to 2. All other parameters were default. Then we prompted the model to classify each *Zhihu* segment according to the same coding scheme. The classifier reached an accuracy of 0.72 and weighted F1 of 0.71.

Eventually, among the 150,996³ text segments coded as "relevant" to the U.S., 76,116 segments were classified as "negative", i.e., anti-American. We further obtained post-level sentiment score by averaging segment-level scores.

PD

Assuming that the titles of official media reports summarize the focus of the article, we directly conducted sentiment analysis on *PD* articles without doing a separate relevance classification.

We used a gpt-4o-2024-05-13 model after experimenting with both gpt-4o and gpt-4-turbo models. We used zero-shot prompting instead of fine-tuning given the satisfactory performance of zero-shot models. Specifically, we set the temperature parameter of the

³ Among the 150,996 relevant segments, 6 did not receive a valid classification from the sentiment classifier and were removed from the final data.

GPT-4o model to 0 and the max numbers of token to 2. Then we prompted the model to classify each *PD* article according to the coding scheme⁴ in Table S3.

Table S3. Coding Scheme for Sentiment Classification, *People's Daily*

Code	Coding Scheme
1 / Positive	This article is about any good or positive events that portray a positive image of the U.S. Or it is a report or commentary that praises the U.S. Or it is about positive diplomatic progress between China and the U.S., or between other countries and the U.S.
0 / Neutral	This article consists of factual descriptions of any events without positive or negative implications for the image of the U.S. Or it is a report or commentary that presents an ambivalent stance or mixed opinions with both positive and negative evaluations of the U.S. Or it is about regular China-U.S. diplomatic communications or factual descriptions of American diplomatic policies and practice.
-1 / Negative	This article is about any bad or negative events that portray a negative image of the U.S., such as natural disasters, social problems, etc. Or it is a report or commentary that criticizes the U.S. Or it is about diplomatic disputes or issues between China and the U.S., or between other countries and the U.S.

The performance of the zero-shot GPT-4o model was evaluated against a set of hand-coded *PD* articles (N = 89) following the same coding scheme. The classifier reached an accuracy of 0.79 and weighted F1 of 0.78. Eventually, 1,909 *PD* articles were identified as negative, i.e., anti-American.

3. LDA topic modeling

After sentiment classification, we zoomed in to anti-American discourses on both platforms. Two corpora were constructed: (1) *Zhihu* anti-American posts, defined as answers with an average negative sentiment score aggregated from segment-level scores (N = 47,489); (2) *PD* anti-American articles, defined as articles with a negative sentiment score (N = 1,909); For each corpus, we then pooled data from different years and used Latent Dirichlet Allocation (LDA), a topic modeling technique (Blei et al., 2003), to identify main topics in official and popular anti-American discourses.

3.1 Methods

LDA takes a number of documents (in this study, *PD* articles or *Zhihu* posts), and it requires researchers to specify the number of topics they expect to find within these documents. Assuming that documents talking about the same topic use similar words,

⁴ The GPT prompts used the Chinese translations of the coding scheme. We also experimented with the English version, but the model performance was slightly worse.

LDA discovers topics by identifying groups of words that frequently co-occur within the documents.⁵

After natural language processing, we tested a set of LDA models with the number of topics ranging from 3 to 10 on both *Zhihu* and *PD* corpus. Based on human judgements and perplexity scores⁶, we found that the LDA model with 8 topics for *Zhihu* and the one with 6 topics for *PD* yield the clearest and most distinct topics. We proceeded with the two selected models.

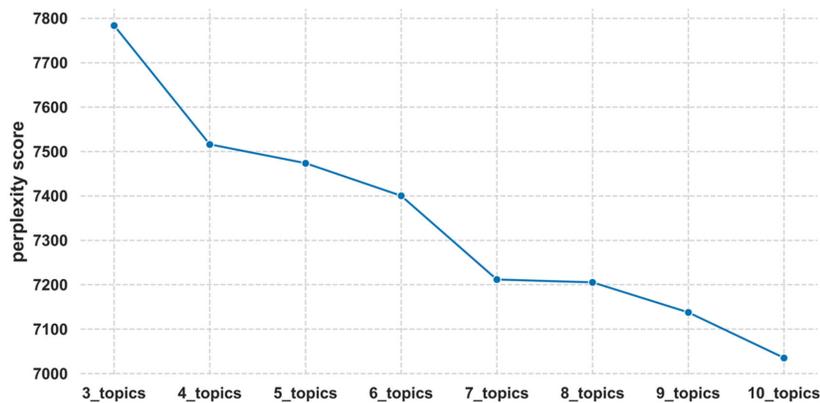


Figure S1. Perplexity score of LDA model by number of topics, *Zhihu* corpus

⁵ LDA assumes that documents are probability distributions over topics, and topics are probability distributions over words. Both of the two probability distributions have a Dirichlet prior to control the sparsity of the distribution. For more details of LDA, see Blei et al. (2003).

⁶ Perplexity is a commonly used quantitative metric to evaluate the performance of topic models. It measures how well the model predicts unseen documents. A lower perplexity score indicates better model performance.

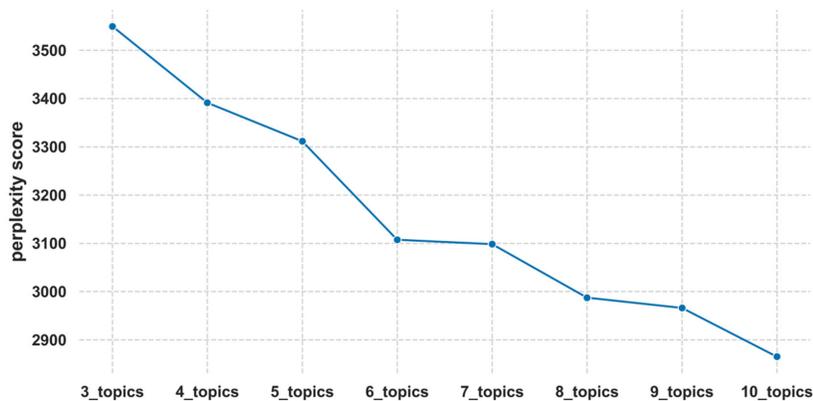


Figure S2. Perplexity score of LDA model by number of topics, *PD* corpus

3.2 Results

The main topics and their associated keywords identified by the LDA models are summarized in Figure 3 in the main text. *Zhihu* and *PD* have three overlapping themes: identity politics, COVID-19, and U.S.-China relations. Identity politics was the most salient topic on both platforms, encompassing discussions around racial, gender and immigration-related issues. This topic is thoroughly discussed in the main text.

In addition to the shared concern on identity politics, *Zhihu* and *PD* also overlaps in their discussion of COVID-19 and U.S.-China relations. However, such topics were smaller, and the most relevant keywords for the two platforms were different, reflecting their varying focuses in the discussions. For example, the discussions about COVID-19 on *Zhihu* primarily contrasted the American government's responses with those of China, suggesting that the Chinese government did a much better job in controlling the pandemic. By contrast, *PD* mainly focused on reporting the cases and the development of the pandemic in China and around the world. In terms of U.S.-China relations, *Zhihu* discussions highlighted the geopolitical conflicts between the U.S. and China, such as Taiwan, while *PD* focused on reporting high-level meetings, dialogues and collaboration strategies between the two countries.

Beside the overlapping topics, *Zhihu* and *PD* also featured unique topics in their anti-American narratives. The second most salient topic on *Zhihu* was about party politics and the political system in the U.S., featuring general discussions about American democracy and its perceived issues. The topic about "elections" further zoomed into the electoral politics and particularly the 2016 and 2020 presidential elections. The topic on "diplomacy" focused on the military and diplomatic practices and the geopolitical

concerns of the U.S. Other topics on *Zhihu* cared more about the economy (economy & consumption and monetary policies).

PD focused on quite different topics. The second most salient topic on *PD* concerned about American trade policies and military operations in Afghanistan, Iraq and Syria, both of which served as evidence of American hegemonism. Another topic focused on American allegations against China, featuring commentaries on how the U.S. interferes with China's domestic affairs and makes false allegations about China's human rights concerns while having poor human rights records at home. The last topic was about the economy and financial policies, particularly about U.S. debts and budgets.

Overall, anti-American discourses on *Zhihu* and *PD* showed both overlaps and divergences. Identity politics was the most significant topic shared by popular and official narratives about the U.S., followed by topics around U.S.-China relations and COVID-19. However, *Zhihu* discourse mainly consisted of general discussions about American political and economic systems with a focus on the social problems of the current American society. By contrast, *PD* focused more on criticizing American political, economic, diplomatic, and military practices, such as foreign trade policies, espionage, and political interference in other country's internal affairs.

4. Qualitative Thematic Analysis

Finally, we delved into *Zhihu* and *PD* discussions around identity politics through qualitative thematic analysis. To identify discussions on identity politics, we matched *Zhihu* posts and *PD* articles with a list of keywords summarized in Table S4. These keywords represent widely debated topics concerning identity politics and progressive movements in the U.S., including Black Lives Matter (BLM) movements for racial topics, feminist and LGBT movements for gender-related topics, and issues around illegal immigrants for immigration-related topics.

Table S4. Keywords for Identifying Discussions Around Identity Politics

Category	Keywords
Race	BLM, <i>heiminggui</i> , ⁷ black lives are also lives (<i>heiren de ming ye shi ming</i>), race (<i>zhongzu</i> , <i>PD</i> only)
Gender	LGBT, feminist rights (<i>nv quan</i>), feminist fists ⁸ (<i>nv quan</i>), gender (<i>xingbie</i> , <i>PD</i> only), female (<i>nv'xing</i> , <i>PD</i> only), women (<i>fu'nv</i> , <i>PD</i> only)
Immigration	Illegal immigrants (<i>feifa yimin</i>)

⁷ "*heiminggui*" and "*heiren de ming ye shi ming*" are two different Chinese translations of "Black Lives Matter". The first term carries a negative undertone as its literal translation is "black lives are precious/expensive," while the second term is a more neutral/positive translation, meaning "black lives are also lives."

⁸ "feminist fists" is a homophone of "feminist rights" in Chinese and carries a negative connotation.

Since very few *PD* articles contained keywords related to LGBT and feminist rights, and they also rarely involved direct references to BLM movements, we supplemented the keywords with a few alternative terms for *PD* only.

403 *PD* articles and 3,721 *Zhihu* segments were matched with at least one of the keywords. The first author then qualitatively analyzed 138 *PD* articles (about 34%) and 607 text segments (about 16%) through thematic analysis.

Thematic analysis is an iterative and reflective process involving six steps as identified by Braun and Clarke (2021). Although this process generally follows a sequential pattern, researchers typically revisit and refine the codes and themes in one step while working on another step as the researchers analyze more data, read additional literature or gain new insights. Following this approach, the first author closely read and coded *Zhihu* segments and *PD* articles to find relevant themes. Multiple iterations of codes and themes were developed during this process until reaching theoretical saturation.

Four main themes were identified for *Zhihu* and *PD* discussions around identity politics, respectively. The two prominent themes in *Zhihu* posts, “anti-baizuo” and “populism”, are thoroughly discussed in the main text. In addition, *Zhihu* discussions also reflected an emphasis on “realism” and “egoism”. Realism highlights the priority of materialistic public goods such as poverty alleviation, economic development, social stability, and campaigns against drugs and crimes. Progressive agenda and identity politics were dismissed as secondary concerns or even distractions to “the more important things.” Egoism blames progressive agenda for neglecting and even hurting Asians for the political and social gains of other minorities, referring to policies like affirmative action. In these narratives, progressive activists were perceived as self-entitled individuals seeking “privileges” and thus jeopardizing rather than promoting social equality.

For *PD* articles, three of four themes are thoroughly discussed in the main text. The remaining theme, “public goods,” is similar to *Zhihu*’s “realism” theme. Despite their opposing rhetoric, *PD* and *Zhihu* shared an emphasis on materialistic public goods. In *Zhihu* discourse, realism emerged as a critique of progressive priorities, lamenting that *baizuo* focused excessively on minorities and identity politics while neglecting realistic problems. By contrast, *PD* articles highlighted the importance of public goods yet argued that the U.S. government fails to adequately address these very issues, especially for minorities and vulnerable populations.

References:

- Blei, David M., Andrew Y. Ng, and Michael I. Jordan. 2003. "Latent Dirichlet Allocation." *Journal of Machine Learning Research* 3: 993-1022.
- Braun, Virginia and Victoria Clarke. 2021. *Thematic Analysis: A Practical Guide*. California: Sage Publications.
- Zhou, Di. 2022. "The Elements of Cultural Power: Novelty, Emotion, Status, and Cultural Capital." *American Sociological Review*, 87 (5): 750–81.
<https://doi.org/10.1177/00031224221123030>.