Journal of Economic Behavior and Organization 220 (2024) 660–674

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Research Paper

Do political conflicts influence daily consumption choices? Evidence from US-China relations [☆]

Celina Proffen [a,*], Lukas Jürgensmeier [b]

[a] *Goethe University Frankfurt, Theodor-W.-Adorno-Platz 3, 60323 Frankfurt am Main, Germany*
[b] *Goethe University Frankfurt, Theodor-W.-Adorno-Platz 4, 60323 Frankfurt am Main, Germany*

A R T I C L E I N F O

JEL classification:
D12
D83
F51
P0

Keywords:
Political conflict
Consumption
Boycotts
Ethnocentrism

A B S T R A C T

Does political conflict with another country influence domestic consumers' daily consumption choices? We exploit the volatile US-China relations in 2018 and 2019 to analyze whether US consumers reduce their visits to Chinese restaurants when bilateral relations deteriorate. We measure the degree of political conflict through negativity in media reports and rely on smartphone location data to measure daily visits to over 190,000 US restaurants. A deterioration in US-China relations induces a significant decline in visits not only to Chinese but also to other foreign ethnic restaurants, while visits to typical American restaurants increase. We identify consumers' age, race, and cultural openness to moderate the strength of this ethnocentric effect.

1. Introduction

Political conflicts between countries influence public opinion. Previous work shows that changes in public opinion impact consumers' general intentions of buying foreign products (Riefler and Diamantopoulos, 2007), trade flows (Michaels and Zhi, 2010), and car purchasing decisions (Sun et al., 2021; Chen and Zhong, 2024). However, there is only limited evidence on whether individuals react to political conflicts by changing their everyday consumption choices, which could, in the aggregate, have sizable economic consequences.

In this paper, we investigate how the political conflict between the US and China during the Trump administration affected American consumers' restaurant choices. More specifically, our revealed preferences approach measures how the heightened tensions between both countries influence visits to Chinese, American, and other ethnic cuisines, relying on daily smartphone location data from around 7% of US smartphone users.

The daily nature of our data allows us to identify the impact of political conflict on consumption choices absent a direct effect of that conflict on the characteristics of the consumption good. While Chinese restaurants' food and service characteristics are malleable,

☆ Both authors contributed equally. The authors thank the co-editor, Joachim Winter, and two anonymous reviewers for excellent feedback during the revision process. We also thank the participants of the EMAC Doctoral Colloquium 2022, the EMAC Annual Conference 2023, the Queen's Workshop on the Economics of Media 2022, the VfS (German Economic Association) Annual Conference 2022, the RGS Doctoral Conference 2023, Doctorissimes 2023, the efl – the Data Science Institute Jour Fixe, and the FQMG Brown Bag Seminar for valuable feedback on presentations. Further, thanks to Davide Difino, Georg Dürnecker, Nicola Fuchs-Schündeln, Alexander Hillert, Felix Holub, Lennart Kraft, Chiara Lacava, and Bernd Skiera for their helpful comments. Funding: This work was supported by efl – the Data Science Institute.

* Corresponding author.
E-mail addresses: proffen@econ.uni-frankfurt.de (C. Proffen), juergensmeier@wiwi.uni-frankfurt.de (L. Jürgensmeier).

https://doi.org/10.1016/j.jebo.2024.02.031
Received 19 September 2023; Received in revised form 15 February 2024; Accepted 19 February 2024
Available online 7 March 2024

they are unlikely to change immediately within a day due to political conflict or new trade regulations regarding China. Notably, it is not entirely clear to which extent such consumption reactions affect the country of conflict (i.e., China), as the Chinese restaurants in question are primarily tax-paying American businesses.

Besides estimating the direction and magnitude of the consumption effect of increased political conflict with a foreign country, this research also investigates heterogeneous responses of different consumer groups and substitution dynamics across comparable consumption goods. This nuanced analysis allows us to make some inferences about the mechanisms activated by political conflicts. Our findings support the notion that deteriorating political relations with a foreign country increase the salience of consumers' national identity (Horowitz, 1985; Kaufmann, 1996), which can manifest in increased consumption of goods perceived as domestic (Pandya and Venkatesan, 2016; Nardotto and Sequeira, 2021).

To conduct our analysis, we exploit a particularly volatile episode in US-China relations during the Trump administration in 2018 and 2019. Using textual analysis methods, we measure the intensity of the political conflict on a daily level based on the share of negative words in US newspaper reports about China. We then investigate how the intensity of political conflict affects the everyday consumption decisions of US citizens by analyzing patterns of visits to more than 190,000 dine-in restaurants across the US, out of which 8.5% are Chinese restaurants.

To quantify the effect of political conflict, we illustrate how a 0.01 increase in our measure of political conflict affects visits to Chinese restaurants. Such an increase corresponds to 1.6 standard deviations or, more intuitively, to one additional negative word in each 100-word long paragraph about China published in US newspapers during the past week. A simple time series regression suggests that this 0.01 increase in political conflict is associated with decreased visits to Chinese restaurants of up to 4.1%. To fully disentangle politically motivated boycotts from the conflict's effect on overall consumption, we exploit the ethnic affiliations of restaurants (e.g., establishments selling Chinese, American, Italian, or other ethnic food) in our main regression. More precisely, we implement a two-way fixed effects model that quantifies the differential effect of increased political conflict with China on visits to Chinese versus non-Chinese restaurants.

We estimate a causal decline of up to 1.3% in daily dine-in visits to Chinese restaurants relative to non-Chinese restaurants following a 0.01 increase in political conflict. Leveraging the geographic variation in our data, we find that the age, racial composition, and cultural openness of the local customer base are significant moderators of our effect. Areas with older, whiter, and more socially cohesive consumers tend to react significantly more strongly to aggravations in political conflict by decreasing their visits to Chinese restaurants. As expected, areas with above-median shares of Asian Americans exhibit significantly weaker reactions than those with below-median shares. Finally, areas with high shares of Republican voters in the 2016 presidential election also tend to frequent Chinese restaurants less often after political conflict with China than areas with high Democratic voter shares.

We then proceed to investigate to which extent the political conflict with China triggered substitution dynamics across different ethnic restaurant types. We find that the conflict also negatively affects visits to other foreign non-Chinese restaurants, while visits to American restaurants increase significantly by up to 0.8%. We interpret these dynamics as an indication that political conflicts trigger ethnocentric consumer behavior in the spirit of Shimp and Sharma (1987). Although typical American restaurants tend to exhibit higher-than-average numbers of daily visitors, the substitution towards them does not fully compensate for the negative effect of the conflict on overall restaurant visits. Our estimates show that overall restaurant visits in the US decrease by up to 2.6% after a 0.01 increase in our measure of political conflict.

To corroborate the external validity of our findings, we expand our analysis to the political relations between the US and Mexico. Focusing on the same period of the Trump administration, we construct a comparable measure of political conflict based on US newspaper articles mentioning Mexico. We find slightly smaller but consistently negative effects on visits to Mexican restaurants in the US after an increase in political conflict. Furthermore, as in the case of Chinese restaurants, increased political conflict with Mexico leads to fewer visits to most foreign cuisines and more visits to restaurants offering American food.

Our research relates to three broad strands of literature. First, we are thematically tied to literature discussing the effects of aggravated political conflict between the US and China. A vast literature analyzes the conflict from different angles, so here we restrict ourselves to mentioning the contributions that are methodologically most related to ours. Lu et al. (2018) use textual analysis techniques to obtain yearly county-level indices of media slant against China. They show that, in the medium term, counties with greater exposure to Chinese imports report more negatively about China, which leads to changes in voting behavior. For an earlier observation period between 1990 and 2010, Ramirez and Rong (2012) show that "China-bashing," as measured by the number of adverse reports about China in US media, increases 3 – 4 months after an unexpected increase in the US trade deficit with China.

Most recently, Cao et al. (2023) find that, in the context of the COVID-19 pandemic, President Trump's "China Virus" tweets increased the frequency of anti-Asian tweets and violent incidents in the US. They also find that counties supporting Trump in the 2016 presidential election exhibited much stronger reactions to the tweets than those supporting Democrats. Huang et al. (2023) also analyze the implications of worsened US-China relations. Focusing on the early days of the COVID-19 pandemic, which induced racist anti-Chinese sentiment in the US, they estimate an 18% decline in Chinese restaurant visits relative to comparable non-Chinese establishments.

Our study extends Huang et al. (2023) to a more general setting, in which the effect of bilateral political conflict on consumption can be clearly distinguished from pandemic-related health concerns regarding restaurants associated with China, where the virus originated. In our study, we purposefully exclude the period of the COVID-19 pandemic to isolate consumption responses within the more general framework of political conflicts, such that our effects are purely mediated through political sentiments. Furthermore, we provide initial insights into which demographic characteristics are most predictive for reacting to conflicts through boycotts of products associated with the country of conflict.

Our paper conceptually relates to research on the role of social identities in human behavior (Akerlof and Kranton, 2000; Charness and Chen, 2020). Evidence from the Middle East, for example, shows that social identities matter strongly for the behavior of judges, medical patients, and exam graders, who exhibit strong in-group biases favoring individuals that belong to the same ethnicity (Shayo and Zussman, 2011; Zussman, 2023; Lavy et al., 2018). Inter-group conflicts tend to increase the salience of humans' social identities (Horowitz, 1985; Kaufmann, 1996), which may in turn affect their behavior. In the context of allocation decisions, for example, Shayo (2020) finds more pronounced in-group biases after exposure to inter-group conflicts. The political conflict between the US and China that we consider could, thus, be interpreted as a trigger for the salience of US consumers' national or cultural identity.

In this paper, we are particularly interested in examining how consumers' national or cultural identity shapes *consumption choices*. In India, where food choices tightly link to religious or ethnic identities, Atkin et al. (2021) rely on a revealed preferences approach to investigate the determinants of identity-conform consumption decisions. They document that group salience, status, and economic costs affect how strongly people purchase identity-conform consumption bundles. Interestingly, inter-group conflicts, which increase the salience of belonging to a group, lead people to behave more strongly in conformity with their identity.

Focusing on situations of political conflict between countries, Pandya and Venkatesan (2016) and Nardotto and Sequeira (2021) investigate subtle changes in consumption patterns. Pandya and Venkatesan (2016) analyze scanner data in the US shortly after France's veto to use military force against Iraq in the UN Security Council. They document that US consumers significantly reduced their consumption of products that have particularly French-sounding brand names. Using more aggregated data in the same context, Michaels and Zhi (2010) show that this political conflict caused a 9% decrease in bilateral trade between the US and France. Analyzing grocery shopping choices around the Brexit referendum, Nardotto and Sequeira (2021) document how UK consumers strongly shift away from products associated with the European Union and towards those that are "Made in the United Kingdom." They also find that media reporting and social media affect how strongly consumers act in line with their national identity.

Finally, we build on marketing literature on politically motivated consumer behavior that mainly relies on surveys to test and analyze consumers' purchase intentions. A concept that is closely tied to our study is consumer ethnocentrism, which Shimp and Sharma (1987) define as "the beliefs held by consumers about the appropriateness, indeed morality, of purchasing foreign-made products" (p. 280). We refer to Balabanis and Siamagka (2022) for an extensive review of the literature. Consumer animosity, a related concept, describes consumers' reluctance to buy products from a *specific* foreign country or region due to past or ongoing conflicts that can be economic, political, or even military (Klein and Ettensoe, 1999).

The research in this domain pays special attention to which consumer characteristics can predict politically motivated behavior. Balabanis et al. (2001) document that ethnocentric tendencies decrease as consumers' income levels, education, and internationalism rise, while they increase with age, nationalism, and patriotism. Klein and Ettensoe (1999) and Fernández-Ferrín et al. (2015) show that individuals' predictive characteristics for consumer animosity differ from those for ethnocentrism. While they find that patriotism is predictive for both constructs, they show that income and education relate to ethnocentrism but not animosity. In contrast, age, race, and union membership strongly correlate with animosity but not ethnocentrism.[1]

Most of the marketing literature discussed above measures how political conflict matters for consumer behavior by measuring purchase intentions through surveys. Recently, however, two studies have used market-level data to document large decreases in actual monthly sales of Japanese firms in the Chinese automobile market following a dispute amongst the countries over a group of islands in the East China Sea. Sun et al. (2021) show that 75% of this decrease was due to consumers switching car brands, with the majority switching to Chinese brands. Chen and Zhong (2024) provide similar evidence but focus on substitution dynamics between Japanese and Chinese car companies. This finding is consistent with Fouka and Voth (2013), who document a shift away from German cars on the Greek automobile market in the context of the Euro crisis. Instead of purchase decisions for durable goods, we investigate restaurant visits, where the social signaling concern regarding consumption is presumably more limited.

The remainder of this paper is structured as follows. We first introduce our estimation approach to quantify the effect of political conflict on daily consumption. Then, in Section 3, we present the data underlying our analysis with key summary statistics. We also describe the textual analysis methods used to quantify political conflict between two countries. Section 4 presents our main results of how political conflict with China affects visits to Chinese restaurants in the US. After analyzing heterogeneous reactions across consumer demographics, we confirm our findings' robustness and external validity. Then, Section 5 investigates substitution dynamics across various ethnic restaurant types and discusses these findings in light of political economy and marketing theories. The final section concludes.

2. Approach

This paper aims to investigate how political conflicts influence daily consumption choices. To do so, we exploit the context of the US-China trade conflict and measure changes in US-Chinese political relations through variations in how negatively the US media reports about China. To measure consumption choices, we use daily data on dine-in restaurant visits in the US. Finally, we assess how aggravations in the political conflict with China affect visits to Chinese restaurants compared to the average non-Chinese restaurant in the US.

Our baseline analysis relies on estimating the following two-way fixed effects Generalized Linear Model (GLM) with log-link (also known as Poisson regression):

[1] A further related concept, political consumerism, focuses on how consumers reward or punish *individual companies* for their political stances through their buying decisions (Copeland and Boulianne, 2022). In a meta-study, the authors document that higher education levels, the use of media, stronger organizational ties, more political interest, and less political trust are positive predictors for individuals engaging in political consumerism.

$$\mathbb{E}[Visits_{it}|Chinese_i, Negativity_t, \gamma_i, \tau_t] = exp\left[\beta\left(Chinese_i \times Negativity_t\right) + \gamma_i + \tau_t\right] . \quad (1)$$

In this equation, $Visits_{it}$ denotes the number of visits to restaurant i on day t. $Chinese_i$ is a dummy variable that takes the value of one if restaurant i is a Chinese restaurant and zero otherwise. β is the coefficient of interest, and indicates how the political conflict, measured by $Negativity_t$ in media reporting about China affects the number of visits to Chinese restaurants on day t relative to the average non-Chinese restaurant in the US. γ_i is a restaurant-fixed effect that accounts for time-constant differences in the number of visits to individual restaurants due to, for example, varying restaurant size, business model, or location. τ_t are day-fixed effects, which capture day-level variation in visits that are common to all restaurants across the US. The day-fixed effects thus account for seasonal effects or national holidays, which affect the number of visits across all restaurant types. We account for possible serial correlation of the error terms by implementing cluster-robust standard errors at the county level.

We interpret the estimated values of β as the effect of political conflict with China on Chinese restaurant visits compared to the average non-Chinese restaurant in our sample. This differential effect is interesting in its own right but not necessarily equivalent to the full effect of political conflict on visits to dine-in Chinese restaurants. The reason is that the stable unit treatment value assumption (SUTVA) is possibly violated if the political conflict with China causes individuals to change their consumption of non-Chinese food as well. We address this concern in Section 4.1 by implementing an alternative estimation strategy based solely on times series variation.

The identification strategy outlined in this section depends on three assumptions. First, in the short term, our measure of political conflict does not affect visits to Chinese restaurants differently than visits to non-Chinese restaurants through any channel other than consumers' sentiment towards the countries of conflict. The high-frequency nature of our data helps us to fulfill this assumption: even if supply chains across restaurant types are different, the effects of trade sanctions impact the supply of inputs in the medium to long term only.

Second, our measure of political conflict on a given day is exogenous to the number of restaurant visits on the same day and in the future, ruling out the possibility of reverse causality. This assumption is reasonable given that our measure relies on newspaper articles, which are typically published before consumers visit dine-in restaurants for lunch or dinner on a given day. Furthermore, our newspaper articles rarely mention restaurants at all, indicating that they are not informative about future plans or promotions in the restaurant industry.

Finally, our identification strategy is only valid if individuals associate the consumption good of interest with the political conflict. If consumers do not associate a given product with China, they cannot act upon the political conflict by altering their consumption, even if they wish to. For Chinese restaurants, this association is clear because most identify with China through their names and marketing strategies.[2]

One might still argue that Americans *know* that boycotting local restaurants does not effectively "punish" the other country but instead harms tax-paying American businesses and, thus, the American economy. Hence, if the goods were, for example, imported technology, the aggregate consumption effect of the political conflict with China on such consumption would likely be considerably larger. Contrary, if in-person consumption is simply substituted by ordering the same good through a delivery service, we may well be overestimating the overall effect on consumption.

3. Data

In this section, we first explain how we obtain and prepare daily data on the consumption of conflict-associated goods, measured by visits to dine-in restaurants in the US. We then describe how we measure political conflict using textual analysis methods. Throughout, we provide key summary statistics for both metrics.

3.1. Daily measure of consumption: restaurant visits

We proxy consumption of conflict-associated goods using daily visits to dine-in restaurants in the US during 2018 and 2019, distinguishing among restaurants serving different types of ethnic food. The underlying data come from SafeGraph, a company that gathers smartphones' GPS location and attributes visits of these devices to specific location polygons corresponding to individual buildings. SafeGraph then attributes a visit to a restaurant once a device resides at the dine-in restaurant for at least four minutes. Ideally, we would also consider data on restaurant visits from earlier years to analyze the effects of negative reporting on China in various political contexts, but the data is only available since 2018. For years after 2020, the COVID-19 pandemic introduced health concerns for dining out and a bias against food produced in establishments associated with China since the virus originated in China (Huang et al., 2023). Hence, we restrict ourselves to the full years 2018 and 2019.

After obtaining metadata on US dine-in restaurants and information on their daily visits registered through tracked smartphones, we use SafeGraph's sampling rate by small geographic areas (census block groups) to remove sampling bias and estimate the actual number of visits. We scale the raw data by monthly sampling rates for each census block group, which we calculate by dividing the number of smartphones tracked in the group by the number of its residents according to the census. We exclude all restaurants with an average daily visitor number smaller than 10 since SafeGraph intentionally randomizes low numbers of visits to preserve visitors' privacy.

[2] The Chinese restaurants in our sample typically have names including the words "China" or "Chinese," or the name of a Chinese city or region.

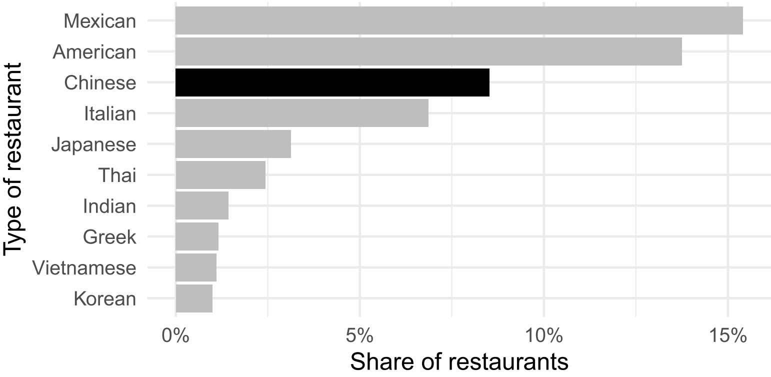

Fig. 1. Share of restaurant types in our sample.

Table 1
Dependent variable—restaurant visits by restaurant type.

Restaurant type	Mean visits	SD visits	Observations	Number of restaurants
All	74.28	85.77	140,978,726	193,926
Non-Chinese	75.86	87.39	128,944,278	177,410
Chinese	57.45	63.61	12,034,448	16,516

Fig. 1 shows the share of restaurants in our sample that fall into each of the 10 largest ethnic restaurant categories as tagged by SafeGraph. The data provider determines category tags by combining proprietary algorithms and fixed tags (e.g., for well-known brands).[3] Restaurants can have multiple tags that do not necessarily correspond to ethnic cuisines. For example, a restaurant can be classified as “Buffet” and “American.” We consider a restaurant affiliated with an ethnic cuisine if one of its tags clearly refers to an ethnic cuisine.[4] Following this rule, we refer to restaurants as Chinese if they are tagged as “Chinese,” which is the case for 8.5% of establishments in our sample.

Table 1 displays our sample’s summary statistics overall and separately for Chinese and non-Chinese restaurants. The mean daily visits are considerably smaller for Chinese restaurants than for non-Chinese ones, although both display a similar standard deviation relative to their mean. The high standard deviation is likely the result of restaurants being closed on some weekdays. For the vast majority of the 193,926 restaurants in our sample, we record the number of visits for each day between January 1, 2018 and December 30, 2019.[5]

3.2. *Measure of political conflict: newspaper negativity*

Our measure of political conflict is based on how negatively the US media reports about China over time. To capture the *affectual* or *subjective* content of communication, we rely on a domain of textual analysis methods often referred to as *sentiment analysis* (for an abundant literature review, see Mohammad, 2021). These methods can be applied to various types of texts, such as newspaper articles, social media posts, financial statements, and political speeches (Loughran and McDonald, 2016). The text type is, however, crucial for choosing the appropriate sentiment analysis method. Shapiro et al. (2022) evaluate the performance of multiple textual analysis methods applied to newspaper articles by benchmarking automated with human-rated sentiment scores and conclude that a customized dictionary[6] approach is most appropriate for the context of newspaper-based sentiment analysis. We build on this finding and implement a similar approach when measuring the degree of political conflict with China from articles published in US newspapers. While transformer-based machine learning models might perform better under certain circumstances, dictionary-based methods appeal through their transparency and interpretability in our context.

To quantify the negativity towards China conveyed in newspaper articles, we construct a negativity dictionary, i.e., a list of negative words. In line with Shapiro et al. (2022), we acknowledge that negativity is highly context-specific and customize our dictionary to fit our context. We start by evaluating the contextual fit of the negative words extracted from five established dictionaries.[7] For each dictionary, we obtain the 20 most frequent negative words appearing in paragraphs with the words “China” or “Chinese.” We

[3] Source: https://community.deweydata.io/t/a-few-general-questions-about-core-and-pattern-data/25685, last accessed on February 1, 2024.

[4] We consider restaurants’ full tags to avoid misattributing a restaurant due to individual words or characters, which ensures we avoid incorrectly attributing “Latin American” restaurants to “American” ones.

[5] 1,514 non-Chinese and 75 Chinese restaurants feature at least one day with missing observations. We decided to keep these restaurants in the sample as there is no clear pattern or documented reason behind these missing values, and the share of affected restaurants is small.

[6] In textual analysis, a dictionary represents a list of words that reflect a specific topic or sentiment.

[7] See Web Appendix 7.1.1 for a detailed description of the baseline dictionaries.

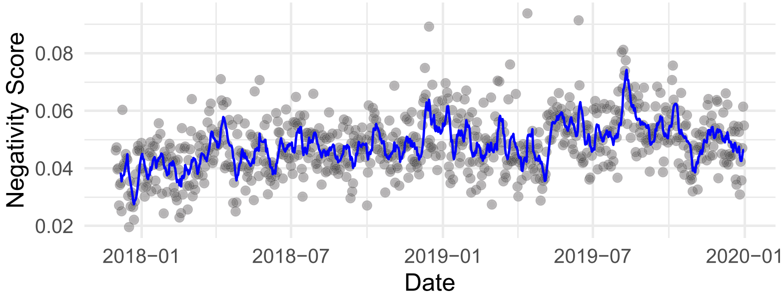

Fig. 2. Text-based negativity measure for US-China relations. **Note:** The dots indicate the daily negativity measure computed by averaging across newspaper-specific negativity. The line indicates its seven-day moving average.

randomly select 20 paragraphs mentioning each of these words and manually verify that the words are indeed negative in the context of US-Chinese political relations.

We then use the union of all five initial dictionaries as the baseline for our negative word list. Subsequently, we exclude words that indicate a negation but are not necessarily negative themselves, such as "not," "none," or "never." Finally, we meticulously go through the list and exclude words that are not negative but rather descriptive in the context of the US-China trade conflict. For example, in our setting, "vice" mostly refers to vice presidents instead of bad habits, and "cloud" refers to cloud technologies and not bad weather. In the end, we obtain a list of 9,464 unique negative words. To give the interested reader a sense of the negative words in our dictionary, we list the most frequent (negative) words in our corpus of articles in Table A1 in Web Appendix 7.1.2. We do not rely on positive word lists to create positivity or net-negativity scores, as these positive words often introduce ambiguity in the interpretation of negative statements and add little incremental information (Loughran and McDonald, 2016).

We then use our customized dictionary to quantify the negativity in China-related reporting of US daily newspapers. We obtain newspaper articles from the global news database Factiva and initially consider all available daily newspapers, excluding only highly industry-specific ones (e.g., "Rubber & Plastic News" or "The Connecticut Law Tribune"). Our final sample consists of 109 newspapers listed in Web Appendix 7.1.3. While we rely on both the digital and print editions of each newspaper, we remove duplicate articles after identifying them through their high levels of Jaccard similarity (Jaccard, 1901). This procedure results in a corpus of 34,022 China-related newspaper articles in 2018 and 2019. Finally, we remove stop words, links, HTML code, and common abbreviations and decapitalize all words in the articles.

Using our customized negative word list, we compute a negativity score for each article that mentions "China" or "Chinese" at least once. The procedure occurs in two steps. First, we compute the negativity score for each paragraph as the fraction of negative words over total words. We do not account for negations, as these tend to make negative words less negative but not positive (Zhu et al., 2014). Second, we compute articles' negativity scores by averaging the paragraph-level scores from all paragraphs containing the words "China" or "Chinese." Basing our measure on the paragraph level is crucial for correctly identifying China-related negativity, given that many articles only partially relate to China. For example, an article analyzing general stock market performance may only refer to the impact of current developments in China in a single paragraph. Since reporting about unrelated issues should arguably not influence readers' views on the country, we do not take such unrelated paragraphs into account.

We then obtain a day-specific negativity score for each newspaper by averaging across all articles' scores published in the relevant newspaper on a given day. The rationale here is that the reader will take away an overall impression from her daily read: the higher the score, the more negative a newspaper reports about China. Finally, we take the average of all newspapers' daily scores to arrive at a measure of negativity in reporting about China on a given day. This choice implies that each newspaper's negativity contributes equally to the national negativity score. In other words, we treat national newspapers with a large readership equal to smaller regional ones whenever the newspaper features any reporting on China on a given day. Since small newspapers typically report less frequently on China than large ones, they still contribute less to our national negativity score.

There are valid alternatives to this approach, with their own advantages and disadvantages. Most intuitively, we also construct an alternative measure that weighs newspapers' negativity by the size of their readership. While this would, in theory, be our preferred approach, the necessary circulation data is only available for approximately one-fourth of our sample's newspapers. We discuss the alternative measures of negativity in Web Appendix 7.1.4 and assess our results' robustness to these choices in Section 4.3.

Fig. 2 displays the equally-weighted daily text-based reporting negativity series, together with its seven-day moving average. Both displayed series share their mean negativity level of 0.049, meaning that, on average, paragraphs mentioning the words "China" or "Chinese" contain 4.9 negative words per 100 words. Both series are subject to significant variation over the two-year observation period. The daily negativity series exhibits a standard deviation of 0.010, which amounts to one additional negative word in each 100-word-long paragraph about China. The standard deviation of the seven-day moving average is lower at 0.0062. During our study's observation period, the measure increases (i.e., reporting becomes more negative), which is in line with the notion of deteriorating political relations between the US and China.

We use the seven-day moving average as our preferred measure of political conflict with China. We intentionally capture the whole past week of reporting to make our measure comparable across time, given that not all newspapers in our sample publish daily and that reporting may vary systematically across days of the week.

Table 2
Main regression estimates (Poisson regression).

Dependent Variable:	Restaurant visits			
Model:	(1)	(2)	(3)	(4)
Variables				
Chinese × Negativity	-0.6439***	-0.6113***	-1.146***	-1.282***
	(0.1863)	(0.1860)	(0.1302)	(0.1628)
Chinese × Time			3.38×10^{-5}***	3.41×10^{-5}***
			(7.98×10^{-6})	(7.94×10^{-6})
Chinese × Reporting quantity				1.34×10^{-5}*
				(7.82×10^{-6})
Fixed-effects				
Date	✓	✓	✓	✓
Restaurant	✓			
Restaurant × Day of week		✓	✓	✓
Fit statistics				
Observations	140,014,391	140,013,653	140,013,653	140,013,653
Squared Correlation	0.64080	0.70404	0.70404	0.70404
Pseudo R^2	0.62775	0.67846	0.67846	0.67846

Notes: Clustered standard errors at the county level displayed in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Estimating regressions with this measure as the independent variable helps us to account for the fact that consumers often plan visits to dine-in restaurants several days in advance. This empirical strategy allows us to identify effects even if the reactions to conflicts measured in restaurant visits do not materialize fully on the same day. At the same time, we implicitly assume that readers' perception of the US-China trade conflict is affected in the same way by newspaper articles today as by those in the past six days, irrespective of the weekday on which they were released.

Our measure of political conflict is subject to two limitations. First, the newspaper articles underlying our measure may not accurately reflect media reporting on the conflict more generally. While the presence of prominent news agencies such as the Associated Press or Reuters increases the similarity of reporting across various channels, there are still differences between the reporting formats of, for example, television, radio, and print media. Suppose the way consumers perceive the conflict is susceptible to the reporting format. Then, the negativity metric may suffer from measurement error whenever there are stark differences in print versus other media content. A second concern relates to how we quantify the negativity of bilateral political relations, i.e., by computing the share of negative words used in related articles. By design, our analysis is thus always concerned with how negative developments in a conflict influence consumption and does not consider positive developments. While we believe the downside is more relevant than the upside, this study provides no evidence regarding the impact of positive news on consumption.

In addition to measuring how negatively the conflict with China evolves, we are also interested in its salience in the media. For this purpose, we compute the average daily number of paragraphs newspapers published about China over the past week. This measure of reporting quantity serves as a first indication of how prominent the conflict is at the moment without directly speaking to the degree of negativity in US-Chinese relations. In some of our regressions, we use this variable to control for the amount of information published about China, irrespective of its content or tone.

4. Effects on visits to Chinese restaurants

4.1. Main results

This section presents the results of estimating Equation (1) using our preferred proxy for political conflict with China, i.e., the average of text-based negativity in China-related reporting over the present day and the past six days. Table 2 displays the main results. The first column presents our estimation of Equation (1), a GLM with log-link, while the three additional columns present more flexible model variations. The estimated coefficient $\hat{\beta}$ on the interaction between the two variables $Chinese \times Negativity$ quantifies the impact of increased political conflict on consumption and is our main coefficient of interest. Across all four specifications, $\hat{\beta}$ is negative and statistically significant, indicating that more political conflict with China leads to a decline in visits to Chinese restaurants relative to the average non-Chinese restaurant in the US. To interpret the magnitude of these results, assume an increase in the week's average reporting negativity of 0.01. This change corresponds to one additional negative word in each 100-word-long paragraph about China during the past week. Absolute changes in our measure of conflict with China of at least this magnitude occur on 4.55% of the days covered in our sample. After such an increase, our estimated coefficient in column (1) indicates that the average Chinese restaurant's daily visits decline by $(\exp(0.01 \times \hat{\beta}) - 1) \times 100 = -0.64\%$.

In columns (2) – (4), we replace the restaurant-fixed effects with "Restaurant × Day of week"-fixed effects to account for restaurant-specific visitor patterns over the course of a week. In this way, the dependent variable is free of any structural patterns due to, for example, a restaurant being closed on Mondays or having special offers and thus systematically more visitors on Thursdays. Column (2) presents the results from estimating Equation (1) with this alternative fixed effect. The coefficient remains similar, estimating 0.61% fewer visits to the average Chinese restaurant after a 0.01 increase in our political conflict measure.

We continue to increase the flexibility of our model and additionally control for differential linear time trends in visits to Chinese and non-Chinese restaurants. The rationale here is to account for potential long-term changes in consumer preferences regarding ethnic cuisines. We present the results in column (3). Including this control increases the magnitude of the estimated effect to 1.14% fewer visits to the average Chinese restaurant following a 0.01 increase in negativity. The coefficient of interest remains significant at the 1% level, which reassures us that our results are not merely driven by long-term trends underlying restaurant preferences and that variations in political relations indeed play an essential role in explaining the observed consumption behavior.

Finally, in column (4), we introduce a measure of reporting quantity about China as an additional control variable. This measure captures the average daily number of paragraphs about China published in our newspaper sample over the past week and is thus related to our measure of political conflict. However, instead of informing about the tone of reporting, it controls for the quantity of new information released each week. After introducing this control, we find that a 0.01 increase in our measure of political conflict leads to a 1.27% reduction in visits. The coefficient remains highly significant at the 1% level. The estimated coefficient on the quantity of China-related reporting is significant at the 10% level, suggesting some relevance for consumers' restaurant consumption choices.[8] We also find some evidence for longer-lasting effects on consumption choices by including lagged values of our independent variable into our regression framework (see Web Appendix 7.2.2). The coefficient on our measure of political conflict is consistently and statistically significantly negative for up to two weeks into the past.

The above effects need to be interpreted as the differential effect on Chinese restaurant visits relative to non-Chinese restaurants. Thus, if the average non-Chinese restaurant also suffers declining visits due to political conflict, the previous approach underestimates the full negative effect on Chinese restaurant visits. To obtain estimates for the full effect on Chinese restaurants, we also estimate a time series regression in which we explain the variation in visits to Chinese restaurants with our measure of political conflict and parametric time trends. Depending on the specification, a 0.01 increase in political conflict leads to a 3.32 – 4.09% reduction in visits to Chinese restaurants. While this approach can estimate the full effect on Chinese restaurant visits, the absence of a control group in the time series regression prevents us from controlling for day-specific events that affect visits to all types of restaurants. We present a complete description of the time series approach and the corresponding results in Web Appendix 7.2.1.

4.2. Heterogeneous effects by consumer characteristics

In this section, we investigate whether and how consumer characteristics moderate the effect of the US-China political conflict on visits to Chinese restaurants in the US. Survey-based literature has conducted extensive analyses on the demographic characteristics that are most important in mediating consumer behavior *vis à vis* products associated with a foreign country. Research on consumer ethnocentrism finds that individuals' tendencies to buy domestic versus foreign products increase with higher levels of consumers' patriotism, conservatism, and collectivism but decreases with their income, education, and cultural openness (Sharma et al., 1995). Klein and Ettensoe (1999) show that age, race, and union membership are highly predictive of consumer animosity (i.e., the likelihood of discriminating against products associated with a *specific* foreign country), albeit they are not significant antecedents for ethnocentric behavior. Motivated by these findings, we investigate how the consumption responses vary by customers' demographic characteristics in our setting of deteriorated political relations with China.

To obtain information about restaurants' customer bases, we rely on the demographic characteristics of the population living in proximity to the restaurants. Using restaurants' geographic coordinates, we match restaurants to information on their surrounding areas obtained from other geo-coded data sets. From the 2019 American Community Survey (ACS), we obtain data on the median income, median age, the population share of various races, and the population share of college graduates at the census block group level. These characteristics are self-explanatory measures for each restaurant's local customer base's income, age, race, and education. Furthermore, we proxy the partisanship of restaurants' clients with the local Republican voter shares in the 2016 presidential elections. We obtain this information from the MIT Election Data and Science Lab (2018), and match it to restaurants at the county level.

To obtain an area-specific measure of cultural openness, we take advantage of the county-level "social capital" data set of Chetty et al. (2022). Specifically, we rely on a measure of social cohesion ("clustering") that captures the degree to which social networks are fragmented into cliques. A higher level of clustering indicates that a higher percentage of one's friends are also friends with each other. Low clustering levels arise in areas where people move or travel more frequently and make friends outside of the usual shared networks. Accordingly, this measure is typically low in internationally oriented cities such as San Francisco or New York City.

To see how these characteristics moderate the effect of political conflict on consumption, we operationalize each demographic variable as a dummy indicating whether an area lies above or below the median value of that variable. We then estimate our preferred model specification, corresponding to column (4) of Table 2, and include additional interaction terms to investigate whether restaurants' local customer base above a variable-specific median exhibits significantly different effects. Table 3 presents the key coefficient estimates of this exercise. In the following, we discuss the estimates of the three-way interaction between the variables $China \times Negativity$ and the respective demographic dummy variable.

Among the standard socioeconomic indicators, neither income nor education seem to matter for the effect of political conflict on consumption. This result contradicts previous findings from the ethnocentrism literature, which would have predicted less ethnocentric behavior for high-education and high-income groups (Sharma et al., 1995). However, our findings are in line with Nardotto and Sequeira (2021), who show that the consumers' socioeconomic characteristics play a limited role in explaining ethnocentric

[8] The average daily reporting quantity on China in our sample period is 317 paragraphs, with a standard deviation of 105 paragraphs.

Table 3
Heterogeneity analysis (Poisson regression).

Heterogeneity dimension	Chinese × Negativity × Above median	Chinese × Negativity	Observations	Pseudo R^2
Share college educated	0.185	-1.395***	136,319,411	0.67552
	(0.2568)	(0.2263)		
Median household income	0.1622	-1.349***	132,044,920	0.67611
	(0.2561)	(0.1900)		
Age	-0.8416***	-0.8300***	136,313,989	0.67550
	(0.2764)	(0.1786)		
Share white	-1.383***	-0.5333***	136,335,995	0.67551
	(0.3106)	(0.2014)		
Share Asian	1.221***	-1.856***	136,335,995	0.67554
	(0.3347)	(0.2794)		
Republican voter share (2016 presidential election)	-0.6527*	-0.6295***	133,621,198	0.67504
	(0.3374)	(0.2403)		
Social cohesion: clustering	-1.433***	-1.126***	139,618,448	0.67850
	(0.5309)	(0.1732)		

Notes: This table shows heterogeneities in effects by whether an area's value of a given heterogeneity dimension lies above or below the median of the country-wide value. All models control for the simple interaction of Chinese × Above median, restaurant-specific day of week and date fixed effects, long-term time trends, and reporting quantity. Clustered standard errors at the county level are displayed in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

changes in UK consumer behavior following the Brexit referendum. Nevertheless, we interpret these findings with caution because the treatment intensity may systematically correlate with some demographic characteristics. For example, suppose consumers with high educational levels are better informed about developments in the US-China conflict because they follow the news media more closely. In that case, they might be able to react to that news, whereas less educated consumers may have perceived the political conflict less strongly.

An above-median share of old or white individuals in the local population significantly exacerbates the overall negative effect of political conflict on Chinese restaurant visits. In contrast, higher shares of Asians are associated with less negative effects. This finding is intuitive, given the hypothesis that political tensions trigger consumers' reactions by strengthening their adherence to their national identity. Areas with higher shares of Asians may exhibit lower identification with the US or even identify with Asian countries. This demographic difference could rationalize why, in these areas, consumption responses are so much less pronounced.

Areas with above-median Republican voter shares in the 2016 presidential election exhibit significantly more negative effects than our baseline estimation, though only at $p < 0.1$. This result aligns with previous findings from the ethnocentrism literature that associate more patriotic and conservative attitudes with stronger tendencies to discriminate against foreign products in general (Sharma et al., 1995). It is also consistent with Huang et al. (2023), who find that during the COVID-19 pandemic, visits to Chinese restaurants decreased more sharply for areas with higher Republican voting shares. Our study extends that insight to pre-pandemic times and shows that it holds even in a political conflict between two countries, which is more common than a once-in-a-century pandemic. In interpreting this finding, we note that differences in reporting styles among newspapers read by Republicans versus Democrats may well be an alternative explanation of our finding: if the reporting style in more Republican areas portrays news on China in a more inflammatory way, the stronger consumption reactions may be due to the stronger signal instead of consumers' partisanship.

Finally, we find that areas with above-median levels of social cohesion ("clustering") exhibit more negative effects than areas that lie below the median. This result confirms the intuition that political conflict leads to a weaker ethnocentric reaction in areas with high levels of cultural openness.[9]

Overall, we find that demographic characteristics moderate the strength of political conflict-induced consumption changes. In particular, age, racial composition, and cultural openness seem to affect how strongly consumers respond to political conflicts. The evidence on partisanship is less clear, with coefficients estimated differently from zero only at the 10% significance level.

4.3. Robustness checks of the main result

The findings presented above suggest that increases in political conflict with China have both a statistically and economically meaningful impact on visits to Chinese restaurants in the US. We test the robustness of this result in two ways: first, we replicate our main analysis using alternative measures of the political conflict with China. In this way, we address concerns about how to best capture the negativity in US media reports about China. Second, we verify that our results are not driven by other economic or sociopolitical dynamics correlated with our measure of political conflict. We do so by extending our main regression with additional controls for potential omitted variables that could be plausible alternative explanations of our main effect.

[9] To account for potentially varying treatment intensities for individuals of different education and income levels, i.e., for people with presumably different exposure to news media, we additionally control for these variables in non-reported regressions. The triple-interaction terms for age, race, voter shares, and social cohesion all keep their sign and statistical significance after adding the controls for areas' education and income levels.

Table 4
Main regression estimates using alternative negativity metrics.

Dependent variable:	Restaurant visits				
Negativity metric:	Heading	Four most-read newspapers	Circulation-weighted	State-level	
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
Chinese × Negativity	-1.005***	-1.132***	-0.6555***	-0.0070***	-0.0048**
	(0.0774)	(0.1915)	(0.1353)	(0.0019)	(0.0019)
Fixed-effects					
Date	✓	✓	✓	✓	
Restaurant × Day of week	✓	✓	✓	✓	✓
Chinese × Date					✓
Fit statistics					
Observations	140,013,653	140,013,653	140,013,653	101,391,140	101,391,140
Squared Correlation	0.70405	0.70404	0.70404	0.70904	0.70921
Pseudo R^2	0.67847	0.67846	0.67846	0.68026	0.68041

Notes: The alternative negativity measures are described in detail in Web Appendix 7.1.4. All models include a control for how intensively newspapers covered the political conflict (*Chinese* × *Reporting quantity*) as in the main analysis. Columns (1) – (4) also continue controlling for a Chinese-restaurant-specific time trend (*Chinese* × *Time*), which is not possible in column (5) due to the *Date* × *Chinese*-fixed effects. Clustered standard errors at the county level are displayed in parentheses. *** p < 0.01, ** p < 0.05, * p < 0.1.

We begin by replicating our analysis using four alternative negativity measures that capture the intensity of the political conflict between the US and China. On a national level, we construct three alternative negativity series that address the some concerns regarding our negativity equally weighted text-based negativity metric. One could argue that (1) newspaper headlines may be more decisive to readers' reactions than text bodies, so we compute a headline-based negativity metric. (2), larger and more international newspapers may trigger different reactions than local ones, so we use only the four largest national newspapers. (3), a concern could be that different newspapers have varying audience sizes and should not be weighted equally. Hence, we use a circulation-weighted negativity metric, which we construct based on the subset of newspapers for which circulation data is available. Lastly, (4) we also construct a state-level negativity series based exclusively on local (as opposed to national or international) newspapers from states with at least two such newspapers. In Web Appendix 7.1.4, we describe the construction of these four alternative negativity measures in detail and present the corresponding summary statistics.

Table 4 shows the results of estimating the preferred model specification from column (4) of Table 2 using all alternative measures for political conflict with China. It is important to note that their means and standard deviations differ by construction, impeding a direct comparison between the coefficient estimates in our main analysis and the ones below. Columns (1) – (3) rely on the above alternative national negativity series as independent variables. As with our baseline negativity series, all three coefficient estimates are statistically significantly negative.

We replicate our analysis using state-specific negativity series in column (4). We are interested in whether these state-specific measures can explain the variation in visits to Chinese restaurants located in this state over time. Again, the coefficients of interest are statistically significantly negative, supporting the evidence from our main analysis. Finally, in column (5), we replace the date-fixed effects with "Chinese × Date"-fixed effects. This change allows us to verify whether our state-specific negativity measures retain their predictive power at the state level once we control for day-specific variations in visits common to all Chinese restaurants in the US. Because the coefficient retains its negative direction and a considerable magnitude, these results reassure us that our measure appropriately captures the degree of the political conflict transmitted by the media.

In our second robustness check, we assess whether our analysis is prone to omitted variable bias. Specifically, one may be concerned that our measure of political conflict with China simply reflects broader economic or sociopolitical dynamics. For example, if our negativity series correlated strongly with the general economic sentiment, the latter could be the underlying reason for the reactions we observe instead of a China-specific sentiment. Shifting away from Chinese restaurants during an economic downturn could be explained by a desire for typical American comfort food, such as burgers, mac and cheese, or fried chicken. To address this concern, we measure economic sentiment and sociopolitical dynamics through three distinct indices described in detail in Web Appendix 7.2.3 and re-estimate our main regression model from column (4) of Table 2 with those additional controls. We present the results in Table 5.

The first two columns control for economic sentiment in the US. The economic sentiment index underlying the regression in column (1) is based on a daily time series published by the *Federal Reserve (Fed)* and relies on a similar sample of US newspapers, which we use to construct our measure of political conflict with China. Similar to our main negativity measure's construction, these sentiment scores are computed using a customized dictionary method. Higher sentiment scores indicate a more positive economic sentiment. The economic sentiment index underlying the regression in column (2) is based on a sample of international newspaper articles included in the full edition of *RavenPack News Analytics (RPNA)* and which are associated with the US through a US place tag. The event sentiment score, which *RPNA* attributes to each article, is higher when experts associate it with increased stock market returns.

Table 5
Regression estimates of alternative independent variables characterizing the US-China conflict.

Dependent Variable:	Restaurant visits			
Model:	(1)	(2)	(3)	(4)
Variables				
Chinese × Negativity	-0.5989***	-1.269***	-1.258***	-0.6552***
	(0.1407)	(0.1592)	(0.1628)	(0.1399)
Chinese × Fed economic sentiment	0.0961***			0.0896***
	(0.0130)			(0.0132)
Chinese × RPNA economic sentiment		8.54×10^{-5}		-0.0002**
		(9.03×10^{-5})		(8.51×10^{-5})
Chinese × Number of Trump tweets			-0.0017***	-0.0014***
			(0.0001)	(0.0001)
Fixed-effects				
Date	✓	✓	✓	✓
Restaurant-Day of week	✓	✓	✓	✓
Fit statistics				
Observations	140,013,653	140,013,653	140,013,653	140,013,653
Squared Correlation	0.70404	0.70404	0.70404	0.70404
Pseudo R^2	0.67847	0.67846	0.67847	0.67847

Notes: All models include controls for the Chinese-restaurant-specific time trend (*Chinese* × *Time*) and how intensively newspapers covered the political conflict (*Chinese* × *Reporting quantity*) from the main analysis. For both economic sentiment measures (Fed and Ravenpack News Analytics, RPNA), more positive values indicate a more positive sentiment. The Fed economic sentiment index ranges from -1 to 1 and has a mean (standard deviation) of -0.0086 (0.1244). The RPNA economic sentiment index ranges from -50 to 50 and has a mean (standard deviation) of -4.5040 (4.8483). Trump tweets indicate the number of daily Trump tweets. In our sample period, the series has a mean (standard deviation) of 7.277 (4.4942). Clustered standard errors at the county level are displayed in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Given the economic nature of the US-Chinese political conflict and the relevance of China for the American economy, we expect our series of political conflicts to be rather strongly correlated with both of these indices. In fact, the correlations of our measure of political conflict on the one hand, and the *Fed*- and the *RPNA*-based measures on the other, are -0.581 and -0.276, respectively. In the case of the *Fed*-based measure, we find that more negative economic sentiment is associated with a significant shift away from Chinese restaurants relative to other restaurant visits. For the *RPNA*-based measure, we do not confirm this finding. Importantly, controlling for both economic sentiment measures does not affect the sign and significance of our main coefficient of interest.

Besides economic sentiment, a concern could be that our measure of political conflict reflects broader sociopolitical dynamics, such as increasing protectionist attitudes or patriotism. While it is not obvious what good measures of such dynamics could be, we attempt to capture them relying on variations in President Trump's Twitter activity, a prominent impulse for sociopolitical dynamics. As expected, given President Trump's critical stance towards China, more tweets by him are associated with a shift away from Chinese restaurants. Reassuringly, however, our main coefficient of interest is robust to controlling for this measure of sociopolitical dynamics and preserves its size and statistical significance in column (3).

Finally, in column (4), we control for all the discussed potential omitted variables. We still find a sizable and highly statistically significant negative effect of our measure of political conflict on visits to Chinese relative to non-Chinese restaurants in the US. Thus, all four models retain effect sizes similar to our main analysis in Table 2.

4.4. External validity of the main result: effects of the US-Mexican conflict

To further test how political conflicts affect the consumption of goods associated with the focal country of conflict, we replicate our primary analysis in a different setting: the political relations between the US and Mexico in 2018 and 2019. Similar to US-China relations, US-Mexico relations were volatile during the Trump administration. Hence, we construct a comparable proxy of political conflict between both countries. Our measure relies on the same set of newspapers, but we only use articles that mention the words "Mexico," "Mexican," or "Mexicans." The negativity dictionary deployed for the main analysis does not reveal any obvious misclassifications applied to Mexico-related articles. Thus, we keep the same dictionary, which ensures comparability between our two measures of conflict.

In Web Appendix 7.1.4, we show summary statistics for our measure of political conflict with Mexico. The negativity measure exhibits a lower mean (0.034), i.e., fewer negative words on average in paragraphs mentioning the country, compared to the China-related metric (0.049). The conflict is also slightly less volatile, exhibiting a standard deviation of 0.0038 (see Table A2 in Web Appendix 7.1 for more details). Analyzing the relative frequency with which the words "trade" and "immigration" occur reveals that the nature of the political conflict between each of these two countries is inherently different. "Trade" accounts for 0.52% of all words in the China-related but only for 0.15% in the Mexico-related articles. In contrast, "immigration" accounts for less than 0.01% of words in articles mentioning China but 0.08% of those mentioning Mexico.

Table 6 replicates our main results (see Table 2), now estimating the effect that the US-Mexico political conflict has on visits to Mexican restaurants in the US. To interpret the size of the effect, we again assume a 0.01 increase in our measure of political conflict,

Table 6
Regression estimates on the effects of the political conflict with Mexico.

Dependent Variable:	Restaurant visits			
Model:	(1)	(2)	(3)	(4)
Variables				
Mexican × Negativity	-0.6392***	-0.5673***	-0.9829***	-1.187***
	(0.1849)	(0.1859)	(0.1834)	(0.2057)
Mexican × Time			$6.87 \times 10^{-5***}$	$6.89 \times 10^{-5***}$
			(9.34×10^{-6})	(9.34×10^{-6})
Mexican × Reporting quantity				$1.94 \times 10^{-5***}$
				(6.18×10^{-6})
Fixed-effects				
Date	✓	✓	✓	✓
Restaurant	✓			
Restaurant × Day of week		✓	✓	✓
Fit statistics				
Observations	140,978,726	140,977,986	140,977,986	140,977,986
Squared Correlation	0.64058	0.70385	0.70388	0.70388
Pseudo R^2	0.62755	0.67836	0.67839	0.67839

Notes: Clustered standard errors at the county level displayed in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

i.e., one additional negative word in each 100-word-long paragraph about Mexico throughout the past week. Depending on the specification, this increase in negativity leads to 0.57 to 1.18% fewer visits to Mexican restaurants relative to the average restaurant in the US. The results of Table 6 align with our main results that political conflicts do indeed influence everyday consumption behavior by decreasing visits to restaurants affiliated with the foreign country of conflict.

5. Effects on visits to other ethnic restaurants

So far, the analysis has investigated how the political conflict with China affects visits to Chinese compared to non-Chinese restaurants in the US. However, our main analysis hides potential effects on other ethnic restaurant types. It seems plausible that changes in the political relations with China also trigger non-trivial substitution dynamics across different ethnic restaurant types, such as Japanese, Italian, or American restaurants.

If a substitution away from Chinese restaurants occurred solely due to the desire to avoid consuming goods associated with the country of conflict, we would expect consumers to replace their visits to Chinese restaurants one-to-one with visits to restaurants that they consider close substitutes.[10] To test this possibility, we check whether the overall number of visits to all restaurants in the US remains the same by running a time series regressions of visits to all types of restaurants in the US on our measure of political conflict with China. If consumers would simply substitute Chinese restaurants one-to-one with visits to non-Chinese restaurants, we should not observe any change in the overall number of restaurant visits in the US following negative reporting about China. We do, however, find that following a 0.01 increase in our measure of political conflict, visits to the average restaurant (including both Chinese and non-Chinese) in the US decrease by 0.77 to 1.52% (see Web Appendix 7.2.1, Table A6). This result suggests that consumers do not fully substitute their Chinese restaurant visits with visits to other restaurant types. Instead, they may replace it by dining at home or ordering food via delivery services.

We proceed by assessing whether and how the political conflict with China triggered shifts in consumption patterns across various ethnic restaurant types. Therefore, we estimate our preferred model specification, corresponding to column (4) of Table 2, separately for the ten most prominent country-related cuisines in our sample. In other words, we replace the *Chinese* indicator with indicators for restaurants labeled as *American*, *Mexican*, etc., and run a regression where we replace the *Chinese* indicator of column (4) in Table 2 with the indicator for each of the ten cuisines. We illustrate the estimated effects of a 0.01 increase in our measure of political conflict with China in Fig. 3.

As before, the estimated effects need to be interpreted relative to the effect that reporting negativity about China has on the average restaurant in our sample, excluding the focal cuisine type. Restaurants selling American food receive significantly more visits when the political conflict with China deteriorates. Following a 0.01 increase in our measure of political conflict, American restaurants experience a 0.75% increase in visits compared to the average non-American restaurant in the US. All other restaurant types exhibit negative effects, albeit these are not statistically different from zero for Japanese, Italian, and Korean establishments.[11] Overall, Fig. 3 illustrates that visits to most foreign restaurants decrease following an aggravation in conflict with China, while visits to American establishments increase.

[10] We cannot distinguish whether the change in visits is due to consumers' active choice against Chinese *food* or because of racist motives, according to which consumers try to avoid interactions with ethnic Chinese staff, who often work for Chinese restaurants (Lee, 2019).

[11] We find similar patterns when we conduct this analysis for the political conflict with Mexico: following a 0.01 increase in the measure of political conflict, Mexican restaurants receive 1.18% fewer daily visits than the average restaurant in the US. Most other foreign ethnic restaurants also experience a decline (on average 0.43%), while American restaurants experience an increase of 1.21%.

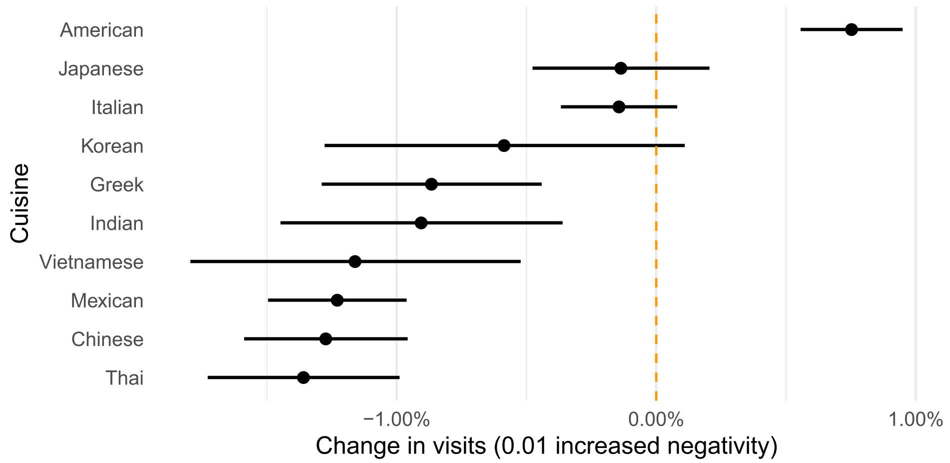

Fig. 3. Effect of political conflict with china on visits to other ethnic restaurants.

The marketing literature refers to consumers' reluctance to purchase foreign products in general as consumer ethnocentrism (Shimp and Sharma, 1987), while the reluctance to buy products associated with a *specific* foreign country is commonly termed consumer animosity (Klein and Ettensoe, 1999). The consumption shifts we observe across various ethnic restaurant types align more with consumer ethnocentrism than animosity. Our results suggest that negative reporting about China induces consumers to refrain from visiting foreign restaurants in general rather than singling out Chinese restaurants. Conversely, an animosity-driven reaction would imply fewer visits to Chinese restaurants, but not necessarily to those serving, for example, Indian, Greek, or Mexican food.[12] The concurrent increase in visits to restaurants offering American food may reflect both a simple substitution effect and a direct behavioral response to the political conflict, consistent with a "rally 'round the flag" effect. Opting for typical American cuisine may serve as a means of expressing support for the national economy during a political adverse context.

Our results suggest that ethnocentric consumption behaviors may be amplified in the context of a political conflict with a foreign country. This behavioral pattern aligns with existing evidence from the political economy literature. This literature suggests that inter-group conflicts, such as the political conflict between the US and China, intensify consumers' attachment to their national identities, which in turn strengthens their in-group bias and adherence to group norms (e.g., Atkin et al., 2021; Shayo, 2020; Pandya and Venkatesan, 2016). Opting to dine at in-group establishments—in this case, American restaurants—is an expression of amplified in-group bias.[13] Moreover, the shift towards choosing typical American food may allow consumers to signal conformity to group norms, a particularly salient behavior during conflict.

It is beyond the scope of this paper to disentangle the relative importance of specific behavioral mechanisms behind the observed ethnocentric reactions to the political conflict, such as increased in-group bias or adherence to group norms. However, one avenue to further investigate these behavioral mechanisms would be to extend our analysis of in-person food consumption to delivery orders of (ethnic) food. If reductions in consumption were primarily driven by heightened adherence to group norms, we might expect to observe smaller decreases in foreign food deliveries compared to dine-in visits, given the lower public visibility of delivery transactions relative to in-person restaurant visits. Conversely, similarly-sized changes across both delivery and in-person food consumption would be indicative of an important role for in-group biases in the reactions to the political conflict.

Finally, exploring whether reactions vary between cheaper and more expensive establishments could provide insights into the price sensitivity of politically motivated consumption behavior. Recent evidence suggests that higher economic costs reduce consumers' adherence to identity-conform food consumption in India, while inter-group conflicts—triggering the salience of consumers' identities—have the opposite effect (Atkin et al., 2021). Research in the context of the Brexit referendum has documented that the willingness and ability to substitute EU-made supermarket products with British products is decreasing with increased prices of the considered consumption goods (Nardotto and Sequeira, 2021). Due to the lack of price data in our restaurant sample, our analysis concerning this important aspect remains limited. However, the insignificant effects of increased political conflict with China on visits to Japanese restaurants, which often sell higher-priced food such as Sushi, is consistent with the notion of less reactive behavior in the case of more expensive cuisines.

6. Conclusion

While anecdotes and survey-based research suggest that political conflicts can affect the consumption of foreign products, it is often difficult to quantify these effects due to accompanying confounders, such as changes in the pricing and availability of products.

[12] It is worth noting that the observed results may arise from a combination of behavioral reactions in line with consumer ethnocentrism and animosity.

[13] Nardotto and Sequeira (2021) examine the impact of the Brexit referendum on British consumers' product choices during grocery shopping. Even though their study focuses on a less visible form of consumption behavior compared to dining choices, they find substantial shifts toward purchasing British produce.

We examine how the political conflict between the US and China affects consumer behavior in the short run. Using smartphone location data, we show that this political conflict significantly reduces visits to Chinese dine-in restaurants. These restaurants are directly associated with China through their branding. However, boycotting visits to any tax-paying American business (including Chinese restaurants) is unlikely to be an effective protectionist action—which could have been consumers' intention for changing their consumption choices away from Chinese restaurants.

Our estimates imply that one additional negative word in each 100-word-long paragraph mentioning China decreases daily visits by up to 1.27% compared to the average non-Chinese restaurant. Notably, demographic characteristics such as age, race, partisanship, and cultural openness moderate the strength of these reactions. Our findings are robust to a wide array of alternative conflict measures and model specifications. Additionally, we highlight the result's external validity by replicating our analysis in the context of varying US-Mexican political relations and their effects on visits to Mexican restaurants in the US.

The political conflict between the US and China also affects other ethnic cuisines. We find that consumers also reduce their consumption of other foreign non-Chinese cuisines while visits to establishments selling American food increase significantly. In this sense, we find empirical support for political conflicts acting as a trigger for consumer ethnocentrism, i.e., a consumption behavior discriminating against foreign products in general (Shimp and Sharma, 1987). Albeit we cannot fully test the mechanisms behind the observed reactions, our findings are in line with the idea that political tensions strengthen American consumers' adherence to their national identity (Shayo, 2020). The shifts in consumption patterns towards goods associated with a particular national identity could be explained by increased in-group bias and adherence to social norms.

Our findings are relevant for firms and policymakers. On the one hand, consumer ethnocentrism is a relevant risk factor for firms: products associated with a political conflict could suffer from sudden decreases in demand, even if they are not the subject of the conflict but merely perceived as foreign. On the other hand, consumer ethnocentrism affects countries' aggregate economic outcomes during international political conflicts. By distorting the competitive positions of local businesses in an unintended way, it makes the allocation of productive resources less efficient.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests:

We would like to declare our cooperation with the following parties:

- *efl – the Data Science Institute*, who supported us through a research grant,
- *SafeGraph*, who provided us with aggregated smartphone location data, and
- *Pollux*, who provided us with access to the newspaper database Factiva.

Data availability

The authors do not have permission to share data.

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.jebo.2024.02.031.

References

Akerlof, G.A., Kranton, R.E., 2000. Economics and identity. Q. J. Econ. 115 (3), 715–753.
Atkin, D., Colson-Sihra, E., Shayo, M., 2021. How do we choose our identity? A revealed preference approach using food consumption. J. Polit. Econ. 129 (4), 1193–1251.
Balabanis, G., Diamantopoulos, A., Mueller, R.D., Melewar, T., 2001. The impact of nationalism, patriotism and internationalism on consumer ethnocentric tendencies. Int. J. Bus. Stud. 32 (1), 157–175.
Balabanis, G., Siamagka, N.T., 2022. A meta-analysis of consumer ethnocentrism across 57 countries. Int. J. Res. Mark. 39 (3), 745–763.
Cao, A., Lindo, J.M., Zhong, J., 2023. Can social media rhetoric incite hate incidents? Evidence from Trump's "Chinese Virus" tweets. J. Urban Econ. 137, 103590.
Charness, G., Chen, Y., 2020. Social identity, group behavior, and teams. Annu. Rev. Econ. 12, 691–713.
Chen, N., Zhong, Z., 2024. History and country-of-origin effects. Mark. Sci. 43 (1), 192–212.
Chetty, R., Jackson, M.O., Kuchler, T., Stroebel, J., Hendren, N., Fluegge, R.B., Gong, S., Gonzalez, F., Grondin, A., Jacob, M., et al., 2022. Social capital I: measurement and associations with economic mobility. Nature 608 (7921), 108–121.
Copeland, L., Boulianne, S., 2022. Political consumerism: a meta-analysis. Int. Polit. Sci. Rev. 43 (1), 3–18.
Fernández-Ferrín, P., Bande-Vilela, B., Klein, J.G., del Río-Araújo, M.L., 2015. Consumer ethnocentrism and consumer animosity: antecedents and consequences. Int. J. Emerg. Mark. 10 (1), 73–88.
Fouka, V., Voth, H.-J., 2013. Reprisals remembered: German-Greek conflict and car sales during the Euro crisis. CEPR Discussion Paper No. DP9704.
Horowitz, D.L., 1985. Ethnic Groups in Conflict. Univ. California Press, Oakland.
Huang, J.T., Krupenkin, M., Rothschild, D., Lee Cunningham, J., 2023. The cost of anti-Asian racism during the COVID-19 pandemic. Nat. Hum. Behav., 1–14.
Jaccard, P., 1901. Distribution florale dans une portion des Alpes et du Jura. Bull. Soc. Vaud. Sci. Nat. 37, 547–579.
Kaufmann, C., 1996. Possible and impossible solutions to ethnic civil wars. Int. Secur. 20 (4), 136–175.
Klein, J.G., Ettensoe, R., 1999. Consumer animosity and consumer ethnocentrism: an analysis of unique antecedents. J. Int. Consum. Mark. 11 (4), 5–24.
Lavy, V., Sand, E., Shayo, M., 2018. The surprisingly small effects of religion-based discrimination in education. Technical report. National Bureau of Economic Research.

Lee, R., 2019. Who does the dishes? Precarious employment and ethnic solidarity among restaurant workers in Los Angeles' Chinese enclave. Ethnicities 19 (2), 433–451.

Loughran, T., McDonald, B., 2016. Textual analysis in accounting and finance: a survey. J. Account. Res. 54 (4), 1187–1230.

Lu, Y., Shao, X., Tao, Z., 2018. Exposure to Chinese imports and media slant: evidence from 147 US local newspapers over 1998–2012. J. Int. Econ. 114, 316–330.

Michaels, G., Zhi, X., 2010. Freedom fries. Am. Econ. J. Appl. Econ. 2 (3), 256–281.

MIT Election Data and Science Lab, 2018. County Presidential Election Returns 2000–2020.

Mohammad, S.M., 2021. Sentiment analysis: automatically detecting valence, emotions, and other affectual states from text. In: Meiselman, H.L. (Ed.), Emotion Measurement. Elsevier, Amsterdam, pp. 323–379.

Nardotto, M., Sequeira, S., 2021. Identity, media and consumer behavior. Working Paper.

Pandya, S.S., Venkatesan, R., 2016. French roast: consumer response to international conflict—evidence from supermarket scanner data. Rev. Econ. Stat. 98 (1), 42–56.

Ramirez, C.D., Rong, R., 2012. China bashing: does trade drive the "bad" news about China in the USA? Rev. Int. Econ. 20 (2), 350–363.

Riefler, P., Diamantopoulos, A., 2007. Consumer animosity: a literature review and a reconsideration of its measurement. Int. Mark. Rev. 24 (1), 87–119.

Shapiro, A.H., Sudhof, M., Wilson, D.J., 2022. Measuring news sentiment. J. Econom. 228 (2), 221–243.

Sharma, S., Shimp, T.A., Shin, J., 1995. Consumer ethnocentrism: a test of antecedents and moderators. J. Acad. Mark. Sci. 23 (1), 26–37.

Shayo, M., 2020. Social identity and economic policy. Annu. Rev. Econ. 12, 355–389.

Shayo, M., Zussman, A., 2011. Judicial ingroup bias in the shadow of terrorism. Q. J. Econ. 126 (3), 1447–1484.

Shimp, T.A., Sharma, S., 1987. Consumer ethnocentrism: construction and validation of the CETSCALE. J. Mark. Res. 24 (3), 280–289.

Sun, Q., Wu, F., Li, S., Grewal, R., 2021. Consumer boycotts, country of origin, and product competition: evidence from China's automobile market. Manag. Sci. 67 (9), 5857–5877.

Zhu, X., Guo, H., Mohammad, S., Kiritchenko, S., 2014. An empirical study on the effect of negation words on sentiment. In: Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics, vol. 1. Association for Computational Linguistics, pp. 304–313.

Zussman, A., 2023. Discrimination in times of crises and the role of the media. Am. Econ. J. Appl. Econ. 15 (4), 422–451.