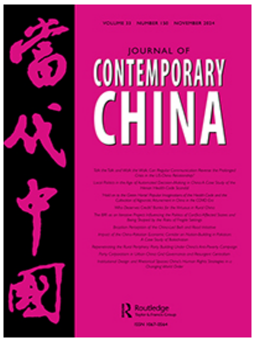

Journal of Contemporary China

ISSN: (Print) (Online) Journal homepage: www.tandfonline.com/journals/cjcc20

How COVID-19 has Impacted American Attitudes Toward China: A Study on Twitter

Gavin G. Cook, Junming Huang & Yu Xie

To cite this article: Gavin G. Cook, Junming Huang & Yu Xie (18 Nov 2024): How COVID-19 has Impacted American Attitudes Toward China: A Study on Twitter, Journal of Contemporary China, DOI: 10.1080/10670564.2024.2427942

To link to this article: https://doi.org/10.1080/10670564.2024.2427942

JOURNAL OF CONTEMPORARY CHINA
https://doi.org/10.1080/10670564.2024.2427942

How COVID-19 has Impacted American Attitudes Toward China: A Study on Twitter

Gavin G. Cook[a], Junming Huang[a] and Yu Xie[a,b]

[a]Princeton University, USA; [b]Peking University, China

ABSTRACT

Past research has studied social determinants of attitudes toward foreign countries. Confounded by potential endogeneity biases due to unobserved factors or reverse causality, the causal impact of these factors on public opinion is usually difficult to establish. Using social media data, we leverage the suddenness of the COVID-19 pandemic to examine whether a major global event has causally changed Americans' views of another country. We collate a database of 297 million posts on the social media platform Twitter about China or COVID-19 up to June 2020, and we take tweeting about COVID-19 as a proxy for an individual's awareness of COVID-19. Using regression discontinuity and difference-in-difference estimation, we reveal that awareness of COVID-19 causes a sharp rise in anti-China attitudes. Our work has implications for understanding how self-interest affects policy preference and how Americans view foreign countries.

KEYWORDS

Twitter; COVID-19; Anti-Chinese Sentiment; Sentiment Analysis

Introduction

The average American does not know a whole lot about the rest of the world.[1] Americans harbor latent, inchoate images about foreign peoples and places that are informed by their own values and other heuristics[2] and only activated in the presence of external primes or triggers.[3] While leaders in other countries may assume that Americans closely follow the intrigues of their home countries, this is not the case.[4] Many Americans simply gloss over news about foreign countries or foreign policies due to a lack of interest.[5]

Regardless of how under- or mis-informed its citizenry may be, the US remains a democracy, and the views that ordinary Americans hold on foreign peoples and places can have strong effects on American foreign policy.[6] This is because public demands affect the policy decisions of elected officials,[7] including in the realm of foreign policy. The opinions of business elites may have a stronger

CONTACT Gavin G. Cook gcook@princeton.edu Paul and Marcia Wythes Center on Contemporary China, 106 Wallace Hal, Princeton University, Princeton, NJ 08544 USA

[1]Holsti O, *Public Opinion and American Foreign Policy, Revised Edition* (University of Michigan Press 2004).

[2]Powlick PJ and Katz AZ, 'Defining the American Public Opinion/Foreign Policy Nexus' [1998] 42 Mershon International Studies Review 29.

[3]Converse PE, 'The Nature of Belief Systems in Mass Publics (1964)' [2006] 18 Critical Review 1.

[4]Anholt S, 'Place Image as a Normative Construct; and Some New Ethical Considerations for the Field' [2010] 6 Place Branding and Public Diplomacy 177.

[5]Graber DA, *Processing the News: How People Tame the Information Tide* (2nd ed, Longman 1988); Gans HJ, *Deciding What's News: A Study of CBS Evening News, NBC Nightly News, Newsweek, and Time/Herbert J. Gans* (Northwestern University Press 2004).

[6]Baum M and Potter PBK, *War and Democratic Constraint: How the Public Influences Foreign Policy* (Princeton University Press 2015).

[7]Shapiro RY, 'Public Opinion and American Democracy' [2011] 75 Public Opinion Quarterly 982; Page BI and Shapiro RY, 'Effects of Public Opinion on Policy' [1983] 77 The American Political Science Review 175.

Supplemental data for this article can be accessed online at https://doi.org/10.1080/10670564.2024.2427942.

impact on foreign policy outcomes than the opinions of the general public, but the attitudes of the general public still have a measurable and significant influence.[8] Examples abound of presidents or officials taking less-than-optimal foreign policy actions to appease constituents, such as then-US President Bill Clinton supposedly advocating for the expansion of NATO to appease voters of Central and Eastern European descent in the 1990s.[9] Public opinion may play a role in maintaining peace between democracies[10] and may be more important now than during the Cold War, when the Soviet Union was as a unifying enemy in the American mind.[11] The views of the American people on China today are important because they constrain and shape the possible actions and stances that the American state may adopt when engaging with China in the future. This is especially significant in an international environment fraught with the potential for armed conflict.[12]

In this study, we seek to understand how the 2020 outbreak of the COVID-19 pandemic may have impacted American views on China by analyzing a large corpus of posts on the social media platform Twitter. The unique exogeneity and magnitude of COVID-19's effects on the US provide a valuable opportunity to contribute to unresolved questions in the literatures on self-interest and foreign policy by drawing on causal inference methods. We use regression discontinuity and difference-in-difference (DID) methods to estimate the causal effect of awareness of COVID-19 on sentiment toward China, controlling for time of exposure. We use posting on Twitter ('tweeting') about COVID-19 as a proxy for awareness of COVID-19 and find that awareness of COVID-19 leads to an increase in anti-China sentiment.

Self-Interest and Foreign Policy

COVID-19 started as a foreign issue but quickly became a domestic crisis in the United States. The virus began to spread rapidly from March 2020 onward in the US. Within four months of the first case of COVID-19 in the US, 100,000 Americans had lost their lives to the disease.[13] As of 7 June 2021, 34 million Americans had been infected, and 612,000 Americans had died as a result of their infection.[14] Further consequences include impairment of mental health, including elevated suicide rates[15] and increased economic anxiety[16] and the persistent neurological symptoms experienced by sufferers of 'long COVID'.[17]

Americans also suffered devastating economic consequences as a direct result of the pandemic. Consumer spending fell in all states.[18] Stay-at-home orders caused a 50% reduction in trips to non-essential businesses and a 19% drop in revenue for small businesses.[19] By April 2020, employment rates had fallen by 21% relative to pre-pandemic levels, and this disproportionately affected small

[8]Jacobs LR and Page BI, 'Who Influences U.S. Foreign Policy?' [2005] 99 The American Political Science Review 107.
[9]Alison Mitchell, 'Clinton Urges NATO Expansion in 1999' *The New York Times* (New York 23 October 1996) 20.
[10]Tomz MR and Weeks JLP, 'Public Opinion and the Democratic Peace' [2013] 107 American Political Science Review 849.
[11]Holsti O, *Public Opinion and American Foreign Policy, Revised Edition* (University of Michigan Press 2004).
[12]Allison G, *Destined for War: Can America and China Escape Thucydides's Trap?* [2018].
[13]Dave Sanders, 'Four Months After First Case, U.S. Death Toll Passes 100,000' *The New York Times* (27 May 2020). https://www.nytimes.com/2020/05/27/us/coronavirus-live-news-updates.html.
[14]Worldometer, 'United States COVID: 34217,262 Cases and 612,512 Deaths—Worldometer' (7 June 2021). https://www.worldometers.info/coronavirus/country/us/.
[15]Sher L, 'The Impact of the COVID-19 Pandemic on Suicide Rates' [2020] 113 QJM: monthly journal of the Association of Physicians 707.
[16]Mann FD, Krueger RF and Vohs KD, 'Personal Economic Anxiety in Response to COVID-19' [2020] 167 Personality and Individual Differences 110,233.
[17]Graham EL et al., 'Persistent Neurologic Symptoms and Cognitive Dysfunction in Non-Hospitalized COVID-19 "Long Haulers"' [2021] 8 Annals of Clinical and Translational Neurology 1073.
[18]Dong D et al., 'Personal Consumption in the United States during the COVID-19 Crisis' [2021] 53 Applied Economics 1311.
[19]Alexander D and Karger E, 'Do Stay-at-Home Orders Cause People to Stay at Home? Effects of Stay-at-Home Orders on Consumer Behavior' (Federal Reserve Bank of Chicago 2020).

businesses and low-wage workers.[20] Overall, the GDP of the US nosedived from 21.5 trillion to 19.5 trillion dollars between the first and second quarter of 2020.[21]

As we will demonstrate later in this paper, there was a sharp increase in negative attitudes toward China when COVID-19 began to affect American lives. Many Americans associated COVID-19 with China in the first months of the pandemic. Then-US President Donald Trump insisted on referring to COVID-19 as the 'China Virus' despite the media and much of the populace calling him a xenophobe. Nonetheless, as the pandemic reached the US and children were pulled from schools and customers kept from businesses, many Americans blamed China for their misfortunes.

The COVID-19 pandemic exposed many Americans to a historic and unprecedented combination of economic and physical hardship that they could blame on a foreign source. As such, it offers an excellent opportunity to examine the question of how self-interest informs political opinions—in this case, views on foreign nation-states. An active literature has been devoted to understanding the relationship between self-interest and political behavior, with two main camps adopting opposing views. On one side of the debate are scholars who argue that self-interest does not strongly influence political behavior. These scholars argue that voters behave with 'sociotropic', loosely meaning societal- or national-level, interests in mind.[22] Findings in this vein include those of Citrin et al.,[23] who show that an individual's economic situation does not predict his or her views on immigration restrictionism but that his or her appraisal of the nation's economic situation does.[24] Research finding that self-interest does matter, however, is much more common. American men drafted to fight in the Vietnam War developed an understandable opposition to war and consequently became more liberal and voted more consistently for the Democratic Party later in life.[25] Self-interest similarly motivates gun owners to oppose bans on firearms.[26] Interest may also operate at the group level. Calculations of group-based self-interest have led Whites to oppose race-based redistributive measures.[27] Wealthier seniors are more politically active than other senior citizens but, because they rely less on social security than their peers, they do not vote with social security in mind and are not mobilized to vote by appeals to social security as an issue.[28] This extends to farmers, who participate in politics more eagerly than their socioeconomic status would otherwise predict[29] and people who lost their jobs during the Great Recession, who were more likely to support welfare programs after experiencing economic hardship.[30]

It is possible to reconcile these conflicting claims with reference to psychological models of political behavior, but this reconciliation does not bring us closer to resolving the two opposing views with an empirical test. Priming subjects with vignettes or questions that invoke respondents to reflect on their own financial circumstances leads to greater influence of self-interest on voting

[20]Cajner T et al., 'The U.S. Labor Market during the Beginning of the Pandemic Recession' (National Bureau of Economic Research 2020) w27159. http://www.nber.org/papers/w27159.pdf.
[21]U.S. Bureau of Economic Analysis, 'Table 1.1.5. Gross Domestic Product'. https://apps.bea.gov/iTable/iTable.cfm?reqid=19&step=2#reqid=19&step=2&isuri=1&1921=survey.
[22]Meehl PE, 'The Selfish Voter Paradox and the Thrown-Away Vote Argument*' [1977] 71 American Political Science Review 11; Kinder DR and Kiewiet DR, 'Sociotropic Politics: The American Case' [1981] 11 British Journal of Political Science 129.
[23]Citrin J et al., 'Public Opinion Toward Immigration Reform: The Role of Economic Motivations' [1997] 59 The Journal of Politics 858.
[24]Sears DO and Funk CL, 'The Limited Effect of Economic Self-Interest on the Political Attitudes of the Mass Public' [1990] 19 Journal of Behavioral Economics 247.
[25]Erikson RS and Stoker L, 'Caught in the Draft: The Effects of Vietnam Draft Lottery Status on Political Attitudes' [2011] 105 American Political Science Review 221.
[26]Wolpert RM and Gimpel JG, 'Self-Interest, Symbolic Politics, and Public Attitudes Toward Gun Control' [1998] 20 Political Behavior 241.
[27]Bobo L and Kluegel JR, 'Opposition to Race-Targeting: Self-Interest, Stratification Ideology, or Racial Attitudes?' [1993] 58 American Sociological Review 443.
[28]Campbell AL, 'Self-Interest, Social Security, and the Distinctive Participation Patterns of Senior Citizens' [2002] 96 The American Political Science Review 565.
[29]Lewis-Beck MS, 'Agrarian Political Behavior in the United States' [1977] 21 American Journal of Political Science 543.
[30]Margalit Y, 'Explaining Social Policy Preferences: Evidence from the Great Recession' [2013] 107 American Political Science Review 80.

behavior.[31] Some voters, however, may find it difficult to figure out the influence a particular policy might have on their daily lives. More-educated voters are more likely to see connections between their own self-interest and policy proposals and then vote with their self-interest in mind,[32] but many voters 'report no connection between government actions and personal experience'.[33] This disconnect may lead to behaviors that appear altruistic but are actually motivated by a lack of understanding of one's own self-interest. This is to say that some voters do not think with self-interest in mind because they do not understand how abstract policies impact their own lives. Uninformed voters, then, do not make voting decisions on the basis of self-interest but instead default to values-based heuristics like party affiliation.[34] This means that it is challenging to craft a research design that tests both the sociotropic and self-interest theories of political behavior or to use techniques of causal inference to examine these topics. We do not seek to answer this debate. We instead aim to use causal inference in a context where the concerns of both of sides of the debate are answered.

The shock of COVID-19 has been so unprecedentedly large that it elides the distinctions between the sociotropic and self-interest models of political behavior. It matters not if the Twitter users in our data weight individual-level interest more heavily than national-level interest in their decision-making calculus, or vice versa; COVID-19 has negatively affected countries and individuals alike. Because the impact of COVID-19 was so sudden, widespread, and deeply felt, we are able to use methods of causal inference on our corpus of Twitter posts ('tweets'). Our outcome of interest, however, is not how COVID-19 affected American views of particular policies but the related question of how COVID-19 changed Americans views on a foreign country.

American Views on China

While America and China have a long history of interaction and exchange, they have butted heads over trade and human rights very dramatically over the past decade. China and the US first came into concerted contact in the 1800s, when Chinese laborers from Taishan migrated to San Francisco to work in California gold mines and were eventually limited from migrating to the United States with the 1882 passage of the Chinese Exclusion Act.[35] This period was marked by anxieties over Chinese immigration and anti-Chinese sentiment, eventually sparking the firing of E. A. Ross from Stanford and in 1900 and provoking nationwide debates about freedom of speech and racial vitriol.[36] While these centuries-old contacts certainly inform contemporary American stereotypes of the Chinese people, we are concerned with American views on the Chinese state. The most relevant period for our work begins with Nixon's visit to China and the ensuing normalization of bilateral US-China relations in 1972.[37] Regular surveys of how the American public view China began shortly after, and they indicate that Americans generally viewed China favorably in the 1970s and into the 1990s. The Tiananmen Square massacre

[31]Chong D, Citrin J and Conley P, 'When Self-Interest Matters' [2001] 22 Political Psychology 541.
[32]Young J et al., 'When Self-Interest Makes a Difference: The Role of Construct Accessibility in Political Reasoning' [1991] 27 Journal of Experimental Social Psychology 271.
[33]Funk CL, 'The Dual Influence of Self-Interest and Societal Interest in Public Opinion' [2000] 53 Political Research Quarterly 37.
[34]Popkin SL, *The Reasoning Voter: Communication and Persuasion in Presidential Campaigns* (2nd ed, University of Chicago Press 1994); Achen CH and Bartels LM, *Democracy for Realists: Why Elections Do Not Produce Responsive Government: With a New Afterword by the Authors* (Princeton University Press 2017).
[35]Lew-Williams B, The Chinese Must Go: Violence, Exclusion, and the Making of the Alien in America (Harvard University Press 2018).
[36]As an example, Ross opined in the preface of his 1913 *The Old World in the New* that 'We [White Americans] could have helped the Chinese a little by letting their surplus millions swarm in upon us a generation ago; but we have helped them infinitely more by protecting our standards and having something worth their copying when the time came'. See: Edward Alsworth Ross, The Old World in the New: The Significance of Past and Present Immigration to the American People (Project Gutenberg 2015). https://www.gutenberg.org/ebooks/47954.
[37]Sutter RG, US-China Relations: Perilous Past, Uncertain Future (Fourth edition, Rowman & Littlefield 2022).

provoked a moderate degradation in the American public's views on China, but it prompted a much larger decline in journalistic views on China.[38] American views on China begin to decline precipitously shortly after China's election to the WTO in 2001, though the media's view of China declined at a slower clip. Trump's election in 2016 marked the entry of politicos with intense anti-China views walking into the halls of government.[39]

At present, Americans generally view China more negatively than Chinese nationals view America, but there are important differences in the drivers of these opinions. More politically conservative Americans have more protectionist and isolationist views on foreign policy,[40] and they also tend to view China more negatively than other do other Americans.[41] While Chinese nationals are mostly concerned with questions of China's national interests, American opinion on China is shaped by perceptions of China's environmental and human rights record.[42] Indeed, most extant survey data reveal that Americans do not hold favorable views of China.[43] Americans are anxious about China's growing economic and military power,[44] and they doubt that the needs of the Chinese people can be adequately served by the Chinese political system.[45] Many Americans appreciate Chinese culture and respect the Chinese people while recognizing a difference between the Chinese people and the Chinese state.[46] It may be the case that Americans dislike the Chinese state because Americans view the Chinese people so favorably and view the Chinese state as the main oppressor of the Chinese people. A survey of English-language social media posts about China in 2011 found that users appreciated China's economic and cultural progress but disliked the repressiveness of China's political system.[47] Given this, it is possible that knowledge of China's strict COVID-19 lockdowns may have negatively affected American attitudes toward China. Some of the tweets in our data support this interpretation (Table S1), but we do not explore this explanation in detail.

There is some evidence that COVID-19 has negatively impacted American views on the Chinese state and the Chinese people. When White Americans are prompted with a vignette of new scientific evidence about the threat of a global pandemic from Asia, their bias toward Asian Americans increases.[48] Regions with higher rates of COVID-19 incidence view the Chinese people and the Chinese state less favorably than regions with lower rates of COVID-19.[49,50] Our work adds to these literatures by clarifying the causal link between COVID-19 and the emergence of negative sentiment toward the Chinese state in the earliest months of the COVID-19 pandemic.

38Huang J, Cook GG and Xie Y, 'Between Reality and Perception: The Mediating Effects of Mass Media on Public Opinion toward China' [2021] 53 Chinese Sociological Review 431; For a briefer treatment, see: Huang J, Cook GG and Xie Y, 'Large-Scale Quantitative Evidence of Media Impact on Public Opinion toward China' [2021] 8 humanities and Social Sciences Communications 181.
39Sutter (n 37).
40Bonikowski B and DiMaggio P, 'Varieties of American Popular Nationalism' [2016] 81 American Sociological Review 949.
41Gries PH, Crowson HM and Cai H, 'God, Guns, and . . . China?' [2012] 12 International Relations of the Asia-Pacific 1.
42Wang D et al., 'In the Eyes of the Beholder: How China and the US See Each Other' [2022] 31 Journal of Contemporary China 232.
43Xie Y and Jin Y, 'Global Attitudes toward China: Trends and Correlates' [2022] 31 Journal of Contemporary China; Cao Y and Xu J, 'The Tibet Problem in the Milieu of a Rising China: Findings from a Survey on Americans' Attitudes toward China' [2015] 24 Journal of Contemporary China 240; Aldrich J, Lu J and Kang L, 'How Do Americans View the Rising China?' [2015] 24 Journal of Contemporary China 203.
44Gries PH and Crowson HM, 'Political Orientation, Party Affiliation, and American Attitudes Towards China' [2010] 15 Journal of Chinese Political Science 219; Yang YE and Liu X, 'The "China Threat" through the Lens of US Print Media: 1992–2006' [2012] 21 Journal of Contemporary China 695.
45Aldrich J, Lu J and Kang L, 'How Do Americans View the Rising China?' [2015] 24 Journal of Contemporary China 203.
46Gries PH and Crowson HM, 'Political Orientation, Party Affiliation, and American Attitudes Towards China' [2010] 15 Journal of Chinese Political Science 219.
47Xiang D, 'China's Image on International English Language Social Media' [2013] 19 Journal of International Communication 252.
48Zhao J, Tinkler JE and Clayton KA, 'Assessing the Causal Link between the COVID-19 Pandemic and Racial Discrimination' [2022] 8 Socius: Sociological Research for a Dynamic World 237,802,312,210,953.
49He Q, Zhang Z and Xie Y, 'The Impact of COVID-19 on Americans' Attitudes toward China: Does Local Incidence Rate Matter?' [2022] 85 Social Psychology Quarterly 84.
50For more survey data on how American views on China changed after COVID-19, see: 'Unfavorable Views of China Reach Historic Highs in Many Countries'. Pew Research Center, Washington, D.C. 2020. https://www.pewresearch.org/global/2020/10/06/unfavorable-views-of-china-reach-historic-highs-in-many-countries/

Identifying the Attitudes of Twitter Users Toward China

We draw upon Twitter as a barometer for how the American public views China. Social media data have been used for similar purposes by a number of other scholars. A sentiment analysis of American tweets found that negatively polarized tweeting behavior was predictive of sharing 'fake news'.[51] Twitter data have been used to track how far-right European politicians use fake bot accounts to gain inflated followings[52] and how the news media influence public discussion of given issues.[53] Twitter has been used extensively in health and medicine research.[54] Large corpora comprised of posts from Chinese social microblogging sites that are very similar to Twitter, such as Weibo, have been used to track censorship in mainland China.[55] Finally, there is precedent for our design of using tweet sentiment in a causal inference framework. Pan and Siegel[56] analyze tweets of Saudi dissidents to examine how tweets about the Saudi regime changed before and after dissidents were arrested.

There is also a growing body of work on Twitter in the context of COVID-19. Public health emergencies create negative sentiment,[57] and COVID-19 vaccination rates have been found to correlate negatively with negative sentiment regarding COVID-19 on Twitter.[58] Tweets from government authorities in the US about COVID-19, however, have improved public attitudes toward the pandemic.[59] There is also a literature on Twitter and COVID-19 specifically in relation to anti-Chinese discrimination.[60] Some of this work has focused specifically on the backgrounds of users who tweet hateful content and has found, for example, that the types of accounts that a user follows are highly predictive of liability to tweet hateful content.[61] Some authors suggest that the increase in anti-Asian tweets, defined in this case as tweets that refer to COVID-19 as the 'Chinese virus', caused many to see Asian-Americans as less American on Implicit Association Tests.[62]

Our work extends this literature by exploring American sentiment not on the pandemic itself but how COVID-19 affected American views on a foreign country that was frequently associated with COVID-19, namely China. This is particularly important because foreign countries and foreign peoples are often maligned by association with plagues.[63] While data

[51]Osmundsen M et al., 'Partisan Polarization Is the Primary Psychological Motivation behind Political Fake News Sharing on Twitter' [2021] American Political Science Review.

[52]Silva BC and Proksch S-O, 'Fake It "Til You Make It: A Natural Experiment to Identify European Politicians" Benefit from Twitter Bots' [2021] 115 American Political Science Review 316.

[53]King G, Schneer B and White A, 'How the News Media Activate Public Expression and Influence National Agendas' [2017] 358 Science 776.

[54]Sinnenberg L et al., 'Twitter as a Tool for Health Research: A Systematic Review' [2017] 107 American Journal of Public Health e1.

[55]King G, Pan J and Roberts ME, 'How Censorship in China Allows Government Criticism but Silences Collective Expression' [2013] 107 American Political Science Review 326; King G, Pan J and Roberts ME, 'How the Chinese Government Fabricates Social Media Posts for Strategic Distraction, Not Engaged Argument' [2017] 111 American Political Science Review 484.

[56]Pan J and Siegel AA, 'How Saudi Crackdowns Fail to Silence Online Dissent' [2020] 114 American Political Science Review 109.

[57]Li T and others, 'Constructing a Multi-Layer Heterogeneous Networks Model to Explore the Public Opinion Evolution Pattern of Key Users in Public Health Emergencies' [2023] Journal of Information Science 016555152311699.

[58]Jun J and others, 'Adverse Mentions, Negative Sentiment, and Emotions in COVID-19 Vaccine Tweets and Their Association with Vaccination Uptake: Global Comparison of 192 Countries' [2022] 10 Vaccines 735.

[59]Xi H and others, 'Public Emotional Diffusion over COVID-19 Related Tweets Posted by Major Public Health Agencies in the United States' [2022] 4 Data Intelligence 66.

[60]Hswen Y and others, 'Association of "#covid19" Versus "#chinesevirus" With Anti-Asian Sentiments on Twitter: March 9–23, 2020' [2021] 111 American Journal of Public Health 956; Toliyat A and others, 'Asian Hate Speech Detection on Twitter during COVID-19' [2022] 5 Frontiers in Artificial Intelligence 932,381.

[61]An J and others, 'Predicting Anti-Asian Hateful Users on Twitter during COVID-19', *Findings of the Association for Computational Linguistics: EMNLP 2021* (Association for Computational Linguistics 2021) https://aclanthology.org/2021.findings-emnlp.398.

[62]Darling-Hammond S and others, 'After "The China Virus" Went Viral: Racially Charged Coronavirus Coverage and Trends in Bias Against Asian Americans' [2020] 47 health Education & Behavior 870.

[63]As an example, the Chinese have been maligned with the disease-adjacent epithet 'Yellow Peril' in English-language publications for centuries. See: Tchen, John Kuo Wei, and Dylan Yeats, eds. *Yellow Peril! An Archive of Anti-Asian Fear* (London New York: Verso 2014).

from other countries in the Anglosphere suggest that US-based users of Twitter and Facebook may be younger, more liberal, and more educated than the general populace,[64] Twitter data is still useful for its shear breadth. Almost 60 million Americans use Twitter,[65] and the platform's reach means that it is the closest thing to a barometer of national-level public opinion.[66]

We searched for tweets that mentioned China and merged the resulting data with a public dataset of tweets that mention COVID-19.[67] We then identified users whose public profiles revealed locations in the US.[68] Together, the corpus consists of 297 million tweets between January 2017 and June 2020. We are aware that this is likely not a representative sample of all Twitter users, for previous work has shown that Twitter users who choose to reveal their locations tend to be wealthier on average than those who do not.[69] Our intent, however, is not to analyze the characteristics of Twitter users but to investigate how Twitter users in the United States respond to a discrete external stressor. By only inspecting the tweets of users who have declared their profile location to be in the United States as opposed to using other methods of location detection that may cast a broader net, such as using all users who have ever tagged individual tweets with US-based locations or inferring location algorithmically,[70] our sample of US-based Twitter users is very conservative and likely provides an accurate representation of the discourses regarding China and COVID-19 on Twitter in the United States from January 2017 to June 2020.[71]

We pointedly do not attempt to create a link between data from surveys on how Americans view China and our Twitter data. Many early papers on Twitter data treated data culled from social media as either equal or superior to that drawn from surveys with an enthusiasm that could be read in the very titles of said papers,[72] but this enthusiasm has long since waned, and we, like most in the field, are cautious about recklessly generalizing our results. Surveys of how Americans viewed China during the early months of the pandemic certainly exist, and they are of high quality.[73] These surveys also generally agree with our conclusions. Methodologically, however, harmonizing Twitter data and survey data is very complex and far outside the scope of this paper. Other teams have provided an illustration of how surveys might be harmonized more

[64]Mellon J and Prosser C, 'Twitter and Facebook Are Not Representative of the General Population: Political Attitudes and Demographics of British Social Media Users' [2017] 4 Research & Politics 205,316,801,772,000.

[65]NCES, 'Twitter: number of users in the United States 2017–2022' [2022] Statista—The Statistics Portal. https://www.statista.com/statistics/232818/active-us-twitter-user-growth/.

[66]Twitter data is the premier data source for social media research for a few reasons. Firstly, as a simple survey of Twitter readily proves, many public opinion elites and news companies use Twitter, implying that the platform has outsize importance for shaping American public opinion. Secondly, it is relatively easy to access: data from platforms owned by Meta, such as Facebook and Instagram, is more difficult to access in the wake of the Cambridge Analytica scandal. While there are other social media platforms beyond those previously mentioned, most of them, such as Gab and 4Chan, are devoted to free speech and hence very loosely moderated. They have become havens for the far-right and are filled with offensive speech. For more information on the Cambridge Analytica scandal, see: Schneble CO, Elger BS and Shaw D, 'The Cambridge Analytica Affair and Internet-mediated Research' (2018) 19 EMBO reports e46579. For more information on Gab, see: Zannettou S and others, 'What Is Gab: A Bastion of Free Speech or an Alt-Right Echo Chamber', Companion of the The Web Conference 2018 on http://dl.acm.org/citation.cfm?doid=3184558.3191531.

[67]Chen E, Lerman K and Ferrara E, 'COVID-19: The First Public Coronavirus Twitter Dataset' [2020] http://arxiv.org/abs/2003.07372.

[68]We recognize that many Twitter users located in the US may not be American citizens, but for the purposes of this paper, we will describe all of the users in this paper as 'Americans'. Our research question is agnostic to concerns of nationality or citizenship; the only individual-level characteristic of interest is whether or not a user tweets from the United States.

[69]Sloan L and Morgan J, 'Who Tweets with Their Location? Understanding the Relationship between Demographic Characteristics and the Use of Geoservices and Geotagging on Twitter' [2015] 10 PLOS ONE e0142209.

[70]Serere HN, Resch B and Havas CR, 'Enhanced Geocoding Precision for Location Inference of Tweet Text Using spaCy, Nominatim and Google Maps. A Comparative Analysis of the Influence of Data Selection' [2023] 18 PLOS ONE e0282942.

[71]For more information on how geotagging rules on Twitter have changed over time, see: Kruspe A and others, 'Changes in Twitter Geolocations: Insights and Suggestions for Future Usage' (arXiv, 22 September 2021) http://arxiv.org/abs/2108.12251.

[72]Bollen J, Mao H and Zeng X, 'Twitter Mood Predicts the Stock Market' [2011] 2 Journal of Computational Science; Tumasjan A and others, 'Election Forecasts With Twitter: How 140 Characters Reflect the Political Landscape' [2011] 29 Social Science Computer Review 402.

[73]Kat Devlin, Laura Silver, and Christine Huang, 'U.S. Views of China Increasingly Negative Amid Coronavirus Outbreak' (*Pew Research Center, Washington, D.C.*, 21 April 2020) https://www.pewresearch.org/global/2020/04/21/u-s-views-of-china-increasingly-negative-amid-coronavirus-outbreak/.

generally[74] and how Twitter data and survey data might be directly compared,[75] and it is worth noting that researchers generally find that data from Twitter and from surveys agree with one another.[76] We invite readers to compare Figure S1, which illustrates changes in Twitter sentiment on China from 2017 to 2020, to Pew data. We stress, however, that we do not formally compare nor otherwise harmonize these sets of data.

After collating our data, we employed and trained eight human coders to manually label 5,000 tweets with sentiment scores ranging from 'most unfavorable' to 'most favorable' (Table S1).[77] We then used the tagged tweets to fine-tune BERT, a deep-learning model for natural language processing,[78] to automatically identify the sentiment scores of all of the tweets in our corpus. Those scores are presented in a range from −100 (most unfavorable) to 100 (most favorable).

To avoid confusing attitudes toward China with attitudes toward COVID-19, we only measure, as our outcome variable, attitudes toward China in tweets that mention China but not COVID-19. To avoid biasing our attitude measurements toward highly active users, we calculate sentiment scores at the user level before we calculate averages between individuals. We report this two-stage average as the average daily sentiment toward China in the US.

Data

Social media data were collected from Twitter public posts and profiles to analyze American attitudes toward China. Starting from an open library of nearly 100 million COVID-19-related tweets, we use an open-source tool to search tweets containing China-related and COVID-19-related keywords. We specifically searched for tweets containing any mention of the words 'China', 'Chinese', 'Wuhan', and 'Beijing'[79] from January 2017 to June 2020 and tweets containing any mention of the words 'covid', 'covid-19', 'covid19', 'ncov', 'ncov19', 'coronavirus', 'corona+virus', 'quarantine', 'lockdown', 'epidemic', 'pandemic', 'pandemia', 'outbreak', 'social+distancing', 'mask', 'confirm+case', 'death+case', 'flatten +curve' from January 2020 to June 2020. After merging the two datasets and removing duplicate tweets, we end up with 297,232,539 unique tweets, 120,250,229 of which mention China-related terms. Our dataset contains tweets from 14 million Twitter users, 9 million of which have mentioned China at least once and 8 million of which have mentioned COVID-19 at least once. We further obtain a subset of 68,979,579 unique English-language tweets mentioning China but not COVID-19. Each tweet bears a time stamp and a unique user identifier, and we anonymize all user identifiers after downloading tweets. We search for Twitter users who have voluntarily revealed geographic information in their profiles and then use Google Maps to identify those located in the United States. We end up with 174,187 users who (1) are living in the US and (2) tweeted at least one tweet about China in every month from January to April 2020. We then analyze the individual-level trends of sentiment toward China of those users.

[74]Wang D, Xie Y and Huang J, 'Trend Analysis with Pooled Data from Different Survey Series: The Latent Attitude Method' [2024] 54 Sociological Methodology 118.

[75]Van Klingeren M, Trilling D and Möller J, 'Public Opinion on Twitter? How Vote Choice and Arguments on Twitter Comply with Patterns in Survey Data, Evidence from the 2016 Ukraine Referendum in the Netherlands' [2021] 56 Acta Politica 436.

[76]Johannes Breuer and others, 'Assessing the Relationship between Survey Data and Twitter Data as Measures of Public Opinion —A Methodological Pilot Study', ESRA 2021 full program [2021]. https://www.europeansurveyresearch.org/conf2021/uploads/219/2790/62/Relationship_between_survey_data_and_Twitter_data_as_measures_of_public_opinion_ESRA2021.pdf.

[77]More specifically, the authors began by annotating roughly 1,000 tweets. Eight human coders were then trained on the specific coding task and told to annotate a larger selection of 5,000 tweets. The annotations from the coders with the highest accuracy were chosen to train the model, and the inter-rater reliability of the chosen annotations was 0.65.

[78]Devlin J and others, 'BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding' (arXiv, 24 May 2019) http://arxiv.org/abs/1810.04805.

[79]We chose these words specifically because we are interested in sentiment toward China at the national level: 'China' and 'Chinese' give us information on the Chinese state and the Chinese people. 'Wuhan' is where many Americans suspect the COVID-19 pandemic began, which makes it important for tracing sentiment on the pandemic, and because the Chinese state is headquartered in Beijing, 'Beijing' is often used a shorthand for the CCP.

Tweet sentiments are quantitatively inferred with BERT, a natural language processing model that utilizes deep neural networks. We first employed research assistants to manually label the expressed China-related sentiment of 5,000 English-language tweets. These tweets were labeled with an integer-valued favorability score from −2 to 2, i.e. −2 (most unfavorable), −1 (somewhat unfavorable), 0 (neutral), 1 (somewhat favorable), 2 (most favorable). Table S1 contains examples of tweets with a range of favorability scores. We then use the labeled tweets to fine-tune a pretrained BERT model. While large language models (LLMs) are considered state of the art for classification tasks, they are very new and are only just beginning to find use in the social sciences. While the literature on LLMs in the social sciences is in its infancy,[80] the usage of BERT and other transformer-based models in the social sciences is prevalent enough that there are multiple pre-trained BERT libraries for social scientific tasks.[81] We elected to use the modeling approach with a longer history of use. Our model was able to map tweet text to integer favorability score with an accuracy of 88.8%.[82] The fine-tuned model assigns favorability scores to all English-language tweets in the entire corpus. We finally proportionally map the scores from [−2, 2] to [−100, 100] for easier presentation, where −100 represents most unfavorable, 0 represents neutral, and 100 represents most favorable.

Attitudes toward China are measured for tweets mentioning China but not COVID-19. This is because a large fraction of tweets that mention China in early 2020 also mention COVID-19, and the attitudes expressed in these tweets may be directed at China, COVID-19, or both China and COVID-19. In fact, 5 in 6 of the Twitter users who talked about China in February 2020 also mentioned COVID-19 in the same tweet, and most tweets about COVID-19 are unsurprisingly negative. Given that the rapid increase in overwhelmingly negative tweets on COVID-19 would negatively bias our measurements of China-related sentiment, we exclude said tweets when measuring attitudes toward China.

Results

A Descriptive Summary of Tweet Volume and Sentiment Toward China on Twitter Before January 2020

Inchoate attitudes are activated by external triggers, and COVID-19 is a significant trigger for attitudes toward China. China was very much on the minds of the American people in the initial stages of the global pandemic beginning in January 2020, with tweet volume on China reaching a three-year high. Figure 1 displays the sentiment and volume of tweets on China. The solid line (main axis) reveals a two-stage reaction, with two distinct dips in sentiment toward China after the initial outbreak of COVID-19. The secondary axis denoted in the figure by dark blue shading shows that the number of users tweeting on China increased tenfold, shooting from less than 6,000 users per day to a peak of greater than 60,000 users per day.

We observe two drops in China-related sentiment among American Twitter users: an initial, smaller drop January 2020 that neatly coincides with the outbreak of COVID-19 in China and a second, larger drop in March 2020 that aligns with the enacting of COVID-19 quarantine measures in a number of American states (Figure 1). The main objective of this paper is to quantify the causal relationship between COVID-19 and anti-China sentiment, and we do not

[80] Jensen JL and others, 'Language Models in Sociological Research: An Application to Classifying Large Administrative Data and Measuring Religiosity' [2022] 52 Sociological Methodology 30.

[81] For a general example, see: Shen S and others, 'SsciBERT: A Pre-Trained Language Model for Social Science Texts' [2023] 128 Scientometrics 1241. For an example specific to economics, see: Lasri K and others, 'EconBERTa: Towards Robust Extraction of Named Entities in Economics', *Findings of the Association for Computational Linguistics: EMNLP 2023* (Association for Computational Linguistics 2023). https://aclanthology.org/2023.findings-emnlp.774.

[82] The F1 score of the model was 0.94, the precision of the model was 0.89, and the recall of the model was 0.98.

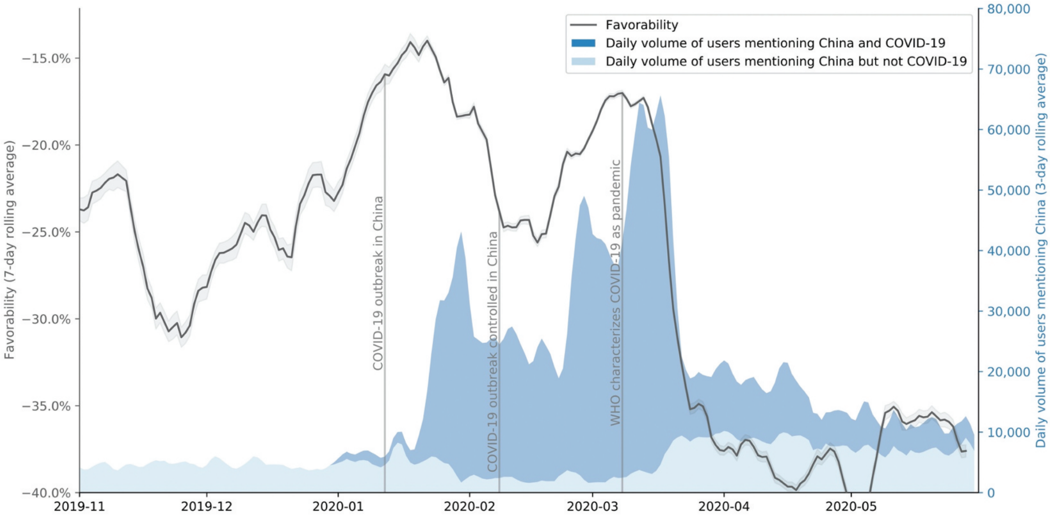

Figure 1. American attitudes toward China declined to a 3-year low of − 38 after the outbreak of COVID-19, on a range from − 100 to + 100. Attitude toward China is measured by averaging sentiment across Twitter users in tweets that mention China but not COVID-19 in a 7-day sliding window. Sudden and steep declines in China-related sentiment occur nationwide. The declines are accompanied by increases in the volume of tweets on China, denoted in the figure with light and dark blue shaded areas. The daily number of users mentioning China on Twitter increased from 5,096 in late December 2019 to 66,535 by the middle of March 2020. See Figure S1 for a full version of this chart.

claim to conclusively elucidate the mechanisms behind either of these declines. The first decline in January 2020 occurred before the material interests of Americans were most threatened by COVID-19. This decline may have been caused by a disgust reaction triggered by the 'behavioral immune system'.[83,84] It may also be the case that the implementing of full-city lockdowns in China irritated freedom-conscious Americans, according with preexisting stereotypes of China as a repressive, authoritarian state (Table S1). The second and more dramatic drop in China-related sentiment in March 2020 is more readily but not conclusively attributable to self-interest.

The sharp drops in China-related sentiment on Twitter after the outbreak of COVID-19 were sudden and unprecedentedly drastic. In the three years before January 2020, the average weekly sentiment of China-related tweets, which we quantify on a range from −100 (most unfavorable) to 100 (most favorable), was between – 30 and – 10 (Figure S1). The drop in sentiment toward China in January 2020 from −15 to −25 was sudden and sharp. This implies that the increasing prevalence of anti-China attitudes on Twitter was not the result of normal swings but was instead a response to an external trigger.

We also see this decline at the state level. Across the 10 states with the most Twitter users, namely California, New York, Texas, Florida, District of Columbia, Kansas, Illinois, Washington, Georgia, and Massachusetts, attitudes toward China declined sharply at almost the same time and to almost the same depth as at the national level (dashed lines in Figure S1). This consistency across states further

[83] Schaller M and Park JH, 'The Behavioral Immune System (and Why It Matters)' [2011] 20 Current Directions in Psychological Science 99.

[84] According to behavioral immune system theory, while the biological immune system fights pathogens once they enter the body, the behavioral immune system conditions us to keep our distance from sickness by provoking emotional reactions, primarily disgust, to induce avoidance. Such behavioral immune system reactions can be operationalized at the society level, inducing groups of humans to intensely discriminate against outgroups. For example, a higher disease burden from communicable illness has been found to be associated with elevated rates of outgroup hostility and avoidance in the US. See: O'Shea BA et al., 'Infectious Disease Prevalence, Not Race Exposure, Predicts Both Implicit and Explicit Racial Prejudice Across the United States' [2020] 11 Social Psychological and Personality Science 345.

Figure 2. Changes in the composition of Twitter users who mention China. Users grouped by sentiment toward China before the COVID-19 outbreak (October—December 2019) are displayed on the left side of the figure. 70.56% of users had negative attitudes, 19.44% had neutral attitudes, and 10.00% had positive attitudes. The right side shows the user distribution over sentiment groups after the COVID-19 outbreak (March—May 2020). 76.91% of users were in the negative group, 15.20% were in the neutral group, and 7.90% were in the positive group. A total of 18.38% of users moved to a more negative group (e.g. neutral to negative) while 11.49% moved to a more positive group; 2.18% of users moved from the negative group to the positive group, and 3.74% of users moved from the positive group to the negative group.

suggests that the decline in China-related sentiment was not the outcome of random fluctuations but was instead caused by a single external trigger.

The decline reflects the changing composition of user sentiment as more users adopted negative attitudes toward China. We segment all users into three categories, negative (sentiment from – 100 to – 20), neutral (sentiment between – 20 and 20), and positive (sentiment from 20 to 100), based on their three-month average attitudes toward China before the outbreak in January 2020. Following the outbreak, 18% of users changed their attitudes toward China to that of a more negative sentiment category while only 11% moved to a more positive category. Users in the negative category increased from 71% before the outbreak to 77% after the outbreak while users in the positive category declined from 10% to 8%. The number of users who moved from the neutral category to the negative category is almost five times the number of users who moved from the neutral to the positive category (Figure 2).

The sharp decline in sentiment on China was accompanied by a rapid uptick in the number of tweets about China. For the three years before January 2020, between 2,000 and 13,000 Twitter users located in the US tweeted or retweeted tweets about China per day. As rumors about what would later be called COVID-19 began to spread in early January 2020, the number of users who mentioned China in tweets began to increase dramatically. Within four weeks, the daily volume of those users reached a peak of 46,792, over six times the usual level. After a brief respite toward the end of February 2020, this volume then reached a second peak with a total of 66,535 users mentioning China on a single day in mid-March, when COVID-19 hit American shores in earnest (dark blue shading in Figure 1). By contrast, the number of users discussing non-COVID-19-related China topics (light blue shading in Figure 1) stayed almost constant and even decreased in the initial stage of the pandemic. The imbalanced composition of users tweeting about both COVID-19 and China and those tweeting about China but not COVID-19 further implicates COVID-19 in both the increasing volume of tweets and the decline in sentiment toward China.

Using Regression Discontinuity to Infer the Impact of Tweeting About COVID-19 on Anti-Chinese Sentiment

The coincidence of the outbreak of COVID-19 and the rise of anti-China sentiment on social media is remarkable. The exogeneity of COVID-19 allows us to employ causal inference methods to investigate whether COVID-19 is a direct cause of the increase in anti-China attitudes. We examine two hypotheses: a causal hypothesis that attributes the rise in anti-China attitudes to the outbreak of COVID-19 and a non-causal hypothesis that posits a mere association between the two. We use regression discontinuity and DID estimation to adjudicate between these hypotheses. For every Twitter user in our data, each of the two strategies analyzes whether COVID-19 changed their attitude toward China. We find that awareness of COVID-19 causes a rise in anti-China sentiment.

Capitalizing on the timestamps in the Twitter data, we use a regression discontinuity design to model the immediate change in the sentiment of an individual's tweets on China after they were 'exposed' to information about COVID-19.[85] If there is no causal effect of exposure to information about COVID-19 on expressed sentiment toward China, we would assume that an individual's sentiment toward China should be stable within a small window before and after their exposure to news about COVID-19. If, on the other hand, there is a causal effect of knowledge about COVID-19 on attitudes toward China, we would expect to see a discontinuity in an individual's sentiment on China before and after exposure. This assumption is reasonable given that both anti-China sentiment and the volume of tweets on COVID-19 appear to rise gradually when viewed at the week level but spike dramatically when viewed at the year level. This means that a discontinuity in a given user's sentiment toward China would stand out as a strong signal in the noise of gradual weekly change.

Information about COVID-19 spread to most Twitter users within the initial six weeks of the outbreak from mid-January to late February, and sentiment on China declined steadily over the same time frame. The outbreak eventually garnered the attention of 90% of Twitter users by late February (Figure S4). Though there is a patch of abrupt change in late January, the curves of decline in sentiment and the cumulative percentage of treated users are both relatively smooth.

We identify the specific day that a Twitter user first posts a COVID-19-related tweet and compare their sentiment toward China before and after the day of exposure. This day is taken to be the treatment day, and we use it as a proxy for the day the user is first exposed to information about COVID-19. In a typical regression discontinuity design, a researcher would model the individual's sentiment as a temporal continuous function in a neighborhood of the treatment day. This approach, however, necessitates parametric assumptions that require data far away from the treatment threshold. We instead use a conservative non-parametric approach that compares a user's average sentiment toward China one week before the treatment day to that user's sentiment toward China one week after the treatment day. It measures the treatment effect ΔRD as:

$$\Delta_{RD} = E_d[y_d(t = d) - y_d(t = d - 1)]$$

where $y_d(t)$denotes the mean attitude in week t averaged over all individuals who are treated in week d and $E_d[\cdot]$ denotes the average over all weeks d. At the individual level, the treatment effect is the difference between sentiment pre- and post-treatment, i.e.$y_d(d) - y_d(d - 1)$.

[85]In the scope of this paper, we use the terms 'treat', 'treatment', 'control", and 'expose' in the context of causal inference. Although this paper is on COVID-19, we stress that these terms are not intended to refer to medicine in any way and that our paper has nothing to do with the medical reality of COVID-19.

There are two potential pitfalls to our approach. The first is that, in theory, there are many discrete pathways that could plausibly have led to the souring of American opinion on China as measured by Twitter data around March 2020. These include but are not limited to through Trump's widely-televised and controversial speeches, from press conferences hosted by other politicians,[86] from workplaces encouraging employees to alter work schedules and/or use respiratory protection, or from personal connections. Because all of these paths, however, are downstream of COVID-10, we make the simplifying assumption that, regardless of how news of COVID-19 reached American ears, knowledge of COVID-19 was the ultimate driver of the unprecedented shift in the opinion of Twitter users on China in March 2024. These multiple paths serve as post-treatment intervening variables and are, in other words, mechanisms through which COVID-19 affected opinion. While we make this simplifying assumption for substantive reasons, it has the additional and fortunate methodological consequence of ensuring that the stable unit treatment value assumption, or SUTVA, of the Neyman-Rubin causal inference model is not violated.[87] Trump's comments on China and COVID-19, while inflammatory, do not merit special treatment in our model because, in addition to the methodological reason fact that Trump's comments on China and COVID-19 are downstream of COVID-19, experimental work in political psychology suggests that exposure to repeated messaging from untrusted politicians does not sway voters from opposing parties.[88] As a result, it is very possible

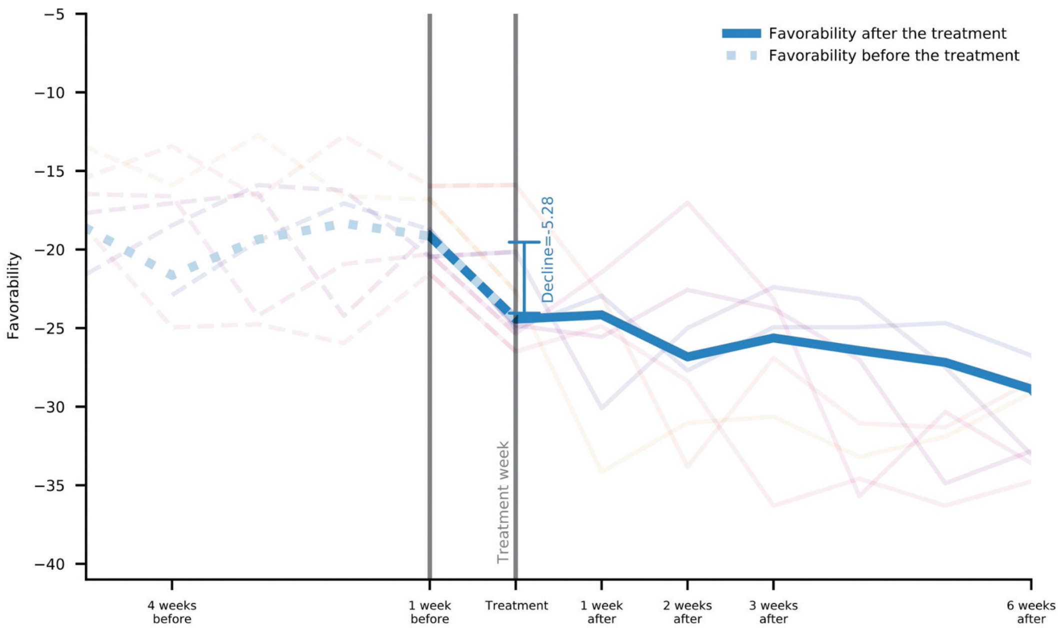

Figure 3. Identifying causality with regression discontinuity. Notes: Plotting the average sentiment of China-related tweets before and after the date on which a Twitter user first posts a tweet mentioning COVID-19 shows that China-related sentiment declines sharply after the users' first post about COVID-19. Individuals are segmented by treatment week, denoted by vertical gray lines.

[86]The Trump administration's messaging on China became much more hardline after COVID-19. For more information on this, see chapters 6 and 7 in Sutter RG, *US-China Relations: Perilous Past, Uncertain Future* (Fourth edition, Rowman & Littlefield 2022).

[87]Rubin DB, 'Randomization Analysis of Experimental Data: The Fisher Randomization Test Comment' [1980] 75 Journal of the American Statistical Association 591.

[88]Orchinik R, Bhui R and Rand DG, 'Repetition Does Not Increase Belief in Claims From Distrusted Politicians' (26 April 2024) https://osf.io/y7pn5 accessed 22 July 2024. For additional evidence showing how Trump's tweets calling for civil disobedience against COVID-19 mandates affected the behavior of White Americans, who generally support Trump, and not Black Americans, who do not, see: Dickson ZP and Hobolt SB, 'Elite Cues and Noncompliance' [2024] American Political Science Review 1.

that Trump did not turn pro-China twitter users to anti-China hawks as much as he stoked extant anti-China sentiment in Twitter users who were already anti-China.[89] We indirectly explore other explanations for this in the Supplementary Information section. Secondly, we recognize that using the date a Twitter user first tweets about COVID-19 as a proxy for the date at which a Twitter user is first exposed to information about COVID-19 is an imperfect proxy accurate. It is, however, the most accurate proxy that we could arrive at given the structure of our data. While it would be possible to integrate survey experiments with Twitter data by polling Twitter users for additional information outside the scope of our data, it is also unclear if a survey experiment would add clarity to our work because survey respondents would most likely not be able to accurately recall when or how they were first informed of COVID-19.

Though using the date of a user's first tweet about COVID-19 as a proxy for exposure to knowledge of COVID-19 is not perfect, both this assumption and the causal hypothesis outlined above are both possible to verify. If the non-causal hypothesis is true, the rise in anti-China sentiment would happen regardless of the presence or absence of information about COVID-19, and the treatment day would not provide any significant information for inferring the sentiment of any given Twitter user on China. If this were the case, we would predict that a user's attitude toward China would be statistically stable on the days before and after treatment. If the causal hypothesis holds, we would expect to observe a marked discontinuity, specifically a sharp dip, in an individual's attitude toward China after treatment.

Our empirical results support the causal hypothesis. Figure 3 shows that, on average, individuals maintain relatively stable and neutral attitudes toward China before treatment but that their attitudes toward China dip abruptly by – 5.28 within one week of treatment. Their attitudes then continue to slowly decline for another six weeks before reaching a saturation point. The step-function-like curve reveals a sharp discontinuity around the treatment day, which supports the causal hypothesis.

We then group individual users by treatment week and report the change in sentiment for each group (vertical gray lines in Figure 3, full version in Figure S2). The sentiment trends for all groups are broadly similar and consist of a stable level before treatment, a sudden and sharp decline immediately after treatment, and a slow and continued decline to a saturation point. The consistency of this pattern across the groups that were exposed at different time points strongly suggests that timing does not impact the treatment effect. Regardless of when an individual is exposed to information about COVID-19, they experience a similar dip in China-related sentiment and eventually settle at a relatively negative opinion toward China. Because individual-level attitudes toward China are stable before treatment, our empirical results validate the assumption that the day a user first posts about COVID-19 is a robust proxy for treatment.

Using Difference-In-Difference to Infer the Impact of Tweeting About COVID-19 on Anti-Chinese Sentiment

The regression discontinuity design we used earlier to identify the causal effect of COVID-19 on attitudes toward China is basically, in the language of Campbell and Stanley,[90] a 'one-group pretest-posttest design'. As such, it lacks a comparison group and thus could suffer from a bias attributable to 'history' – concomitant events that overlap with treatment. To guard against such biases, we supplement the above regression discontinuity analyses with a DID strategy.

The DID strategy compares the change in attitudes toward China over time for the treated group and the change in attitudes over the same time for the control (i.e. untreated) group. Because COVID-

[89]Other work has found that tweets with the hashtag '#chinavirus', a phrase infamously associated with Trump, contained more anti-Chinese sentiment than tweets with '#covid19', but this does not control for the prior tweeting histories of users. See Hswen et al. (n 55).

[90]Campbell DT and Stanley JC, *Experimental and Quasi-Experimental Designs for Research* (Wadsworth 2011).

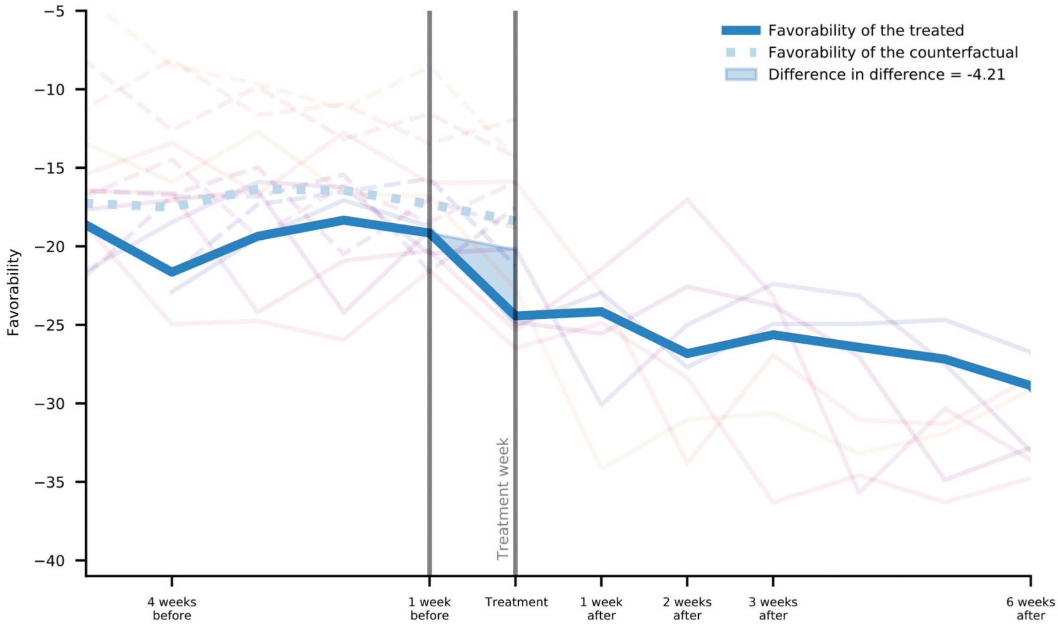

Figure 4. Identifying causality with difference-in-difference estimation. Notes: The sentiment of Twitter users toward China declines suddenly in the treatment week, i.e. the week after they post their first tweets about COVID-19 (solid lines). In contrast, the counterfactual group, who had not yet been treated by a given treatment week, maintain stable sentiment toward China (dotted lines). The treatment effect is estimated by first taking the difference in sentiment on China before and after the treatment within groups and then taking the difference of these within-group differences between groups. This suggests that tweeting about COVID-19 causes an immediate decline of –4.21 in sentiment toward China.

19 was a very hot topic on social media in early 2020, almost 90% of Twitter users had mentioned COVID-19 at least once by March 2020 (Figure S4). If we were to define the control group as users who had not ever mentioned COVID-19, we would be left with a control group so vanishingly small that it would be difficult to use any techniques of statistical inference to compare it to the treated group. We therefore create a DID comparison with moving windows and treat those who had not yet mentioned COVID by the end of each window as eligible controls. For every week in January and February 2020, we create a treated group with all individuals who have mentioned COVID-19 before a given week and a control group with all individuals who have not mentioned COVID-19 by the same week. The overall treatment effect ΔDID is then measured as the macro-level mean over all windows:

$$\Delta_{\mathrm{DID}} = E_d[(y_d(t=d) - y_d(t=d-1)) - (y_{d'>d}(t=d) - y_{d'>d}(t=d-1))],$$

Where $y_d(t)$ measures the macro-level average attitude in week t among individuals who are treated in week d, and $y_{d'>d}(t)$ represents the average attitude in week t among individuals who are treated in any week d′ after week d. In every moving window d, we compute $y_d(t=d) - y_d(t=d-1)$, the change in attitudes over time for the treated group, and $y_{d'>d}(t=d) - y_{d'>d}(t=d-1)$, the change in attitudes over time for the control group. The DID estimator Δ_{DID} should tell us whether knowledge of COVID-19 triggers a Twitter user to tweet negatively about China compared to a counterfactual situation in which he or she had no knowledge of COVID-19. If the non-causal hypothesis is true, the treated group and the control group would be statistically indistinguishable, and we would predict an expectation of zero for the DID estimator, i.e. $\Delta_{\mathrm{DID}} = 0$. If the causal hypothesis holds, the drop in the treated group's attitudes toward China would be significant relative to the 'natural' change in the control group's attitudes toward China. That is, we would we predict Δ_{DID} to be non-zero.

Our empirical DID results again support the causal hypothesis. As shown in Figure 4, we see an average decline of $\Delta DID = -4.21$ in attitudes toward China in the week following a Twitter user's treatment. The treated group and control group have very similar attitudes toward China until the treatment week, upon which the treated group's attitudes toward China degrade suddenly while the control group's attitudes toward China remain almost constant. A more detailed analysis (Figure S3) shows that this pattern is consistent across all of the moving windows we construct.

Conclusions and Discussion

In the wake of the outbreak of COVID-19, American attitudes toward China sank to a three-year low while tweet volume on China soared to a three-year high. Based on a corpus of 297 million tweets on COVID-19 and/or China, we identify the causal impact of the COVID-19 outbreak on anti-China sentiment. Using regression discontinuity and DID estimation methods, we discover that awareness of COVID-19, measured by proxy as tweeting about COVID-19, led to a sudden and sharp decline in American attitudes on China. Moreover, we find that the relationship between this awareness and this decline is causal.

We acknowledge that there are limitations to using Twitter as a data source. While this study has treated Twitter as roughly representative of the American population, Twitter itself can artificially exacerbate negativity. For American users of social media platforms, reading opposing views may actually polarize users of the platforms even further.[91] Contact with like-minded counterparts may also breed hostility toward outgroups. This finding has been replicated for face-to-face conversation. Discussions in groups of individuals with similar ideologies can lead to extreme speech.[92] Conversations in small, bounded groups, offline or online, have a tendency to foment polarization. Because many discussions on Twitter take place in 'filter bubbles'[93] that mimic these small groups and because many Twitter users are infrequently but regularly exposed to opposing views, Twitter is an almost-perfect environment for polarization. Further research is needed to understand how inimical sentiments spread on Twitter and to what extent this impacts the broader applicability of our findings. For our purposes, these concerns are at least somewhat attenuated. While the filter bubble phenomenon may have artificially polarized tweets on China and inflated the increase in anti-China sentiment on Twitter, the rise in anti-China sentiment was so large that some inflation does not change our results. Additionally, survey evidence[94] supports the connection between COVID-19 and shifts in sentiment toward China. This preponderance of evidence indicates that we can consider Twitter to be representative of public opinion and allows us to leverage the main benefits of social media data, which are scale and scope. While it is not feasible to survey nor experimentally manipulate the opinions of millions of Americans, it is feasible to mine their tweets, and our work reveals a causal connection between COVID-19 and anti-China sentiment at scale across America that cannot be demonstrated with other sources of data.

The COVID-19 pandemic and the unique data we have collated offer powerful methodological and theoretical insight. Firstly, the suddenness of the COVID-19 pandemic makes it a uniquely exogenous trigger without precedent in the social media era for shifts in attitudes on foreign countries. This creates an opportunity for the assumptions of causal inference to hold in situations where they may have previously been untenable. Adopting COVID-19 as

[91]Bail CA et al., 'Exposure to Opposing Views on Social Media Can Increase Political Polarization' [2018] 115 Proceedings of the National Academy of Sciences 9216.
[92]Gabbay M et al., 'Frame-Induced Group Polarization in Small Discussion Networks' [2018] 81 Social Psychology Quarterly 248.
[93]Bozdag E et al., 'Does Offline Political Segregation Affect the Filter Bubble? An Empirical Analysis of Information Diversity for Dutch and Turkish Twitter Users' [2014] 41 Computers in Human Behavior 405.
[94]He Q, Zhang Z and Xie Y, 'The Impact of COVID-19 on Americans' Attitudes toward China: Does Local Incidence Rate Matter?' [2022] 85 Social Psychology Quarterly 84.

a natural experiment specifically allows us to use causal inference methods on a large corpus of social media data. Our analysis reveals that awareness of COVID-19 leads to a marked rise in anti-China sentiment on Twitter. Secondly, COVID-19 offers a rare situation where foreign affairs become domestic: when the pandemic struck, many Americans perceived the pandemic as a very foreign threat to their very local interests and physical safety. Because Americans generally remain safely isolated behind the natural walls of two oceans, the American heartland has been far removed from the military dramas of other continents. Similarly, though a few rare events, notably the 1973 oil embargo, have impacted the pocketbooks of American consumers, the economy of the US has generally been somewhat insulated from economic shocks originating overseas. The COVID-19 pandemic, however, exposed many Americans to a historic and unprecedented combination of economic and physical hardship that they could blame on a foreign source. We combine insight from literatures on self-interest and policy preference and on how Americans view foreign countries, and we find that Americans were considerably less China-focused prior to the outbreak of the COVID-19 pandemic, but as the livelihoods, lives, and lifestyles of all Americans began to be affected we see a sharp spike in online animus toward China. While we do not conclusively prove the proposed link between self-interest and political behavior in the realm of foreign policy, our findings are suggestive of its existence. We have shown that COVID-19, a pandemic that negatively affected American livelihoods and that many Americans associated with China, directly caused a profound souring of American attitudes toward China. When a nation's circumstances are worsened, and a foreign power can be scapegoated or otherwise maligned for this worsening, it follows that citizens of said nation will view the foreign power more negatively. Our findings may be used for further investigations of how self-interest affects views on foreign places and peoples using big data from social media platforms.

Acknowledgments

The authors thank Susan T. Fiske, Xuechunzi Bai, and Gandalf Nicolas Ferreira at Princeton University for valuable suggestions on the Stereotype Content Model. We thank Phillip Rush, Karla M. Perez-Gazca, Rebecca J. Han, Lucas J. Irwin, Hamna Khurram, Louison M. Sall, Ryan Sung, Chris Venturella, and Sandoval E. Wood at Princeton University for coding tweets.

Disclosure Statement

No potential conflict of interest was reported by the author(s).

ORCID

Gavin G. Cook http://orcid.org/0000-0003-4617-6469
Junming Huang http://orcid.org/0000-0002-2532-4090