How Racial Animus Forms and Spreads: Evidence from the Coronavirus Pandemic

Runjing Lu   Sophie Yanying Sheng*

October 4, 2021

Abstract

This paper studies the formation and the spread of crisis-driven racial animus during the coronavirus pandemic. Exploiting plausibly exogenous variation in the timing of the first COVID-19 diagnosis across US areas, we find that the first local case leads to an immediate increase in local anti-Asian animus, as measured by Google searches and Twitter posts that include a commonly used derogatory racial epithet. The majority of racist tweets come from users who post the epithet for the first time, suggesting that the pandemic affected racial animus at the extensive margin. Racist tweets are predicted by users' interactions with other anti-Asian individuals, highlighting the role of social network in the spread of racial animus. Moreover, 75 percent of the anti-Asian tweets do not explicitly mention the virus, suggesting that the racial animus may persist beyond the duration of the pandemic. Finally, online animosity and offline hate incidents against Asians both increase with the salience of the connection between China and COVID-19; while the increase in racial animus is not associated with the local economic impact of the pandemic.

Keywords: COVID-19, Racial animus, Xenophobia, Social media, Hate crime
JEL Classification Numbers: J15, D74, I18

*Lu: University of Alberta (runjing1@ualberta.ca); Sheng: UC San Diego (yas107@ucsd.edu). An earlier version of this paper circulated under the title "From Fear to Hate: How A Pandemic Sparks Racial Animus in the United States." We are grateful for the support and guidance from Eli Berman, Julie Cullen, and Gordon Dahl. We would also like to thank David Arnold, Julian Betts, Georgy Egorov, Itzik Fadlon, Ruixue Jia, William Mullins, Jacob Shapiro, Jesse Shapiro, and the participants at the UCSD applied lunch seminar for their constructive feedback.

1. Introduction

Racial animus can affect welfare in measurable ways, as economists have noted since the seminal work of Becker (2010). Recent papers have shown that racial animus can hinder economic development, affect political institutions, and induce social unrest.[1] To curb racial animus at the outset and to mitigate its consequences, a crucial first step is to understand how it forms and grows.

In this paper, we shed light on what factors motivate racial animus, which individuals are more susceptible to such factors, and how racial animus spreads, using the coronavirus (COVID-19) pandemic as a natural experiment. The Centers for Disease Control and Prevention (CDC) has emphasized that people of Asian descent are at no greater risk of spreading the virus than other Americans. Nonetheless, since the outbreak of the virus, news reports of hate crimes against Asian Americans have increased (Mullis and Glenn, 2020).[2] The unexpected nature and regional variation of the pandemic provide a valuable opportunity to study the rise and spread of racial animus – in this case, against Asians.

To proxy for an area's racial animus against Asians, we use the percentage of Google searches and Twitter posts (tweets) that include the words "chink" or "chinks" (hereafter, the ch-word).[3] Google searches can capture *private* racial animus given others cannot view one's searches. Past papers have documented a clear relationship between Google searches of racial slurs and racial animus against minorities (Anderson et al., 2020, Depetris-Chauvin, 2015, Stephens-Davidowitz, 2014). Furthermore, as we will show below, an area's monthly Google searches for the epithet is positively correlated with monthly anti-Asian hate crimes and is negatively correlated with monthly visits to Chinese restaurants. Our second proxy is based on tweets, which has been used to measure *public* displays of racial animus (Nguyen et al., 2018). These two proxies are valuable alternatives to more traditional measures, such as offline hate crimes which may only capture the most extreme hatred and may not fully reflect the levels of racial animus due to

[1]For instance, see Charles and Guryan (2008), Card et al. (2008), Stephens-Davidowitz (2014), and Healy and Searcey (2020).

[2]For example, see NBC News, New York Times, and USA Today.

[3]We focus on the use of the ch-word because it is the most salient and unambiguously pejorative racial slur against Asians. According to the Philadelphia Bar Association, the epithet "is now widely used throughout the United States as a racial slur against people of Asian descent" (Association, 2014). Importantly, it has not been reclaimed by the Asian American community (Anderson and Lepore, 2013).

blanket stay-at-home orders during the pandemic. In addition, use of racial slurs online is an important outcome in and of itself, as researchers have shown a strong relationship between exposure to racial discrimination online and depression and anxiety measured offline (Tynes et al., 2008).

To motivate, we exploit the timeline of COVID-19 developments in the United States to understand the general evolution of anti-Asian animus during the pandemic. We find little increase in the national racial animus upon the first US COVID-19 case and only a small uptick in the week when WHO declared COVID-19 a pandemic. In contrast, we observe a clear jump in the week when President Trump tweeted "Chinese virus."

In order to causally identify the effects of COVID-19 on racial animus against Asians, we use a difference-in-differences (DID) event study design exploiting the variation in the timing of the *first local* COVID-19 diagnosis across areas. Specifically, we compare the change in racial animus following the first diagnosis in an area to the change in other areas during the same period. First local diagnoses are likely to increase the salience of the virus, and the salience of diseases has been shown to induce xenophobia in lab experiments (Faulkner et al., 2004). The identifying assumption is that the *precise* timing of the first diagnosis in an area is plausibly exogenous; whether an area has its first diagnosis this week (day) or the next is largely unpredictable and unlikely to correlate with other factors that simultaneously change local racial animus.[4]

Our DID event study reveals that, in the week after the first local COVID-19 diagnosis, an area's Google search rate of the ch-word increases by 22.6 percent of the area's maximum search rate during the sample period, and an area's Twitter post rate of the epithet increases by 118.6 percent of the average post rate across all areas during the sample period. These effects persist for six weeks after the first local case. Given the correlation based on historical data, the increase in Google search rate of the ch-word would be associated with a 6.5 percent increase in anti-Asian hate crimes.[5] The results, where applicable, are quantitatively unchanged under a *dynamic* event study design which allows for varying treatment effects across event periods (Sun

[4]Papers like Egorov et al. (2020) have noted that areas with larger population sizes or better medical systems tend to have first diagnoses earlier. We include area fixed effects to control for these time-invariant characteristics.

[5]Note, given the wide presence of stay-at-home orders during this time of the pandemic, we would not expect to find an effect on offline hate crimes.

and Abraham, 2020). Our results are also robust to using alternative racial animus measures based on tweets which include other anti-Asian slurs and are not counter-hate; to excluding early- and hard-hit states; and to controlling for severity of local infection, existence of stay-at-home orders, general local attention to Asians, and area and year-month fixed effects.

When we examine the content of ch-word tweets, we find that the share showing emotions of anger and disgust increases from 23.3 to 40.8 percent after the first local diagnosis. This shift in sentiment suggests that the increase in these racially charged tweets represents a real change in attitude towards Asians.

Moreover, we show that the increase in racial animus is directed *only* at Asians and not at other minority groups. The singling out of Asians implies that the increase is likely not due to an overall rise of ethnic distrust or tensions from general uncertainty about cross-group differences in health status or risk-taking behavior. Rather, it is targeted at a specific group associated with the geographical origin of the virus. In addition, 75 percent of ch-word tweets posted following the first local case do not explicitly mention COVID-19, implying that the pandemic-induced racial animus towards Asians extends to broader topics and may persist beyond the duration of the pandemic.

We also leverage the rich information in historical tweets and Twitter user network to study *which individuals* are more likely to start expressing hate because of the pandemic. We find that the surge in ch-word tweets is driven primarily by the extensive margin (i.e., existing Twitter users who post the term for the first time) rather than the intensive margin (i.e., in increase in tweets from users who have previously used the term). The pandemic-induced first-time ch-word users have frequently interacted with other anti-Asian users in the past, and roughly two out of ten of them had tweeted racial slurs against other minorities before the pandemic. This finding implies that the pandemic may have redirected these users' anti-minority sentiments towards Asians.

Finally, we turn our attention to the factors fueling the spread of racial animus among individuals. Exposure to anti-Asian users is one such factor. We find that interacting with anti-Asian users in a day predicts a higher likelihood of tweeting the ch-word the next day. The salience of the connection between COVID-19 and the Asian population is another major factor. We proxy for this salience by using the number of President Trump's tweets that mention China and COVID-19 simultaneously. We find that one additional such China-and-COVID tweet in a day corresponds to an eight percent increase in anti-Asian hate incidents and an increase in national ch-word tweets on the same day, equivalent to 14 percent of the daily average. An event study using *hourly* tweet data also reveals an immediate increase in ch-word tweets following the president's China-and-COVID tweets but not before. In contrast, we find little evidence that negative economic impacts from the pandemic motivates the initial rise of racial animus. Areas with a more severe economic damage from the pandemic do not exhibit a higher increase in racial animus than areas with a less severe impact.

This paper contributes to the literature studying the causes of animus toward minorities. This body of work has shown that negative shocks such as terrorist attacks and deterioration of economic conditions induce animus against racial or religious minorities. For example, Kaushal et al. (2007), Hanes and Machin (2014), and Ivandic et al. (2019) document that 9/11 and jihadi terror attacks lead to increases in anti-Muslim hate crimes. Anderson et al. (2017) and Anderson et al. (2020) find that the Great Recession and negative shocks to agricultural income in historical Europe contribute to animus against minorities. In addition, desire to avoid health threats has also been postulated to motivate racial bias (Schaller and Neuberg, 2012). Lab experiments have shown that exposing subjects to disease-related primes leads to increased xenophobia (Faulkner et al., 2004, O'Shea et al., 2020, Bartos et al., 2020). However, the causal evidence on whether infectious diseases lead to racial animus in the field is still lacking.[6] An exception is Jedwab et al. (2019), which documents that the black death caused an increase in anti-Jewish pogroms in medieval Europe.

Our contribution is to provide causal evidence on how negative shocks, such as pandemics, trigger racial animus and shed light on who are more susceptible to such shocks and how racial animus spreads. Our findings have implications for mitigating animus amid future crises. We find that the rise in racial animus is specific to Asians who are associated with the geographical origin of the virus and that the salience of this association amplifies animus against the group. Therefore, careful naming of a disease (e.g., COVID-19 and Delta variant as opposed to Chinese virus and Indian variant) and debunking claims of a purported connection between a disease and

[6] More recent papers on the prevalence of hate during the COVID-19 pandemic are mostly descriptive (e.g., Schild et al. (2020), Ziems et al. (2020), Lyu et al. (2020), Croucher et al. (2020)).

a group could be helpful in curbing animus. Additionally, our findings reveal that the extensive margin and social network play an important role in spreading racial animus, suggesting that moderating racist individuals and their interaction with others on social media could help constrain racial animus in the future.

Finally, our paper speaks to the literature on political rhetoric. Political rhetoric has been shown to influence public opinion and behavior, such as presidential approval (Druckman and Holmes, 2004), public perception of a foreign country (Silver, 2016), and anti-minority hate crimes (Müller and Schwarz, 2019). We add to this literature by providing another example of how the rhetoric of political figures regarding a public crisis influences racial animus at the national level. On the flip side, harnessing these public figures' opinion-shaping power could be useful in curbing animus.

2. Measures of Racial Animus

2.1 Google and Twitter Proxies

We use two measures to proxy for an area's racial animus against Asians: the percentage of Google searches and the percentage of tweets that include the words "chink" or "chinks." The ch-word is not uncommon in Google searches or tweets. Between June 2019 and June 2020, this racial epithet was included in more than a quarter million searches and 60,000 tweets.[7] Google searches and tweets that include the epithet are mostly negative. For instance, "chink eye" and "chink virus" are common terms in such Google queries and Twitter posts. People may search the epithet to look for jokes or memes about Asians or to look for like-minded others with whom they can share anti-Asian sentiments.

We use Google Trends to obtain weekly Google search data for the ch-word at the media market level between July 2019 and April 2020. The data are not the raw number of searches but the weekly percentage of searches that include the term (*search rate*), taken from a random sample of total searches representative at the media market and time levels and scaled by the highest weekly search rate in the same market during the entire extraction period – in our case,

[7] The number of Google searches is an approximation from https://searchvolume.io/. The data are only available for the 12-month period before our query on June 8, 2020.

between July 2019 and April 2020. In particular, the racially charged Google search index for media market m at time t extracted over period T is

$$\text{Search Index}_{mt,T} = 100 \times \frac{\frac{\text{Searches including "chink(s)"}_{mt}}{\text{Total searches}_{mt}}}{\max\limits_{t \in T}\{\frac{\text{Searches including "chink(s)"}_{mt}}{\text{Total searches}_{mt}}\}} \tag{1}$$

Note that Google returns a zero value when the racially charged search index for a given area and time falls below an unreported threshold. To alleviate concerns of selecting our sample based on outcomes, we focus on media markets that have valid racially charged Google search index in the baseline period (2014-2018). This leaves us with 60 of 210 media markets, covering approximately 74 percent of the US population and 78 percent of the US GDP in 2019 across 33 states. Compared to other media markets, the ones in our sample tend to have a larger population, higher percentage of Asians, slightly lower baseline anti-Asian hate crime rate, and more enplanements of international airports (Table A1 column (1)). Shaded areas in Figure A1 panel A indicate the media markets in our sample.

The above metric can capture the timing but not the level of a change in an area's search index. As an alternative, we rescale the search index so that the search rate in different media markets is normalized by *one* base search rate. We try three different bases: Huntsville-Decatur (Florence)'s search rate on March 15, 2020; Wilkes Barre-Scranton on March 29, 2020; and Buffalo on April 5, 2020. We choose these bases to obtain rescaled indexes for as many media markets as possible, i.e., 35, 29, and 29, respectively. As detailed in Appendix 1, rescaling drops many media markets whose search rate is zero on the date when the base search rate occurs (*benchmark date*). For this reason, we only use the rescaled version as a robustness check.

We obtain Twitter data from Crimson Hexagon, which houses all public tweets through a direct partnership with Twitter. We downloaded all geo-located tweets that include the ch-word between November 1, 2019, and May 2, 2020. Crimson Hexagon does not provide the total number of tweets posted in a given area and time. We thus extract the number of all public tweets that include the word "the," the most common word on Twitter, in a given area and time as a substitute. Assuming that the proportion of tweets that include "the" is stable across areas, the number of tweets that include "the" can approximate Twitter activity. We define the racially charged Twitter post index for a given area and time as the number of tweets including the ch-word per 100,000 tweets including the word "the."

We calculate the Twitter post index for 658 counties across 50 states and Washington D.C., encompassing 60 percent of the US population and nearly 70 percent of the US GDP in 2019. Counties are included if their residents ever posted "the" tweets between 2014 and 2018.[8] Counties with Twitter data tend to have a larger population, higher support for the Democratic Party, and higher enplanements of international airports, but show no difference in baseline anti-Asian hate crime rate compared to other counties (Table A1 column (2)). Shaded areas in Figure A1 panel B are counties with Twitter data. Analyses using Twitter data are conducted at the county level unless noted otherwise.

The fact that Google and Twitter data do not cover the full of the US should not affect internal validity of our study, but it could pose threat to external validity. Therefore, we use *both* data sources, which could alleviate concerns about the external validity of our findings.

2.2 Relationship between Racial Animus, Hate Crimes, and Restaurant Visits

For the racially charged Google search index and Twitter post index to be meaningful proxies for racial animus, the only assumption we need is that an increase in racial animus makes a person more likely to use the ch-word. Under this assumption, higher racial animus results in a higher percentage of Google searches and tweets that include the racial epithet. Existing papers that use a similar proxy for racial animus suggest that the assumption is likely to hold (Anderson et al., 2020, Depetris-Chauvin, 2015, Stephens-Davidowitz, 2014).

To better understand the above proxies, we check how they predict anti-Asian hate crimes and visits to Chinese restaurants. Hate crime data come from the FBI Uniform Crime Reports (UCR) and are available up to 2018. A majority of these hate crimes are simple or aggravated assault (30 percent) and in-person intimidation (34 percent). Table 1, panel A, columns (1) through (4) report the media market-level correlation between the monthly racially charged search index and the monthly number of anti-Asian hate crimes between January 2014 and December 2018, controlling for local population size, unemployment rate, year-month fixed

[8]About half of the tweets in the sample lack geo-identifiers and hence cannot be associated with a certain county.

effects, and media market fixed effects. On average, a one-standard-deviation increase in the racially charged search index corresponds to an increase in the anti-Asian hate crimes in the same month, amounting to 8.9 percent of the monthly average.[9] The correlation is robust to controlling for the search index for "Asian(s)," which is related to the ch-word but neutral in connotation, as shown in column (2). In columns (3) and (4), we include both the index in the current month and the index in the prior month. The relationship between the racially charged search index and hate crimes is mainly contemporaneous.

Next, we change the dependent variable to monthly visits to Chinese restaurants in each media market between January 2018 and December 2019 while additionally controlling for the monthly visits to all local restaurants. The visit data are from Safegraph and are available starting in 2018.[10] Table 1, panel A, columns (5) and (6) show that a one-standard-deviation increase in the racially charged search index is linked to 484 fewer monthly visits to Chinese restaurants, equaling 0.5 percent of the monthly average. The relationship between the search index and visit rate is also contemporaneous.

Finally, we replicate the above correlations using Twitter data in Table 1, panel B. We aggregate hate crimes to the media market level due to their low occurrences at the county level. To maintain consistency, we also aggregate restaurant visits to the media market level. Overall, the racially charged Twitter post index does not correlate with anti-Asian hate crimes or visits to Chinese restaurants. One potential explanation is that Twitter data represent public displays of racial animus and undergo more social censoring. We may only see a change on Twitter when the shift in racial animus is substantially large.

3. Evolution of Racial Animus in US amid the Pandemic

To motivate our causal analysis, we study the general evolution of anti-Asian animus as the pandemic develops. An ideal experiment would be to contrast rates of racially charged Twitter posts and Google searches in the U.S. during the pandemic to counterfactual rates absent the pandemic. However, a perfect counterfactual does not exist because all individuals and areas

[9]The percent increase is calculated by multiplying the standard deviation of the index (23.07) with the coefficient and dividing the product with the outcome mean.

[10]Safegraph provides data on foot traffic to roughly 4.1 million points of interest in the United States.

were more or less impacted by the pandemic. For this reason, we use racially charged Twitter posts and Google searches in 2019 as controls. The assumption is that racially charged Twitter posts and Google searches in 2020 would have been the same as in 2019 absent the pandemic.[11]

We first compare an individual's weekly likelihood of tweeting the ch-word during the first 16 full weeks in 2020 and the *same* person's likelihood of doing so in the corresponding weeks in 2019. An advantage of this analysis is that it does not require geo-identifiers, so we can include all 26,065 Twitter users who ever tweeted the ch-word during our baseline period (2014-2018).[12] We use the following specification:

$$\mathrm{Y}_{iyw} = \sum_{w=2}^{16} \beta_w \times 1\{y = 2020\} + \alpha_i + \alpha_w + \epsilon_{iyw} \tag{2}$$

where Y_{iyw} is a binary variable which equals one if individual i tweets the ch-word in week w of year y. We use $w = 1$, the first full week of a year, as the comparison period. Our treatment variable is $1\{y = 2020\}$, which equals one if the year is 2020, and 0 if the year is 2019. We include person fixed effects α_i and week-of-year fixed effects α_w to absorb individuals' baseline propensity to tweet the racial epithet and the seasonality in such tweets. We cluster standard errors by individual.

The individual-level analysis reveals that the likelihood of tweeting the ch-word co-moves with important developments of COVID-19. In Figure A2 panel A, we plot β_w from equation 2. While we find little to no increase in the likelihood of tweeting the term following the first US COVID-19 case or declarations of health emergency and only a small uptick in the week when WHO declared COVID-19 a pandemic, we observe a clear jump in the week when President Trump first tweeted "Chinese virus."

Lacking individual-level search data, we compare media market-level weekly racially charged Google search index in 2020 and the index in the corresponding markets and weeks in 2019. Specification is the same as equation 2 except that Y_{iyw} is now the Google search index in media market i in week w of year y. We plot β_w in Figure A2 panel B. While we also see a spike in

[11]This assumption could be violated if there are other contemporaneous shocks affecting racial animus. Our strategy in the next section avoids this issue.

[12]We make this sample selection because we do not have the universe of tweets. Crimson Hexagon only allows tweet extraction based on keywords.

Google search index in the week when President Trump tweeted "Chinese virus," we cannot draw definitive conclusions for other weeks.

4. Evidence from DID Event Study

We now turn to our main empirical strategy, exploiting the precise timing of the first COVID-19 diagnoses across the United States. We compare the changes in racially charged Google search index (Twitter post index) in the weeks before and after the first local case to the changes in other media markets (counties) during the same period. This design allows us to avoid concerns about contemporaneous shocks that influence racial animus at the same time as the pandemic develops.

4.1 Timing of the First Local Case and Empirical Strategy

We download the data on US COVID-19 cases and deaths between January 21 and May 2, 2020, from the Johns Hopkins University Coronavirus Resource Center. We match the date of the first case in each media market and county to those with valid Google and Twitter data. Table A4 displays the number of media markets and counties by the timing of their first local diagnoses. All media markets have their first diagnoses in the sample period and have Google data for at least six weeks after the diagnosis. These media markets make up the Google sample. Seventeen counties with Twitter data are excluded because they did not have diagnoses in the sample period.[13] The remaining 641 counties make up the Twitter sample and have data for at least one week after the first local diagnosis; the number of counties decreases to 636, 629, 613, 555, 416 in weeks two to six.[14] Therefore, the Google (Twitter) sample is a panel of media markets (counties) from six weeks before to six weeks after the first local diagnosis. Table A2 reports summary statistics for each of the samples.

To understand predictors of the diagnosis timing, we regress the week of first local diagnosis on a battery of local characteristics in Table A3. The analysis reveals that larger population size predicts earlier diagnoses for both the Google and the Twitter samples, while enplanements of international airports predict slightly later diagnoses for the Google sample. However, propor-

[13] Results are unaffected when we include these counties.

[14] Crimson Hexagon discontinued in July, 2020, so we cannot extend Twitter data.

tion of Asians does not have predictive power for the timing, consistent with CDC's statement that Asians are at no greater risk of spreading the virus. More importantly, pre-pandemic anti-Asian hate crime rate does not predict the timing either, suggesting that the treatment timing is orthogonal to baseline racial animus.

We then estimate the following regression:

$$\mathrm{Y}_{it} = \sum_{t=-6, t\neq -1}^{6} \beta_t + \gamma' X_{it} + \alpha_i + \alpha_{ym(t)} + \epsilon_{it} \tag{3}$$

where Y_{it} is the racially charged Google search index (Twitter post index) in media market (county) i in event time t, which is the number of time periods relative to the first local diagnosis. β_t represents event dummies for six weeks before to six weeks after the first local diagnosis, excluding our comparison period $t = -1$. X_{it} is a vector of area-specific characteristics such as the local number of COVID-19 diagnoses or deaths, an indicator for a state-level stay-at-home order, and the Google search index or Twitter post index for "Asian(s)". We include county or media market fixed effects α_i and year-month fixed effects $\alpha_{ym(t)}$ to control for an area's baseline racial animus and national trends in racial animus.[15] We cluster standard errors by media market for Google data and by county for Twitter data. We also estimate equation 3 at the daily level, where we include event dummies from 14 days before to 21 days after the first local diagnosis while omitting the dummy for day -4 and additionally control for day-of-week fixed effects.

If the trends of racially charged Google search index (or Twitter post index) across media markets (or counties) are parallel in the absence of local COVID-19 cases, and treatment effect of the first local case does not vary across event times, β_t identifies the weighted average treatment effect across treatment areas in time t on local searches or posts of the ch-word. Testing for parallel pre-trends can shed light on the first identifying assumption. As we will show, this assumption appears to hold. The second assumption is harder to test, and its violation could bias the estimates in unknown directions. For example, if earlier treated areas experience an

[15] Although the Google search index is a normalized search rate so that the maximum search rate in a media market is equivalent to a search index of 100, there is still considerable variation in the sample mean of this normalized index varying between eight and 50. For example, 50 means that the average search rate in a media market is half of its maximum search rate during the extraction period.

increasing (or decreasing) treatment effect over time due to evolving local pandemic situation, using these areas as controls for later treated places could bias the average treatment effect downward (upward).

To alleviate concerns about varying treatment effects, we use a *dynamic* DID event study comparing areas with a first case before and after the case, using areas that *have not had* any cases as controls. To implement the dynamic event study, we follow Novgorodsky and Setzler (2019) and stack our data as a series of 2×2 matrices (treated/not-yet-treated × pre/post). We define areas which have their first cases in calendar week g as cohort g, and cohort-specific event time in calendar month m as $e_g = m - g$. The treatment effect on cohort g in event time e_g is labeled as β_{e_g}. Following Sun and Abraham (2020), we define the average treatment effect for event time e among all cohorts G as:

$$\beta_e = \sum_{g \in G} \beta_{e_g} \times w_g \tag{4}$$

where w_g (the aggregation weight) is the population in areas belonging to cohort g. We calculate clustered standard errors at the area level for β_e via the delta method.

A limitation of the dynamic event study is that it requires enough not-yet treated areas in event time e_g to estimate β_{e_g}. Since counties are smaller than media markets, the former have more variation in the timing of first local diagnoses. Most counties in the Twitter sample have control counties for multiple post periods while most media markets in the Google sample do not have controls after event 0. Therefore, we only apply the dynamic event study to Twitter data and use this approach as a robustness check.

4.2 Effects of the First Local Case on Racial Animus

4.2.1 Main Findings

We start by examining how an area's Google searches for the ch-word respond to the first COVID-19 case in the local area. Figure 1, panel A plots β_w from equation 3 using an area's racially charged Google search index as the outcome. The index jumps markedly in the week after the first local case and persists at high levels in the following weeks. The pre-trends are flat and statistically insignificant, suggesting that the parallel trend assumption is likely to hold. Regression results corresponding to this figure are found in Table 2, column (1). For example, consider the $+1w$ coefficient: compared to the week before the first local case, an area's racially charged search rate increases by 22.6 percent of the area's maximum search rate over the sample period in the first week after the diagnosis. The treatment effects remain mostly above 17 percent for the following five weeks. Given our findings of the correlation between Google searches and hate crimes (see Table 1), the increase in the search index in the month after the first local diagnosis translates to an increase of 0.0095 anti-Asian hate crimes or 6.5 percent of the average monthly anti-Asian hate crimes between 2014 and 2018.[16]

Figure A3 shows the event time plot when we replace the original search index with the indexes rescaled using three different bases. The patterns are qualitatively similar. Note that the magnitude of the estimates are roughly half of those from the main specification. This is explained by the fact that base search rates for the rescaled indexes are higher than search rates in most media markets. Note also that the standard errors are much larger than those from the main specification. This is explained by the fact that rescalling forces us to drop nearly half of the media markets (see Appendix 1 for detail). We present the results using the original search index in the rest of the paper.

We next turn to Twitter to understand how the first local case affects *public* use of the ch-word. In Figure 1, panel B, we plot the effect of the first local case on the racially charged Twitter post index. Similar to the Google search index, the Twitter post index also jumps in the week after the first case. Specifically, relative to the week before the case, racially charged tweets increase by 0.7 per 100,000 "the" tweets in the week after, amounting to 118.6 percent of the weekly average during the sample period. The effects remain high in weeks 2-6. Table 3, column (1) reports the regression results.

To confirm that our results are not driven by the functional form of the index or the specific racial epithet we choose, we use alternative functional forms and other ways of identifying anti-Asian content. Raw number of ch-word tweets and number of ch-word tweets per million population reveal similar patterns as the original Twitter post index, as shown in Table 3 columns (2) and (3). Additionally, we construct a new index using COVID-related tweets

[16] The level increase in anti-Asian hate crimes is calculated by multiplying the average of coefficients between weeks one and four with the number of monthly anti-Asian hate crimes per unit increase in the Google search index.

posted between January 15 and April 17, 2020 that are classified as anti-Asian via machine learning.[17] Column (4) shows that the effects estimated with this new index share a similar pattern to the ones in column (1) but are *seven* times as large. Our original Twitter post index is thus likely a conservative measure of racial animus.

An evolving local pandemic situation may produce time-varying treatment effects, which could bias results of a regular DID event study. To alleviate this concern, in column (5), we re-estimate the effect using a dynamic DID event study, following Sun and Abraham (2020). This design compares counties with a first COVID-19 case before and after the case, using counties that have not had any cases as controls. The estimates are quantitatively similar to those in column (1), implying that time-varying treatment effect is likely not an issue.

Figure A4 presents results using *daily* search and post indexes. The indexes increase starting on the second or the third day after the first local case, suggesting that it may take time for the information about new cases to disseminate and for local residents to react.

4.2.2 Discussion and Robustness Check

In this subsection we discuss potential alternative explanations for the rise in the ch-word use in an area after the first local COVID-19 case and explore the robustness of our main findings.

Increased ch-word usage may result from a general rise in online activities due to blanket stay-at-home orders rather than a change in racial animus. However, our search index and post index already account for an overall change in online activities because they are normalized by the total searches and tweets in a given area and time. In addition, when we include an indicator for state-level stay-at-home order in Table 2 column (2) and Table 3 column (6), results are quantitatively similar to those from our main specification.

Alternatively, an increase in general attention to Chinese or Asians may lead to higher ch-word usage. In Table 2 column (3) and Table 3 column (7), we control for searches or tweets of terms that capture such general attention but are neutral in connotation, i.e., "Asian(s)." Results are unaffected.

[17]We thank Ziems et al. (2020) for providing the data. These anti-Asian tweets include phrases like "Chinese Virus," "Wuhan Virus," and "Kung Flu" and exclude counter-hate tweets that may have racist keywords in them. We include the counties that had their first diagnoses between February 16 and March 22, 2020.

Our results are also robust to excluding early- and hard-hit states like New York, Washington, and California, as shown in Table 2 column (4) and Table 3 column (8). Our findings thus represent a general phenomenon across the United States rather than only in a few states particularly impacted by the pandemic.

One may also worry that "Twitter bots" rather than actual users are responsible for the rise in ch-word use on Twitter. However, only 10.4 percent of users who post anti-Asian tweets between January 2020 and April 2020 are potential bots (Ziems et al., 2020). Moreover, our results are quantitatively unchanged when we exclude users who are more likely to be bots, i.e., those who tweeted the ch-word more than five times (99 percentile) during the sample period, in Table 3 column (9).

The increase in searches and tweets including the ch-word could also come from the seasonality in ch-word use and may exist absent the pandemic. To test this possibility, we generate a placebo diagnosis date for each area using the same calendar day and month of its actual diagnosis date but changing the year from 2020 to 2019. We reestimate equation 3 using the placebo dates and plot the effects in Figure A7. Reassuringly, the racially charged search index and post index do not change around the placebo dates, suggesting that seasonality cannot explain our findings.

Finally, the increase in ch-word use on Twitter could reflect a change in the social cost of publicly expressing racial animus rather than a shift in attitudes towards Asians. However, this would not explain the increase in racist Google searches, which are done in private. Several other pieces of evidence also support a shift in attitudes. First, the proportion of ch-word tweets showing emotions of anger and disgust increases from 23.3 percent between November 2020 and the first local diagnosis to 40.8 percent in the six weeks after the first diagnosis.[18] Second, data on self-reported hate incidents from Asian Pacific Policy and Planning Council (AP3CON) Stop Hate Reporting System show that the daily average of anti-Asian hate incidents nationwide was alarmingly 70 in late March 2020 and 13 between April and May 2020 (Figure A5). Third, Pew Research Center's Global Attitudes Survey, conducted in June to August 2020, shows that unfavorable views of China have reached historic high (Silver, Delvin, and Huang, 2020). All

[18]Crimson Hexagon assigns each tweet emotion tag(s) generated via a natural language processing algorithm. Please refer to https://www.brandwatch.com/blog/understanding-sentiment-analysis for more details.

the above suggests that the rise in ch-word usage represents a real change in animus against Asians and not just a lower cost of publicly expressing it.

4.3 What Motivates Racial Animus and Who Responds the Most

Thus far, we have provided evidence that animus against Asians, as measured by Google searches and Twitter posts including the ch-word, surges immediately following the first diagnosis in an area. As a next step, we explore *what* motivates individuals to increase animus in response to the pandemic and *who* respond the most.

As a first step towards understanding the rise in racial animus, we test whether the rise is specific to Asians. If the racial animus is motivated by an overall increase in ethnic distrust or tensions from general uncertainty about cross-group differences in health status or risk-taking behavior, we expect to see an increase in animus against other minorities too. By contrast, if the racial animus is targeted *only* at Asians, it is more likely to be motivated by the association between Asians and the geographical origin of the virus.

To proxy for racial animus against other minorities, we construct Google search and Twitter post indexes for common racial epithets against major minority groups in the United States, such as the n-word (both singular and plural) against African Americans, "wetback(s)" against Hispanics, and "kike(s)" against the Jewish population.[19] We estimate equation 3 using racially charged searches and tweets against these minorities as outcomes.[20] The coefficients on the event dummies are plotted in Figure A6. None of the examined racial epithets experience an increase in Google searches following the first local diagnosis. A similar pattern is found for tweets using the w-word and the k-word.[21] The lack of response in the use of racial epithets against other minorities suggests that the pandemic-induced racial animus is mainly driven by

[19] We do not use "spic(s)" as the epithet against Hispanics because the cleaner brand "Spic and Span" experienced growing interest during the pandemic. We do not include "redskin(s)" because the corresponding queries and tweets are about an American football team formerly called "Washington Redskins".

[20] When using the n-word as the outcome, we include an indicator for the week of January 26, 2020 because the word's use spiked due to an MSNBC anchor using the n-word when broadcasting Kobe Bryant's death. When using the k-word as the outcome, we include an indicator for the week of February 23, 2020 because Los Angeles Dodgers player Enrique ("Kiké") Hernandez led to a spike in the word's use.

[21] We present the result for tweets using the n-word in Figure A6 panel C. N-word tweets may not be a valid proxy for racial animus against African Americans on Twitter because of Black Lives Matter protests, Black History Month in February, and seasonality which is evident when comparing the n-word usage between 2019 and 2020 in panel D. Note that we include an indicator for the week of February 9, 2020 in panel A because a viral n-word tweet unrelated to COVID-19 contributed to 95 percent of the n-word tweets on that day.

the connection between Asians and the geographical origin of the virus.

Although the anti-Asian animus is motivated by the potential geographical origin of the virus, we find that the racially charged tweets extend to broader topics than just the virus. Figure 2 demonstrates that the increase in ch-word tweets mostly comes from those that *do not* explicitly mention COVID-19, i.e., no mention of "virus", "COVID", "pandemic" or "epidemic". As a result, pandemic-induced racial animus may persist beyond the duration of the pandemic.

We next study which type of individuals are more susceptible to the pandemic shock. We begin by examining whether the increase in ch-word usage comes from users who start to harbor animus after the pandemic hits or from existing racists who step up their animosity. We define *existing* ch-word users as individuals who tweeted the ch-word at least once between 2014 and the sixth week before the first local COVID-19 diagnosis. We define *first-time* ch-word users as individuals who never tweeted the ch-word between 2014 and the sixth week before the first local diagnosis and who posted at least 10 tweets before their first ch-word tweet. Note, this definition avoids counting newly registered Twitter users as first-time ch-word users.

Figure 3 plots the breakdown in effects by the first-time versus existing ch-word user status. The increase in ch-word tweets from first-time users is roughly 4.5 times of that from existing users in the first two weeks after the first local diagnosis. This breakdown suggests that the extensive margin could play an more important role than the intensive margin in driving racial animus during crises. After the first local diagnosis, 4,621 Twitter users started to use the racial epithet, potentially exposing their combined 13 million followers to racially charged content and creating a multiplier effect on racial animus.

To better understand the type of individuals whose anti-Asian sentiment is easily influenced by the pandemic, we analyze the characteristics of the first-time ch-word users based on their user profiles and historical tweets extracted in August, 2021.[22] Table 4 reports the results. On average, 67.6 percent of these users are men, which is slightly higher than the proportion of US Twitter users who are men in 2021 (62 percent) (Statistia, 2021). In addition, they are seasoned Twitter users: their average account age is over six years, and their average number of followers and followings are over 4,100 and nearly 1,500, respectively.

[22] We downloaded historical tweets and user profiles for 3,033 of these users in August 2021. We cannot download the rest because their accounts are private, suspended, or no longer exist on Twitter.

Besides the information from user profiles, we also examine users' historical tweets to understand what they tweeted and who they interacted with before and during the pandemic. First we note that less than two percent of first-time ch-word users mentioned COVID-related conspiracies during the pandemic, suggesting that such conspiracies may not be the main cause of their racial animus. Over 16 percent have tweeted racial epithets against non-Asian minorities before the pandemic, implying that the crisis may have redirected anti-minority sentiments of these users towards Asians. Nearly 87 percent of these users were connected with at least one other ch-word user before the pandemic. As we will see in the next section, this connection is a key element in spreading animus against Asians.

Public opinion about coronavirus is more politically divided in the United States than many other advanced economies (Mordecai and Connaughton, 2020). We investigate whether first-time ch-word users have any party leaning. We find that 56 percent of these users mentioned President Trump in their tweets before January 2020, while 23 percent mentioned Senator Nancy Pelosi, followed by 18 percent for Senator Chuck Schumer, 14 percent for Senator Kevin McCarthy, and 2 percent for Senator Mitch McConnell. To place these numbers in perspective, roughly 43 percent of Twitter users follow President Trump in January, 2021 while only 3.5, 1.6, 0.6, and 1 percent users follow Senators Pelosi, Schumer, McCarthy, and McConnell, respectively.[23] Our finding suggests that the first-time ch-word users may be more politically active than average Twitter users. Given that media consumption is polarized along party lines (Jurkowitz et al., 2020), we also look at the party leaning of the Twitter news accounts these users paid attention to. We find that 24 percent of the users mentioned *only* Republican-leaning Twitter news accounts before the pandemic while 1 percent mentioned *only* Democratic-leaning ones.

5. Factors Fueling Racial Animus

In this section, we explore factors that may have helped propagate anti-Asian animus during the pandemic. Understanding these factors is crucial to stopping the spread of animus from the outset amid future crises.

[23]These percentages are calculated by multiplying the follower number of each Twitter account with the proportion of US Twitter users and dividing the product with the number of US Twitter users.

We know from the previous section that first-time ch-word users are the main driving force behind the rise of ch-word usage on Twitter during the pandemic. In Table 5, we zoom into these users and their Twitter activity between the date of the first local COVID-19 case and the end of sample (May 2, 2020) to understand what predicts their ch-word tweets. Doing so could shed light on factors fueling the spread of racial animus. We regress a user's likelihood of tweeting the ch-word in a given day on a series of indicators for whom they interacted with in the day before and a vector of user characteristics. We control for county, year-of-week, and day-of-week fixed effects to absorb the average propensity to tweet the ch-word in a county and the national trend of such tweets.

Exposure to anti-Asian individuals. Interaction with anti-Asian users (i.e. users who have previously used the ch-word) in a given day is associated with a 0.26 percentage point increase in the likelihood of tweeting the epithet in the following day, or 21 percent of the average likelihood in the sample period. This finding highlights the importance of social media in spreading racial animus and is consistent with papers which document how social media influence real outcomes like voting behaviors (e.g., Fujiwara et al., 2021). Our finding suggests that moderating racist individuals and their interaction with others on social media could constrain the spread of animus.

Opinions of public figures. The only other positive predictor is a user's interaction with President Trump. (Re)tweeting, replying, or mentioning the president in a day is associated with a 0.30 percentage point increase in the likelihood of tweeting the ch-word on the next day (24 percent of average). This finding is consistent with Müller and Schwarz (2019) which shows that President Trump's tweets affect public behavior such as hate crimes. In contrast, mentioning other prominent politicians of either parties or national news accounts has little to no predictive power or even predict a lower likelihood of tweeting the epithet. When we additionally control for the number of new COVID-19 cases or deaths in the local area in column (2), the results are unaffected. Taken together, certain public figures play a key role in shaping public opinions of a subject matter. Harnessing their opinion-shaping power could be useful in curbing animus in the future.

Salience of Asian-COVID connection. One potential factor mediating the relationship between ch-word use and interaction with President Trump is the salience of the connection between COVID-19 and the Asian population. It is possible that President Trump's tweets that mention COVID-19 and China simultaneously (hereafter, China-and-COVID tweet) may increase the salience of the connection and influence racial animus. We categorize all President Trump's tweets between January 1, 2020 and May 2, 2020 that contain any of the words "china", "chinese", "huawei", "xi", "covid", "covid-19", "corona", "coronavirus", "virus", "epidemic", or "pandemic" into three categories: those mentioning only China (China-only), only COVID-19 (COVID-only), and both China and COVID-19 (China-and-COVID). Table A5 presents examples of President Trump's tweets. Figure A8 plots the daily frequency of his tweets.

In Table 6, we regress the daily racially charged Twitter post index at the national level on the number of the president's tweets in each of the three categories while controlling for year-week and day-of-week fixed effects. Column (1) shows that one additional China-and-COVID tweet of President Trump in a day corresponds to roughly five more racially charged tweets per million "the" tweets nationwide on the same day. This increase is non-trivial and is equivalent to 14 percent of the national daily average. Importantly, the Twitter post index does not co-move with the president's China-only or COVID-only tweets, highlighting that the *connection* between China and COVID-19 is what matters. Results are unaffected when we control for the daily number of new COVID-19 cases and deaths nationwide in column (2).

The time-series correlation may be confounded by contemporaneous shocks unrelated to the president's tweets. To alleviate this concern, we conduct an event study comparing nationwide racially charged Twitter post index in the *hours* before and after President Trump's China-and-COVID tweets, using the index during the same hours-of-day on days without such tweets as controls. Figure 4 shows that the index in the four hours leading up to the China-and-COVID tweets is no different from other times, but it jumps in the first hour after such tweets and continues to grow. The immediacy of the change suggests a causal interpretation of the relationship between the salience of the China-and-COVID connection and the anti-Asian sentiment at the national level.[24]

In addition, we study whether the salience of the connection has translated into hate incidents against Asians. We obtain self-reported anti-Asian hate incidents from AP3CON Stop

[24]An alternative explanation is that there are other events that simultaneously affect President Trump's likelihood of China-and-COVID tweets and national racial animus.

AAPI Hate Reporting Center, a hate incident self-reporting website that went online on March 17, 2020. This is the best hate-tracking organization specialized in anti-Asian hate incidents in the United States (CBS News, 2020). In Table 6, column (3), we regress the log of daily hate incidents at the national level on the number of the president's tweets in each of the aforementioned categories while controlling for year-week and day-of-week fixed effects. We find that one additional China-and-COVID tweet from the president in a day corresponds to a roughly eight percent increase in self-reported hate incidents against Asians nationwide on the same day.[25] When we control for daily number of new COVID-19 cases and deaths nationwide in column (4), results are unchanged.[26]

In contrast to the clear relationship between the anti-Asian sentiment and the president's tweets, we find little evidence that the sentiment co-moves with tweets from other prominent politicians or national news outlets (Table A6).[27] The difference is likely due to the large gap in the number of Twitter followers between the president and the others. President Trump amassed 88.7 million followers before Twitter suspended his account in January 2021, while the follower number as of October, 2021 for the prominent politicians and national news outlets are mostly below 10 millions with only Fox and CNN reaching 20.2 and 54.7 millions, respectively.

Economic downturn. The COVID-19 pandemic poses risks on both lives and livelihoods. Existing work has documented that a deterioration of economic conditions can fuel animus towards minorities (Anderson et al., 2017, 2020, Sharma, 2015). We thus study the heterogeneity in racial animus increases by the level of the pandemic's negative impact on the local economy. We partition the main regression samples by whether the proportion of an area's annual average employment in "leisure and hospitality" and "education and health services," the two hardest-hit industries in employment according to the Bureau of Labor Statistics (BLS), is above or below the sample median (32 percent in Google data and 35 percent in Twitter data). We also partition the samples by whether the percent change in net revenue between January and March, 2020 among local small businesses is above or below the sample median

[25] We conduct the analysis at the daily level because the exact hour of the incidents is not available. We cannot estimate equation 3 with AP3CON data due to lack of pre-periods given the late start date of the data.

[26] In unreported results, we find little relationship between the racially charged Google search index and President Trump's tweets.

[27] Some coefficients are omitted because the Twitter account in question does not have the corresponding type of tweets.

(-39 percent in the Google sample and -37 percent in the Twitter sample) using data built by Chetty et al. (2020). Figure 5 shows that the areas that experience high versus low negative economic impact do not respond differently to the first local COVID-19 diagnosis. In other words, the negative economic impact of the disease plays a relatively weaker role in motivating the initial rise of racial animus. One potential reason is that the long-term impact was not well understood at the beginning of the pandemic.

6. Conclusion

This paper studies the formation and the spread of crisis-driven racial animus using evidence from the COVID-19 pandemic. We exploit variation in the timing of the first COVID-19 diagnosis across US regions and find that the first local case leads to an immediate increase in racial animus in an area. This rise in animus specifically targets Asians, implying that the association between this group and the potential geographical origin of the virus likely motivates the animosity. The majority of racist tweets come from users who post the epithet for the first time, and their interaction with other anti-Asian individuals predict such tweets, highlighting the role of extensive margin and social network in the spread of hatred. Moreover, the anti-Asian sentiment spills over to topics broader than the pandemic, as most of the racist tweets do not explicitly mention the virus. This suggest that the animus could persist beyond the duration of the pandemic. Finally, online animosity and offline hate incidents against Asians both increase when President Trump more frequently links China and COVID-19 in his tweets. The increase in racial animus does not differ by the economic impact of the pandemic.

Growing racial tension is a serious challenge facing society. Our analysis sheds light on how negative shocks, such as pandemics, motivate the animus, the types of individuals susceptible to such shocks, and how the animus spreads. From a policy perspective, our findings have practical implications. Careful naming of a shock, debunking claims of any alleged connection between a shock and a group, moderating racist individuals and their interaction with others on social media could all be helpful in curbing animus amid future crises.

This paper opens up several avenues for future research, such as predicting and stopping the spread of racist content online. While we estimate the effect of pandemics on racial animus, it would also be interesting to know the downstream consequences of such crisis-driven animus, for example, on labor market, geographical sorting, and immigration.

References

Anderson, D. M., Crost, B., and Rees, D. I. (2020). Do economic downturns fuel racial animus? *Journal of Economic Behavior & Organization*, 175:9–18.

Anderson, L. and Lepore, E. (2013). Slurring words. *Noûs*, 47(1):25–48.

Anderson, R. W., Johnson, N. D., and Koyama, M. (2017). Jewish persecutions and weather shocks: 1100–1800. *The Economic Journal*, 127(602):924–958.

Association, P. B. (2014). Resolution opposing use of racial slurs in business names.

Bartos, V., Bauer, M., Cahlíková, J., and Chytilová, J. (2020). Covid-19 crisis fuels hostility against foreigners.

Becker, G. S. (2010). *The economics of discrimination*. University of Chicago press.

Card, D., Mas, A., and Rothstein, J. (2008). Tipping and the dynamics of segregation. *The Quarterly Journal of Economics*, 123(1):177–218.

Charles, K. K. and Guryan, J. (2008). Prejudice and wages: an empirical assessment of becker's the economics of discrimination. *Journal of political economy*, 116(5):773–809.

Chetty, R., Friedman, J. N., Hendren, N., Stepner, M., et al. (2020). How did covid-19 and stabilization policies affect spending and employment? a new real-time economic tracker based on private sector data. Technical report, National Bureau of Economic Research.

Croucher, S. M., Nguyen, T., and Rahmani, D. (2020). Prejudice toward asian americans in the covid-19 pandemic: the effects of social media use in the united states. *Frontiers in Communication*, 5:39.

Depetris-Chauvin, E. (2015). Fear of obama: An empirical study of the demand for guns and the us 2008 presidential election. *Journal of Public Economics*, 130:66–79.

Druckman, J. N. and Holmes, J. W. (2004). Does presidential rhetoric matter? priming and presidential approval. *Presidential Studies Quarterly*, 34(4):755–778.

Egorov, G., Enikolopov, R., Makarin, A., and Petrova, M. (2020). Divided we stay home: Social distancing and ethnic diversity. Technical report, National Bureau of Economic Research.

Faulkner, J., Schaller, M., Park, J. H., and Duncan, L. A. (2004). Evolved disease-avoidance mechanisms and contemporary xenophobic attitudes. *Group Processes & Intergroup Relations*, 7(4):333–353.

Fujiwara, T., Müller, K., and Schwarz, C. (2021). The effect of social media on elections: Evidence from the united states. Technical report, National Bureau of Economic Research.

Hanes, E. and Machin, S. (2014). Hate crime in the wake of terror attacks: Evidence from 7/7 and 9/11. *Journal of Contemporary Criminal Justice*, 30(3):247–267.

Healy, J. and Searcey, D. (2020).

Ivandic, R., Kirchmaier, T., and Machin, S. J. (2019). Jihadi attacks, media and local hate crime.

Jedwab, R., Johnson, N. D., and Koyama, M. (2019). Negative shocks and mass persecutions: evidence from the black death. *Journal of Economic Growth*, 24(4):345–395.

Jurkowitz, M., Mitchell, A., Shearer, E., and Walker, M. (2020). 2. americans are divided by party in the sources they turn to for political news. *Pew Research Center*.

Kaushal, N., Kaestner, R., and Reimers, C. (2007). Labor market effects of september 11th on arab and muslim residents of the united states. *Journal of Human Resources*, 42(2):275–308.

Lyu, H., Chen, L., Wang, Y., and Luo, J. (2020). Sense and sensibility: Characterizing social media users regarding the use of controversial terms for covid-19. *IEEE Transactions on Big Data*.

Mordecai, M. and Connaughton, A. (2020). Public opinion about coronavirus is more politically divided in us than in other advanced economies. *Pew Res. Cent.*

Müller, K. and Schwarz, C. (2019). From hashtag to hate crime: Twitter and anti-minority sentiment. *Available at SSRN 3149103*.

Mullis, S. and Glenn, H. (2020). New site collects reports of racism against asian americans amid coronavirus pandemic.

Nguyen, T. T., Meng, H.-W., Sandeep, S., McCullough, M., Yu, W., Lau, Y., Huang, D., and Nguyen, Q. C. (2018). Twitter-derived measures of sentiment towards minorities (2015–2016) and associations with low birth weight and preterm birth in the united states. *Computers in human behavior*, 89:308–315.

Novgorodsky, D. and Setzler, B. (2019). Practical guide to event studies.

O'Shea, B. A., Watson, D. G., Brown, G. D., and Fincher, C. L. (2020). Infectious disease prevalence, not race exposure, predicts both implicit and explicit racial prejudice across the united states. *Social Psychological and Personality Science*, 11(3):345–355.

Schaller, M. and Neuberg, S. L. (2012). Danger, disease, and the nature of prejudice (s). In *Advances in experimental social psychology*, volume 46, pages 1–54. Elsevier.

Schild, L., Ling, C., Blackburn, J., Stringhini, G., Zhang, Y., and Zannettou, S. (2020). " go eat a bat, chang!": An early look on the emergence of sinophobic behavior on web communities in the face of covid-19. *arXiv preprint arXiv:2004.04046*.

Sharma, S. (2015). Caste-based crimes and economic status: Evidence from india. *Journal of comparative economics*, 43(1):204–226.

Silver, L. R. (2016). China in the media: Effects on american opinion.

Stephens-Davidowitz, S. (2014). The cost of racial animus on a black candidate: Evidence using google search data. *Journal of Public Economics*, 118:26–40.

Sun, L. and Abraham, S. (2020). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*.

Tynes, B. M., Giang, M. T., Williams, D. R., and Thompson, G. N. (2008). Online racial discrimination and psychological adjustment among adolescents. *Journal of adolescent health*, 43(6):565–569.

Ziems, C., He, B., Soni, S., and Kumar, S. (2020). Racism is a virus: Anti-asian hate and counterhate in social media during the covid-19 crisis. *arXiv preprint arXiv:2005.12423*.

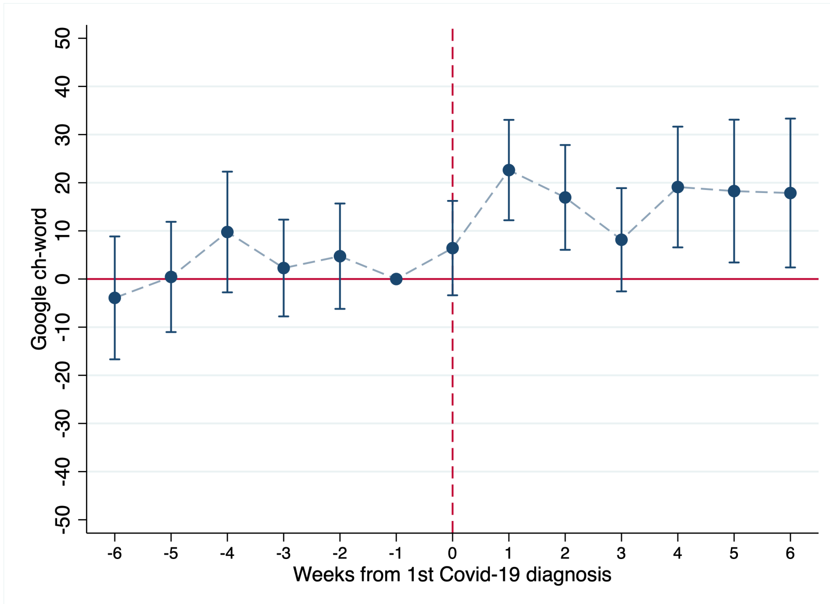

A: Google search index

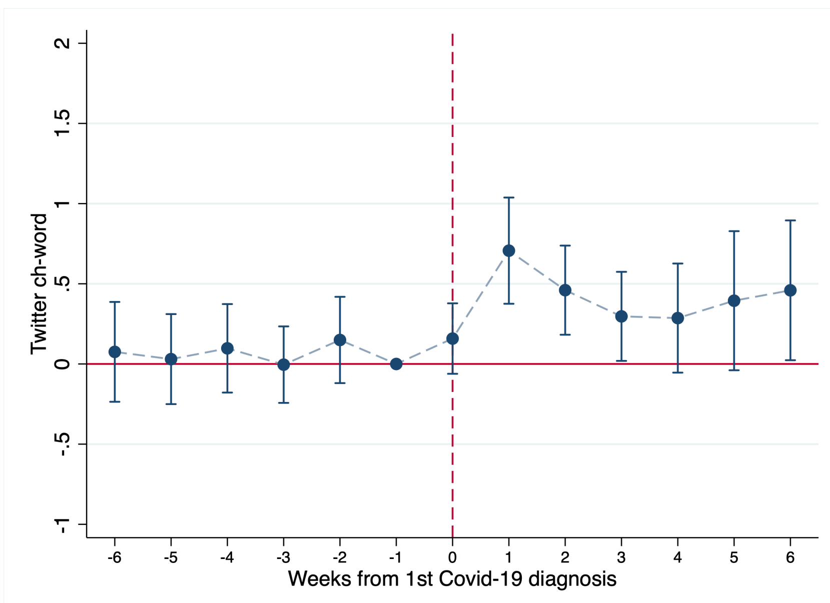

B: Twitter post index

Figure 1: The Effect of the First Local COVID-19 Diagnosis on Racial Animus

Note: The figure presents the effect of the first local COVID-19 diagnosis on the racially charged Google search index and Twitter post index. See section 2.1 for the definitions of the indexes. Panels A and B plot the estimates and 95 percent confidence intervals of the coefficients on the event dummies in equation 3 using the Google search index and the Twitter post index as the outcome, respectively. The estimates in panel A and panel B correspond to column (1) of Table 2 and column (1) of Table 3. Regressions control for year-month fixed effects and media market or county fixed effects. Standard errors are clustered by media market (panel A) or by county (panel B).

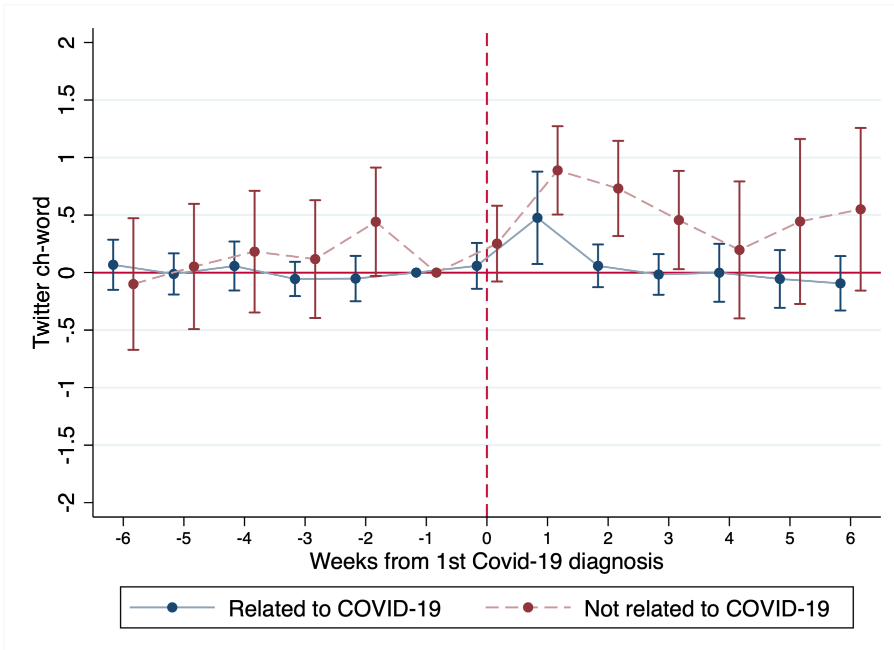

Figure 2: The Effect of the First Local COVID-19 Diagnosis on Racially Charged Tweets COVID-Related vs Non-COVID-Related Tweets

Note: The figure presents the effect of the first local COVID-19 diagnosis on the racially charged Twitter post index by whether or not the tweets are related to COVID-19. COVID-related racially charged Twitter post index are defined as the number of ch-word tweets including keywords: "COVID-19", "COVID", "virus", "pandemic", "epidemic", per 100,000 "the" tweets. The blue (red) line plots the estimates and 95 percent confidence intervals of the coefficients on the event dummies in equation 3 using (non-)COVID-related racially charged Twitter post index as the outcome. All regressions control for year-month fixed effects and county fixed effects. Standard errors are clustered by county.

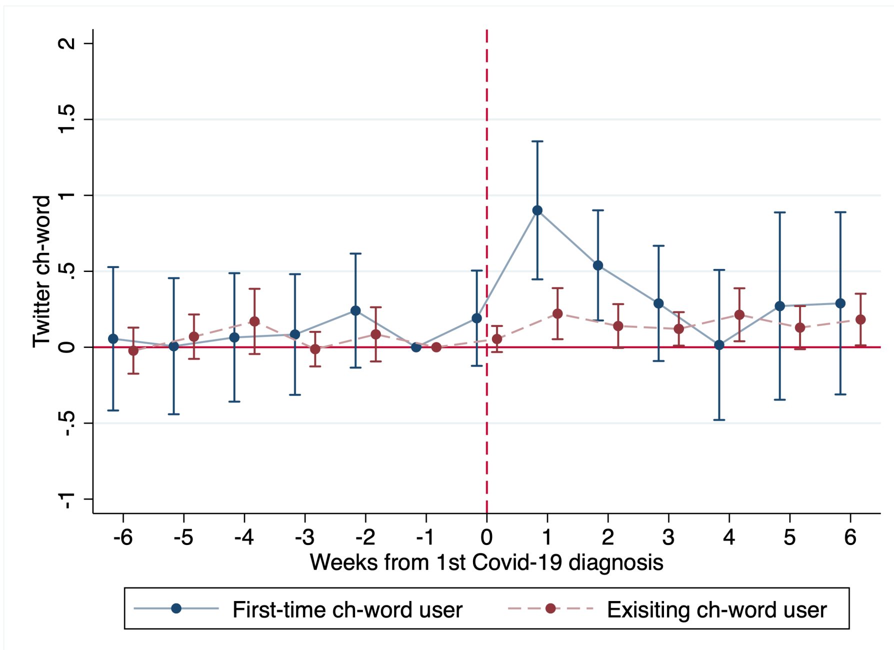

Figure 3: The Effect of the First Local COVID-19 Diagnosis on Racially Charged Tweets
First-time vs Existing Ch-word Users

Note: The figure presents the effect of the first local COVID-19 diagnosis on the racially charged Twitter post index by whether the user is a first-time or an existing ch-word user. See section 4.3 for definitions of first-time and existing ch-word users. The blue (red) line plots the estimates and 95 percent confidence intervals of the coefficients on the event dummies in equation 3 using the racially charged Twitter post index based on first-time (existing) ch-word users as the outcome. All regressions control for year-month fixed effects and county fixed effects. Standard errors are clustered by county.

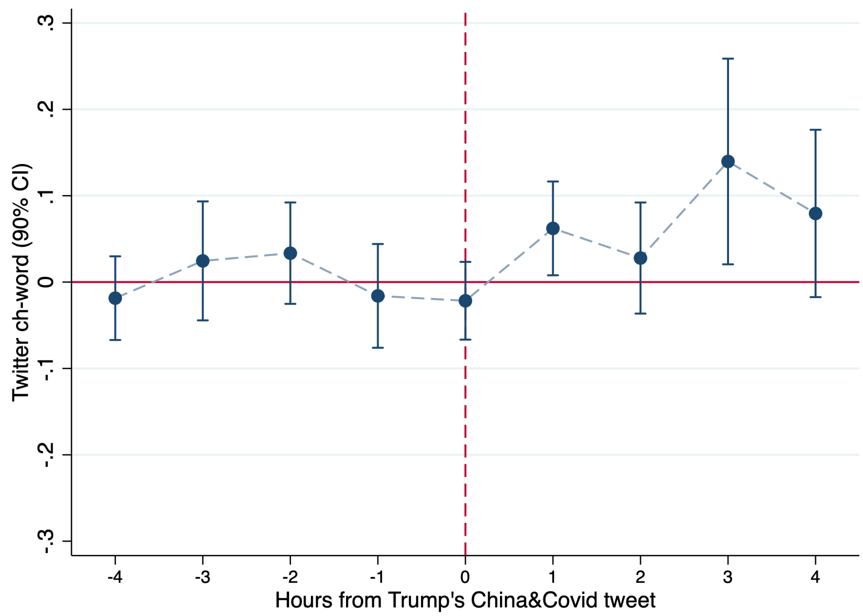

Figure 4: Relationship between Racially Charged Tweets and Trump Tweets

Note: The figure presents the relationship between the number of President Trump's tweets that mention both Covid-19 and China (China-and-COVID tweets) in an hour and the number of ch-word tweets per 100,000 "the" tweets nationwide in the four hours before and the four hours after the president's tweets. It plots the estimates and 90 percent confidence intervals of the coefficients on the interactions between hourly event dummies and the number of Trump's China-and-COVID tweets at hour zero. Event dummy for the hours outside of those being plotted are omitted. We control for year-week fixed effects, day of week fixed effects, and hour fixed effects. Standard errors are clustered by date.

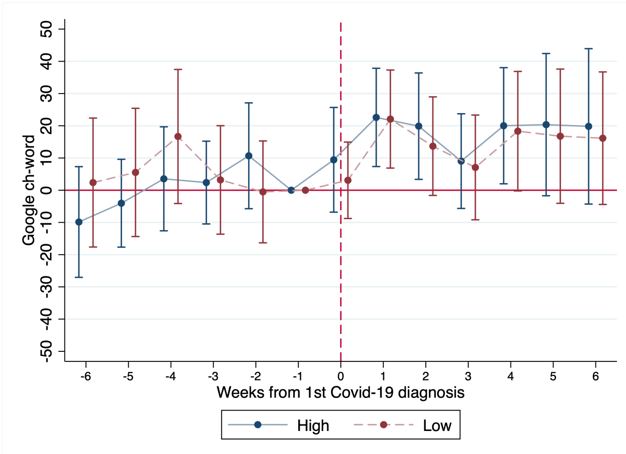

A: Google - Employment shock

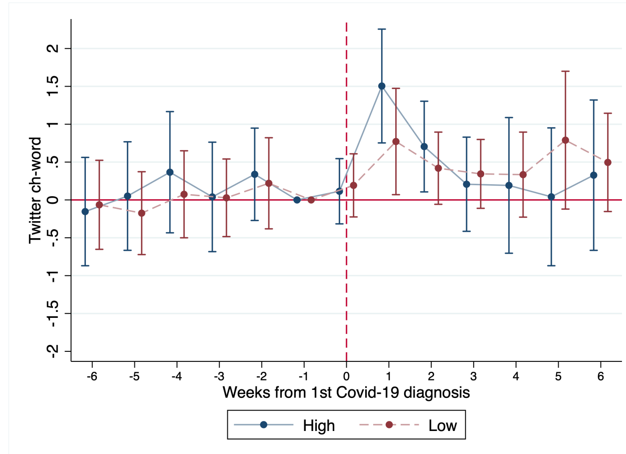

B: Twitter - Employment shock

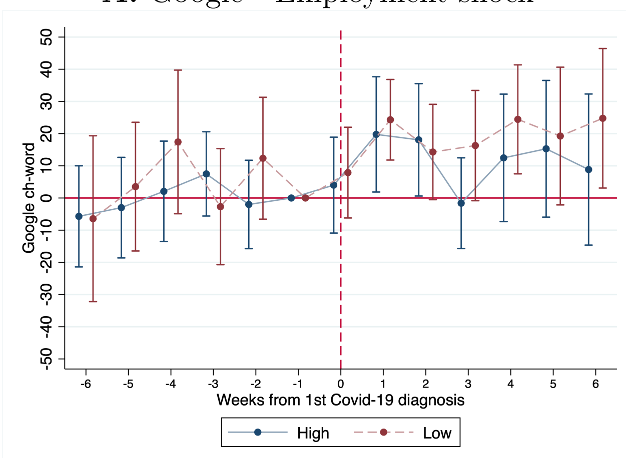

C: Google - Revenue shock

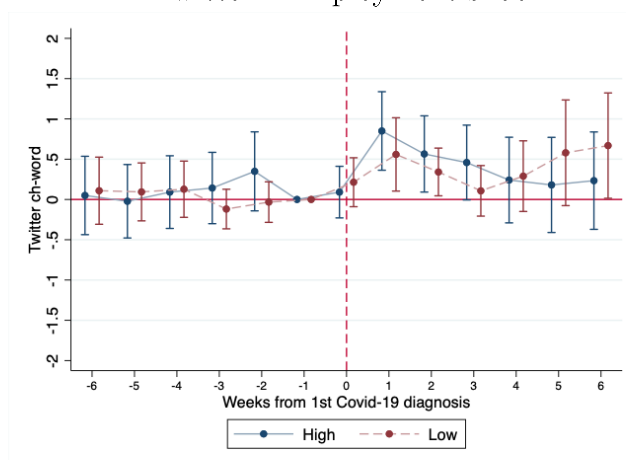

D: Twitter - Revenue shock

Figure 5: The Effect of the First Local COVID-19 Diagnosis on Racial Animus by the Negative Economic Impact of COVID-19

Note: The figures present the heterogeneous effect of the first local COVID-19 diagnosis on the racially charged Google search index and Twitter post index by the negative economic impact of the pandemic. Panels A and B partition the regression sample by whether the proportion of an area's annual average employment in "leisure and hospitality" and "education and health services", the two hardest-hit industries in employment according to BLS, is above or below the sample median (i.e., 32 percent in the Google sample and 35 percent in the Twitter sample). Panels C and D partition the regression sample by whether the percent change in net revenue among local small businesses between January and March is above or below the sample median (i.e., -39 percent in the Google sample and -37 percent in the Twitter sample). Panels A and C (B and D) plot the estimates and 95 percent confidence intervals of the coefficients on the event dummies in equation 3 using the racially charged Google search index (Twitter post index) as the outcome. Specifications in panels A and C mirror those in column (1) of Table 2, and specifications in panels B and D mirror those in column (1) of Table 3.

Table 1: Relationship between Racial Animus, Hate Crimes, and Chinese Restaurant Visits

VARIABLES	(1) Incidents	(2) Incidents	(3) Incidents	(4) Incidents	(5) Visits	(6) Visits	(7) Visits	(8) Visits
Panel A: Google search index								
Google ch-word(t)	0.00057*	0.00057*	0.00057*	0.00057*	-21.069*	-22.017*	-23.489**	-24.381**
	(0.00031)	(0.00031)	(0.00031)	(0.00031)	(12.629)	(12.704)	(11.632)	(11.697)
Google ch-word(t-1)			-0.00019	-0.00018			-13.934	-14.900
			(0.00035)	(0.00035)			(11.467)	(11.594)
Total visits					0.052***	0.052***	0.048***	0.048***
					(0.005)	(0.005)	(0.006)	(0.006)
Population(m)	0.51926***	0.51907***	0.52042***	0.52287***				
	(0.19939)	(0.19942)	(0.19936)	(0.19991)				
Google Asian(s)(t)		-0.00005		-0.00018		-138.636**		-120.800**
		(0.00109)		(0.00111)		(60.169)		(54.317)
Google Asian(s)(t-1)				0.00075				10.816
				(0.00163)				(59.669)
Observations	3,600	3,600	3,600	3,600	1,440	1,440	1,380	1,380
R-squared	0.309	0.309	0.309	0.309	0.996	0.996	0.997	0.997
Outcome mean	.147	.147	.147	.147	104962.736	104962.736	104962.736	104962.736
Panel B: Twitter post index								
Twitter ch-word	-0.00041	-0.00037	-0.00048	-0.00047	-60.249	-61.190	-30.883	-30.178
	(0.00099)	(0.00096)	(0.00099)	(0.00100)	(49.798)	(50.150)	(42.641)	(42.737)
Twitter ch-word (t-1)			-0.00063	-0.00064			-10.216	-10.089
			(0.00065)	(0.00065)			(36.060)	(36.286)
Total visits					0.047***	0.047***	0.044***	0.044***
					(0.005)	(0.005)	(0.006)	(0.006)
Population(m)	0.56488***	0.56517***	0.58097***	0.58136***				
	(0.17387)	(0.17391)	(0.17623)	(0.17630)				
Twitter Asian(s)(t)		-0.00003		-0.00002		0.294		-0.279
		(0.00003)		(0.00002)		(0.682)		(0.646)
Twitter Asian(s)(t-1)				-0.00001				-0.266
				(0.00002)				(0.693)
Observations	11,116	11,116	10,921	10,921	4,493	4,493	4,300	4,300
R-squared	0.220	0.220	0.220	0.230	0.996	0.996	0.997	0.997
Outcome mean	.057	.057	.057	.057	41065.784	41065.784	41065.784	41065.784

Notes: The table presents the relationship between the racially charged Google search index and the Twitter post index, anti-Asian hate crimes, and visits to Chinese restaurants. Hate crime data are from the FBI UCR, visit data are from Safegraph, and all data are at the media market×year-month level. Outcome variables are the monthly number of anti-Asian hate crimes between January 2014 and December 2018 (columns (1)-(4)) and the monthly number of visits to Chinese restaurants between January 2018 and December 2019 (columns (5)-(8)). *Google Asian(s)* is the Google search index for the word "Asian(s)." *Twitter Asian(s)* is the number of tweets including "Asian(s)" per 100,000 "the" tweets. All regressions control for local unemployment rate, year-month fixed effects, and media market fixed effects. *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table 2: The Effect of the First Local COVID-19 Diagnosis on Racial Animus Google Search Index

VARIABLES	(1) Ch-word index	(2) Severity control	(3) Asian control	(4) Exclude states
-6w	-3.920	-2.694	-4.265	-8.979
	(6.379)	(6.620)	(6.404)	(8.341)
-5w	0.431	1.100	-0.198	-2.575
	(5.722)	(5.820)	(5.699)	(7.083)
-4w	9.764	10.088	9.419	9.205
	(6.263)	(6.316)	(6.233)	(7.649)
-3w	2.282	2.503	2.247	2.458
	(5.023)	(5.085)	(5.020)	(5.912)
-2w	4.739	4.899	4.771	2.564
	(5.469)	(5.535)	(5.467)	(6.150)
+0w	6.421	6.326	6.274	6.574
	(4.898)	(4.911)	(4.864)	(5.127)
+1w	22.628***	22.442***	22.030***	22.771***
	(5.210)	(5.246)	(5.280)	(5.721)
+2w	16.945***	15.936***	16.727***	18.104***
	(5.439)	(5.443)	(5.407)	(5.621)
+3w	8.155	5.702	7.894	8.614
	(5.359)	(5.907)	(5.403)	(5.829)
+4w	19.106***	15.972**	18.873***	19.527**
	(6.265)	(6.999)	(6.253)	(7.461)
+5w	18.263**	15.375*	18.041**	14.709*
	(7.411)	(8.113)	(7.428)	(8.679)
+6w	17.861**	15.002*	18.125**	18.017*
	(7.726)	(8.046)	(7.751)	(9.267)
Observations	780	780	780	663
R-squared	0.190	0.192	0.193	0.180
Outcome mean	30.03	30.03	30.03	30.03

Notes: The table presents the effect of the first local COVID-19 diagnosis on the racially charged Google search index. All columns report the estimates of coefficients on the event dummies in equation 3. Column (1) corresponds to Figure A4, panel A. Column (2) controls for the number of COVID-related new cases and deaths and whether the state has any stay-at-home orders in place. Column (3) controls for the Google search index for "Asian(s)." Column (4) excludes Washington, New York, and California. All regressions control for media market fixed effects and year-month fixed effects. Standard errors are clustered by media market. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: The Effect of the First Local COVID-19 Diagnosis on Racial Animus
Twitter Post Index

VARIABLES	(1) Ch-word index	(2) Ch-word level	(3) Ch-word per capita	(4) Exclude counter-hate	(5) Dynamic DID	(6) Severity control	(7) Asian control	(8) Exclude states	(9) Exclude bots
-6w	0.075	-0.022	0.061		-0.037	0.070	0.053	0.127	0.098
	(0.159)	(0.276)	(0.307)		(0.097)	(0.159)	(0.255)	(0.165)	(0.151)
-5w	0.030	-0.069	-0.036	-0.801	-0.085	0.027	0.091	0.056	0.039
	(0.143)	(0.158)	(0.267)	(1.252)	(0.091)	(0.143)	(0.242)	(0.142)	(0.153)
-4w	0.098	-0.128	-0.117	-0.328	-0.025	0.095	0.248	0.113	0.075
	(0.140)	(0.165)	(0.240)	(1.181)	(0.107)	(0.140)	(0.239)	(0.141)	(0.144)
-3w	-0.004	0.024	-0.100	0.450	-0.082	-0.006	0.095	0.018	0.014
	(0.121)	(0.091)	(0.195)	(1.152)	(0.081)	(0.121)	(0.213)	(0.129)	(0.138)
-2w	0.150	0.065	0.412	-0.361	0.120	0.149	0.331	0.136	0.242
	(0.137)	(0.050)	(0.308)	(0.967)	(0.094)	(0.137)	(0.212)	(0.146)	(0.180)
+0w	0.158	0.012	0.390**	5.154***	0.120***	0.163	0.168	0.169	0.203
	(0.112)	(0.069)	(0.170)	(1.005)	(0.094)	(0.159)	(0.171)	(0.122)	(0.142)
+1w	0.707***	0.227**	1.037***	5.075***	0.689***	0.718***	1.077***	0.572***	0.952***
	(0.169)	(0.105)	(0.197)	(1.046)	(0.159)	(0.166)	(0.238)	(0.162)	(0.228)
+2w	0.460***	0.348***	1.140***	2.855***	0.428***	0.478***	0.763***	0.389**	0.538***
	(0.142)	(0.109)	(0.252)	(1.039)	(0.111)	(0.145)	(0.199)	(0.151)	(0.173)
+3w	0.297**	0.631***	1.331***	2.688***	0.181*	0.315**	0.526**	0.300*	0.255*
	(0.141)	(0.193)	(0.396)	(0.842)	(0.095)	(0.152)	(0.204)	(0.154)	(0.137)
+4w	0.286*	0.789**	1.947**	1.521	0.122	0.307*	0.361	0.273	0.132
	(0.173)	(0.310)	(0.771)	(1.257)	(0.103)	(0.184)	(0.269)	(0.187)	(0.157)
+5w	0.394*	0.683***	1.650***	1.158	0.240	0.421*	0.535*	0.385	0.144
	(0.221)	(0.201)	(0.466)	(1.396)	(0.154)	(0.248)	(0.323)	(0.240)	(0.178)
+6w	0.459**	0.696***	1.664***	2.264	0.340**	0.489*	0.533*	0.479**	0.373*
	(0.222)	(0.223)	(0.469)	(1.566)	(0.150)	(0.252)	(0.315)	(0.243)	(0.198)
Observations	7,930	7,976	7,976	3,141	103,694	7,930	5,578	7,188	11,811
R-squared	0.121	0.809	0.270	0.611	0.112	0.121	0.142	0.123	0.060
Outcome mean	0.591	0.681	1.075	6.779	0.591	0.591	0.591	0.591	0.569

Notes: The table presents the effect of the first local COVID-19 diagnosis the racially charged tweets. All columns report the estimates of coefficients on the event dummies in equation 3, except for column (5). Column (1) corresponds to Figure A4, panel B. The outcome variable in column (2) is the number of ch-word tweets, and the regression additionally controls for the number of "the" tweets. The outcome variable in column (3) is the number of ch-word tweets per one million county population. Column (4) uses an alternative Twitter post index, which removes counter-hate tweets (see section 4.2.1). Column (5) presents the estimates from our dynamic DID event study. Column (6) controls for the number of COVID-related new cases and deaths and whether the state has any stay-at-home orders in place. Column (7) controls for the Twitter post index for "Asian(s)." Column (8) excludes Washington, New York, and California. Column (9) excludes tweets from users who are likely Twitter bots. All regressions control for county fixed effects and year-month fixed effects. Standard errors are clustered by county. *** p <0.01, ** p <0.05, * p <0.1.

Table 4: Characteristics of First-Time Ch-word Users

	Mean	SD	Max
Panel A: User characteristics			
Male (excl. unknown)	.676	0.468	1
Account years	6.309	3.652	13.8
Followers	4,104	57,537	2,853,415
Followings	1,441	3,509	92,927
Panel B: Prob. (re)tweet/reply/mention			
During pandemic:			
COVID conspiracy	0.017	0.131	1
Before pandemic:			
Anti-minority content	0.164	0.371	1
Anti-Asian user	0.866	0.341	1
Trump	0.561	0.496	1
McCarthy	0.137	0.344	1
McConnell	0.024	0.154	1
Pelosi	0.227	0.419	1
Schumer	0.178	0.383	1
Fox	0.282	0.450	1
CNN	0.388	0.487	1
CBS	0.115	0.319	1
Rep. news only	0.239	0.427	1
Dem. news only	0.014	0.118	1
Mixed news	0.306	0.461	1
N users	3,033		

Notes: This table presents the characteristics of first-time ch-word users (see section 4.3). Panel A reports information from Twitter user profiles. Panel B reports a user's likelihood of mentioning certain keywords in their tweets or interacting with certain users. "During pandemic" and "Before pandemic" refer to tweets posted before May 2, 2020 and before January 21, 2020, respectively. "COVID conspiracy" uses keywords related to COVID-19 conspiracies: plandemic, fakepandemic, scamdemic, film your hospital, 5gcoronavirus, or coronavirustruth. "Anti-minority content" uses racial epithets against non-Asian minorities as keywords: the n-word, the w-word, and the k-word. "Anti-Asian user" is one if an user has interacted with other ch-word users, and zero otherwise. "Trump" is one if an user has ever mentioned #trump or @realDonaldTrump, and zero otherwise; "McCarthy", "McConnell", "Pelosi", "Schumer", "Fox", "CNN", and "CBS" are similarly defined using @kevinomccarthy, @McConnellPress (or @LeaderMcConnell), @SpeakerPelosi, @SenSchumer, @cnn, @foxnews, @cnn, and @cbsnews as keywords, respectively. "Rep. news only" or "Dem. news only" are one if an user has mentioned only Republican-leaning news accounts or only Democrat-leaning ones in their tweets, while "Mixed news" is one if they have mentioned both.

Table 5: Predictors of Tweeting Ch-word among First-Time Ch-word Users after the First Local COVID-19 Diagnosis

VARIABLES	(1) P(ch-word) (t+1)	(2) P(ch-word) (t+1)
Anti-Asian user(t)	0.260***	0.260***
	(0.072)	(0.072)
COVID consp.(t)	1.180	1.185
	(1.847)	(1.846)
Trump(t)	0.302**	0.302**
	(0.122)	(0.122)
McCarthy(t)	-0.098	-0.100
	(0.457)	(0.457)
McConnell(t)	-2.077***	-2.066***
	(0.442)	(0.443)
Pelosi(t)	-0.423	-0.422
	(0.284)	(0.284)
Schumer(t)	-0.107	-0.106
	(0.379)	(0.379)
CBS(t)	-0.696	-0.696
	(0.825)	(0.825)
CNN(t)	0.170	0.171
	(0.277)	(0.277)
Fox(t)	-0.384	-0.384
	(0.346)	(0.346)
Male	-0.005	-0.005
	(0.049)	(0.049)
Gender unknown	0.068	0.068
	(0.045)	(0.045)
Verified account	-0.118	-0.118
	(0.106)	(0.106)
Log(followers)	-0.028**	-0.028**
	(0.014)	(0.014)
Log(followings)	-0.012	-0.012
	(0.020)	(0.020)
New diagnoses		-0.000
		(0.000)
New deaths		0.000
		(0.000)
Observations	174,164	174,164
R-squared	0.004	0.004
Outcome mean	1.251	1.251

Notes: This table presents the relationship between first-time ch-word users' likelihood of tweeting the ch-word in day and their Twitter activity in the day before as well as their baseline characteristics. See note to Table 4 for the definitions of the independent variables. All regressions control for year-week fixed effects and day-of-week fixed effects. Standard errors are clustered by user. *** p <0.01, ** p <0.05, * p <0.1.

Table 6: Relationship between Racial Animus and Trump Tweets

	(1)	(2)	(3)	(4)
VARIABLES	Twitter ch-word	Twitter ch-word	Log(Incidents)	Log(Incidents)
China&Covid(t)	0.0482**	0.0493**	0.0799*	0.0888**
	(0.0234)	(0.0246)	(0.0453)	(0.0398)
China only(t)	-0.0126	-0.0130	-0.0592	-0.0332
	(0.0131)	(0.0133)	(0.0815)	(0.0844)
Covid only(t)	0.0008	0.0006	-0.0014	0.0004
	(0.0038)	(0.0040)	(0.0146)	(0.0143)
New diagnoses		-0.0000		0.0000*
		(0.0000)		(0.0000)
New deaths		0.0001		0.0001
		(0.0001)		(0.0003)
Observations	123	123	45	45
R-squared	0.519	0.522	0.812	0.829
Outcome mean	.344	.344	3.1932	3.1932

Notes: The table presents the relationship between the number of President Trump's tweets about COVID-19 or China and racial animus in the United States. The outcome variable in columns (1)-(3) is the daily number of ch-word tweets per 100,000 "the" tweets nationwide between January 1, 2020 and May 2, 2020. The outcome variable in columns (4)-(6) is the natural log of the daily number of anti-Asian hate incidents nationwide from AP3CON Stop AAPI Hate Reporting system between March 19 and May 2, 2020. We categorize the president's tweets between January 1, 2020 and May 2, 2020 that include "china", "chinese", "huawei", "xi", "COVID", "COVID-19", "corona", "coronavirus", "virus", "epidemic", or "pandemic" into three categories: those mentioning only China (China only), only COVID-19 (COVID only), and both China and COVID-19 (China-and-COVID). "China-and-COVID" is the daily number of the president's China-and-COVID tweets; "China only" and "COVID only" are similarly defined. "New diagnoses" and "New deaths" are the daily number of COVID-related new cases and deaths in the United States. All regressions control for year-week fixed effects and day-of-week fixed effects. Standard errors are clustered by date. *** p <0.01, ** p <0.05, * p <0.1.

Online Appendix for
"How Racial Animus Forms and Spreads"

Runjing Lu and Sophie Yanying Sheng

Appendix 1. Rescaled Google Search Index

Google Trends reports the search index in either a time series or a cross-sectional format. To construct a panel data set for each media market and time, we need to extract the search index in each media market separately. However, the search index reported by Google Trends is the search rate normalized by the maximum search rate in an extraction and is not comparable across extractions. To build a panel of search indexes that are normalized by the same base, we rescale the search index using the following method.

In a time series extraction of the search index in media market m over period T, the search index in median market m at time t is approximately

$$\text{Search Index}_{mt,T} = 100 \times \frac{\frac{\text{Searches including "chink(s)"}_{mt}}{\text{Total searches}_{mt}}}{\max\limits_{t \in T}\{\frac{\text{Searches including "chink(s)"}_{mt}}{\text{Total searches}_{mt}}\}} \tag{5}$$

Meanwhile, in a cross-sectional extraction of the search index at time t for all media markets $m \in M$, the search index in media market m at time t is approximately

$$\text{Search Index}_{mt,M} = 100 \times \frac{\frac{\text{Searches including "chink(s)"}_{mt}}{\text{Total searches}_{mt}}}{\max\limits_{m \in M}\{\frac{\text{Searches including "chink(s)"}_{mt}}{\text{Total searches}_{mt}}\}} \tag{6}$$

If we are willing to assume that the numerators in equations 5 and 6 are the same, then we can calculate the ratio of the two denominators as

$$Ratio_{m.MT} = \frac{\max\limits_{t \in T}\{\frac{\text{Searches including "chink(s)"}_{mt}}{\text{Total searches}_{mt}}\}}{\max\limits_{m \in M}\{\frac{\text{Searches including "chink(s)"}_{mt}}{\text{Total searches}_{mt}}\}} = \frac{\text{Search Index}_{mt,M}}{\text{Search Index}_{mt,T}} \tag{7}$$

when both search indexes are non-zero. We can scale the time series search index over period T in each media market $m \in M$ by multiplying it with the corresponding $Ratio_{m,MT}$. The resulting time series are normalized by the same $\max\limits_{m \in M}\{\frac{\text{Searches including "chink(s)"}_{mt}}{\text{Total searches}_{mt}}\}$. However, Google Trends returns a zero value when the absolute level of search in a given media market and time is below an unreported threshold, under which the rescaling does not work. After extracting cross-sectional search indexes on all possible weeks in the sample period, we can at best back out the rescaled search index for 35 media markets using Huntsville-Decatur (Florence) media market's search rate on March 15, 2020, as the base. Alternatively, we can back out 29 media markets using Wilkes Barre-Scranton media market's search rate on March 29, 2020, and 29 media markets using Buffalo media market's search rate on April 5, 2020, as the base. When combined, these three measures cover 50 media markets.

Note that Google calculates the search index using a random sample of searches, which can be different across extractions. As a result, the numerators in equations 5 and 6 are similar but may not be exactly the same. To the extent that these two numerators are not the same, we may be introducing measurement errors to the dependent variable and attenuating the main effects.

Appendix 2. Additional Figures & Tables

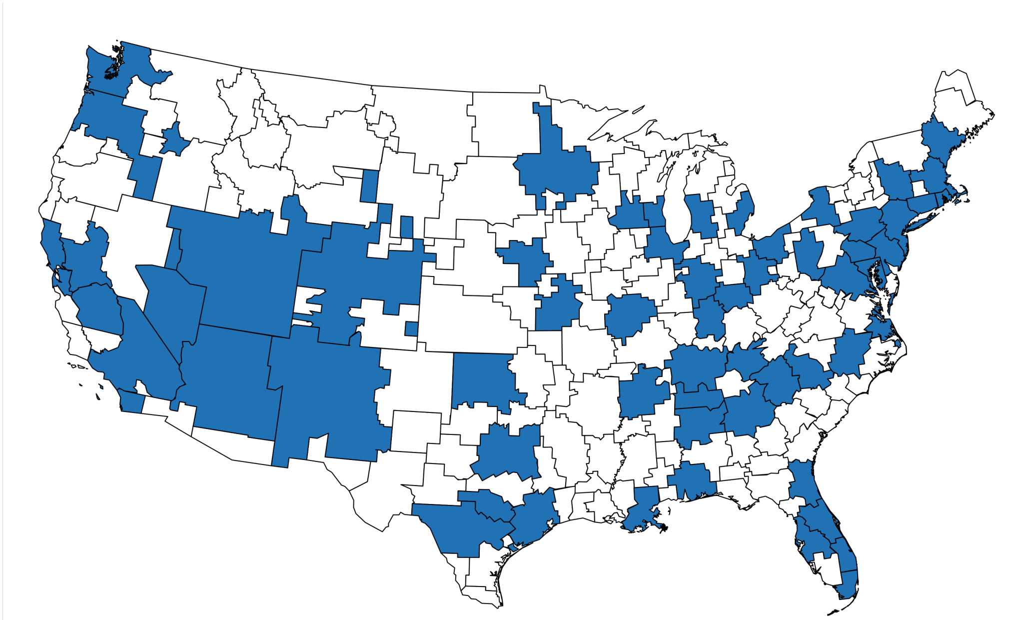

A: Google media market

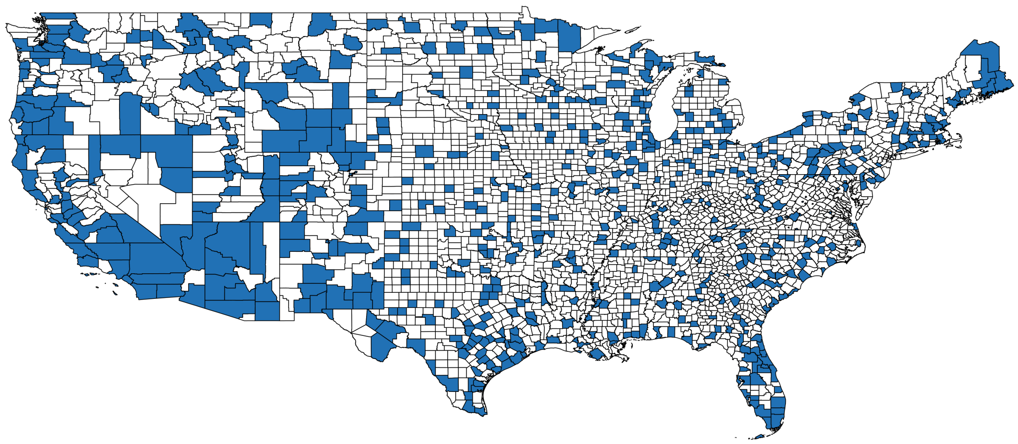

B: Twitter county

Figure A1: Location of Media Markets and Counties with Data

Note: The figure plots the locations of media markets with Google data in panel A and counties with Twitter data in panel B.

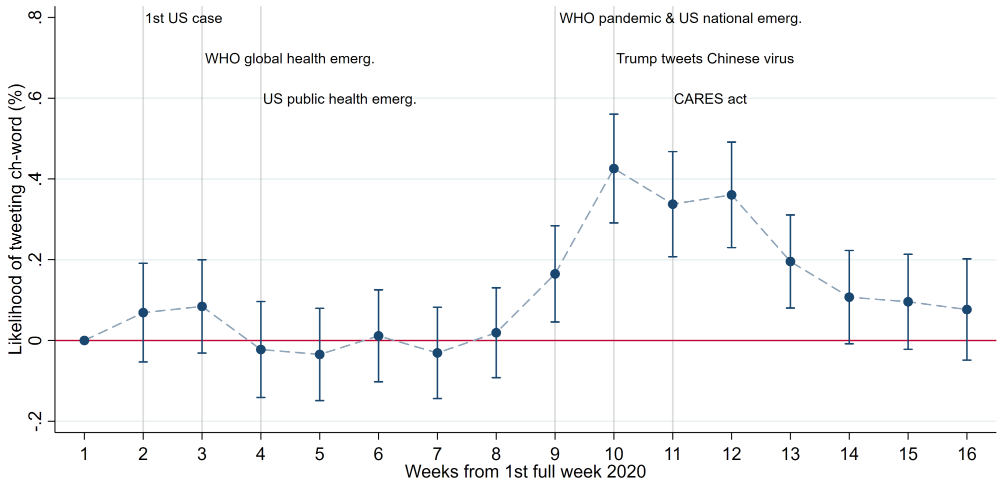

A: Likelihood of tweeting the ch-word

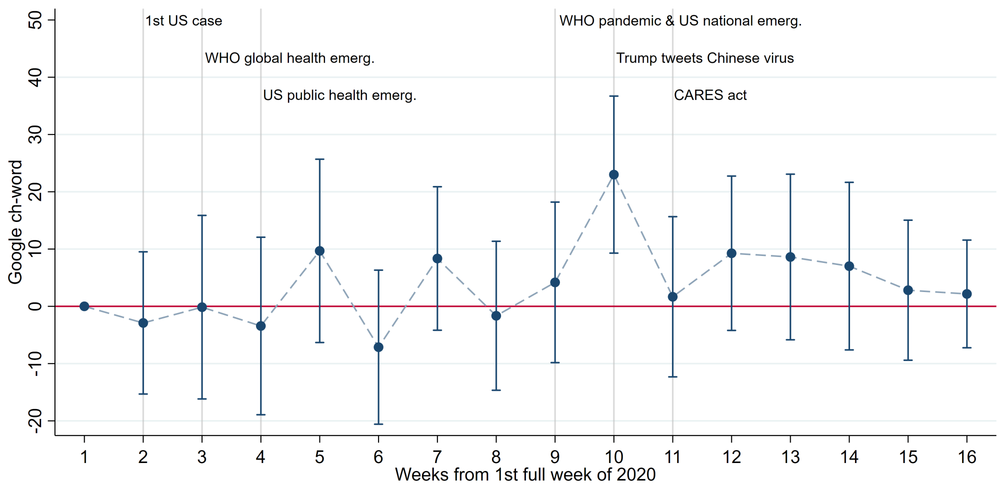

B: Google search index

Figure A2: Timeline of COVID-19 Developments and Evolution of Racial Animus

Note: The figure presents the relationship between the timeline of important COVID-19 developments and the evolution of racial animus in the United States. Panels A and B plot the estimates and 95 percent confidence intervals of the coefficients on the event dummies in equation 2 using user's likelihood of tweeting the ch-word and the racially charged Google search index as the outcome, respectively. Regressions control for week-of-year fixed effects and user fixed effects (panel A) or media market fixed effects (panel B). Standard errors are clustered by user (panel A) or by media market (panel B).

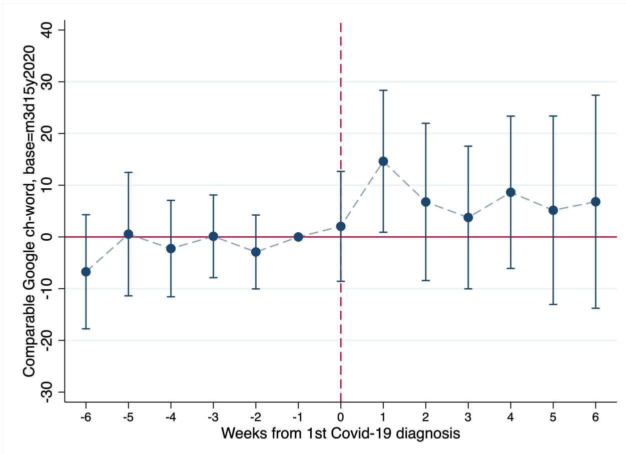

A: Benchmark 3/15/2020

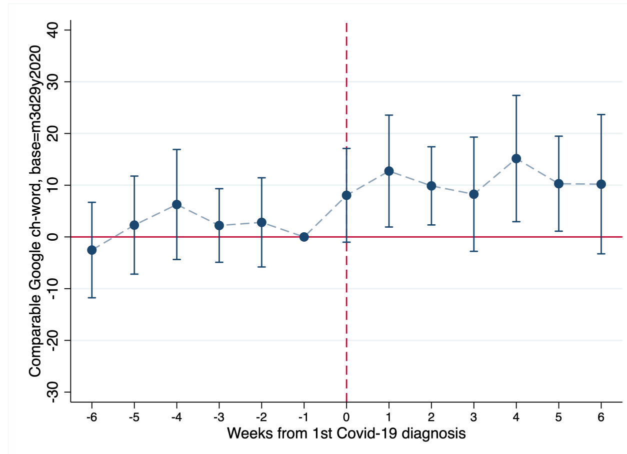

B: Benchmark 3/29/2020

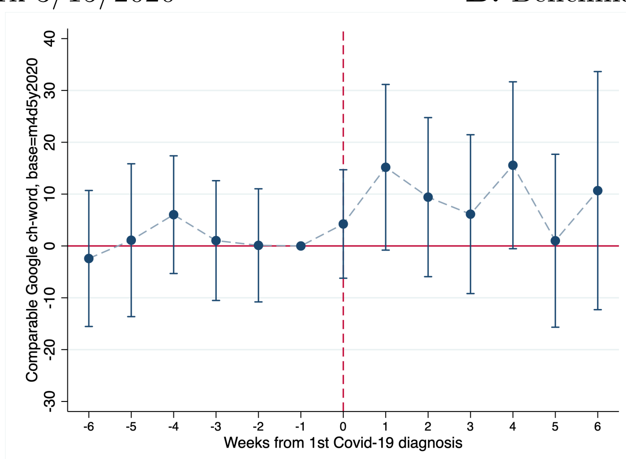

C: Benchmark 4/5/2020

Figure A3: The Effect of the First Local COVID-19 Diagnosis on Racial Animus
Rescaled Google Search Index

Notes: The figure presents the effect of the first local COVID-19 diagnosis on various *rescaled* racially charged Google search indexes. Panels A, B, and C plot the estimates and 95 percent confidence intervals of the coefficients on the event dummies in equation 3 using an area's racially charged Google search rate scaled by Huntsville-Decatur (Florence) media market's search rate on March 15, 2020, by Wilkes Barre-Scranton media market's search rate on March 29, 2020, and by Buffalo media market's search rate on April 5, 2020 as the outcome, respectively. See Appendix 1 for the definitions of these indexes. Specifications mirror those in column (1) of Table 2.

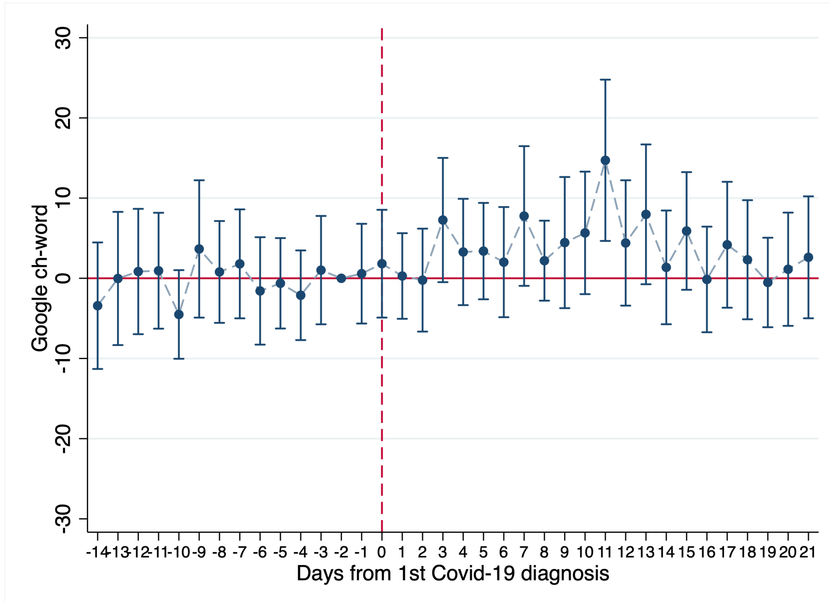

A: Google search index

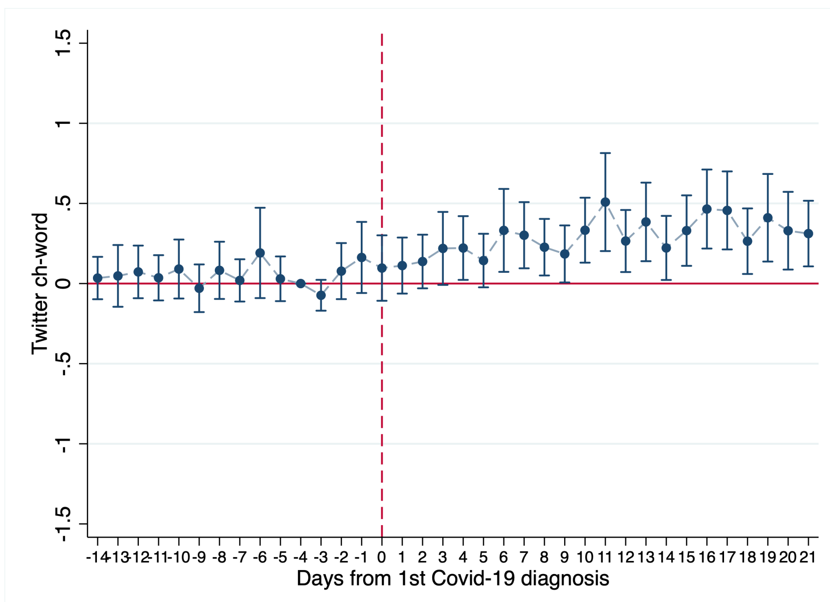

B: Twitter post index

Figure A4: The Effect of the First Local COVID-19 Diagnosis on *Daily* Racial Animus

Note: The figure presents the effect of the first local COVID-19 diagnosis on the *daily* racially charged Google search index and Twitter post index. Panels A and B plot the estimates and 95 percent confidence intervals of the coefficients on the event dummies in equation 3 using the Google search index and the Twitter post index as the outcome, respectively. Regressions control for year-month fixed effects, day-of-week fixed effects, and media market fixed effects (panel A) or county fixed effects (panel B). Standard errors are clustered by media market (panel A) or by county (panel B).

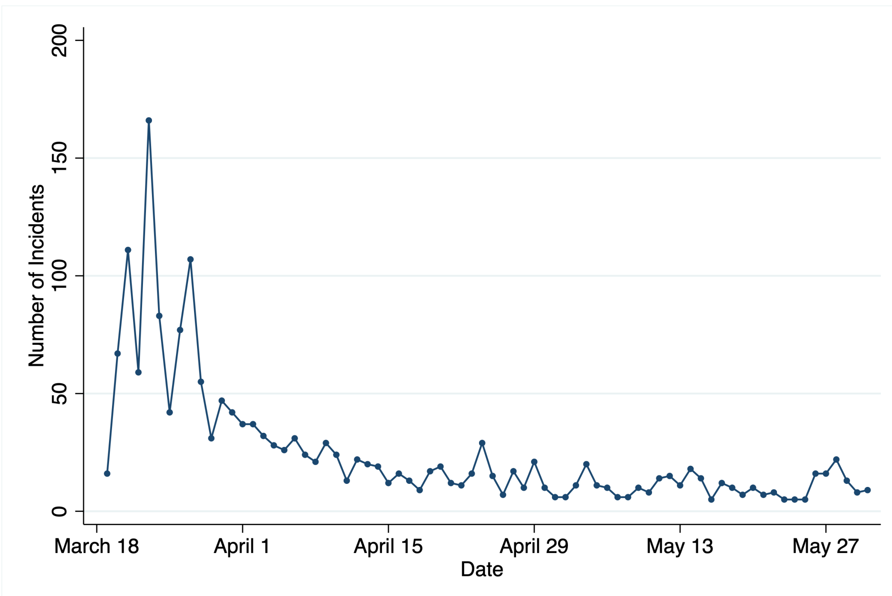

Figure A5: Self-Reported Hate Incidents in the United States

Notes: This figure presents the daily number of hate incidents from AP3CON Stop AAPI Hate Reporting system between March 19, 2020 (start of the data) and September, 2020.

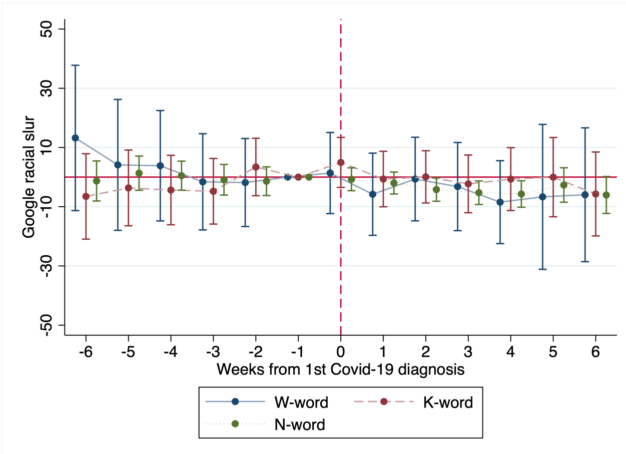

A: Google search index

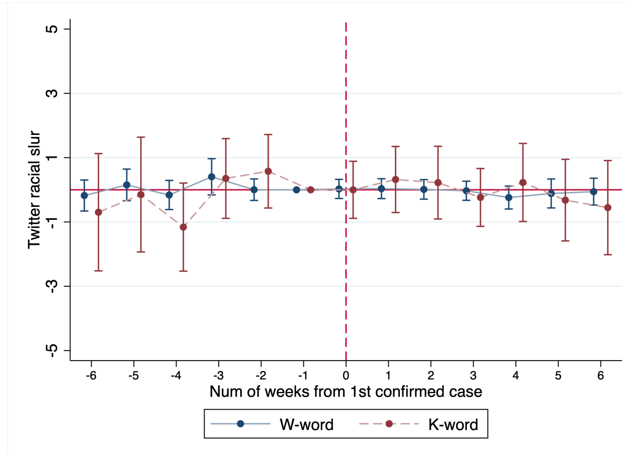

B: Twitter post index

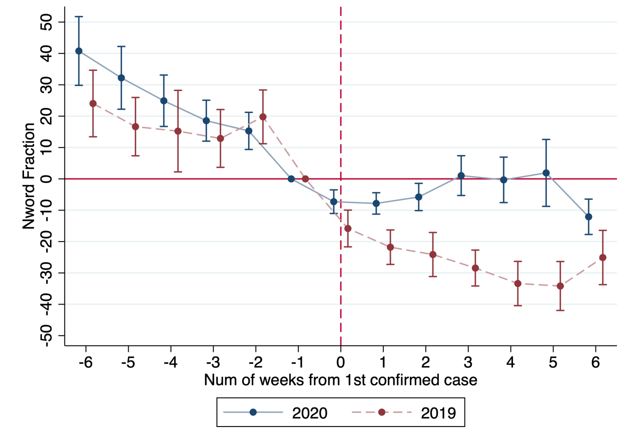

C: Twitter N-word

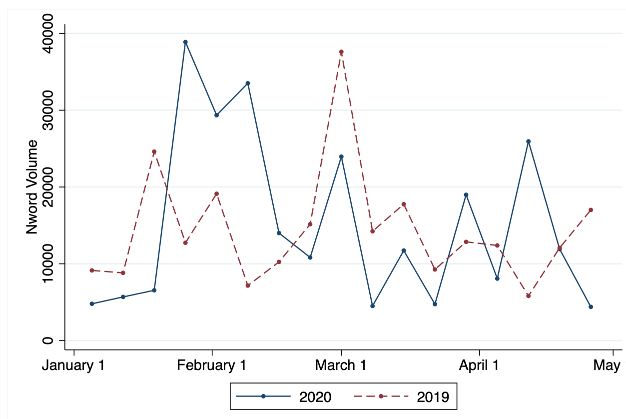

D: N-word trend

Figure A6: The Effect of the First Local COVID-19 Diagnosis on Racial Animus against *Non-Asian* Minorities

Notes: The figure presents the effect of the first local COVID-19 diagnosis on racial animus against the Hispanic, Jewish, and African American population, using the Google search index and Twitter post index for "wetback(s)", "kike(s)", and the n-word as proxies. The indexes are defined following the method outlined in section 2.1. Regression samples for the n-word, k-word, and w-word Google search indexes contain 203, 78, and 27 media markets (panel A). Regression samples for the w-word and k-word Twitter post indexes contain 599 counties (panel B). Estimates of the coefficients and the 95 percent confidence intervals of the event dummies are from estimating equation 3 using the above indexes as outcomes. We include an indicator for the week of January 26, 2020 in the regression for the n-word to control for a spike in its use due to Kobe Bryant's death and MSNBC's anchor using the n-word while reporting the news. We include an indicator for the week of February 23, 2020 in the regression for the k-word to control for a spike in its use due to the Los Angeles Dodges player Enrique ("Kiké") Hernandez's performance in that week. All other specifications in panels A and B mirror those in column (1) of Table 2 and column (1) of Table 3, respectively. Panel C plots the estimates and 95 percent confidence intervals of the coefficients on event dummies in equation 3 using the Twitter post index for the n-word between November 2019 and April 2020 (blue line) and that between November 2018 and April 2019 (red line) as the outcomes. For the regression using the 2018-2019 data, we replace the date of the first local COVID-19 diagnosis with a placebo date which shares the same day and month as the actual date in 2020 but with the year as 2019. For the regression using the 2019-2020 data, we include an indicator for the week of January 26, 2020 to control for Kobe Bryant's death on January 26, 2020 and an indicator for the week of February 9, 2020 to control for an extremely viral video tweet *unrelated* to COVID-19 but mentioning the n-word on February 10, 2020. Panel D plots time trends for the Twitter post index for the n-word in 2020 (blue line) and in 2019 (red line).

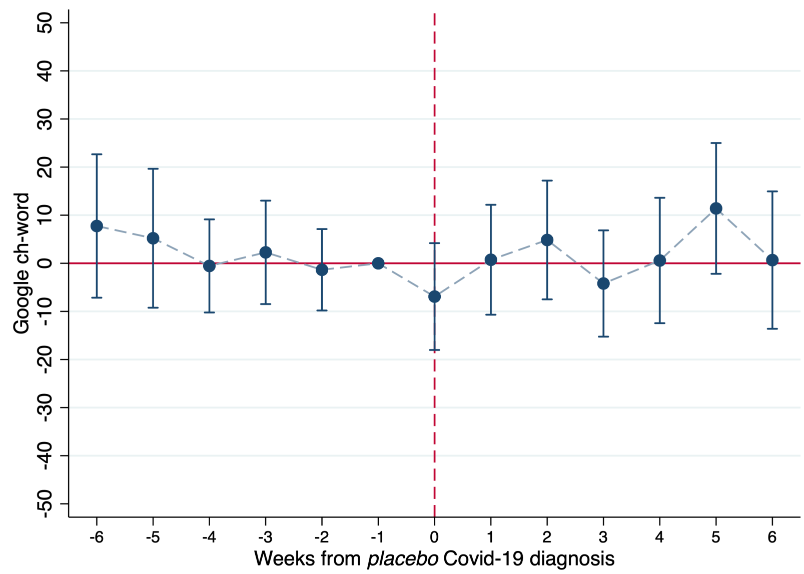

A: Google search index

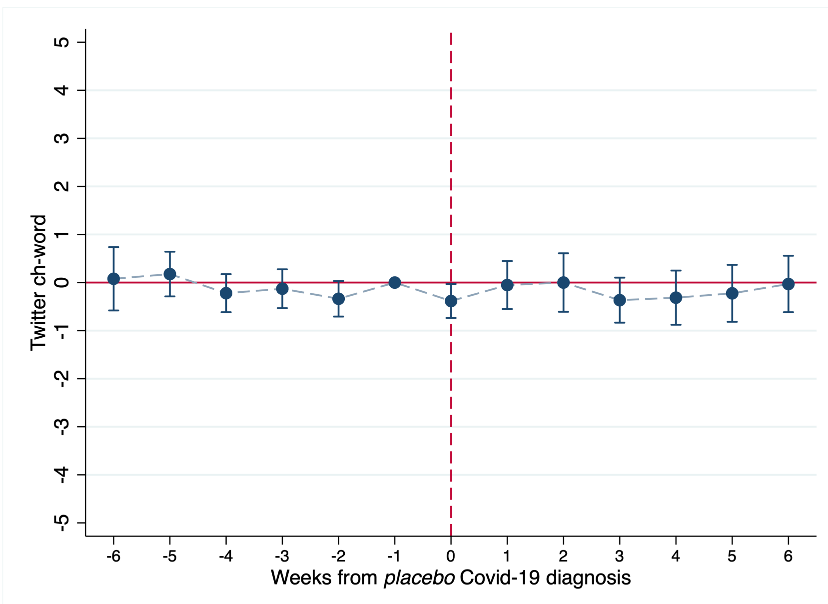

B: Twitter post index

Figure A7: The Effect of the First Local COVID-19 Diagnosis on Racial Animus Placebo Test

Note: The figure presents a placebo test for the effect of the first local COVID-19 diagnosis on the racially charged Google search index and Twitter post index. We replace the date of the first local COVID-19 diagnosis with a placebo date using the same calendar day and month of the actual diagnosis date but changing the year from 2020 to 2019. Panels A and B plot the estimates and 95 percent confidence intervals of the coefficients on the event dummies in equation 3 using the Google search index and the Twitter post index as the outcome, respectively. Specifications in panels A and B mirror those in column (1) of Table 2 and column (1) of Table 3, respectively.

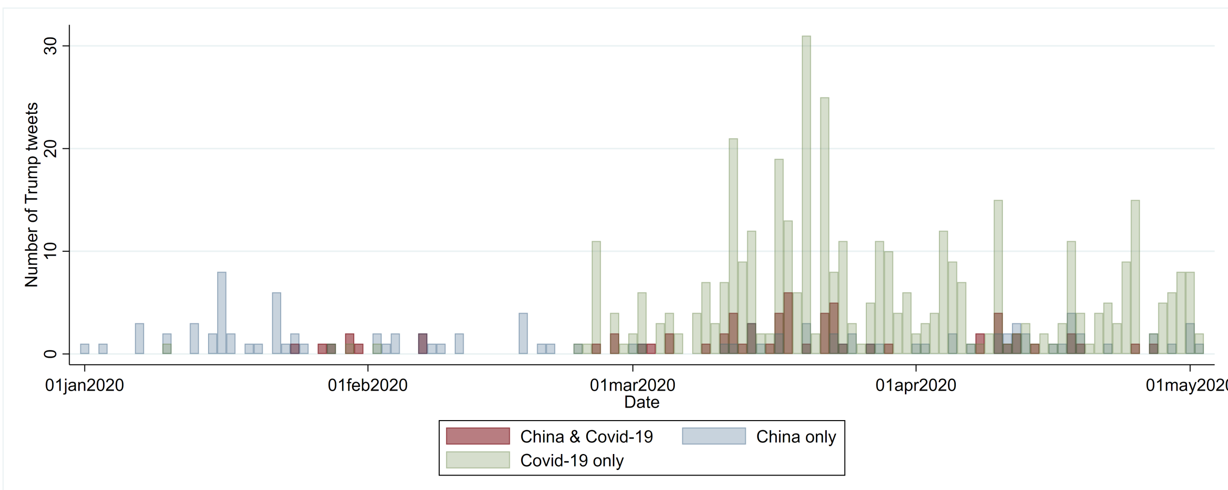

Figure A8: Number of President Trump's Tweets about China or COVID-19

Notes: This figure plots the number of President Trump's tweets by category on each day between January 1, 2020 and May 2, 2020. We categorize the president's tweets that include "china", "chinese", "huawei", "xi", "COVID", "COVID-19", "corona", "coronavirus", "virus", "epidemic", or "pandemic" into three categories: those mentioning only China (China only), only COVID-19 (COVID only), and both China and COVID-19 (China-and-COVID).

Table A1: Sample Selection - Media Markets and Counties with Google and Twitter data

	(1)	(2)	(3)
VARIABLES	Google sample	Twitter data	Twitter sample
Log(pop)	0.193***	0.142***	0.142***
	(0.036)	(0.007)	(0.007)
% Asian	0.059**	0.007	0.007
	(0.030)	(0.009)	(0.009)
% Asian^2	-0.002**	-0.001**	-0.001**
	(0.001)	(0.000)	(0.000)
% Male	0.001	-0.001	-0.001
	(0.037)	(0.002)	(0.002)
% 65+	-0.013	-0.002	-0.003
	(0.015)	(0.002)	(0.002)
% BA+	0.009	0.002	0.002*
	(0.007)	(0.001)	(0.001)
% Unemp	-0.003	0.002	0.002
	(0.020)	(0.005)	(0.005)
% VS dem-rep	-0.001	0.001***	0.001***
	(0.002)	(0.000)	(0.000)
Hate crime/1m	-0.087***	0.001	0.001
	(0.026)	(0.005)	(0.005)
Intl airport enplanement	0.006*	0.006**	0.006**
	(0.003)	(0.002)	(0.002)
Observations	205	3,111	3,111
R-squared	0.695	0.341	0.343
Outcome mean	.292	.205	.202

Notes: The table presents the sample selection in Google and Twitter data. The data are at the media market level in column (1) and at the county level in columns (2) and (3). The outcome is an indicator of having Google data in column (1), an indicator of having Twitter data in column (2), and an indicator of being in the Twitter regression sample in column (3). Note that all media markets with Google data are in the Google regression sample. "Log(pop)" is the natural log of local population estimates in 2018 from Census Bureau. "%Asian", "% Male", "% 65+", and "% BA+" are the percentage of Asians, males, population 65 years old or over, and population with Bachelor's or above degree in the local area from American Community Survey 2014-2018 five-year average. "%Unemp" is the average monthly local unemployment rate between 2014 and 2018 from the Bureau of Labor Statistics. "% VS Dem-Rep" is the difference between the Democratic and the Republican vote shares in the 2012 Presidential election from Harvard Dataverse. "Hate crime/1m" is the annual average number of anti-Asian hate crimes per million population between 2014-2018 from UCR. "Intl airport enplanement" is the international airport enplanements in 2016 according to the Federal Aviation Administration. The number of media markets and counties is less than 210 and 3141 due to missing covariates. All regressions control for state fixed effects. Standard errors in parentheses are clustered by media market in column (1) and by county in columns (2) and (3). *** p <0.01, ** p <0.05, * p <0.1.

Table A2: Summary Statistics

VARIABLE	(1) Google Sample	(2) Twitter Sample
Ch-word index		
Weekly	25.213	0.591
	(29.008)	(2.572)
Daily	6.965	0.309
	(20.316)	(2.215)
Hourly	-	0.357
	-	(0.676)
Other indexes (weekly)		
N-word	31.781	33.448
	(25.042)	(65.104)
K-word	34.183	2.970
	(26.750)	(12.705)
W-word	29.452	0.436
	(29.702)	(2.708)
Asian(s)	79.305	138.026
	(10.815)	(199.423)
Other animus measures		
Anti-Asian hate crime/1m	0.037	0.003
	(0.099)	(0.036)
Chinese restaurant visits/1m	26353	23846
	(13148)	(12328)
Total restaurant visits/1m	698341	606620
	(189770)	(220331)
Geographic unit	Media market	County
Unique geo-units	60	641

Notes: The table presents summary statistics for our main regression samples. See section 2.1 for the definitions of Google search index and Twitter post index. "Anti-Asian hate crime/1m" is the monthly anti-Asian hate crimes per million population in a media market between January 2014 and December 2018. "Chinese (or total) restaurant visits/1m" is the monthly visits to Chinese (or all) restaurants per million population in a median market between January 2018 and December 2019. All other variables are measured at the media market×time level in column (1) and at the county×time level in column (2).

Table A3: Timing of the First Local COVID-19 Diagnosis - Weeks from Jan 19, 2020

VARIABLES	(1) Google sample Weeks from Jan192020	(2) Twitter sample Weeks from Jan192020
Log(pop)	-1.499***	-0.673***
	(0.474)	(0.053)
% Asian	0.156	-0.018
	(0.212)	(0.038)
% Asian^2	-0.004	-0.001
	(0.006)	(0.001)
% Male	-1.158***	0.019
	(0.414)	(0.033)
% 65+	-0.031	0.004
	(0.067)	(0.014)
% BA+	0.048	-0.005
	(0.040)	(0.007)
% Unemp	0.623***	0.061
	(0.212)	(0.044)
% VS dem-rep	-0.022**	-0.000
	(0.009)	(0.002)
Hate crime/1m	-0.850	-0.023
	(0.790)	(0.034)
Intl airport enplanement	0.035**	-0.028
	(0.015)	(0.017)
Observations	60	630
R-squared	0.984	0.646
Outcome mean	5.983	8.12

Notes: The table presents the relationship between the timing of the first local COVID-19 diagnosis and the characteristics of the local area. The data are at the media market level in column and at the county level in column (2). The outcome is the number of weeks from the week of the first diagnosis in the US, i.e., the week of January 19, 2020. See note to Table A1 for variable definitions. The number of observations in columns (1) and (2) is smaller than 60 and 641 due to missing covariates. All regressions control for state fixed effects. Standard errors are clustered by media market in column (1) or by county in column (2). *** $p<0.01$, ** $p<0.05$, * $p<0.1$.

Table A4: Number of Media Markets and Counties by Week of the First COVID-19 Local Diagnosis

Date of Sunday	(1) Media markets	(2) Counties
Jan 19, 2020	2	
Jan 26, 2020	4	5
Feb 9, 2020	1	1
Feb 16, 2020	1	2
Feb 23, 2020	1	1
Mar 1, 2020	20	28
Mar 8, 2020	31	148
Mar 15, 2020		229
Mar 22, 2020		139
Mar 29, 2020		58
Apr 05, 2020		16
Apr 12, 2020		7
Apr 19, 2020		4
Total	60	641

Notes: The table presents the number of media markets and counties in our main regression samples by the week of their first local COVID-19 diagnoses.

Table A5: Examples of President Trump's Tweets about China or COVID-19

Category	Post	Date
China-only Tweet	"Years from now, when we look back at this day, nobody's going to remember nancy's cheap theatrics, they will remember though how president trump brought the Chinese to the bargaining table and delivered achievements few ever thought were possible." @ingrahamangle @foxnews	1/17/20
China-only Tweet	The Wall Street Journal editorial board doesn't have a clue on how to fight and win. Their views on tariffs & trade are losers for the U.S., but winners for other countries, including China. If we followed their standards, we'd have no country left. They should love sleepy joe!	4/11/20
COVID-only Tweet	The coronavirus is very much under control in the USA. we are in contact with everyone and all relevant countries. CDC & World Health have been working hard and very smart. Stock market starting to look very good to me!	2/24/20
COVID-only Tweet	I am fully prepared to use the full power of the federal government to deal with our current challenge of the coronavirus!	3/11/20
China-COVID Tweet	Just received a briefing on the Coronavirus in china from all of our great agencies, who are also working closely with china. we will continue to monitor the ongoing developments. We have the best experts anywhere in the world, and they are on top of it 24/7!	1/30/20
China-COVID Tweet	I will be having a news conference today to discuss very important news from the FDA concerning the Chinese Virus!	3/18/20
China-COVID Tweet	Just finished a very good conversation with President Xi of China. Discussed in great detail the Coronavirus that is ravaging large parts of our planet. China has been through much & has developed a strong understanding of the virus. We are working closely together. Much respect!	3/22/20

Notes: This table presents examples of President Trump's tweets mentioning China or COVID-19. We manually categorize all President Trump's tweets between January 1, 2020 and May 2, 2020 that contain any of the words "china", "chinese", "huawei", "xi", "covid", "covid-19", "corona", "coronavirus", "virus", "epidemic", or "pandemic" into three categories: those mentioning only China (China-only), only COVID-19 (COVID-only), and both China and COVID-19 (China-and-COVID).

Table A6: Relationship between Racial Animus and Tweets from Politicians and National News Outlets

VARIABLES	(1) McCarthy	(2) McConnell	(3) Pelosi	(4) Schumer	(5) CBS	(6) CNN	(7) Fox
Panel A: Twitter post index							
China-and-Covid(t)	0.0328				-0.0131	0.0278*	0.0283
	(0.0594)				(0.0237)	(0.0155)	(0.1858)
China only(t)	0.0076			0.0446	-0.0203	0.0027	
	(0.0201)			(0.0391)	(0.0205)	(0.0132)	
Covid only(t)	0.0170	-0.0213*	-0.0087	-0.0304	0.0063***	-0.0012	0.0235
	(0.0109)	(0.0128)	(0.0226)	(0.0236)	(0.0018)	(0.0015)	(0.0223)
Observations	123	123	123	123	123	123	123
R-squared	0.5118	0.5088	0.4916	0.5057	0.5383	0.5090	0.5074
Outcome mean	.3445	.3445	.3445	.3445	.3445	.3445	.3445
Panel B: Log(hate incidents)							
China-and-Covid(t)	-0.5740*				-0.0175	-0.1768	0.2794
	(0.3287)				(0.1326)	(0.1365)	(0.2438)
China only(t)	-0.0090				0.2764	0.4743**	
	(0.0734)				(0.1714)	(0.2041)	
Covid only(t)	0.0426	-0.0081	0.0077	-0.0745	0.0031	0.0075	-0.0310
	(0.0553)	(0.0265)	(0.1218)	(0.0932)	(0.0052)	(0.0059)	(0.0690)
Observations	45	45	45	45	45	45	45
R-squared	0.8522	0.8198	0.8196	0.8236	0.8348	0.8611	0.8264
Outcome mean	3.1932	3.1932	3.1932	3.1932	3.1932	3.1932	3.1932

Notes: Panel A correlates the racially charged Twitter post index with major politician and news outlets' tweets about COVID-19 or China. The data are at the media market×daily level (columns (1)-(3)) or county×daily level (columns (4)-(6)) between January 1, 2020 and May 2, 2020. Panel B correlates the natural log of the daily number of anti-Asian hate incidents nationwide from AP3CON Stop AAPI Hate Reporting system between March 19 and May 2, 2020 with major politician and news outlets' tweets about COVID-19 or China. Outcome variables are the racially charged Twitter post index (columns (1)-(3)) and the racially charged Google search index (columns (4)-(6)). A tweet is defined to be about China if it contains any of "China", "Chinese", "Huawei", or "Xi" and about Covid-19 if it contains any of "covid", "covid-19", "corona", "coronavirus", "virus", "epidemic", or "pandemic." "New diagnoses" and "New deaths" are the total daily number of new COVID-19 diagnoses and deaths in the United States calculated using the data from Johns Hopkins University Covronavirus Resource Center. Standard errors in parentheses are clustered by county (columns (1)-(3)) or by media market (columns (4)-(6)). *** p <0.01, ** p <0.05, * p <0.1.