

# Economic WARNings:

## The Impact of Negative News on Racial Animus

Adam Bestenbostel<sup>1</sup>

Abigail Peralta<sup>2</sup>

<sup>1</sup>*Texas A&M University*

<sup>2</sup>*Louisiana State University*

October 2021

### **Abstract**

Racially charged rhetoric often surrounds layoff events, with specific minorities often blamed for the loss of “American jobs.” We examine whether information about impending mass layoffs causes racial animus. Our data consist of information on mass layoff notices linked to Google Search Trends and FBI Hate Crime Statistics. We compare outcomes across areas that vary in the timing of news of impending layoffs. Results indicate an increase in both racist internet searches (1.3 percent) and hate crimes (23 percent) following layoff notices.

JEL Codes: J15, J63, K42

Keywords: animus, racism, layoffs

Declaration of interest: None

# 1 Introduction

Economic pressure brought by events such as the Great Recession (Yagan, 2019), NAFTA (Hakobyan and McLaren, 2016), rising import competition from China (Acemoglu, Autor, Dorn, Hanson, and Price, 2016), and globalization in general (Ebenstein, Harrison, McMillan, and Phillips, 2014) have resulted in some U.S. companies are reorganizing their operations, or shutting down entirely. These decisions affect a large fraction of their workforce, often resulting in mass layoffs and plant closings. News of the sudden unemployment of a large number of workers presents a significant disruption to the fabric of affected local communities. In addition, racially charged rhetoric surrounding such events often points to the shifting of “American jobs” to foreign workers (e.g. Preston, 2015; May, 2020).

In this context, linking information about impending mass layoffs to racial animus seems straightforward: as communities learn of workers losing their source of income and social status, local tensions and animosity toward certain groups rise (Falk, Kuhn, and Zweimüller, 2011). Similar responses to news have been found in criminal sentencing (e.g., Philippe and Ouss, 2018; Eren and Mocan, 2018). However, more closely related work on negative economic shocks arising from differential impacts of the Great Recession (Anderson, Crost, and Rees, 2020), and trade competition (Ortega, Di Fruscia, and Louise, 2021 and DiLorenzo, 2021), find mixed results. Anderson, Crost, and Rees (2020) and Ortega, Di Fruscia, and Louise (2021) both find that layoffs increase animus, as measured by either internet searches for racial slurs, or hate crimes against Blacks. On the other hand, DiLorenzo (2021) only finds effects on the number of hate crime groups.

The purpose of this paper is to determine the extent to which information about impending mass layoffs and plant closings lead to animosity toward racial minorities. We overcome several obstacles to answering this research question. First, aggregate economic data, such as county-level employment rates, is often not granular enough to reflect job loss shocks. In addition, disentangling the information content from the experienced job loss is difficult, because local employment data typically reflect changes in current or past local labor market situations. Second, racial animosity is also difficult to measure, largely because people are understandably hesitant to reveal “socially unacceptable attitudes” (Stephens-Davidowitz, 2014). Lastly, causal identification can be hampered by endogenously different labor market conditions across communities and by the difficulty of attributing crimes and other actions to racial animosity. To address these issues, we assemble a novel dataset containing a broad information set of notices about impending mass layoffs, and link this dataset to two different measures of racial animus.

Specifically, we leverage the timing and distribution of Worker Adjustment and Retraining Notification (WARN) Act Notices of mass layoffs and plant closings to identify their impact on animus. Using these

data has several advantages. First, we are able to determine the location and timing of impending layoff events with great detail. This is useful because it allows us to separate the impact of the news from the impact of the layoff itself, as done in Carlson (2015) in his study of the health effects of WARN Notices. Second, because some mass layoff events may affect only a small fraction of the local labor force, our data can pick up smaller labor shocks than may be seen in county-level unemployment rates. Lastly, Krolikowski and Lunsford (2020) show that WARN Notices are a strong leading indicator of local labor market indicators, such as unemployment insurance claims. We assemble data on all available WARN Notices of mass layoffs and plant closings for 23 states.<sup>1</sup> Our compiled data contains over 75,000 notices of mass layoff events that happened between the years 2004 and 2020. Although there is some dispersion, the average layoff event affects 100 workers.

Measuring racial animus is complicated by issues of underreporting and attribution. We draw on two data sources to address these difficulties: Google Trends data on searches for racial slurs and FBI Hate Crime Statistics. Google Trends data were first used to measure racial animus by Stephens-Davidowitz (2014), who showed that internet searches for the most commonly used anti-Black slur were negatively correlated with voting for President Obama in 2008. In this paper, we exploit the ability of Google Trends to also report data on internet searches for other racial slurs, specifically those commonly used to refer to Asians, Hispanics, and Whites. Google Trends can report an index of search activity for specified search terms over time at the Designated Market Area (DMA) level, which is useful because that geographic designation is likely where information about mass layoff notices also propagate.<sup>2</sup>

Although we have reason to believe that online behavior can be both harmful on its own and also reflect behavior in other spheres (Müller and Schwarz, 2020), we also corroborate our findings using data from the FBI Hate Crime Statistics program. The data contain incident-level information on crimes that are determined to have been motivated by bias against certain groups. Using these data allows us to provide further evidence of the increased harm associated with racial animus, and capture a broad picture of any potential increase in animus resulting from mass layoffs. We view these animosity measures to be linked, as recent studies suggest that anti-minority sentiment online is predictive of hate crimes against minorities (Müller and Schwarz, 2021, and Müller and Schwarz, 2020). In using both internet searches and hate crimes, we aim to measure both racially targeted harmful rhetoric and criminal behavior.

Our empirical strategy is to compare how the values of our animosity measures change in counties where mass layoff notices occur relative to other counties. Because mass layoff notices can occur several times for

---

<sup>1</sup>These are 60-day advance notices that large firms (>100 employees) are required to give to workers ahead of a layoff event affecting more than 50 workers at a single employment site.

<sup>2</sup>We have made the editorial decision to not write out the racial slurs used in this paper. The slurs we use come from Anderson and Lepore (2013).

each county during our study period, the main effect we focus on is the contemporaneous effect of the mass layoff notice on animosity. Our approach is similar in spirit to that taken by other studies that examine the effect of news shocks on a variety of outcomes (e.g. Eren and Mocan, 2018; Philippe and Ouss, 2018 – criminal sentencing, and Carlson, 2015 – birth outcomes). We also show that the contemporaneous effect captures most of the treatment effect, as effects dissipate within two months. The identifying assumption is that absent mass layoffs, counties would have experienced changes in these measures of animosity similar to what other counties experienced. We estimate falsification exercises to show evidence for this assumption. In other specifications, we also account for state-month shocks and county-specific time trends.

Estimates indicate a significant increase in racial animosity following WARN Notices of impending mass layoffs. In particular, we estimate a 1.4 percent increase in the search rate for racial slurs. While this outcome does not necessarily imply animosity targeted at some group, we also estimate a 23 percent increase in the number of hate crimes committed within a month of the WARN Notice. Thus, we interpret our results to mean that there is a broad increase in racial animosity with real, harmful effects.

Importantly, we do not find any evidence of reverse causality – future layoffs affecting animosity measures today. We also do not find any evidence of increased internet searches for placebo terms like “weather,” or a slur commonly used to refer to rural Whites. Taken together, these reinforce our interpretation that it is the mass layoff notices that cause the increase in racial animosity toward minorities, and that we are not simply picking up increases in general animosity across groups. Our results are all the more noteworthy because the data we leverage exists, by definition, to soften the blow of mass layoffs. Yet we still estimate an increase in internet searches for racial slurs. Combined with our hate crimes estimates, we conclude that information about impending mass layoffs cause an increase in harmful rhetoric and criminal behavior toward minorities. Our estimates are robust to including controls for state-month shocks, time-varying demographic characteristics, and to allowing counties to follow different trends over time.

Our results contribute to a nascent literature that seek to causally identify the effects of various local labor shocks on racial animus in the US (Anderson, Crost, and Rees, 2020; DiLorenzo, 2021; Ortega, Di Fruscia, and Louise, 2021). These papers each focus on different reasons for layoffs. Anderson, Crost, and Rees (2020) focus on layoffs due to the Great Recession, exploiting cross-state variation in pre-Recession sector composition. Ortega, Di Fruscia, and Louise (2021), who focus on import exposure to China, and DiLorenzo (2021), who uses the universe of layoffs reported to the Trade Adjustment Assistance Database, both focus on layoffs due to trade. Our paper complements these studies by drawing on a wider and more inclusive set of mass layoffs and by separating the timing of the information about impending layoffs from the layoff events.

Through our use of internet searches from Google Trends to corroborate our hate crime results, we also join an emerging literature that leverage novel internet and social media data to measure individual attitudes, and how they relate to observed behavior. Anderson, Crost, and Rees (2020) use internet searches for a commonly used anti-Black slur, from Google Trends, to measure animus. Similarly, a number of papers have used data from Twitter to examine the impacts of social media personalities and access on political behavior (e.g., Fujiwara, Müller, and Schwarz, 2021; Giavazzi, Iglhaut, Lemoli, and Rubera, 2020; Müller and Schwarz, 2020) and Facebook access on hate crimes (Müller and Schwarz, 2021).

This paper proceeds as follows. Section 2 presents our data on layoff notices and racial animus. Section 3 lays out our empirical strategy and robustness checks. In Section 4, we present our estimates of the effect of mass layoff notices on our two measures of racial animus. In Section 5, we show that our results are qualitatively robust to using alternative specifications. Section 6 concludes.

## 2 Data

### 2.1 Layoff Data from WARN Notices

The Worker Adjustment and Retraining Notification (WARN) Act of 1988 provides protections against sudden termination for employees of large firms. Its main provision is the requirement of a 60-day notice to workers ahead of a planned mass layoff or plant closing. The objective of the advance notice is to give workers time to find another job or to enter retraining programs (Bartell, 2001). Firms typically file a report with the State Dislocated Worker Unit.<sup>3</sup>

We construct our data from the WARN Notices that firms file with the relevant State Dislocated Worker Units. The notices include the employment site address, layoff or plant closing date, WARN Notice date, and the number of affected workers. Other information is also sometimes included, but data availability is inconsistent, so we focus on these specific variables. The assembled data contain WARN notices for 75,574 mass layoff events, affecting an average of 105 workers each event. On average, counties experience around 6 mass layoff events during the study period, with the affected workers representing a small fraction of the local labor force, 0.3 percent. We do not observe whether the mass layoffs actually take place, or whether they involve fewer workers than originally stated in the WARN Notice. We take 60 days after the WARN Notice to be the layoff date, and consider the number of affected workers stated in the WARN Notice to be the number of laid off workers.

### 2.2 Google Trends

Google Trends is a tool that provides an index of search activity for specified search terms. We follow (Stephens-Davidowitz and Varian, 2014) and (Anderson, Crost, and Rees, 2020) in using searches for a

---

<sup>3</sup>A list of coordinators for these offices can be found at: <https://www.dol.gov/agencies/eta/layoffs/contact>

commonly used anti-Black slur as a measure of racial animus. We also obtain internet search data for the slurs used to refer to Asians, Hispanics, Whites, and the word “weather.” The last two are placebo search terms where we do not expect to see an effect from mass layoffs. Since the most smallest geographic level available from Google Trends is the DMA (Designated Market Area), we extract data at the DMA-month level. We then link the DMA-month internet searches data to the county-level layoffs data. Because DMAs are larger than counties, we assign the internet search activities observed at the DMA level to all counties within that DMA, and then weight by county-level population.

The data that we obtain from Google Trends are an index – for a chosen geography, it measures the fraction of searches that include the specified search term relative to the total search volume at that time. In essence, the raw data from Google Trends is a relative, rather than absolute, measure of search activity. To be able to use this data in our analysis, we follow previous studies Anderson et al. (2020) and use the logarithm of the reported search rate.

The data have a few limitations, which we account for in this paper. First, Google Trends data are drawn from a sample of all Google searches, using searches that are cached each day. The reported search index are also averaged to the nearest integer. We account for these potential issues by pulling data from Google Trends five times on separate days, and then taking the average of the reported search indices. Second, Google Trends data have a privacy threshold – if search volume for a given area during the specified time period is below the threshold, Google Trends will report a zero. This is one reason we limit to the more commonly-used racial slurs. Finally, Google Trends reports the search activity for racial slurs, but we do not directly observe the intent behind the searches. Using FBI data on hate crimes allow us to speak to race-motivated behavior.

### **2.3 FBI Hate Crime Statistics**

Our second measure of animus is hate crime data from the FBI Hate Crime Statistics, part of the Uniform Crime Reporting program. Hate crimes are defined according to the Hate Crime Statistics act as crimes motivated by bias based on “race, gender or gender identity, religion, disability, sexual orientation, or ethnicity.” Hate crimes reported to the FBI include both violent and property crimes. The most prevalent hate crimes tend to be destruction/damage/vandalism, intimidation, and simple assault (Masucci and Langton, 2017).<sup>4</sup>

Incident-level hate crime data, reported by law enforcement agencies to the FBI, are available for offenses that occurred from 1991 onward. The data contain the date of the incident, the nature of the offense, and

---

<sup>4</sup>More information about hate crimes, including the types of biases considered, can be found at: <https://www.fbi.gov/services/cjis/ucr/hate-crime>

information about the offender and victim. We use the Law Enforcement Agency Identifiers Crosswalk to link the hate-crime data to the county-level layoffs data.<sup>5</sup>

### 3 Empirical Strategy

We estimate the effect of mass layoffs on the various measures of animosity using the following general specification:

$$y_{it} = \alpha_i + \alpha_t + \beta \cdot X_{it} + \gamma_0 \cdot Treated_{it} + \gamma_1 \cdot Treated_{i,t-1} + \gamma_2 \cdot Treated_{i,t-2} + \epsilon_{it} \quad (1)$$

Where  $y_{it}$  is the value of the animosity measure for county  $i$  in month  $t$ ,  $\alpha_i$  and  $\alpha_t$  are county and month fixed effects, respectively,  $X_{it}$  are other control variables. The coefficient of interest is  $\gamma_0$ , the coefficient on the indicator variable  $Treated_{it}$ , which takes on a value of 1 if area  $i$  had a WARN Notice in month  $t$ , and 0 otherwise. The coefficients  $\gamma_1$ , and  $\gamma_2$ , on the indicator variables  $Treated_{i,t-1}$  and  $Treated_{i,t-2}$ , are also of interest. The indicator variable  $Treated_{i,t-1}$  takes on a value of 1 if area  $i$  had a WARN Notice in month  $t - 1$ , while  $Treated_{i,t-2}$ , takes on a value of 1 if area  $i$  had a WARN Notice in month  $t - 2$ . Including these indicator variables allows us to determine whether a WARN Notice that occurred in period  $t$  also had lagged effects over the next two months.

As discussed in Section 2, we measure animosity using both internet searches for racial slurs and hate crimes. For the analysis using internet searches for racial slurs, we define  $y_{it}$  to be the log of the search index for any racial slur.<sup>6</sup> For the analysis using hate crimes, we define  $y_{it}$  to be the number of hate crimes in county  $i$  in that month.

The county fixed effects  $\alpha_i$  account for unobserved, time-invariant, differences across counties that may drive differences in hate crimes. The month fixed effects  $\alpha_t$  account for unobserved, time-varying shocks that affect all counties similarly. In most specifications, we also account for state-month shocks using a set of state by month fixed effects, which subsume the month fixed effects. The addition of state by month fixed effects account for time-varying shocks that may be affecting counties within each state similarly. When all these fixed effects are included in the estimation, our identifying assumption is that the change in animosity measures observed in the unaffected counties in the same state provide a valid counterfactual for the change that would have been experienced in the counties affected by the mass layoff, had the mass layoff not occurred. All standard errors are clustered at the county level. For analyses using internet searches from Google Trends, where we spread DMA-level outcome values to the component counties, we also weight observations by county population.

---

<sup>5</sup><https://www.icpsr.umich.edu/web/NACJD/studies/35158>

<sup>6</sup>The set of racial slurs we consider are a common slur for Blacks, a slur for Asians, and two slurs for Hispanics or Latinos.

We also conduct a falsification exercise to demonstrate that future layoffs do not affect past animosity measures. The equation we estimate is:

$$y_{it} = \alpha_i + \alpha_t + \beta \cdot X_{it} + \gamma_0 \cdot Treated_{it} + \gamma_1 \cdot Treated_{i,t-1} + \gamma_2 \cdot Treated_{i,t-2} + \sum_{n=1}^6 \zeta_n \cdot Treated_{i,t+n} + \epsilon_{it} \quad (2)$$

where the coefficients  $\zeta_n$  measure the effect of future layoffs, in periods  $t + n$ , on animosity measures in period  $t$ . Our empirical strategy requires that our estimated  $\zeta_n$ s be zero.

Finally, because our main estimating equation, Equation 1, implicitly assumes that the effects of mass layoffs only occur in the immediate aftermath of the WARN Notice, we also present estimates from specifications where we allow for lagged treatment effects beyond two months. This is to examine whether treatment effects indeed decay within the first two months of the WARN Notice date. We estimate the equation below, where the coefficients  $\gamma_n$  would indicate lagged effects beyond the first two months, and each additional lagged effect changes the reference period:

$$y_{it} = \alpha_i + \alpha_t + \beta \cdot X_{it} + \gamma_0 \cdot Treated_{it} + \gamma_1 \cdot Treated_{i,t-1} + \gamma_2 \cdot Treated_{i,t-2} + \sum_{n=3}^6 \gamma_n \cdot Treated_{i,t-n} \epsilon_{it} \quad (3)$$

## 4 Results

We first estimate a falsification-type exercise. Evidence against the validity of our approach would be if we see that future layoffs are predictive of animosity “today.” In Figures 1 and 2, we report the estimated coefficients from this falsification exercise, on internet searches and hate crimes, respectively. Figure 1 shows, in particular, layoffs that will occur 6 months, or even as soon as 1 month from “today” have no effect on searches containing racial slurs “today.” We show a similar pattern in Figure 2, where the animosity measure is hate crimes.

We then move to estimating the effect of mass layoffs on animosity measures. Table 1 reports the estimated coefficients  $\gamma_0$ , and the lagged effects  $\gamma_1$  and  $\gamma_2$ , where the animosity measure used is the log of the search index for any racial slur. Column 1 reports estimates from a base specification that includes only state and month fixed effects. In Column 2, we include state-by-month fixed effects, which subsume the month fixed effects. In Columns 3 and 4, we add county linear time trends and time-varying demographic controls to the specification in Column 2. Our preferred specification is this final specification, which allows for state-month shocks as well as for counties to follow different trends over time. Across all four columns, we



consistently estimate approximately a 1.3 percent increase in internet searches for racial slurs in the same month as the mass layoff notice. The effect does not persist beyond the first month.

In Columns 5 and 6 of Table 1, we estimate Equation 1 again but use placebo animosity measures as dependent variables. Specifically, we examine whether mass layoffs affect internet searches for the placebo search terms “weather” and the racial slur for Whites. Since the search term “weather” is very unlikely to be affected by news of mass layoffs, we should not expect to find effects on this outcome. Similarly, given that much of the animus in the US is targeted at racial minorities, we should also not expect to find effects on the rate of searches for a slur typically used to refer to Whites. That we report precisely estimated null effects on the searches for both of these words provides further evidence that our main results identify the effects of mass layoffs.

In Table 2, we break down the treatment effect by the specific racial slur used. Column 1 reports the effect of mass layoff notices on searches for the common anti-Black slur, Column 2 reports the effect on searches for an anti-Asian slur, and Columns 3 and 4 report the effects on searches for two anti-Hispanic slurs. Table 2 tells us that most of the effects we see in Table 1 are driven by increases in animus against Black and Hispanic people. We do not find evidence of increased animosity toward Asians, as the main effect is small and precisely estimated.

To corroborate our findings on internet searches for racial slurs, we also estimate Equation 1 using hate crimes as the outcome variable. Table 3 reports the estimates of  $\gamma_0$ , and the lagged effects  $\gamma_1$  and  $\gamma_2$  from Equation 1, where the animosity measure used is the number of hate crimes. In Table 3, Columns 1 to 4 use the same specifications as in Table 1. Estimates from our preferred specification in Column 4 indicate that the notice of mass layoffs leads to a statistically significant 0.05 increase in the number of hate crimes in that county and month, and that this effect persists for 2 months. With a baseline number of hate crime incidents per county-month of 0.22 (approximately one every five months), this represents a 23 percent increase.

Taken together, our estimates show that mass layoff notices lead to immediate increases in racial animus, that then dissipate within two months. Moreover, the increase in animus causes direct harm to minorities, as evidenced by the estimated increases in hate crime. We also show that the increase in racial animus is largely directed toward Blacks and Hispanics.

## 5 Robustness

The estimates presented in Tables 1, 2, and 3 all implicitly assume that lagged treatment effects only last a maximum of two months. In this section, we relax this assumption and allow for effects to last for 3 to 6

months. Table 4 presents the estimates of Equation 3, for both internet searches (Panel A) and hate crimes (Panel B). Each row  $n$  presents an estimate of  $\gamma_n$  from Equation 3. The first Column in each panel reproduces the estimates from Column 4 of Tables 1 and 3.

The estimates presented in Panel A consistently show that the main treatment effect of the mass layoff notice is an increase in searches for racist slurs in the same month as the notice. The estimates for lagged effects two months or more after the month of the layoff notice are small and precisely estimated. Notably, allowing for these lagged effects does not change the magnitude of our estimated treatment effect: Column 5, which allows for lagged effects up to 6 months, shows that a WARN Notice in month  $t$  increases racist internet searches 1.5 percent in the same month, which is comparable to the main estimate of 1.3 percent.

Similarly, Panel B also shows that the effects of mass layoff notices on hate crimes do not last beyond 2 months. By and large, the estimates of  $\gamma_n$  for  $n > 3$  are significantly smaller than the main contemporaneous effect, with the estimates of  $\gamma_5$  and  $\gamma_6$  being precisely estimated at near zero. As in Panel A, the main effect is also not affected by the inclusion of lagged effects up to 6 months: the estimate in Column 5 is a 0.0465 contemporaneous increase in hate crime incidents, comparable to the main estimate of 0.05.

Finally, in Columns 6 to 8 of Panel B, we also explore different functional forms of Equation 1. In Column 6, we use the logarithm of  $1 +$  the number of hate crime incidents as the dependent variable, while in Column 7, we use the inverse hyperbolic sine of the number of hate crime incidents. Doing so allows us to interpret the estimate of  $\gamma$  as the percent change in the number of hate crimes resulting from the mass layoff notice. The estimated coefficients indicate that mass layoff notices caused an increase in hate crime of 1 to 1.3 percent, and are significant at the 10 percent level. In Column 8, we use a Poisson model on the number of hate crime incidents and find that mass layoff notices caused an increase in hate crime of 5 percent. However, for the estimation to converge successfully, this column uses approximately two thirds the observations as in previous result, so this sample may not be as representative. Still, the Poisson estimate is positive and significant at the 10% level.

## 6 Discussion and Conclusion

This paper presents evidence of sharp increases in racial animosity following mass layoffs notices. By exploiting the advance notice that the WARN Act requires firms to provide to affected workers, we are able to disentangle the effect of the notice of the mass layoff from the layoff event itself. Our findings show that mass layoff notices lead to an increase in internet searches for racial slurs, indicating heightened racial animus. We put this finding in context by showing a similar effect on hate crime incidents. Thus, while some of the estimated effect on online behavior might just reflect people expressing harmless racist


sentiments, or looking up racial slurs, our estimates of the effect on hate crimes indicate that at least part of what we pick up is harmful racial animus directed at minorities. With mass layoff events often a small fraction of a locality’s labor force, we interpret our results to mean that the information content in mass layoff notices has its own effect on racial animus, separate from the effect of the layoff event itself.

We explicitly do not attempt to conduct a cost-benefit analysis of providing the type of advance notice that the WARN Act requires, nor do we focus on its unintended consequences. Rather, we only exploit the timing and information content in the WARN Act layoff notices in our empirical strategy. Other studies document effects of WARN Act notices on educational choice (Acton, 2020) and health outcomes Carlson (2015). We note that although the WARN Notices allow us to separate the information component from the actual layoff event, it is possible that in a counterfactual setting where mass layoffs occur without warning, racial animosity will increase around the time of the layoff event, as would seem to be implied by related studies on economic downturns (Anderson, Crost, and Rees, 2020; Ortega, Di Fruscia, and Louise, 2021; DiLorenzo, 2021).

Our contribution is that we show that information about impending mass layoffs by itself is enough to generate increases in racial animosity. We show that this increase in racial animosity manifests itself both in broad online behavior, and in observed criminal behavior. In an era where information is transmitted much faster than before, and through ever-changing networks and media, it is important to be able to anticipate second and third-order effects of negative economic shocks. This knowledge, combined with the advance notice that WARN Notices provide, can allow local communities to prepare safeguards for their minority populations during times of economic distress.


## Exhibits

Figure 1: The Effect of Mass Layoff Notices in the “Future” on Internet Searches for Racial Slurs “Today”



This figure shows a falsification exercise estimating the impact of layoffs in periods  $t + 1, \dots, t + 6$  on outcomes in period  $t$ . I.e. the effect of layoffs ‘tomorrow’ on hate crimes ‘today.’ The dashed line represents the estimated treatment main effect (from layoffs in period  $t$  on outcomes in period  $t$ ) from column 4 of the corresponding table, and estimates in each month are shown with 95% confidence intervals.

Figure 2: The Effect of Mass Layoff Notices in the “Future” on Hate Crimes “Today”



This figure shows a falsification exercise estimating the impact of layoffs in periods  $t + 1, \dots, t + 6$  on outcomes in period  $t$ . I.e. the effect of layoffs ‘tomorrow’ on hate crimes ‘today.’ The dashed line represents the estimated treatment main effect (from layoffs in period  $t$  on outcomes in period  $t$ ) from column 4 of the corresponding table, and estimates in each month are shown with 95% confidence intervals.

Table 1: The Impact of Mass Layoff Notices on Internet Searches for Racial Slurs and Placebo Words

	Any racial slur				Placebo searches	
	1	2	3	4	'redneck'	'weather'
Main effect	0.0162*** (0.0060)	0.0136*** (0.0048)	0.0132*** (0.0043)	0.0133*** (0.0044)	0.0060 (0.0054)	0.0033 (0.0020)
1 mo. lagged	0.0046 (0.0063)	0.0020 (0.0058)	0.0017 (0.0061)	0.0018 (0.0062)	-0.0055 (0.0058)	0.0015 (0.0021)
2 mo. lagged	-0.0054 (0.0055)	-0.0052 (0.0051)	-0.0053 (0.0052)	-0.0052 (0.0052)	0.0077 (0.0060)	0.0006 (0.0022)
Observations	220,158	220,158	220,158	220,158	220,158	220,158
State-by-month FE		Y	Y	Y	Y	Y
County linear time trend			Y	Y	Y	Y
Demographic controls				Y	Y	Y

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

This table shows estimates for the effect of a WARN notice on searches for any racist term, including a slur for blacks known colloquially as the n word, a term for Asians, and two terms for Hispanics. Standard errors clustered at the county level where treatment is assigned.

Table 2: The Impact of Mass Layoff Notices on Internet Searches for Specific Racial Slurs

	Slur for Blacks 1	Slur for Asians 2	Slur for Hispanics 3	Alternative Slur for Hispanics 4
Main effect	0.0197*** (0.0064)	0.0058 (0.0100)	0.0257*** (0.0095)	0.0021 (0.0104)
1 mo. lagged	-0.0012 (0.0066)	-0.0149 (0.0094)	0.0030 (0.0108)	0.0204* (0.0121)
2 mo. lagged	0.0058 (0.0074)	-0.0206** (0.0094)	-0.0017 (0.0069)	-0.0043 (0.0111)
Observations	220,158	220,158	220,158	220,158
State-by-month FE	Y	Y	Y	Y
County linear time trend	Y	Y	Y	Y
Demographic controls	Y	Y	Y	Y

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

This table shows estimates for the effect of a WARN notice on searches for any racist term, including a slur for blacks known colloquially as the n word, a term for Asians, and two terms for Hispanics. Standard errors clustered at the county level where treatment is assigned.

Table 3: The Impact of Mass Layoff Notices on Hate Crime

	1	2	3	4
Main effect	0.0478 (0.0403)	0.0425 (0.0349)	0.0505** (0.0229)	0.0505** (0.0230)
1 mo. lagged	0.0346 (0.0397)	0.0308 (0.0332)	0.0392** (0.0189)	0.0393** (0.0190)
2 mo. lagged	0.0205 (0.0430)	0.0142 (0.0373)	0.0224 (0.0195)	0.0224 (0.0196)
Observations	220,158	220,158	220,158	220,158
Outcome mean	.221	.221	.221	.221
State-by-month FE		Y	Y	Y
County linear time trend			Y	Y
Demographic controls				Y

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

Standard errors clustered at the county level where treatment is assigned.



Table 4: Robustness to Extending the Treatment Window, and Alternative Specifications

	1	2	3	4	5	6	7	8
Panel A: Search trends								
Main effect	0.0133*** (0.0044)	0.0133*** (0.0042)	0.0134*** (0.0043)	0.0140*** (0.0042)	0.0145*** (0.0042)			
1 mo. lagged	0.0018 (0.0062)	0.0019 (0.0061)	0.0028 (0.0062)	0.0031 (0.0062)	0.0031 (0.0061)			
2 mo. lagged	-0.0052 (0.0052)	-0.0049 (0.0052)	-0.0045 (0.0052)	-0.0042 (0.0051)	-0.0031 (0.0050)			
3 mo. lagged		0.0065 (0.0052)	0.0061 (0.0051)	0.0062 (0.0051)	0.0076 (0.0052)			
4 mo. lagged			-0.0012 (0.0051)	-0.0006 (0.0051)	-0.0001 (0.0050)			
5 mo. lagged				-0.0015 (0.0053)	-0.0009 (0.0053)			
6 mo. lagged					-0.0085 (0.0052)			
Observations	220,158	218,399	216,640	214,881	213,122			
State-by-month FE	Y	Y	Y	Y	Y			
County linear time trend	Y	Y	Y	Y	Y			
Demographic controls	Y	Y	Y	Y	Y			
Panel B: Hate crime						log(y+1)	asinh(y)	Poisson
Main effect	0.0505** (0.0230)	0.0482** (0.0225)	0.0463** (0.0225)	0.0458** (0.0225)	0.0465** (0.0227)	0.0100* (0.0056)	0.0126* (0.0072)	1.0533** (0.0264)
1 mo. lagged	0.0393** (0.0190)	0.0351* (0.0189)	0.0363* (0.0189)	0.0341* (0.0193)	0.0322* (0.0193)	0.0078* (0.0045)	0.0097* (0.0058)	1.0325 (0.0219)
2 mo. lagged	0.0224 (0.0196)	0.0194 (0.0194)	0.0203 (0.0188)	0.0205 (0.0188)	0.0183 (0.0190)	0.0054 (0.0045)	0.0069 (0.0058)	1.0246 (0.0202)
3 mo. lagged		0.0252 (0.0191)	0.0238 (0.0192)	0.0233 (0.0196)	0.0219 (0.0196)			
4 mo. lagged			0.0454** (0.0184)	0.0450** (0.0187)	0.0442** (0.0188)			
5 mo. lagged				0.0016 (0.0206)	0.0013 (0.0204)			
6 mo. lagged					0.0035 (0.0185)			
Observations	220,158	218,399	216,640	214,881	213,122	220,158	220,158	146,609
Outcome mean	.221	.22	.22	.219	.219			
State-by-month FE	Y	Y	Y	Y	Y	Y	Y	Y
County linear time trend	Y	Y	Y	Y	Y	Y	Y	Y
Demographic controls	Y	Y	Y	Y	Y	Y	Y	Y
Incidence-rate Ratio								Y

\*  $p < 0.10$ , \*\*  $p < 0.05$ , \*\*\*  $p < 0.01$

This table shows estimates for the effect of a WARN Notice on searches for any racist term (Panel A), and hate crimes (Panel B). Each column presents results from separate regressions. Columns 1 to 5 expand the time window considered as “treated”, while Columns 6 to 8 explore different functional forms for estimating the effect of WARN Notices on hate crimes. Note: the Poisson estimate in the last column is presented as an incidence-rate ratio.

## References

- Daron Acemoglu, David Autor, David Dorn, Gordon H Hanson, and Brendan Price. Import competition and the great us employment sag of the 2000s. *Journal of Labor Economics*, 34(S1):S141–S198, 2016.
- Riley Acton. Community college program choices in the wake of local job losses. 2020.
- D Mark Anderson, Benjamin Crost, and Daniel I Rees. Do economic downturns fuel racial animus? *Journal of Economic Behavior & Organization*, 175:9–18, 2020.
- Luvell Anderson and Ernie Lepore. Slurring words. *Noûs*, 47(1):25–48, 2013.
- Laura B Bartell. Why warn—the worker adjustment and retraining notification act in bankruptcy. *Bankr. Dev. J.*, 18:243, 2001.
- Kyle Carlson. Fear itself: The effects of distressing economic news on birth outcomes. *Journal of health economics*, 41:117–132, 2015.
- Matthew DiLorenzo. Trade layoffs and hate in the united states. *Social Science Quarterly*, 102(2):771–785, 2021.
- Avraham Ebenstein, Ann Harrison, Margaret McMillan, and Shannon Phillips. Estimating the impact of trade and offshoring on american workers using the current population surveys. *Review of Economics and Statistics*, 96(4):581–595, 2014.
- Ozkan Eren and Naci Mocan. Emotional judges and unlucky juveniles. *American Economic Journal: Applied Economics*, 10(3):171–205, 2018.
- Armin Falk, Andreas Kuhn, and Josef Zweimüller. Unemployment and right-wing extremist crime. *Scandinavian Journal of Economics*, 113(2):260–285, 2011.
- Thomas Fujiwara, Karsten Müller, and Carlo Schwarz. The effect of social media on elections: Evidence from the united states. Technical report, National Bureau of Economic Research, 2021.
- Francesco Giavazzi, Felix Iglhaut, Giacomo Lemoli, and Gaia Rubera. Terrorist attacks, cultural incidents and the vote for radical parties: Analyzing text from twitter. Technical report, National Bureau of Economic Research, 2020.
- Shushanik Hakobyan and John McLaren. Looking for local labor market effects of nafta. *Review of Economics and Statistics*, 98(4):728–741, 2016.
- Pawel Michal Krolikowski and Kurt G Lunsford. Advance layoff notices and labor market forecasting. 2020.
- Madeline Masucci and Lynn Langton. Hate crime victimization, 2004-2015. *NCJ*, 250653, 2017.
- Caroline May. U.s. graduates expect mass layoffs as companies keep hiring h-1b visa workers. *Breitbart*, Mar 2020. URL <https://www.breitbart.com/politics/2020/03/23/u-s-graduates-expect-mass-layoffs-as-companies-hire-more-h-1b-workers/>.
- Karsten Müller and Carlo Schwarz. From hashtag to hate crime: Twitter and anti-minority sentiment. *Available at SSRN 3149103*, 2020.
- Karsten Müller and Carlo Schwarz. Fanning the flames of hate: Social media and hate crime. *Journal of the European Economic Association*, 2021.
- Alberto Ortega, Ema Di Fruscia, and Bryn Louise. Trade liberalization and racial animus. *Contemporary Economic Policy*, 39(1):194–204, 2021.
- Arnaud Philippe and Aurélie Ouss. “no hatred or malice, fear or affection”: Media and sentencing. *Journal of Political Economy*, 126(5):2134–2178, 2018.

Julia Preston. Pink slips at disney. but first, training foreign replacements. *The New York Times*, Jun 2015. URL <https://www.nytimes.com/2015/06/04/us/last-task-after-layoff-at-disney-train-foreign-replacements.html?ref=topics&r=0>.

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

Seth Stephens-Davidowitz and Hal Varian. A hands-on guide to google data. *further details on the construction can be found on the Google Trends page*, 2014.

Danny Yagan. Employment hysteresis from the great recession. *Journal of Political Economy*, 127(5): 2505–2558, 2019.