

University of Chicago Law School

Chicago Unbound

Coase-Sandor Working Paper Series in Law and Economics Coase-Sandor Institute for Law and Economics

2019

Imputing Unreported Hate Crimes Using Google Search Data

Dharmika Dharmapala

Aziz Z. Huq

Follow this and additional works at: https://chicagounbound.uchicago.edu/law_and_economics



Part of the **Law Commons**

Imputing Unreported Hate Crimes using Google Search Data

Dharmika Dharmapala
dharmap@uchicago.edu
University of Chicago Law School

Aziz Huq
huq@uchicago.edu
University of Chicago Law School

Abstract

U.S. law requires the Attorney General to collect data hate crime victimization from states and municipalities. But states and localities are under no obligation to cooperate. Data production hence varies considerably across jurisdictions. This Article addresses the ensuing “missing data” problem by imputing unreported hate crimes using Google search rates for racial epithets. As a benchmark, it uses two alternative definitions of which jurisdictions more effectively collect hate crime data: all states that were not part of the erstwhile Confederacy, and states with statutory provisions relating to hate crime reporting. We regress hate crime rates for racially-motivated hate crimes with African-American victims on Google searches and other relevant variables over 2004-2015 at the state-year level for each group of benchmark states. Google search rates substantially enhance the capacity of such models to predict hate crime rates among benchmark states. We use the results of these regressions to impute hate crime rates, out-of-sample, to non-benchmark jurisdictions that do not robustly report hate crimes. The results imply a substantial number of unreported hate crimes, concentrated in particular jurisdictions. It also illustrates how internet search rates can be a source of data for hard-to-measure victimization patterns.

Key words: hate crimes; victimization surveys; internet search.

Acknowledgments: We thank Rafeh Qureshi and Bartosz Woda of the Coase-Sandor Institute for Law and Economics at the University of Chicago Law School for their outstanding work in constructing the dataset used in this paper. We also thank Ken Schwenke, Scott Klein, and John Keefe for their advice on the hate crime data, and numerous engineers at Google for useful conversations. We thank Nick Stephanopoulos and workshop participants at the University of Chicago for helpful comments. Dharmapala acknowledges the financial support of the Lee and Brena Freeman Faculty Research Fund at the University of Chicago Law School; Huq acknowledges the Frank J. Cicero Fund. Any remaining errors or omissions are, of course, our own.

I. Introduction

Under the 1990 Hate Crimes Statistics Act (“HCSA” or “the Act”), the Attorney General of the United States is mandated to collect and publish annual data on the occurrence of hate crimes (Pub. L. 101-275). The HCSA defines hate crimes as “crimes that manifest evidence of prejudice based on race, gender and gender identity, disability, sexual orientation, or ethnicity.”¹ It also enumerates a non-exclusive class of covered offenses, including “murder, non-negligent manslaughter; forcible rape; aggravated assault, simple assault, intimidation; arson; and destruction, damage or vandalism of property.” The Attorney General has delegated statutory responsibility under the HCSA to the Federal Bureau of Investigation (“FBI”) as part of its obligations under the Uniform Crime Reporting (“UCR”) system (Ahuja 2016).

The FBI solicits information from law enforcement agencies in the fifty states. Federal law, however, imposes no mandate on states to collect or report hate crime data: States’ and municipalities’ cooperation under HCSA is voluntary. Nor does it provide funding. In 2016, 1,776 (11.6 percent) of the 15,254 participating state and local law enforcement agencies reported 6,121 incidents. The remaining 88.4 percent of agencies reported none (FBI 2017).

To what extent does this overwhelming pattern of non-reporting rest on the absence of relevant crimes, and to what extent does it reflect a failure on the part of participating state and local agencies to acquire and pass on information? One analysis estimates that a third of zero-report agencies had investigated at least one relevant bias-motivated crime within a calendar year (Cronin et al. 2007). No comprehensive study, however, estimates the extent (if any) or the geographic distribution of missing hate crimes data. Data from the National Crime Victimization Survey (“NCVS”) on bias-motivated offenses reported by victims represent a partial exception, although (as discussed below) a limited one.

We present in this article a new method for estimating state-specific rates of hate crimes to address this missing data problem. The method allows us to calculate (at best lower-bound) estimates for the rate of predicted anti-Black hate crimes for all states, and

¹ Congress amended the HCSA in 2009 to include the term “gender and gender identity” after “race.” Matthew Shepherd and James Byrd, Jr., Hate Crimes Prevention Act, Pub. L. 111-84, 123 Stat. 2190, 2385, Div. E, § 4708(a).

in particular states for where underreporting seems acute. Our method represents an important step toward filling present gaps in victimization data. In practical terms, it provides guidance as to where and how resources to prevent and address hate crimes might be best employed. In addition, the method also has more general application to impute aggregate rates of otherwise unobserved victimization. The article's contribution is hence both substantive and methodological.

Our method of imputation relies on data generated by the web search engine Google. The Google Trends app, available publicly since 2009, provides a time-series index of the queries that users enter into the Google search engine in a given geographic area in a given time period. Google Trends calculates an index value for the relative prevalence of a given search query within a particular state during the relevant time period (Choi and Varian 2012). Following studies in financial economics, public health, and social sciences, we employ Google search data derived from the Trends app as a proxy for the roughly contemporaneous preferences and behaviors of internet users (see, e.g., Choi and Varian 2012; Jun et al. 2018; Stephens-Davidowitz and Varian 2015). Hate speech and hate crimes, of course, do not necessarily correlate: Not everyone who engages in the former acts violently. The purpose of this study is rather to determine, at an aggregate population level, *whether the frequency of search for hateful terms is correlated with—and hence can be used as a proxy for unobserved hate-based victimization.*

We also use a measure of reported hate crimes from the FBI data. We obtain this from Inter-university Consortium for Political and Social Research at the University of Michigan, at the state-year level for 2004-2015. While this is the best available data on victimization rates for hate crimes, there is substantial variation in reported hate crime rates across states and longitudinally. Several states in the south and southeast that have historically been the locus of intense racial conflict and violence report exceptionally low rates of hate crime. These states also have Google search rates for racial epithets that are substantially above the national average.

There are two interpretations of this difference. A first is that jurisdictions in which online expressions of animus are higher have lower levels of actual hate crimes. This hypothesis is undermined by evidence that, at the national level and in jurisdictions with robust reporting mechanisms, the rate of Google searches for racial epithets and the rate of

reported hate crimes are strongly and positively correlated. A second hypothesis is that where history and current search activity suggest higher levels of racial animus, this difference reflects gaps in reporting hate-crime incidents. For reasons explained below, we think our data broadly supports this second hypothesis.

Our analysis aims to generate estimates of reported rate of hate crimes for those jurisdictions in which underreporting is likely acute. To fill this gap, we hypothesize in the vein of Stephens-Davidowitz (2014) and Chetty et al. (2018) that racial animus, as measured by the Google search rates, is a strong predictor of offline behavior, here actual hate-crime rates. To estimate the relationship between search rates and reported crime rates, however, we need to identify on *a priori* grounds a subset of states that we have reason to believe are *relatively* comprehensive in their reporting. To do this, we follow two alternative approaches. These ultimately yield similar results.

First, as many of the anomalously low-reporting states were historically part of the erstwhile Confederacy, we use all other states (i.e. all states and the District of Columbia (DC), *apart from* the 11 formerly Confederate states) as a first set of benchmark ‘robust reporting’ jurisdictions. Second, and alternatively, we use the presence of training programs for law enforcement officers on hate crimes and other relevant factors to identify 13 “good reporters” (consisting of 12 states and DC).

Using regression analysis, we demonstrate that search rates from Google Trends are strongly and significantly correlated with reported hate crime rates within each of our baseline sets of states.² This correlation persists when a variety of other variables, all arguably predictive of hate crime rates, are included in the regression specification. Critically, adding Google search rates to the model substantially enhances its ability to predict hate crime rates relative to a model consisting only of these other variables.

We then use regression results for our benchmark jurisdictions over 2004-2015 to generate predicted values of the hate crime rate for all states. The out-of-sample portion of this analysis comprises the states excluded from the benchmarking regression (i.e. the

² We use both linear and Poisson maximum-likelihood models in our regressions; these lead to very similar findings. The correlation is stronger cross-sectionally (i.e. across states) than longitudinally (i.e. within a given state over time). But our approach is intrinsically cross-sectional in nature, in that it seeks to extrapolate from a benchmark set of jurisdictions to other states. Thus, the existence of strong cross-state correlation is important, indeed essential, to our approach.

formerly Confederate states for our first analysis, and the states that are not identified as “good reporters” for our second analysis). We use these predicted values to impute hate crime rates for all states, including states excluded from the benchmarking regression. The underlying intuition is that if, for instance, Florida is subject to the same relationship between Google search rates (and other independent variables) and hate crimes as the baseline states (such as California), then the predicted hate crime rate we calculate for Florida tells us how many reported hate crimes we would counterfactually have observed in Florida if its reporting had been as robust as that in the benchmarking states.

The aim here involves imputing values for missing data, not causal inference. We are not making the causal claim that Google searches for racial epithets *cause* hate crimes to be committed. Nor do we claim that the people carrying out such searches are necessarily involved in committing hate crimes. Rather, we hypothesize that the Google search rate serves as a proxy for the *ecological* prevalence of racial animus within a particular jurisdiction at a particular time. We further postulate that it separately correlates to the extent to which residents of a particular state in a given year commit hate crimes targeting African-Americans.

One important limitation of our method is that we are unable to control for cross-state variation in victims’ reporting behavior. If victims of hate crimes were less willing to report such crimes in our baseline states than in the former Confederate states, this would generate inflated estimates of reported hate crime rates. We cannot rule this out. But it seems to us more likely that reporting levels will be higher in states with more robust approaches to tracking and measuring hate crimes than in states where anti-black hate crime is not a priority concern.

The limited nature of the resulting estimates also needs to be stressed. They are not estimates of the “true” rate of hate crimes. Rather, they are likely to be low estimates of the expected reported rate of hate crimes in a jurisdiction assuming high rather than low quality reporting. They primarily illuminate the problem of relative underreporting.³ They

³ We have calibrated our two benchmarks using coarse categories that we cannot manipulate (i.e., membership in the former Confederacy, and enactment of reporting-related legislation). This means that our benchmark jurisdictions are likely heterogeneous: They are unlikely to devote equal amounts of resources to identifying and gathering information about hate crimes. A benchmark that encompassed only very high quality reporters would be likely to generate higher predicted rates for out-of-sample jurisdictions.

do not address the distinct problem of underreporting by victims to local law enforcement agencies: Reported rates, given currently available data from victimization surveys, appear to be around one half of experienced crimes. They also do not address generalized failures on the part of law enforcement agencies to recognize certain acts as hate crimes even in benchmark jurisdictions.

Nevertheless, the reported hate-crime rates that we impute beyond our benchmark sample suggest substantial underreporting among many states, especially (but not exclusively) those in the southeast. Nationally, the mean number of African-American victims of hate crimes reported at the state-year level is about 57 per year in our sample of data from 2004-15. Florida reports a number very similar to this national average: 56. Yet, using the methodology described above, we find that the imputed number of reported hate crimes to be on average about 340 per year between 2004 and 2015. Thus, we impute an average of 284 unreported hate crime victims in Florida per year over 2004-2015.

We also impute substantial numbers of unreported hate crime victims in other states – on average, for example, 259 per year in Texas and 146 per year in Georgia over the same time period. Lest it is thought this analysis merely reiterates conventional wisdom about a north-south racial divide, we note that several non-southern states appear to also suffer from substantial underreporting. For instance, we impute that Illinois had about 89 unreported hate crime victims per year, while Pennsylvania had about 95.

Part II provides background on the definition and measurement of hate crimes, identifying the contours of the missing data problem addressed here. Part III describes the data used in the analysis. It focuses primarily on the Google Trends data, identifying the assumptions upon which reliance on that data is predicated. Part IV presents our central analysis and imputed lower-bound estimates. It also identifies limitations of the analysis. Part V concludes with a discussion of implications and possible extensions.

II. Defining, Measuring, and Analyzing Hate Crime

A. Defining Hate Crimes

There is no single, unified definition of hate crime in the U.S. The HCSA's definition encompasses "crimes that manifest evidence of prejudice based on race, gender and gender identity, disability, sexual orientation, or ethnicity," but does not precisely

define which substantive offenses count. In respect to relevant motivation, the scope of federal practice has changed subtly over time.⁴ Statutory amendments have extended the boundaries of the textual definition of hate crimes, and may have altered, at least at the margins, the kind of data collected.

The HCSA definition, moreover, is not adopted by all fifty states in their own criminal law. Within constraints imposed by the First Amendment,⁵ states use different terminology, such as “bias-related crimes” (D.C. Code §22-3703 (2012)) or “discrimination in public places” (N.D. Cent. Code § 12.1-14-04 (2017)). Some local law extends to crime motivated by the “political affiliation” of a person (D.C. Code §22-3701 (2012)), while other states do not enumerate protected classes (Utah Code Ann. § 76-3-203.4)). States’ regulatory approaches to hate crimes also vary. Sentencing enhancements for bias-related crimes were adopted by most states starting in the 1970s (Grattet et al. 1998). Today, in addition, twenty-nine states and the District of Columbia have a reporting mandate of some kind (ADL 2017). States with no reporting mandate can nevertheless participate in federal reporting efforts under the HCSA (Gillis 2013).

B. Measuring the Frequency of Hate Crimes

Variance in state-law understandings of the term “hate crime” is relevant because it is state and local actors, not federal ones, who collect data on-the-ground. The HCSA does not empower the Department of Justice to collect information directly. It rather allows the Attorney General to solicit such information from the states. But the federal government cannot command state officials in their official capacity to take actions, and HCSA appropriates no funds to support state data collection efforts. The Act instead relies on states’ voluntary efforts to gather information on hate crimes.

Hate crime data is reported by states as part of their submission to the Uniform Crime Reports (“UCR”) data. In 2016, agencies in 49 states⁶ and the District of Columbia, covering almost 89.7 percent of the nation’s population, formally participated in federal

⁴ In addition to a statutory amendment adding the terms “gender” and “gender identity” to the Act in 2009, the Department of Justice has revised the subcategories for race and ethnicity in 2012; altered its definition of rape in 2013 (dropping a force requirement); and in the same year started to allow agencies to report a wider variety of religion-based crimes, including “anti-Buddhist, anti-Eastern Orthodox (Greek, Russian, etc.), anti-Hindu, anti-Jehovah’s Witness, anti-Mormon, anti-Other Christian, and anti-Sikh” acts (FBI 2016).

⁵ *Virginia v. Black*, 538 U.S. 343 (2003).

⁶ In 2016 Hawaii was listed as a non-reporter to the FBI.

efforts to gather hate crimes data through the UCR (FBI 2016). But many agencies within states either did not report anything or reported zero hate crimes for a given year. One journalistic estimate is that 88 percent of participating agencies did not report any hate crimes (Schwencke 2017).

The NCVS also collects data on crimes “motivated by an offender’s bias against them for belonging to or being associated with a group largely identified by [enumerated characteristics]” (DOJ 2017b).⁷ Unlike the UCR, which focuses on state and local law enforcement, the NCVS collects information from a nationally representative sample of households on a nationwide basis. In 2016, the sample was redesigned to produce reliable estimates for the 22 most populous states and metropolitan areas (Department of Justice 2017a). This more granular data, however, will never be available for the historical period we study here. Currently available NCVS data permits estimation of aggregate undercounting of hate crime reporting, but does not allow estimation of state- or locality-level variation in undercounting (Gillis 2012). It is possible, though, to obtain self-reported NCVS victimization data for large regional groupings of states (such as the South), and we discuss how these relate to our estimates in Section IV below.

C. *Undercounting and Variance in Federal Hate Crime Data*

There are several reasons to believe first, that UCR estimates of hate crimes undercount actual rates, and second, that the undercount is not evenly distributed across (state) jurisdictions..

To begin, net undercounting occurs primarily because the UCR captures only crimes reported to police or other responsible state entities (e.g., state universities). In 2016, however, the NCVS estimated that 48.7 percent of serious violent crimes and 57.9 percent of violent crimes were not reported to police (Department of Justice 2017a).⁸ Between 2011 and 2015, approximately 54 percent of hate crimes (including 54 percent of violent hate crimes) identified in the NCVS sample were not reported. Even this might underestimate the hate-crime undercount because some victims of hate crimes may be unwilling

⁷ Question 162 of the present instrument asks respondents about being “targeted ... because of” race, religion, ethnic background or national origin, disability, gender, or sexual orientation.

⁸ Zaykowski (2010) analyzes New York data to suggest that racial minorities are less likely to report a hate crime to police for either race-based or non-race-based reasons, with the suppression effect greatest for the former.

to provide details of their experiences even in the context of a survey such as the NCVS, or may themselves not classify a qualifying experience as a hate crime.

Larger variation in the quality of state and local data may arise because of differences in the institutional arrangements whereby local police and other state agencies collect hate-crime data. Even in jurisdictions where police have an explicit reporting mandate under state law—and hence are not merely unfunded contributors to a federal project—the zeal with which they pursue their task will fluctuate depending upon institutional leadership and culture. Grattet and Jenness (2005a and b) analyzed the practices of 397 Californian police and sheriffs' offices. They found local practices varied greatly. Different departments relied on different definitions of hate crimes and engaged in different practices to acquire information. Different departments also had varying priorities depending on the local ecology of violent crime. Other studies find similar variation in other jurisdictions (Boyd et al. 1996). Agencies with community policing policies also appeared to be correlated with increased compliance with hate-crime mandates (Grattet and Jenness 2008; Jenness and Grattet 2005). Jurisdictions in which there is a history of lynching and a large African-American population are less likely to be characterized by cooperation with federal data-collection efforts and less likely to prosecute anti-black hate crimes (King et al. 2009).

The absence of any reliable way to estimate the degree of underreporting by jurisdiction creates a problem of missing data: How can we know how many hate crimes occur in a given state that are either not reported to authorities, or alternatively are reported but not recorded in federal data for one reason or another? The NCVS provides one answer—but only at the national level. Given that rates of hate crimes vary between jurisdictions, and that therefore any policy response must also be localized, this fails to provide needful guidance for public policy.

D. Prior Literature

Notwithstanding these limitations, there is a small but significant empirical literature on measuring the frequency of hate crimes.⁹ For instance, Medoff (1999) uses an

⁹ There is also a wider empirical literature globally on the phenomenon of hate crime (e.g. Krueger and Pischke, 1997), and a theoretical literature on hate crimes using a variety of different approaches (e.g. Dharmapala and Garoupa, 2004; Gan, Williams, and Wiseman, 2004; Dharmapala and McAdams, 2005; Klump and Mialon, 2013).

economic framework to analyze the determinants of hate crime, using cross-sectional data. He finds support for market wage and law enforcement activity as determinants of hate crime. Gale, Heath and Ressler (2002) use state-level UCR data to analyze the determinants of hate crimes. Their analysis employs a panel dataset, and they control for unobserved state effects. They too find that economic variables such as unemployment are significantly associated with the incidence of hate crimes.

Particularly noteworthy here is Chan, Ghose and Seamans (2016). Using UCR hate crime data, they construct a panel dataset at the county-year level over 2001-2008 to analyze the impact of the spread of internet usage on racially-motivated hate crimes. With instrumental variables based on geographical features of the local terrain that affect the introduction of broadband internet service, they find evidence that broadband availability increases the incidence of hate crimes. They also use Google searches for racially offensive terms as a measure of racial animus. Their aim (unlike ours) is to show that the positive impact of broadband availability on hate crimes exists only in areas with higher levels of racial animus.

III. Data

A. Google Search Data

This study employs Google Trends data as a proxy for animus against specific groups. We use the frequency of Google searches of the word “n_____” in a jurisdiction as a *proxy* for the extent of racial animus directed towards African-Americans within that jurisdiction.

Google Trends generates an index of search activity for a certain search term within a defined geographic area and specific time period. More formally, an ideal search rate measure r_{st} for a particular search term in state s in year t might be defined as follows:

$$r_{st} = \frac{S_{st}^w}{S_{st}} \quad (1)$$

where S_{st}^w is the number of searches for the word “w” in state s in year t and S_{st} is the total number of searches (for all terms) in state s in year t . This can be understood as a measure of the frequency of a search term’s usage relative to the volume of Google searches in that jurisdiction. It hence provides a measure of the relative frequency with which a term is

sought in relation to other searches, notwithstanding changes in the penetration of Internet and search engine usage.

It is not possible to obtain from Google Trends the absolute number of searches S_{st}^w or the change in that absolute number over time. Instead, we can obtain from Google Trends search data for the word “n_____” for each state s in each year t ,¹⁰ and national (US-wide) search data for this term in each year t , relative to a benchmark term. These observed Google Trends numbers can be denoted by G_{st} and G_{USt} , respectively:

$$G_{st} = \frac{r_{st}}{B} * 100 \quad (2)$$

$$G_{USt} = \frac{r_{USt}}{B} * 100 \quad (3)$$

Here, B is the search rate for an arbitrary benchmark term. Typically, this is the maximum monthly value of G_{USt} within year t (however, as described below, B plays no role in the analysis as it drops out of our animus measure).

We construct a measure of racial animus towards African-Americans, denoted by A_{st} , as follows:

$$A_{st} = \frac{G_{st}}{G_{USt}} * 100 = \frac{r_{st}}{r_{USt}} * 100 \quad (4)$$

This measure has a simple interpretation: it is the search rate for the word “n_____” in state s in year t relative to the search rate for the same word in year t in the US as a whole. Note that if Google users in state s search for the word more frequently than US Google users on average in a particular year t , then $A_{st} > 100$ for that state in that year. Note also that A_{st} is not identical to the ideal measure r_{st} defined in Equation (1); the former is a ratio of absolute numbers of searches (i.e. the search rate in a given state in a given year), while the latter is the ratio of the search rate in that state to the national search rate. Our regression specifications, however, includes year fixed effects that, in essence, absorb the r_{USt} term that is common to all states in a given year. Hence, exogenous shocks that affect all fifty states will be controlled for in our analysis. The resulting variation in A_{st} is essentially equivalent to the variation in r_{st} .

¹⁰ More precisely, we obtain this data at that the monthly level, but aggregate to the annual level for each state.

Previous articles use Google Trends to search for evidence of the geographic distribution of racial animus (Stephens-Davidowitz 2014) and to identify its effect on racial minorities' life-course outcomes (Chetty et al. 2018). In the former study, Stephens-Davidowitz demonstrates a correlation between search trends and voting behavior, and identified a relationship between increases in searches for the term “n ___” and depressed vote share for an African-American presidential candidate. In the second of these studies, Chetty and his co-authors use Google searches as an index of “explicit racial animus,” to show that African-American boys and men who grow up in areas of higher “animus” (so defined and measured) have worse life-course outcomes, including lower wages and a greater probability of incarceration (Chetty et al. 2018).

In this study, we use only one search term (and variants). Like Stephens-Davidowitz (2014), we measure searches for the word “n ___,” as a measure of racial animus. According to Stephens-Davidowitz (2014), the most common searches using that term were for ‘n ___ jokes’ and ‘I hate n ___.’ The term is also known to be used in hate crimes (Parks and Jones 2005), although in practice its usages have changed subtly over time (Kennedy 2003). Of course, the use of this search term allows us to analyze only anti-black animus. According to NCVS data from 2011-15, though, almost half (48.1 percent) of all bias-motivated crimes were race related (Department of Justice 2017b, 2).

As shown in Table 1, the Google search rate measure of animus defined in Equation (4) has a mean of about 74 across state-years, with some substantial differences across states. To emphasize again, we do not assume that every person who searches for n ___ (or a variant thereon) is motivated by racial animus. We accept that population level data contains a great deal of noise if those searches are used as a proxy for animus. We hence theorize that the *variation* in the volume of searches for n ___ is a rough proxy for *variation* in racial animus

It is important to identify four reasons for handling such data with caution. First, there is an unreported privacy threshold. When the use of a term falls below a certain proportion of total searches, Google Trends will report zeros. Second, the index reported by Google Trends is calculated based on a sample of searches from a particular jurisdiction

and time period.¹¹ That sample is changed each day. Hence, multiple samples can be drawn, but one must wait a calendar day before drawing a new sample (Stephens-Davidowitz 2014; Stephens-Davidowitz and Varian 2015). Unfortunately, this sampling process implies that it is impossible for the data obtained from Google Trends by one team of scholars to be precisely replicated by another team; even if the latter were to enter identical search terms for the same jurisdictions and time periods, it would obtain slightly different samples. Nevertheless, we verified that the Google search rates we obtain for our terms of interest (averaged by state over 2004-2015) are very highly correlated with the state averages (over 2004-2007) reported by Stephens-Davidowitz (2014, Appendix A): the correlation coefficient is 0.87, suggesting that this measure reliably captures a fairly stable attribute.

Third, because the Google Trends index provides a measure of relative frequency of different searches, it is necessarily sensitive to changes in the demographic composition of Internet (and Google) users over time.¹² This suggests that time trends should be inferred with extreme caution from Google search data. In our regression analysis, we take account of these types of changes at the national level by the inclusion of year fixed effects. Fourth, search behavior is not just exogenously determined. It also reflects endogenous changes by Google's search algorithm (Lazer et al. 2014). Those changes might influence the relative frequency of searches over time.

Accordingly, while Google search data provides insight into attitudes likely to be correlated with related off-line behavior, this inference is unreliable as to any given Internet user. This study, like earlier ones,¹³ leverages the fact that the large aggregation of data

¹¹ It is also reported to the nearest integer. Hence, a search term with a low enough rate of searches in comparison to the most popular search term may yield an index of zero because of rounding, independent of the privacy threshold.

¹² Between 2000 and 2016, the rate of Internet penetration in the United States grew from about 50 per cent to almost 90 percent (PRC 2017).

¹³ Recent studies that employ an analytic strategy roughly akin to ours include studies of trading in financial markets (Preis et al. 2013), stock market fluctuations (Curme et al. 2013; Moat et al. 2013; Preis et al. 2014), product market share (Utsuro et al. 2014), private consumption behavior (Vosen and Schmidt 2012), and unemployment benefits (Choi and Varian 2009). Another study uses online search data to generate estimates of *future* behavior, including weekend box-office revenue for feature films, first-month sales of video games, and the rank of songs on the Billboard Hot 100 chart (Goel et al. 2010; Bollen et al. 2011).

captured by Google Trends can illuminate population-level behavioral trends as a simple consequence of the law of large numbers.

B. UCR and ICPSR Data

Available data on reported hate crimes used in our analysis is ultimately derived from the reports of state and local law enforcement, voluntarily provided to the FBI under the terms of the HCSA (as described above). This data is publicly available through the FBI's Uniform Crime Reporting (UCR) system for each year since 1996.¹⁴ It divides reported hate crimes into five categories based on the type of hate motivation – namely, by bias relating to race, ethnicity, religion, sexual orientation, and disability.

Reports under the HCSA are also available through the Inter-university Consortium for Political and Social Research (ICPSR).¹⁵ In addition to the type of motivation, this records the race or other relevant characteristics of the victim. Thus, it enables us to calculate the number of African-American hate crime victims (within the category of victims of racially-motivated hate crimes), aggregated to the state-year level. As some incidents have multiple victims, this variable differs slightly from the number of hate crime incidents at the state-year level with African-American victims. Nevertheless, using the latter variable instead leads to very similar results. While the hate crime data is available (albeit with the limitations that have already been discussed) since the 1990's, we only use data for 2004-2015, as the Google Trends data is only available from 2004.

As shown in Table 1, there are about 57 African-American victims of reported racially-motivated hate crimes in the typical state-year, though there is also a quite large standard error (and by implication considerable variation across states and years). In our benchmarking analysis, we scale the number of African-American hate crime victims by the total state population in that state-year, to obtain a reported hate crime rate (for African-American victims) per 100,000 state population. State population is obtained from the Census Bureau's intercensal estimates of population for states,¹⁶ and is measured in millions (with a mean of about 6 million people in our sample). The mean of the hate crime

¹⁴ For instance, see <https://ucr.fbi.gov/hate-crime/2015> for 2015 data.

¹⁵ See <https://www.icpsr.umich.edu/icpsrweb/>

¹⁶ See <https://www.census.gov/programs-surveys/popest.html>

rate is about 0.97 African-American victims of racially-motivated hate crimes per 100,000 state population.

An alternative measure would scale by the African-American population rather than by the total state population. This leads, however, to a highly skewed variable due to the small African-American population in some states. Thus, we use for our primary analysis a hate crime rate that is scaled by state population (as described above). We also control in the regression analysis for a polynomial function of the African-American fraction of the state population. This takes account in a flexible manner of any effects of racial contact, competition, or other factors (such as the presence of potential victims, the influence of interracial contact on attitudes, and the effect of racial threat (Bobo and Hutchins 1996)) that may influence the rate at which hate crimes are committed and reported.

The UCR system also reports more general crime statistics, by law enforcement agency and year.¹⁷ We collect from the UCR database a large number of crime rate variables for the period 2004-2015. In our baseline analysis, we use an aggregate measure of all offenses as a proxy for the general crime rate.¹⁸ The mean rate of such offenses is 2.36 per 100,000 state population. We also use a number of other variables, summary statistics for which are provided in Table 1. Unemployment rates at the state level are obtained from the Bureau of Labor Statistics.¹⁹ Mean income (in thousands of dollars) at the state-year level, the percentage of the state population that is college-educated, and the fraction of African-American residents are all obtained from the Census Bureau's Current Population Survey (CPS).²⁰

IV. Empirical Approach and Results

A. Partitioning States Based on Expected Reporting Behavior: Benchmarking Jurisdictions

¹⁷ See <https://www.bjs.gov/ucrdata/abouttheucr.cfm>

¹⁸ The offenses included in this measure are: murder and nonnegligent manslaughter, rape, robbery, assault, burglary, larceny, and motor vehicle theft.

¹⁹ See <https://www.bls.gov/data/#unemployment>

²⁰ See <https://www.census.gov/programs-surveys/cps.html>

We start by identifying on an *a priori* basis a class of jurisdictions to generate an initial estimate of the relationship between Google search rates and reported hate crime rates. This assumes that there exists an identifiable subset of jurisdictions in which reporting practices are reasonably robust. But there is no single objective way to draw this distinction because there is no reliable data comparing the quality of reporting between states (and because actual hate crime rates are obviously unobserved).

We develop two sets of jurisdictions to serve as benchmarks using proxies that are relatively simple but robust to manipulation. We first propose a distinction between former Confederate and all other states as a first approximation to identify states in which the quality of reporting may be low but the level of racial animus high. Historical members of the Confederacy (Alabama, Arkansas, Florida, Georgia, Louisiana, Mississippi, North Carolina, South Carolina, Tennessee, Texas, and Virginia) account for about 22% of our observations. We also use the Census definition of the South²¹ to construct a broader category consisting of the former Confederate states plus five additional (“border”) states (Delaware, Kentucky, Maryland, Oklahoma, and West Virginia). Southern states, defined in this way, comprise about a third of our observations.

We readily concede that we could begin with a number of other distinctions, such as black shares of the voting-age population to jurisdictions that have not passed any regulation of hate crimes. We rely on former Confederacy status because it is not a quantitative measure. It therefore does not require us to make contestable judgments about what ‘counts’ as evidence of likely underreporting. Its insensitivity to manipulation binds our hands.²²

A second and alternative set of benchmarking jurisdictions is a category of states for which there exist *a priori* reasons to believe that they are better reporters of hate crimes to the FBI under the HCSA. The primary criterion that we use is the existence of a state statute mandating training for police on hate crime reporting. Such legislation is directly relevant to the quality of reporting, and again, is a measure that is not amenable to our

²¹ See <https://www.census.gov/geo/reference/webatlas/regions.html>

²² There is, in addition, evidence to suggest that the historical status of a geographic unit continues to have influence today for two reasons. First, Acharya, Blackwell, and Sen (2016) demonstrate that the density of slave-holding within a county in 1860 correlates to levels of racial resentment within that county today. Second, King et al. (2009) demonstrate that the historical rate of lynching influences various metrics of reporting and prosecuting hate crimes.

manipulation. According to a 2016 survey conducted by the Anti-Defamation League (ADL 2017), such laws exist in the following states: California, Connecticut, Illinois, Kentucky, Massachusetts, Minnesota, New Jersey, New Mexico, Oregon, Rhode Island, and Washington. Notably, this category includes one north/south border state (Kentucky). In addition, scholarly commentary suggests that New York has a good reporting system (Gilles 2012). Finally, DC is the only jurisdiction in our data with an African-American majority, and so reporting of hate crimes directed at African-Americans may be particularly salient there. These 12 states and DC constitute a second and alternative benchmarking subset of jurisdictions that we use in some of the analysis below. As shown in Table 1, they constitute about a quarter of our observations.

B. Anomalous Patterns in Reported Hate Crimes: Descriptive Statistics

We begin our analysis by presenting descriptive statistics in graphical form.

To begin with, Figures 1 and 2, both presented below, summarize the catalyzing impulse of the analysis. Figure 1 depicts the average reported hate crime rate by state. Here, the annual data on the number of African-American victims of racially-motivated hate crimes (scaled by state population) is averaged by state over the entire sample period (2004-2015) to obtain an average reported hate crime rate for each state. This map shows considerable variation across states. One particularly striking pattern is that several states in the southeast that have historically been the locus for racial conflict and violence report among the lowest rates of hate crime. Given what is known about the history of these states, this is surprising.

The intuition that underreporting rates may systematically vary across states is supported by Figure 2. This depicts the average value for each state of the racial animus measure A_{st} defined in Equation (4). Again, this measure is averaged by state over the entire sample period (2004-2015) to obtain an average level of racial animus for each state. There is considerable variation in these average levels across states. The southeastern states that tend to report very low rates of hate crime are noteworthy for their relatively high levels of racial animus, at least as measured by Google searches for “n ____”. This pattern might suggest that the low reported hate crime rates in these states are unlikely to reflect low actual rates of hate crime (at least under the assumption that actual hate crime rates are related to racial animus, which in turn is reflected in comparative rates of internet searches

for disparaging terms). Further, the existence of a strong, positive correlation between search rates for racial epithets and reported crime rates, which we describe below, makes that interpretation especially likely.

A more targeted comparison is provided in Figure 3. This figure compares reported hate crime rate for California (again, for anti-black crimes only) for each year in our sample period (2004-2015) to the corresponding rate for Mississippi. California has a relatively low level of racial animus, as measured by the prevalence of Google searches for “n_____”; its value of A_{st} (averaged over 2004-2015) is about 55, relative to the sample mean of about 74. In contrast, Mississippi’s value of A_{st} (averaged over 2004-2015) is about 109, the highest for any state apart from West Virginia. Yet, California reports much higher rates of hate crime (on average, a rate of 1.29 per year, relative to the sample mean of 0.97). Indeed, Mississippi reported zero hate crimes for most years in our sample; its average reported rate is 0.03.

A third, more general comparison suggests that there are similar predictable differences between our first set of benchmarking jurisdictions on the one hand, and the 11 former Confederate states on the other. Figure 4 shows a comparison between reported hate crime rates against African-Americans for the 11 former Confederate states and the corresponding rates for the other 39 states and DC. The reported hate crime rate for each year is averaged over states separately for Confederate and non-Confederate states. Reported hate crime rates are substantially and consistently higher for non-Confederate states. At the same time, the two groups of states seem to follow similar time trends. In both categories, hate crime rates decline over the 2004-2015 period, with a sharp increase against the declining trend in 2008. More recently, reported rates appear to have stabilized or started to increase at the end of this period.²³

Figure 5 shows the analogous time-series variation for A_{st} for each year averaged over states (separately for Confederate and non-Confederate states). Over the earlier part of our sample period, there is once again a declining trend for both categories of states. Moreover, there is an increase in the rate of Google searches for “n_____” in 2008 for both

²³ Media reports suggested that there may have been an increase in hate crimes during and following the 2016 election campaign (Bursch 2017). Because our sample period ends prior to this, in 2015, our results neither support nor disprove such reports.

categories; strikingly, this corresponds to the sharp increase in reported hate crimes for both groups of states in the same year. Since about 2010, the racial animus measure appears to have stabilized overall, but increased slightly in former Confederate states. The regression analysis described below uses year fixed effects to take account of such changes.

C. Bivariate relationships between reported hate crime data and internet search data

An assumption that is crucial for our analysis is that Google searches for “n_____” are predictive of reported hate crime rates (at least among “good” reporter states). To validate this in part, Figure 6 shows the (non-causal) relationship between A_{st} and reported hate crime rates at the state-year level, using only data for non-Confederate states. This graph suggests a positive relationship. The line of best fit shown in Figure 6 is a simple bivariate relationship. That is, it does not control for any potential determinants of reported hate crime rates. That said, when we include such additional variables in our regression analysis, we continue to find a positive relationship between A_{st} and reported hate crime rates. This also suggests that it would be inappropriate to interpret the low rate of A_{st} in southeastern states as a signal of low hate crime rates.

Figure 7 shows the fitted line from Figure 6, estimated only using data for non-Confederate states, together with its 95% confidence interval. It also plots the hate crime rate and A_{st} for the 11 former Confederate states (each averaged by state over 2004-2015). While two of these states lie above the fitted line, most former Confederate states lie well below the 95% confidence interval. This implies that the former Confederate states are very different from other states in terms of the relationship between A_{st} and reported hate crime rates. In particular, the reported hate crime rate for these states is mostly below (and typically far below) that which would be predicted from their observed levels of racial animus as measured by internet search usage. Given the evidence that search rates provide a meaningful proxy for hate crime rates, we think it is most reasonable to interpret Figure 7 as evidence of underreporting by a subset of jurisdictions.

D. Empirical Specification and Regression Results

We now describe the regression models that we use to predict and impute hate crime rates. For our primary analysis, we use a linear model with the following specification:

$$H_{st} = \beta A_{st} + \gamma \mathbf{X}_{st} + \delta_t + \epsilon_{st} \quad (5)$$

Here, H_{st} is the reported hate crime rate for state s in year t , and A_{st} is the measure of racial animus based on Google searches for “n_____” (as defined in Equation (4)). \mathbf{X}_{st} is a vector of additional variables that we hypothesize to be predictive of H_{st} . In the basic analysis, these are the general crime rate (i.e. all offenses), the unemployment rate, mean income, the college-educated percentage of the population, and the state population. In addition, we include a quadratic function of the African-American fraction of the state population to account for the complex effects of racial contact and racial threat. In some specifications, we also include historical variables, such as indicators for former Confederate states and border states. The term δ_t is a year fixed effect, while ϵ_{st} is the error term.

While we include year fixed effects (which capture national trends in H_{st} that are common across states), we do not include state fixed effects. This is because we are *not* seeking to estimate the causal effect of A_{st} on H_{st} (an exercise that would, among other things, require controlling for unobserved heterogeneity across states). Rather, we are seeking to *impute* H_{st} using A_{st} and other relevant variables for out-of-sample jurisdictions. Moreover, state fixed effects are precluded by the intrinsically cross-sectional nature of the underlying approach. Given that we are seeking to extrapolate from the observed relationship between H_{st} and other variables (including A_{st}) in some states to impute (missing or unreliable) values of H_{st} in other states, state fixed effects (or, equivalently, de-meaning H_{st} by state) would create an arbitrary parameter (the state fixed effect or state-level mean) that would make the extrapolation indeterminate. To be sure, the corollary to this is that we must impose the assumption that states are sufficiently similar (in terms of the relationship between the true H_{st} and the other variables) to permit valid extrapolation across states. As discussed below, the most likely sources of error in this regard would tend to bias against our findings.

Column 1 of Table 2 reports the results of a regression for H_{st} that includes only year effects and indicators for former Confederate states and border states (estimated using the entire dataset, with 612 state-year-level observations on 51 jurisdictions). This demonstrates that former Confederate states have significantly lower reported hate crime rates; the difference of 0.45 is quite large relative to the sample mean of 0.97. In contrast, border states tend to have higher reported hate crime rates than the excluded category (i.e. non-Southern states and DC). Column 2 repeats this exercise, adding the various variables

described earlier. The patterns in Column 1 are robust to their inclusion, and the magnitude of the difference between predicted hate crime rates for the former Confederate states and for on-Confederate states is very similar.

Column 3 of Table 2 reports a regression in which H_{st} is replaced as the dependent variable by A_{st} . Without controls, it appears that Southern states - both former Confederate states and border states - have much higher levels of racial animus than the excluded category (non-Southern states). Controlling for the general crime rate and for economic and demographic variables (in Column 4), however, the difference for former Confederate states becomes much smaller and is statistically insignificant (the difference for border states also becomes smaller, although it remains statistically significant). It is worth noting that our claims in this paper do not depend on racial animus necessarily being higher in Confederate states. Rather, we argue only that their reported hate crime rates are unusually low, *given their level of racial animus*. Thus, Table 2 establishes more formally what we have already shown through Figures 1 and 2 – i.e. that the former Confederate states report low rates of hate crime, while their observed racial animus is at least as high as elsewhere.

Column 1 of Table 3 reports results from the regression specification in Equation (5), estimated using only data on non-Confederate states. This sample includes 480 state-year-level observations on 39 states and DC.²⁴ It is evident from the table that A_{st} and H_{st} are very strongly related for these benchmarking jurisdictions. Higher general levels of violent crime are also strongly predictive of H_{st} . It may seem surprising that a higher percentage of college-educated residents is associated with higher reported hate crime rates. But it should be borne in mind that reported hate crime rates reflect both actual hate crime rates and reporting quality. Higher levels of education may be associated with higher levels of reporting on the part of law enforcement agencies or individual victims, even if the college-educated themselves are less likely to commit hate-motivated (or other) crimes.

Column 2 of Table 3 reports results from the regression specification in Equation (5), estimated using only data on our second set of benchmarking jurisdictions. As explained above, this set is defined primarily in terms of legislatively mandated training on

²⁴ The tables report robust (heteroscedasticity-corrected) standard errors. The patterns of statistical significance are fairly similar when standard errors are clustered at the state level (and the predicted values are of course identical).

reporters or its analog. This sample includes 156 state-year-level observations on 12 states and DC. The results are generally very similar to those in column 1. In particular, A_{st} and H_{st} are again very strongly related, and the magnitude of the coefficient (β in Equation (5)) is substantially larger. This may reflect the potential dilution of the estimate in Column 1 due to the over-inclusiveness of the non-Confederate category – that is, including states that are not very good reporters in the baseline sample would be expected to bias the estimated β downwards.

We now consider the extent to which the inclusion of A_{st} in an analysis enhances a model's predictive power. What boost, that is, does search data give to hate-crime imputation? Column 3 of Table 3 reports a specification that includes year fixed effects and the social and economic characteristics described earlier, but excludes A_{st} . R^2 , the basic measure of the model's goodness-of-fit, falls from 0.221 to 0.167, while adjusted R^2 (which adjusts for the number of regressors in the model) falls from 0.112 to 0.058. Put differently, adding A_{st} to a model that initially includes only the other regressors increases R^2 by about one third. By way of comparison, if we were to drop the race variables, R^2 would fall from 0.221 to 0.205, if we were to drop the unemployment rate, R^2 would fall from 0.221 to 0.218, and if we were to drop income, R^2 would be virtually unaffected.²⁵ This suggests that A_{st} has substantial predictive power with respect to hate crimes, relative to the other variables.²⁶

Equation (5) uses a linear specification. The number of African-American hate crime victims in a given state-year takes on only non-negative integer values, and thus is an example of “count” data. The hate crime rate (scaled by state population) that we compute and use has a more continuous distribution, but also necessarily only takes on non-negative values. Moreover, it potentially includes more zero observations than would be expected with a standard normal distribution of the error term. Ultimately, we wish to generate predicted values of hate crime rates. But the linear model does not constrain

²⁵ In addition, a Wald test yields an F-statistic of 6.39 and a p-value of 0.0127, implying that the Google search variable adds significantly (at the 5% level, and almost at the 1% level) to the explanatory power of the model.

²⁶ The impact of adding A_{st} is even greater in relative terms if year fixed effects are excluded from the model. If the non-Confederate baseline is used instead, the increase in R^2 from adding A_{st} is smaller. This is arguably unsurprising, as A_{st} would be expected to be a weaker predictor of reported hate crime rates when the sample includes relatively poorer reporters.

predicted values to be non-negative, and so it is in principle possible to predict negative hate crime rates (although this does not turn out to be a major problem in our sample).

For these reasons, we check the robustness of our results using a specification that better accommodates non-negative count data, the Poisson maximum-likelihood model:

$$H_{st} = \exp(\beta A_{st} + \gamma \mathbf{X}_{st} + \delta_t) \epsilon_{it} \quad (6)$$

where the variables are as defined above. Column 3 of Table 3 reports the results from the regression specification in Equation (6), estimated using only data on non-Confederate states. Column 4 of Table 3 reports results from the regression specification in Equation (6), estimated using only data on “good reporter” states. In each case, it is apparent that the estimates are very similar to those in Columns 1 and 2 of Table 3.

E. Imputing Unreported Hate Crimes

For each of the baseline samples, we use the estimates reported in Table 3 to generate predicted values of hate crime rates, denoted \hat{H}_{st} . For the linear regression in Equation (5), these predicted values are calculated as follows (where $\hat{\beta}$ denotes the estimated value of the coefficient β in Equation (5), $\hat{\gamma}$ denotes the estimated value of γ , and $\hat{\delta}_t$ denotes the estimated year effects for each year t):²⁷

$$\hat{H}_{st} = \hat{\beta} A_{st} + \hat{\gamma} \mathbf{X}_{st} + \hat{\delta}_t \quad (7)$$

Critically, the predicted values of hate crime rates \hat{H}_{st} can be generated for states outside the benchmark set for which Equation (5) is estimated. For example, Florida is excluded from the set of baseline states, but its values of A_{st} and \mathbf{X}_{st} are available in the dataset. We can combine that data with the coefficients $\hat{\beta}$ and $\hat{\gamma}$ estimated in Table 3, along with the estimated year effects $\hat{\delta}_t$ for each year, to construct a predicted \hat{H}_{st} for Florida for each year of our sample period (2004-2015).

Table 4 compares the reported hate crimes rates to the predicted values generated using Equation (7) for all Confederate states and for selected non-Confederate states. The baseline set of states is all non-Confederate states. As discussed below, the results are similar when using the ‘good reporters’ as the benchmark, which we discuss below. The first two columns contrast the reported and predicted mean annual hate crime rates, each

²⁷ Predicted values from the Poisson model in Equation (6) can be derived in an analogous manner; these lead to fairly similar results.

averaged over the years 2004-2015. For example, Florida has a mean annual reported hate crime rate of about 0.3 (substantially below the sample mean of 0.97). But based on its Google search rate for “n _____” and its other observed characteristics, its predicted mean annual hate crime rate is about 1.68 (substantially above the sample mean).

The third column of Table 4 reports the mean annual number of African-American hate crime victims from the ICPSR data, averaged over the years 2004-2015. The fourth column converts the predicted hate crime rate in column 2 into the implied number of African-American hate crime victims per year (by multiplying the rate by the state population). For instance, Florida law enforcement agencies reported an average of 56 African-American victims of racially-motivated hate crimes per year over 2004-2015, a number that is very close to the sample mean of 57. Still, the predicted hate crime rate for Florida implies an average of 340 African-American victims of racially-motivated hate crimes per year over 2004-2015. This in turn implies an average of 284 African-American victims of *unreported* racially-motivated hate crimes per year over 2004-2015, as shown in the final column of Table 4 (which reports the difference between the predicted number of victims per year from column four and the reported number from column three).

The imputed number of unreported hate crimes is large for several other former Confederate states. For instance, the predicted hate crime rate for Texas implies an average of 259 African-American victims of unreported racially-motivated hate crimes per year over 2004-2015; the predicted hate crime rate for Georgia implies an average of 146 African-American victims of unreported racially-motivated hate crimes per year over 2004-2015.

Underreporting is by no means limited to former Confederate states. Even among the benchmark states, there is variation in whether states reported hate crime rates fall above or below the predicted value that represents an average among all states in the benchmark set. For example, the predicted hate crime rate for Pennsylvania implies an average of 95 African-American victims of unreported racially-motivated hate crimes per year over 2004-2015; the predicted hate crime rate for Illinois implies an average of 89 African-American victims of unreported racially-motivated hate crimes per year over 2004-2015.

Figure 8 presents a map of the states showing predicted mean annual hate crime rates averaged over 2004-2015. It is notable that states in the southeast tend to have relatively large predicted values in Figure 8, in sharp contrast to the pattern of reported hate crime rates in Figure 1. The reported and imputed hate crime rates for each of the former Confederate states are represented in Figure 9. This shows that our procedure implies an upward adjustment for each of these states, though the magnitude of the gap between the reported and imputed rates differs substantially across states.

E. Limitations

Our analysis is constrained by a number of important caveats and limitations. It uses *reported* hate crimes in some jurisdictions as a benchmark to estimate the rate at which hate crimes *would be reported* in other jurisdictions. The analysis thus rests on the premise that the benchmark jurisdictions constitute a reasonable comparator for other jurisdictions. In particular, it requires that racial animus translates into hate crimes to a roughly similar degree across states.

One partial check on this is to use the NCVS data based on victim self-reporting (available only at the national and regional, not state, levels). Using the NCVS data over 2003-2018, we compute that about 26% of African-Americans who report being victims of hate crimes reside in the South, compared to about 38% of the total population in our dataset in 2015.²⁸

This suggests that actual hate crime victimization, as measured by in the NCVS data, is somewhat lower in the South compared to other regions. Even if this were true, however, it cannot possibly account for the phenomenon of several Southern states' authorities reporting hate crime numbers of zero or close to zero to the FBI. Even adjusting our estimates of imputed hate crimes downwards to account for putatively lower actual hate crime victimization in the South would leave us with imputed numbers for most Southern states that are dramatically larger than those reported to the FBI.²⁹

²⁸ The NCVS data is available at: <https://bjs.ojp.gov/data-collection/ncvs>. To perform this calculation, we use affirmative answers to question V4526A: "Do you suspect the offender targeted you because of your race?" by respondents whose race is reported as African-American, and classify respondents residence by whether it is reported to be in the South or elsewhere.

²⁹ For example, scaling Georgia's imputed mean number of imputed hate crimes of 145.99 (shown in Table 4) by a factor of 0.7 (the ratio of 26% and 38%) would result in the imputed number of about 102, as compared to a mean reported to the FBI of 11.33. An alternative approach to scaling uses the fraction of the African-

Moreover, there are reasons to be cautious in interpreting the regional NCVS data. One relates to sampling variability, as the sample size for specific subgroups (such as African-American victims of hate crimes in the South) are quite small in the survey. In addition, some victims of hate crimes may be unwilling to provide details of their experiences even to a survey such as the NCVS, or may themselves not classify or recognize a qualifying experience as a hate crime in terms that fit the survey's framework. Importantly, the reluctance to report or to understand an experience in hate-related terms may differ across states and regions based on their cultural environment and history. Hate crimes are intended to create or exacerbate a sense of social vulnerability (Waldron 2012), and often succeed in doing so, possibly to a differential degree across regions. The NCVS employs a two-step process to measure and collect data on sexual assault in recognition of the likelihood that respondents will be reluctant to discuss such events (Department of Justice 2017a, 17). It uses no such protocol for hate crimes, although this may help equalize reporting propensities across regions.

Finally, the wider lesson of the NCVS data is perhaps that the rate at which hate crimes are reported is likely to be substantially lower than the rate at which they occur (Department of Justice 2017b). In addition, reporting practice in 'good' states is far from perfect (Boyd et al. 1996; Grattet and Jenness 2005a, b), suggesting another source of downward bias in our estimates.³⁰

V. Conclusion

Google search data has been identified as a proxy for otherwise unobservable attitudes on the ground that trends in search rates are correlated with overall attitudinal shifts or differences at the population level (see, e.g., Chetty et al. 2018; Choi and Varian 2012; Jun et al. 2018; Stephens-Davidowitz and Varian 2015). This study extends that insight to a field in which data in respect to an important question of law enforcement and public policy has been of inconsistent quality. We have hypothesized that there would be

American population in the South as a baseline. This is about 68% in our dataset in 2015, the imputed number for Georgia would become 55, still about five times the number reported to the FBI.

³⁰ Despite our efforts to clarify the algorithm's operation with engineers at Google, we remain uncertain whether the privacy lower-bound varies with the choice of search term, or the extent to which daily sampling introduces error into any given estimate. Better and more transparent access to search data might increase predictive capacity in respect to hate crimes.

a correlation between population-level Google search trends for racially derogatory terms in a jurisdiction and rates of reported anti-black hate crimes in the same area. In jurisdictions for which our prior would suggest that state and municipal data collection efforts are more rather than less robust, we have estimated that correlation statistically, and demonstrated that these Google search rates add substantial predictive power relative to a model that includes only state-year socioeconomic and demographic characteristics. This relationship can be used to extrapolate beyond the benchmarking sample and thereby impute hate crime rates for jurisdictions that are, by hypothesis, poor reporters. Our estimates, to be clear, are only a lower bound of the expected reported hate crime rate. In light of NCVS data, we think that this is likely about one half of the actual rate of such crimes.

As we have emphasized, this is not the first study to postulate a correlation between population-level search trend data. Nevertheless, we are not aware of any other efforts to deploy population-level search trends as a proxy for an unobservable parameter. There are many other possible applications of this same method, if it is found to be generally reliable. An obvious extension of our work here is to other kinds of hate crimes, or to offenses that are in expectation under-reported. We can imagine, for instance, extending this method to impute rates of rape or sexual assault. More generally, it is possible that Google search trends can be validated as proxies for other attitudinal and behavioral trends with which the law is concerned. Our contribution here, in short, should be understood to be as much methodological as substantive: It entails the introduction of a new method of empirical research for a legal scholarly audience.

Bibliography

Acharya A, Blackwell M, and Sen M (2016) The Political Legacy of American Slavery. *Journal of Politics*, 78(3), 621–641.

Ahuja H (2015) The Vicious Cycle of Hate: Systemic Flaws in Hate Crime Documentation in the United States and the Impact on Minority Communities. *Cardozo Law Review*, 37, 1867–1906.

Anti-Defamation League (2017) #50 Against Hate. <https://www.adl.org/50statesagainsthate> (accessed 19 January 2018).

Bobo L and Hutchings V (1996) Perceptions of racial group competition: Extending Blumer's theory of group position to a multiracial social context. *American Sociological Review*, 61(4): 951-972.

Boyd EA, Berk RA, & Hamner, KM (1996), 'Motivated by Hatred or Prejudice': Categorization of Hate-motivated Crimes in Two Police Divisions. *Law & Society Review*, 30(4), 819–850.

Bursch A (2017) Spread of Hate Crimes has Lawmakers Seeking Harsher Penalties. *New York Times*. April 30.

Chae DH, Clouston S, Hatzenbuehler ML, Kramer MR, Cooper HLF, Wilson SM, Link BG (2015) Association between an Internet-Based Measure of Area Racism and Black Mortality. *PLOS ONE*, 10(4), e0122963.

Chan J, Ghose A, & Seamans R (2016) The internet and racial hate crime: Offline spillovers from online access. *MIS Quarterly*, 40(2), 381-403.

Chetty R, Hendren N, Jones M, and Porter S (2018) Race and Opportunity in the United States: An Intergenerational Perspective. NBER Working Paper, http://www.equality-of-opportunity.org/assets/documents/race_paper.pdf

Choi H, and Varian H (2009) Predicting initial claims for unemployment benefits. *Google Inc*, 1-5.

Choi H, and Varian H (2012) Predicting the Present with Google Trends. *The Economic Record*, 88(s1), 2–9.

Cronin SW, McDevitt J, Farrell A, and Nolan JII (2007) Bias-Crime Reporting: Organizational Responses to Ambiguity, Uncertainty, and Infrequency in Eight Police Departments. *American Behavioral Scientist*, 51(2), 213–231.

Curme C, Preis T, Stanley HE, and Moat HS (2014) Quantifying the semantics of search behavior before stock market moves. *Proceedings of the National Academy of Sciences*, 111(32), 11600–11605. <https://doi.org/10.1073/pnas.1324054111>

- Dharmapala D and Garoupa, N (2004) Penalty enhancement for hate crimes: An economic analysis. *American Law and Economics Review*, 6(1), 185-207.
- Dharmapala D and McAdams RH (2005). Words that kill? An economic model of the influence of speech on behavior (with particular reference to hate speech). *Journal of Legal Studies*, 34(1), 93-136.
- Federal Bureau of Investigation (2016) About Hate Crime Statistics. <https://ucr.fbi.gov/hate-crime/2016/resource-pages/about-hate-crime> (accessed 19 January 2018).
- Federal Bureau of Investigation (2017) Hate Crime by Jurisdiction. <https://ucr.fbi.gov/hate-crime/2016/topic-pages/jurisdiction> (accessed 19 January 2018).
- Gale L, Heath W, and Ressler R (2002). An Economic Analysis of Hate Crime. *Eastern Economic Journal*, 28, 2002, 203–16.
- Gan L, Williams RC and Wiseman T (2011). A simple model of optimal hate crime legislation. *Economic Inquiry*, 49(3), 674-684.
- Gillis B (2012). Understanding hate crime statutes and building towards a better system in Texas. *American Journal of Criminal Law* 40(3), 197-226.
- Grattet R, and Jenness V (2005) The Reconstitution of Law in Local Settings: Agency Discretion, Ambiguity, and a Surplus of Law in the Policing of Hate Crime. *Law & Society Review*, 39(4), 893–941.
- Grattet R and Jenness V (2008) Transforming Symbolic Law into Organizational Action: Hate Crime Policy and Law Enforcement Practice. *Social Forces*, 87(1), 501–527.
- Grattet R, Jenness V, and Curry TR (1998) The Homogenization and Differentiation of Hate Crime Law in the United States, 1978 to 1995: Innovation and Diffusion in the Criminalization of Bigotry. *American Sociological Review*, 63(2), 286–307.
- Jenness V, and Grattet R (2005). The Law-In-Between: The Effects of Organizational Perviousness on the Policing of Hate Crime. *Social Problems*, 52(3), 337–359.
- Jun SP, Yoo HS, & Choi S (2017). Ten years of research change using Google Trends: From the perspective of big data utilizations and applications. *Technological Forecasting and Social Change*, 130: 69-87
- King RD, Messner SF, & Baller, RD (2009) Contemporary Hate Crimes, Law Enforcement, and the Legacy of Racial Violence. *American Sociological Review*, 74(2), 291–315
- Klumpp T and Mialon HM, (2013) On Hatred. *American Law and Economics Review*, 15(1), 39-72.

Lazer D, Kennedy R, King G, and Vespignani A (2014) The parable of Google Flu: traps in big data analysis. *Science*, 343(6176), pp.1203-1205.

Krueger AB and Pischke JS (1997) A statistical analysis of crime against foreigners in unified Germany. *Journal of Human Resources*, 182-209.

Medoff M (1999) Allocation of Time and Hateful Behavior: A Theoretical Analysis of Hate and Hate Crimes. *American Journal of Economics and Sociology*, 58, 959–73.

Parks GS and Jones SE (2008) ‘Nigger: A Critical Race Realist Analysis of the N-Word Within Hate Crime Law. *Journal of Criminal Law and Criminology* 98(4), 1305–1352.

Pew Research Centers (PRC) (2017) Internet Usage in the United States.
<http://www.pewinternet.org/fact-sheet/internet-broadband/>

Preis T, Moat HS, & Stanley HE (2013) Quantifying Trading Behavior in Financial Markets Using Google Trends. *Scientific Reports*, 3.

Schwencke K (2018) Why America Fails at Gathering Hate Crime Statistics. ProPublica, January 19, 2018, <https://www.propublica.org/article/why-america-fails-at-gathering-hate-crime-statistics>

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

Stephens-Davidowitz SI (2017) Everybody lies: Big Data, New Data, and What the Internet Can Tell Us About Who We Really Are. New York: Harper Collins, 2017.

Stephens-Davidowitz SI, and Varian H (2014) A Hands-on Guide to Google Data.
<http://people.ischool.berkeley.edu/~hal/Papers/2015/primer.pdf>.

Stovel K (2001) Local Sequential Patterns: The Structure of Lynching in the Deep South, 1882–1930. *Social Forces*, 79(3), 843–880.

U.S. Department of Justice (2017a) Criminal Victimization, 2016. Office of Justice Programs, <https://www.bjs.gov/content/pub/pdf/cv16.pdf> (accessed 19 January 2018).

U.S. Department of Justice (2017b) Hate Crime Victimization, 2004-15. Office of Justice Programs, <https://www.bjs.gov/content/pub/pdf/hcv0415.pdf> (accessed 19 January 2018).

Utsuro T, Zhao C, Xu L, Li J, and Kawada Y (2017) An Empirical Analysis on Comparing Market Share with Concerns on Companies Measured Through Search Engine Suggests. *Global Journal of Flexible Systems Management*, 18(1), 3–19.

Vosen S, and Schmidt T (2011) Forecasting private consumption: survey-based indicators vs. Google trends. *Journal of Forecasting*, 30(6), 565–578.

Waldron J (2012) *The Harm in Hate Speech*. Cambridge: Harvard University Press.

Zaykowski H (2010) Racial Disparities in Hate Crime Reporting. *Violence and Victims; New York*, 25(3), 378–394.

Cases

Virginia v. Black, 538 US 343 (2003).

State Statutes

D.C. Code §22-3703 (2012)

N.D. Cent. Code § 12.1-14-04 (2017)).

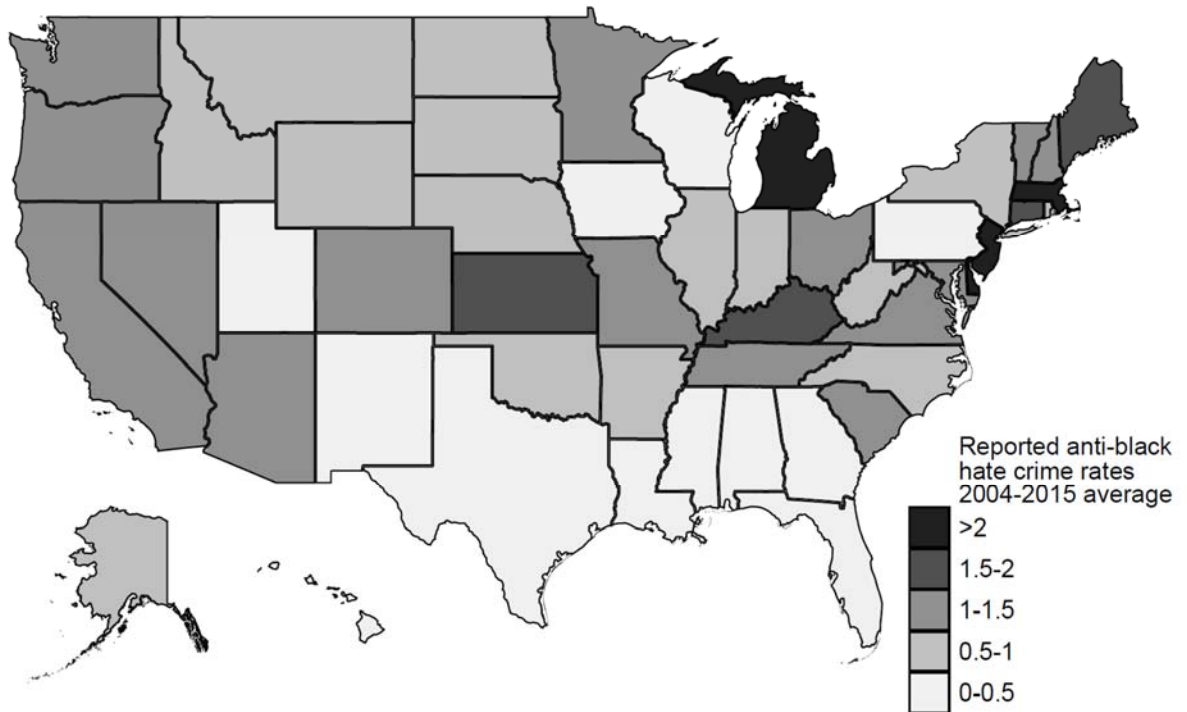
Utah Code Ann. § 76-3-203.4 (2017)

Federal Statutes

Hate Crime Statistics Act, 34 USC § 534 note (2012).


Matthew Shepherd and James Byrd, Jr., Hate Crimes Prevention Act, Pub L 111-84 (2009).

Figure 1: Rates of Reported Hate Crime (African-American Victims), Averaged by State over 2004-2015




Note: This map depicts the average reported hate crime rate for each state. Annual data on the number of African-American victims of racially-motivated hate crimes (scaled by state population) is averaged by state over the 2004-2015 period to obtain the average reported hate crime rate for each state. The data on reported hate crimes is from the Inter-university Consortium for Political and Social Research (ICPSR).

Figure 2: Google Search Rates for “N_____”, Averaged by State over 2004-2015




Note: This map depicts the average value for each state of the racial animus measure defined in Equation (4). Annual data on this measure is averaged by state over the 2004-2015 period to obtain an average level of racial animus for each state. The racial animus measure uses data on Google searches for “n_____,” obtained using the Google Trends app.

Figure 3: Rates of Reported Hate Crime (African-American Victims) in California and Mississippi, 2004-2015




Note: This figure depicts the reported hate crime rates for two states - California and Mississippi. In particular, it uses annual data for each of these states on the number of African-American victims of racially-motivated hate crimes (scaled by state population) for each year in our sample period (2004-2015). The data on reported hate crimes is from the Inter-university Consortium for Political and Social Research (ICPSR).

Figure 4: Rates of Reported Hate Crime (African-American Victims), 2004-2015




Note: This figure depicts reported hate crime rates for the 11 former Confederate states and for the other 39 states and DC. Data on the number of African-American victims of racially-motivated hate crimes (scaled by state population) is averaged over states separately for Confederate and non-Confederate states, for each year in our sample period (2004-2015). The data on reported hate crimes is from the Inter-university Consortium for Political and Social Research (ICPSR).

Figure 5: Google Search Rates for “N_____” 2004-2015




Note: This figure depicts the racial animus measure defined in Equation (4) for the 11 former Confederate states and for the other 39 states and DC. Data on this measure is averaged over states separately for Confederate and non-Confederate states, for each year in our sample period (2004-2015). The racial animus measure uses data on Google searches for “n_____,” obtained using the Google Trends app.

Figure 6: Hate Crime Rates and Google Search Rates for “N_____”, 2004-2015




Note: This figure depicts a scatterplot of observations at the state-year level (for non-Confederate states only) on the racial animus measure defined in Equation (4) – which is on the horizontal axis – and the number of African-American victims of racially-motivated hate crimes (scaled by state population) – which is on the vertical axis. The racial animus measure uses data on Google searches for “n_____,” obtained using the Google Trends app. The data on reported hate crimes is from the Inter-university Consortium for Political and Social Research (ICPSR).

Figure 7: Mean Hate Crime Rates and Google Search Rates for Confederate States (2004-2015), Relative to the Fitted Line for Non-Confederate States




Note: This figure depicts the fitted line from a bivariate regression of the number of African-American victims of racially-motivated hate crimes (scaled by state population) on the racial animus measure defined in Equation (4). Observations are at the state-year level, for non-Confederate states only. The 95% confidence interval is also shown. In addition, the figure plots the average hate crime rate and the average racial animus measure for each of the 11 former Confederate states (averaged by state over 2004-2015). The racial animus measure uses data on Google searches for “n_____,” obtained using the Google Trends app. The data on reported hate crimes is from the Inter-university Consortium for Political and Social Research (ICPSR).

Figure 8: Predicted Mean Hate Crime Rates, Averaged by State over 2004-2015 (Using All Non-Confederate States as the Benchmark)




Note: This map depicts the predicted hate crime rate for each state, using all non-Confederate states as the benchmark jurisdictions, generated using the procedure described in the text.

Figure 9: Reported and Predicted Mean Hate Crime Rates for Confederate States, 2004-2015 (Using All Non-Confederate States as the Benchmark)



Note: This figure depicts the average reported and predicted hate crime rate for each of the 11 former Confederate states. Annual data on the number of African-American victims of racially-motivated hate crimes (scaled by state population) is averaged by state over the 2004-2015 period to obtain the average reported hate crime rate for each state. The predicted hate crime rate for each state, using all non-Confederate states as the benchmark jurisdictions, generated using the procedure described in the text. The data on reported hate crimes is from the Inter-university Consortium for Political and Social Research (ICPSR).

Figure 10: Reported and Predicted Mean Hate Crime Rates for Selected Non-Confederate States, 2004-2015 (Using All Non-Confederate States as the Benchmark)



Note: This figure depicts the average reported and predicted hate crime rate for selected non-Confederate states. Annual data on the number of African-American victims of racially-motivated hate crimes (scaled by state population) is averaged by state over the 2004-2015 period to obtain the average reported hate crime rate for each state. The predicted hate crime rate for each state, using all non-Confederate states as the benchmark jurisdictions, generated using the procedure described in the text. The data on reported hate crimes is from the Inter-university Consortium for Political and Social Research (ICPSR).

Table 1: Summary Statistics

Variable	Observations	Mean	Standard Deviation
Number of African-American Victims	612	56.70915	83.1503
Hate Crime Rate (African-American Victims)	612	0.9654077	0.7868782
Google Search Rate for “N_____”	612	74.06093	23.80281
Confederate=1	612	0.2156863	0.411634
South=1	612	0.3333333	0.4717901
“Good Reporter” States=1	612	0.254902	0.4361626
General Crime Rate	612	2.364942	2.953855
Unemployment Rate (%)	612	6.226634	2.113752
Mean Income (thousands of \$)	612	23.97207	2.398346
College-educated %	612	38.3388	4.285081
African-American Fraction	612	0.1185008	0.120905
State Population (millions)	612	6.028551	6.754069

Note: This table reports summary statistics for each of the variables used in the analysis. Observations are at the state-year level. The number of African-American victims is obtained from the Inter-university Consortium for Political and Social Research (ICPSR). The hate crime rate variable scales this number per 100,000 state population. The racial animus measure – defined in Equation (4) - uses data on Google searches for “n_____,” obtained using the Google Trends app. Confederate and Southern states are defined using historical sources and the Census Bureau definition of the Southern region. “Good reporter” states are based primarily on statutes regarding hate crime reporting, as described in the text. The general crime rate refers to all offenses, from the UCR database. The unemployment rate is from the Bureau of Labor Statistics, while mean income, the percentage college-educated, the fraction of residents who are African-American, and the state population are all from the Census Bureau’s Current Population survey.

Table 2: Hate Crime Rates (African-American Victims) and Google Search Rates for “N_____”

	(1)	(2)	(3)	(4)
	Dependent Variable: Hate Crime Rate (African-American Victims)		Dependent Variable: Google Search Rate for “N_____”	
Confederate=1	-0.45296*** (0.059)	-0.45449*** (0.116)	17.56620*** (1.718)	1.84890 (2.281)
Border=1	0.24102** (0.122)	0.38683*** (0.129)	19.09177*** (3.164)	5.92022** (2.371)
General Crime Rate		0.02844 (0.019)		-1.50010*** (0.402)
Unemployment Rate		0.07348*** (0.024)		-0.01213 (0.511)
Income		0.00639 (0.014)		-2.21973*** (0.377)
College-educated %		0.05291*** (0.008)		-2.66186*** (0.226)
State Population		-0.01909** (0.008)		-0.09098 (0.161)
Quadratic Function of African-American %?	No	Yes	No	Yes
Year Fixed Effects?	Yes	Yes	Yes	Yes
Observations	612	612	612	612
States	51	51	51	51
R-squared	0.136	0.239	0.442	0.638

Note: This table reports regressions for the hate crime rate (for African-American victims) and for the racial animus measure defined in Equation (4). Observations are at the state-year level. The hate crime data is obtained from the Inter-university Consortium for Political and Social Research (ICPSR). The hate crime rate variable scales this number per 100,000 state population. The racial animus measure – defined in Equation (4) - uses data on Google searches for “n_____,” obtained using the Google Trends app. Confederate and Southern states are defined using historical sources and the Census Bureau definition of the Southern region. The general crime rate refers to all offenses, from the UCR database. The unemployment rate is from the Bureau of Labor Statistics, while mean income, the percentage college-educated, the fraction of residents who are African-American, and the state population are all from the Census Bureau’s Current Population survey. Robust standard errors are in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Table 3: Regressions Predicting Hate Crime Rates

	(1)	(2)	(3)	(4)	(5)
	Dependent Variable: Hate Crime Rate (African-American Victims)				
	Linear Specification			Poisson Specification	
	All Non-Confederate States	“Good Reporter” States	“Good Reporter” States	All Non-Confederate States	“Good Reporter” States
Google Search Rate for “N _____”	0.00702*** (0.002)	0.01740** (0.007)		0.00652*** (0.002)	0.01270*** (0.005)
General Crime Rate	0.11309** (0.048)	0.09501** (0.046)	0.10778* (0.061)	0.13069*** (0.039)	0.11293** (0.057)
Unemployment Rate	0.05737** (0.027)	-0.04432 (0.055)	-0.04610 (0.060)	0.05343** (0.024)	-0.03271 (0.044)
Income	0.01494 (0.017)	-0.00480 (0.051)	-0.00124 (0.055)	0.00453 (0.019)	-0.00603 (0.042)
College-educated %	0.05911*** (0.011)	0.08981*** (0.023)	0.04688** (0.021)	0.05719*** (0.010)	0.07213*** (0.021)
State Population	-0.04252** (0.019)	-0.02321 (0.020)	-0.03861 (0.025)	-0.05029*** (0.016)	-0.03442 (0.025)
Quadratic Function of African-American %?	Yes	Yes	Yes	Yes	Yes
Year Fixed Effects?	Yes	Yes	Yes	Yes	Yes
Observations	480	156	156	480	156
States	40	13	13	40	13
R-squared	0.195	0.221	0.167		
Adjusted R-squared	0.162	0.112	0.058		

Note: This table reports regressions for the hate crime rate (for African-American victims). Observations are at the state-year level. The hate crime data is obtained from the Inter-university Consortium for Political and Social Research (ICPSR). The hate crime rate variable scales this number per 100,000 state population. The racial animus measure – defined in Equation (4) - uses data on Google searches for “n _____,” obtained using the Google Trends app. Confederate states are defined using historical sources. “Good reporter” states are based primarily on statutes regarding hate crime reporting, as described in the text. The general crime rate refers to all offenses, from the UCR database. The unemployment rate is from the Bureau of Labor Statistics, while mean income, the percentage college-educated, the fraction of residents who are African-American, and the state population are all from the Census Bureau’s Current Population survey. Robust standard errors are in parentheses; *** p<0.01, ** p<0.05, * p<0.1.

Table 4: Mean Annual Imputed Hate Crimes and Hate Crime Rates for Selected States, 2004-2015

State	Mean Annual Reported Hate Crime Rate	Mean Annual Predicted Hate Crime Rate	Mean Annual Reported African-American Hate Crime Victims	Mean Annual Predicted African-American Hate Crime Victims	Mean Annual Imputed Number of Unreported Hate Crimes (African-American Victims)
<i>Confederate</i>					
Alabama	0.16	1.41	7.58	68.59	61.00
Arkansas	0.85	1.21	24.25	36.17	11.92
Florida	0.30	1.68	56.33	340.37	284.03
Georgia	0.12	1.54	11.33	157.32	145.99
Louisiana	0.22	1.28	9.67	59.78	50.11
Mississippi	0.03	1.26	1	37.58	36.58
North Carolina	0.75	1.56	70	156.07	86.07
South Carolina	1.04	1.58	46.92	77.38	30.46
Tennessee	1.26	1.39	79.08	91.69	12.60
Texas	0.38	1.28	92.17	351.60	259.44
Virginia	1.30	1.36	102.42	114.13	11.72
<i>Other</i>					
Iowa	0.41	0.85	12.42	26.53	14.11
Illinois	0.57	1.26	72.67	161.39	88.73
Indiana	0.53	1.01	33.92	66.68	32.76
New York	0.64	1.15	124.58	226.24	101.66
Pennsylvania	0.29	1.02	36.17	130.89	94.73
Rhode Island	0.50	1.22	5.33	12.88	7.55
Wisconsin	0.44	1.01	25.08	57.97	32.89

Note: This table reports the average reported and predicted hate crime rate (per 100,000 state population) for each of the 11 former Confederate states and for selected non-Confederate states (rounded to 2 decimal places). Annual data on the number of African-American victims of racially-motivated hate crimes (scaled by state population) is averaged by state over the 2004-2015 period to obtain the average reported hate crime rate for each state. The predicted hate crime rate for each state, using all non-Confederate states as the benchmark jurisdictions, is generated using the procedure described in the text. The data on reported hate crimes is from the Inter-university Consortium for Political and Social Research (ICPSR).