DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD
DASH.HARVARD.EDU

HARVARD LIBRARY
Office for Scholarly Communication

Essays Using Google Data

Citation
Stephens-Davidowitz, Seth Isaac. 2013. Essays Using Google Data. Doctoral dissertation, Harvard University.

Permanent link
http://nrs.harvard.edu/urn-3:HUL.InstRepos:10984881

Terms of Use
This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA

Share Your Story
The Harvard community has made this article openly available.
Please share how this access benefits you. Submit a story .

Accessibility

Essays Using Google Data

A dissertation presented

by

Seth Stephens-Davidowitz

to

The Department of Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Economics

Harvard University

Cambridge, Massachusetts

April 2013

© 2013 Seth Stephens-Davidowitz
All rights reserved.

Contents

List of Tables

List of Figures

Acknowledgments

Alberto Alesina was my first advisor and guided me throughout my PhD program. He was instrumental in my choosing research topics (racism and political economy). Alberto's world-class ability to understand people and how to talk to them is unheard of in a top economics program and is something that I have attempted, without much success, to emulate.

David Cutler, despite being a health economist, agreed to help on all my projects. Although the subject matter was often not his main area of study, he had an uncanny ability to point me in the right direction – where to search for the perfect quote to start my job market paper, for example.

Edward Glaeser is superhuman. He achieves mastery at anything he does, including writing papers, giving talks, and advising. Hearing Ed talk and explore ideas – about cities, politics, religion, media, and marriage – was one of the great thrills of graduate school.

Lawrence Katz is the most logical person that I have ever met. If Larry thinks what you are saying is true and interesting, he will nod yes; otherwise, he will not nod. Larry's beliefs about what are true and interesting correlate almost perfectly with reality. I realized, near the end of the graduate program, that I could give a presentation and, just by watching Larry's nods, know what to cut and what to keep.

Other faculty who helped me include Raj Chetty, Claudia Goldin, and Lawrence Summers.

Hal Varian, upon learning of my research, invited me to give a talk at Google. Later, he offered me a position. I found, shortly after arriving at Google, that a huge number of people count Hal as a mentor. Pretty soon, I began to learn why. His clear thinking has taught me much in a short time.

My brother, Noah Stephens-Davidowitz, read all my papers, discussed numerous ideas, and often calmed me down. Lauren Stephens-Davidowitz and other family members also gave me much love and support.

Rebecca Kahane dated me through much of my PhD program. She is an angel of compassion.

I thank all my friends, and especially Sam Asher, Zinhle Dlamini, Joshua Gottlieb, Nathaniel Hilger, Pan-Pan Jiang, Shira Lazarus, Jacob Leshno, Noah Popp, Ramon Roullard, and Greg Sobolski.

To my parents, Esther Davidowitz and Mitchell Stephens

Introduction

A little more than one decade ago, Google was founded with an ambitious goal: to organize the world's information.

Its success is well-documented.

There is a little-discussed side effect of Google's staggering success: Google not only finds and organizes existing information; its existence has created important new information. How people interact with the world's great information database is itself immensely informative.

Most Americans, and increasing numbers from the rest of the world, turn to Google for information on health, politicians, and jobs. They look for pornography, jokes, and other entertainment. They shop for cars, nannies, and shoes. They learn of ways to cheat on their spouse and ways to catch their cheating spouse. It does not seem extreme to call Google search queries the most impressive dataset ever collected on the human psyche.

Yet, this data has only been used in a handful of papers focused on its immediate availability (Varian and Choi, 2010; Ginsberg et al., 2009). The first Google papers, while undeniably important and trailblazing, only scratch the surface of the dataset's potential. The goal of this dissertation is to add another scratch or two.

Chapter 1 suggests that Google can be used to measure racial animus against African-Americans. The Google measure shows that racial animus cost Barack Obama significantly more votes than found by surveys.

Chapter 2 suggests that Google can be used to predict election turnout. These predictions prove stronger than other available indicators.

Chapter 3 suggests that Google can be used to measure child maltreatment. I argue that, contrary to evidence using official data, the recent economic downturn significantly increased child maltreatment in the United States.

Since completing this dissertation, I have continued scratching away at this dataset, as an intern at Google. It is hoped that each of these chapters encourages more scratches from more scratchers. There is so much more to be learned.

Chapter 1

The Cost of Racial Animus on a Black Candidate: Evidence Using Google Data

I Introduction

Does racial animus cost a black candidate a substantial number of votes in contemporary America? The most recent review of the literature is inconclusive: "Despite considerable effort by numerous researchers over several decades, there is still no widely accepted answer as to whether or not prejudice against blacks remains a potent factor within American politics" (Huddy and Feldman, 2009).

There are two main reasons the answer to this question is of interest to scholars: first, it would help us better understand the extent of contemporary prejudice[1]; second, it would increase our understanding of the determinants of voting.[2] There is one main reason the question has proven so difficult: individuals' tendency to withhold socially unacceptable attitudes, such as negative feelings

[1]Charles and Guryan (2011) surveys some of the voluminous literature studying modern discrimination. Creative field environments used to study discrimination include NBA referees (Price and Wolfers, 2010); baseball umpires (Parsons et al., 2011); baseball card sales (List, 2004); motor vehicle searches (Knowles et al., 2001); and employers receiving manipulated resumes (Bertrand and Mullainathan, 2004).

[2]Rational choice theory says that economic impacts of outcomes fully determine voting. A number of scholars have previously found important deviations from an extreme interpretation of this model (Benjamin and Shapiro, 2009; Alesina and Rosenthal, 1995; Berggren et al., 2010; Wolfers, 2002).

towards blacks, from surveys (Tourangeau and Ting, 2007; Berinsky, 1999; Berinsky, 2002; Gilens et al., 1998; Kuklinski et al., 1997).

This paper uses non-survey-based methodology. I suggest a data source not previously used to study prejudice. I proxy an area's racial animus based on the percent of Google search queries that include racially charged language. I compare the proxy to Barack Obama's presidential vote shares, controlling for the previous Democratic candidate, John Kerry's, presidential vote share. This empirical specification is most similar to that of Mas and Moretti (2009). They use a survey measure of support for a law banning interracial marriage from the General Social Survey (GSS) as their state-level proxy for racial attitudes. They do not find evidence that racial attitudes affected Obama's 2008 vote share.

Google data, evidence suggests, are unlikely to suffer from major social censoring: Google searchers are online and likely alone, both of which make it easier to express socially taboo thoughts (Kreuter et al., 2009). Individuals, indeed, note that they are unusually forthcoming with Google (Conti and Sobiesk, 2007). The large number of searches for pornography and sensitive health information adds additional evidence that Google searchers express interests not easily elicited by other means. Furthermore, aggregating information from millions of searches, Google can meaningfully reveal social patterns. The percent of Google searches that include the word "God," for example, explains more than 60 percent of areas' variation in belief in God.

I define an area's racially charged search rate as the percent of Google searches, from 2004-2007, that included the word "nigger" or "niggers." I choose the most salient word to constrain data-mining.[3] I do not include data after 2007 to avoid capturing reverse causation, with dislike for Obama causing individuals to use racially charged language on Google.

The epithet is searched for with some frequency on Google. From 2004-2007, the word "nigger(s)" was included in roughly the same number of Google searches as words and phrases such as "migraine(s)," "economist," "sweater," "Daily Show," and "Lakers." The most common searches that include the epithet, such as "nigger jokes" and "I hate niggers," return websites with derogatory

[3] Kennedy (2003, p.22) says this is "the best known of the American language's many racial insults ... the paradigmatic slur."

material about African-Americans. From 2004-2007, the searches were most popular in West Virginia; upstate New York; rural Illinois; eastern Ohio; southern Mississippi; western Pennsylvania; and southern Oklahoma.

Racially charged search rate is a significant, negative predictor of Obama's 2008 and 2012 vote shares, controlling for Kerry's 2004 vote share. The result is robust to controls for changes in unemployment rates; home-state candidate preference; Census division fixed effects; demographic controls; and long-term trends in Democratic voting. The estimated effect is somewhat larger when adding controls for an area's Google search rates for other terms that are moderately correlated with search rate for "nigger" but are not evidence for racial animus. In particular, I control for search rates for "African American," "nigga," (the alternate spelling used in nearly all rap songs that include the word), and profane language.

A non-racial explanation for the results might be that areas with higher racially charged search rates became less likely, during this time period, to support Democratic candidates, more generally. This, though, does not fit the evidence. There is not a significant relationship between an area's racially charged search rate and changes in either House Democratic vote shares or measured liberalism over the same time period.

The preferred point estimates imply that, relative to the most racially tolerant areas in the United States, prejudice cost Obama 4.2 percentage points of the national popular vote in 2008 and 4.0 percentage points in 2012. These numbers imply that, among white voters who would have supported a white Democratic presidential candidate in 2008 (2012), 9.1 (9.5) percent did not support a black Democratic presidential candidate.

Obama lost substantially more votes from racial animus, I argue, than he gained from his race. Back-of-the-envelope calculations suggest Obama gained at most only about one percentage point of the popular vote from increased African-American support. The effect was limited by African-Americans constituting less than 13 percent of the population and overwhelmingly supporting every Democratic candidate. Evidence from other research, as well as some new analysis in this paper, suggest that few white voters swung in Obama's favor in the 2008 or 2012 *general* elections due to his race.[4]

This paper builds on and contributes to the large literature, reviewed by Huddy and Feldman (2009), testing for the effects of racial attitudes on black candidates.[5] In addition, the new proxy of area-level prejudice might be useful to literatures in social economics (Alesina et al., 2001; Alesina and La Ferrara, 2002), labor economics (Charles and Guryan, 2008), and urban economics (Cutler et al., 1999; Card et al., 2008).

More generally, this paper adds further support for a potentially large role for Google data in the social sciences. Previous papers using the data source have tended to note correlations between Google searches and other data (Ginsberg et al., 2009; Seifter et al., 2010; Varian and Choi, 2010; Scheitle, 2011). This paper shows clearly that Google search query data can do more than correlate with existing measures; on socially sensitive topics, they can give better data and open new research on old questions. If I am correct that the Google database contains the best evidence on such a well-examined question, that the Google database might contain the best evidence on many important questions does not seem such a large leap.

II Google-Search Proxy For an Area's Racial Animus

II.A Motivation

Before discussing the proxy for racial animus, I motivate using Google data to proxy a socially sensitive attitude. In 2007, nearly 70 percent of Americans had access to the internet at home (CPS, 2007). More than half of searches in 2007 were performed on Google (Burns, 2007). Google searchers are somewhat more likely to be affluent, though large numbers of all demographics use

[4]The effect of race on the overall probability of being elected president would also have to consider the effects of race on primary voting and on fundraising. These questions are beyond the scope of this paper.

[5]More recent papers compare individuals' self-reported racial attitudes near the time of the election to self-reported decision to support Obama and various controls (Piston, 2010; Pasek et al., 2010; Schaffner, 2011; Lewis-Beck et al., 2010; Kinder and Dale-Riddle, 2012; Tesler and Sears, 2010a). They generally find smaller effects than the effects found here, suggesting that individual surveys fail to fully capture the effects of racial attitudes. In addition, these papers may be open to the critique of Schuman (2000) and Feldman and Huddy (2005), discussed in Huddy and Feldman (2009): surveys' measures of prejudice, such as saying that African-Americans would be more successful if they tried harder, may capture omitted conservative ideology. The unambiguous proxy of racial animus and the fact that I control for administrative vote data in the 2004 presidential election greatly limit this argument. This is discussed in more detail in Section IV.A

the service (Hopkins, 2008).

Aggregating millions of searches, Google search data consistently correlate strongly with demographics of those one might most expect to perform the searches (See Table 1.1). Search rate for the word "God" explains 65 percent of the variation in percent of a state's residents believing in God. Search rate for "gun" explains 62 percent of the variation in a state's gun ownership rate. These high signal-to-noise ratios hold despite some searchers typing the words for reasons unrelated to religion or firearms and not all religious individuals or gun owners actually including the term in a Google search (The 'top search' for "God" is "God of War," a video game. The 'top search' for "gun" is "Smoking Gun," a website that reveals sensational, crime-related documents.) If a certain group is more likely to use a term on Google, aggregating millions of searches and dividing by total searches will give a good proxy for that group's area-level population.

Furthermore, evidence strongly suggests that Google elicits socially sensitive attitudes. As mentioned in the Introduction, the conditions under which people search – online, likely alone, and not participating in an official survey – limit concern of social censoring. The popularity of search terms related to sensitive topics further supports this use. The word "porn," for example, is included in more searches in the United States than the word "weather."[6]

II.B Proxy

I define an area's racially charged search rate as the percentage of its Google searches, from 2004-2007, that included the word "nigger" or its plural.[7]

---

[6]Only about 20 percent of Americans admit to the GSS that they have watched a pornographic movie within the past year.

[7]As mentioned in the Introduction, data prior to 2008 are used to avoid capturing reverse causation. About five percent of searches including "nigger" in 2008 also included the word "Obama," suggesting feelings towards Obama were a factor in racially charged search in 2008. Searches including both the epithet and "Obama" were virtually non-existent in 2007. It is also worth noting that area-level search rates for the racial epithet are highly correlated through time, and any choice of dates will yield roughly similar results. For example, the correlation between 2004-2007 and 2008-present state-level racially charged search rate is 0.94. Using just one word or phrase, even one that can be used for different reasons, to proxy an underlying attitude builds on the work of scholars who have conducted text analysis of newspapers. For example, Saiz and Simonsohn (2008) argue that news stories about a city that include the word "corruption" can proxy a city's corruption. And Gentzkow et al. (2011) show that, historically, Republican (Democratic) newspapers include significantly more mentions of Republican (Democratic) presidential candidates.

Table 1.1: *Signal-to-Noise Ratio in Google Search Terms*

Term	*Underlying Variable*	*t-stat*	R^2
God	Percent Believe in God	8.45	0.65
Gun	Percent Own Gun	8.94	0.62
African American(s)	Percent Black	13.15	0.78
Hispanic	Percent Hispanic	8.71	0.61
Jewish	Percent Jewish	17.08	0.86

Notes: The t-stat and R^2 are from a regression with the normalized search volume of the word(s) in the first column as the independent variable and measures of the value in the second column as the dependent variable. The normalized search volume for all terms are from 2004-2007. All data are at the state level. Percent Black are Percent Hispanic are from the American Community Survey, for 2008; the Jewish population is from 2002, gun ownership from 2001, and belief in God from 2007. Jewish data are missing one observation (South Dakota); belief in God data are missing for 10 states. The data for belief in God, percent Jewish, and percent owning guns can be found at `http://pewforum.org/how-religious-is-your-state-.aspx`, `http://www.jewishvirtuallibrary.org/jsource/US-Israel/usjewpop.html`, and `http://www.washingtonpost.com/wp-srv/health/interactives/guns/ownership.html`, respectively.

$$Racially\ Charged\ Search\ Rate_j = 100 \cdot \frac{\left[\frac{Google\ searches\ including\ the\ word\ "nigger(s)"}{Total\ Google\ searches}\right]_j}{\left[\frac{Google\ Searches\ including\ the\ word\ "nigger(s)"}{Total\ Google\ searches}\right]_{max}} \tag{1.1}$$

The racial epithet is a fairly common word used in Google search queries: It is now included in more than 7 million searches annually.[8] Figure 1.1 shows terms included in a similar number of searches, from 2004-2007, as the racial epithet.[9]

The word "migraine" was included in about 30 percent fewer searches. The word "Lakers" and the phrase "Daily Show" were each included in about five percent more searches than the racial epithet.[10] While these words and phrases were chosen rather arbitrarily as benchmarks, the number

[8]These are approximations calculated using AdWords. It combines searches on 'Desktop and Laptops' and 'Mobile devices.'

[9]The percentage of Google searches including the racial epithet was roughly constant from 2004 through 2008. There were, though, notable spikes in the days following Hurricane Katrina and in early November 2008, particularly on Election Day. The percentage of Google searches including the term dropped after the 2008 election and has consistently been about 20 percent *lower* during Obama's presidency than prior to his presidency. An emerging literature is examining how Obama's presidency has affected racial attitudes (DellaVigna, 2010; Valentino and Brader, 2011; Tesler, 2012; Tesler and Sears, 2010b).

[10]Google data are case-insensitive. So I am comparing the racial epithet to searches that include either "lakers" or

Figure 1.1: *Selected Words and Phrases Included in Google Searches Roughly as Frequently as "nigger(s)," 2004-2007*

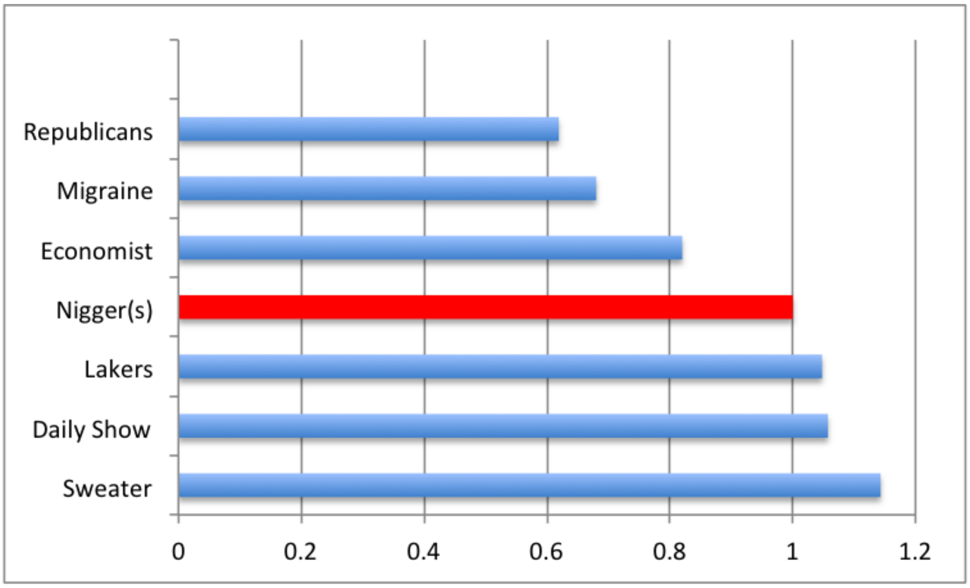

Notes: This figure shows selected words and phrases included in a similar number of searches, from 2004-2007, as "nigger(s)." The number corresponds to the ratio of total Google searches that include that word to total Google searches that include the racial epithet. "Daily Show," for example, was included in about 6 % more searches than the racial epithet. "Economist" was included in about 20 % fewer searches. It is worth emphasizing again that this counts any searches including the word or phrase. So searches such as "The Daily Show" and "Daily Show clips" will be counted in the search total for "Daily Show." And Google considers searches case-insensitive. So "daily show" and "daily show clips" would also count. While the words included were rather arbitrarily selected, another benchmark to use is "weather." "Weather" was included in only about 81 times more searches than "nigger(s)" during this time period. All numbers presented were estimated using Google Trends.

of searches including the term can also be compared to the number of searches including one of the most common words, "weather." Search volume including the racial epithet, from 2004-2007, was within two orders of magnitude of search volume including "weather."

What are searchers looking for? About one quarter of the searches including the epithet, from 2004-2007, also included the word "jokes," searches that yield derogatory entertainment based on harsh African-American stereotypes. These same joke sites, with derogatory depictions of African-Americans, are also among the top returns for a Google search of just the epithet or its plural,

"Lakers."

representing about 10 percent of total searches that included the epithet.[11] More information on the searches can also be gleaned from the 'top searches,' the most common searches before or after searches including the word (See Table 1.2). Searchers are consistently looking for entertainment featuring derogatory depictions of African-Americans. The top hits for the top racially charged searches, in fact, are nearly all textbook examples of antilocution, a majority group's sharing stereotype-based jokes using coarse language outside a minority group's presence. This was determined as the first stage of prejudice in Allport's (1979) classic treatise.

Table 1.2: *Top Searches for "nigger(s)"*

Rank	*'04-'07 Search* DATA USED	*'08-'11 Search* DATA NOT USED
1.	jokes	jokes
2.	nigger jokes	nigger jokes
3.	white nigger	obama nigger
4.	nigga	nigga
5.	hate niggers	black nigger
6.	i hate niggers	funny nigger
7.	black jokes	nigger song
8.	the word nigger	the word nigger
9.	racist jokes	nas nigger
10.	kkk	i hate niggers

Notes: This table shows the 'top searches' for "nigger(s)." 2004-2007 is the time period for the search volume used in the regressions and figures to limit reverse causation. Results would be similar regardless of time period selected, as the state-level correlation between the two periods is 0.94. Depending on the draw, the 'top searches' might be slightly different. Top searches, according to Google, 'are related to the term,' as determined 'by examining searches that have been conducted by a large group of users preceding the search term you've entered, as well as after,' as well as by automatic categorization.

I obtain data for all 51 states and 196 of 210 media markets, encompassing more than 99 percent of American voters.[12] I use media-market-level regressions when other data sources are available at

[11] I do not know the order of sites prior to my beginning this project, in June 2011. The ordering of sites for searches of just the epithet has changed slightly, from June 2011-April 2012 . For example, while joke sites were the second, third, and fourth returns for a search for "niggers" in June 2011, these sites were passed by an Urban Dictionary discussion of the word by April 2012.

[12] Google Trends says that the media market data corresponds to measures of Arbitron. I have confirmed that they

the media-market level and state data when such data are not available.[13]

Racially charged search rates, for the 50 states and the District of Columbia, are shown in Table 1.3. Racially charged search rates for media markets are shown in Figure 1.2. The search rate was highest in West Virginia; upstate New York; rural Illinois; eastern Ohio; southern Mississippi; western Pennsylvania; and southern Oklahoma. The search rate was lowest in Laredo, TX – a largely Hispanic media market; Hawaii; parts of California; Utah; and urban Colorado.

II.C Predictors of Racially Charged Search Rate

Comparisons with GSS

Figure 1.3 compares the Google-based proxy to the GSS measure of Mas and Moretti (2009). Since the GSS only includes data for 44 states plus the District of Columbia, the figures and regressions only include 45 observations. The Google measure has a correlation of 0.6 with the measure of Mas and Moretti (2009), support for a law banning interracial marriage from 1990 to 2004.[14]

Some of the outliers are likely due to small samples for some states using GSS data. For example, Wyoming ranks as significantly more racially prejudiced using the Mas and Moretti (2009) proxy than the Google proxy. However, only 8 white individuals living in Wyoming were asked this question by the GSS. (Two, or twenty-five percent, said they supported a law banning interracial marriage.)

The GSS and Google proxies for racial prejudice noticeably differ in their relationship with ideology. The GSS supports some popular wisdom that racial prejudice is now a larger factor among Republicans than Democrats: The higher Kerry's 2004 vote share in a state, the lower the percentage of whites admitting opposition to interracial marriage. In contrast, there is no statistically

---

actually correspond to designated media markets, as defined by Nielsen. I match other data to the media markets using Gentzkow and Shapiro (2008).

[13] It should be noted that some of this data are not easily obtained. If total number of searches, for a given area and time period, is below an unreported, but clearly high, threshold, Google does not report the data. In the Appendix, I show what I think is the first algorithm for obtaining data that does not cross the threshold.

[14] The Google measure has a correlation of 0.66 with the measure of Charles and Guryan (2008), average prejudice from 1972 to 2004. I thank the authors for providing their data.

Table 1.3: *Racially Charged Search Rate, State*

Rank	*State*	*Racially Charged Search Rate*
1.	West Virginia	100
2.	Louisiana	86
3.	Pennsylvania	85
4.	Mississippi	83
5.	Kentucky	82
6.	Michigan	78
7.	Ohio	78
8.	South Carolina	76
9.	Alabama	76
10.	New Jersey	74
11.	Tennessee	73
12.	Florida	71
13.	New York	71
14.	Rhode Island	70
15.	Arkansas	70
16.	North Carolina	69
17.	Georgia	69
18.	Connecticut	68
19.	Missouri	68
20.	Nevada	67
21.	Illinois	65
22.	Delaware	65
23.	Oklahoma	65
24.	Maryland	64
25.	Indiana	63
26.	Wisconsin	63
27.	Kansas	62
28.	Texas	62
29.	Virginia	59
30.	Vermont	59
31.	California	57
32.	Maine	56
33.	Nebraska	55
34.	New Hampshire	54
35.	North Dakota	54
36.	Iowa	53
37.	Massachusetts	52
38.	Arizona	51
39.	Washington	50
40.	South Dakota	50
41.	Alaska	50
42.	Wyoming	48
43.	Montana	48
44.	Oregon	47
45.	Minnesota	46
46.	District of Columbia	44
47.	Idaho	39
48.	New Mexico	39
49.	Colorado	39
50.	Hawaii	34
51.	Utah	30

Notes: *Racially Charged Search Rate* is Web Search, from January 2004-December 2007, for either "nigger" or "niggers." This data can be found here: `http://www.google.com/insights/search/#q=nigger%2Bniggers%2C%2C%2C%2C&geo=US&date=1%2F2004%2048m&cmpt=q`.

Figure 1.2: *Racially Charged Search Rate, Media Market*

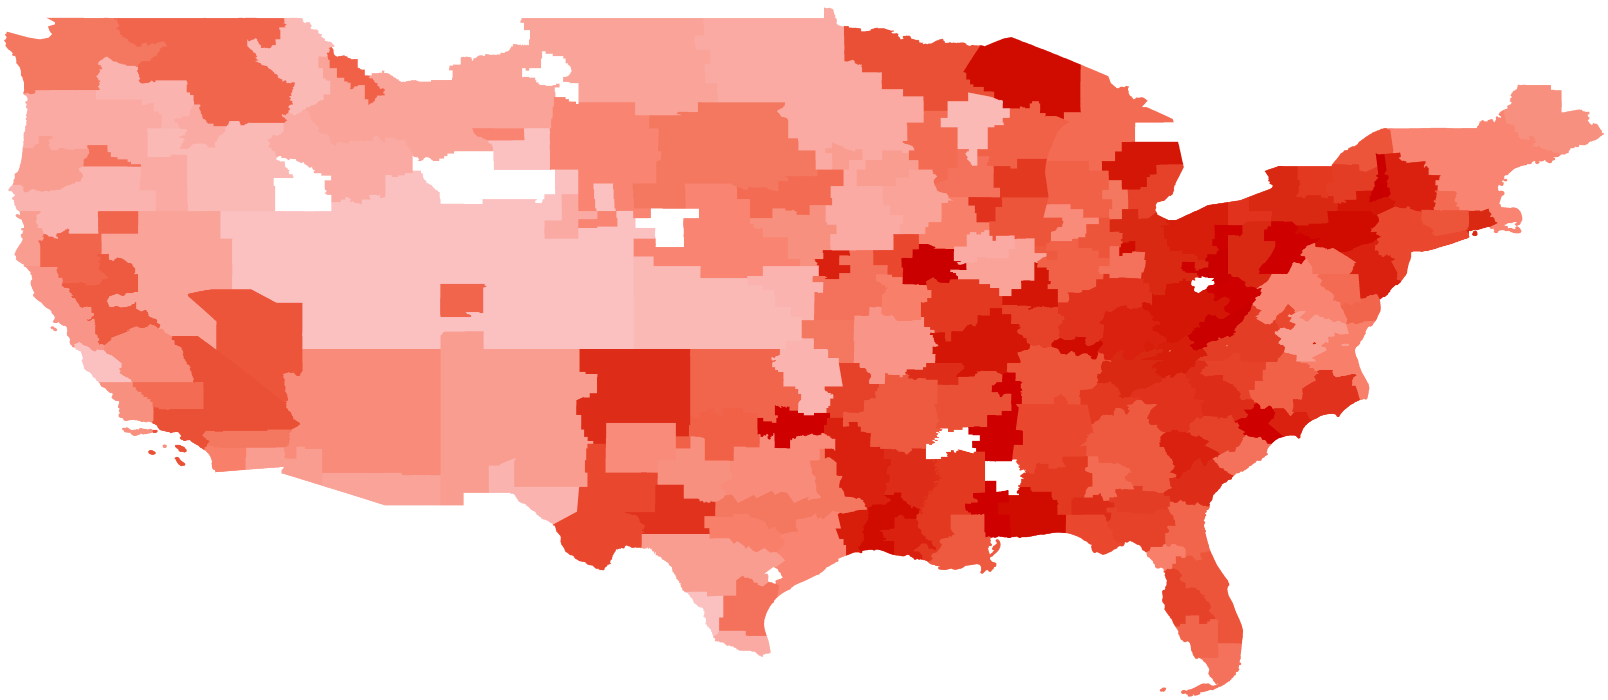

Notes: This maps search volume for "nigger(s)," from 2004-2007, at the media market level. Darker areas signify higher search volume. White areas signify media markets without data. Alaska and Hawaii, for which data are available, are not shown.

Figure 1.3: *Google Racially Charged Search Compared to GSS Opposition to Interracial Marriage*

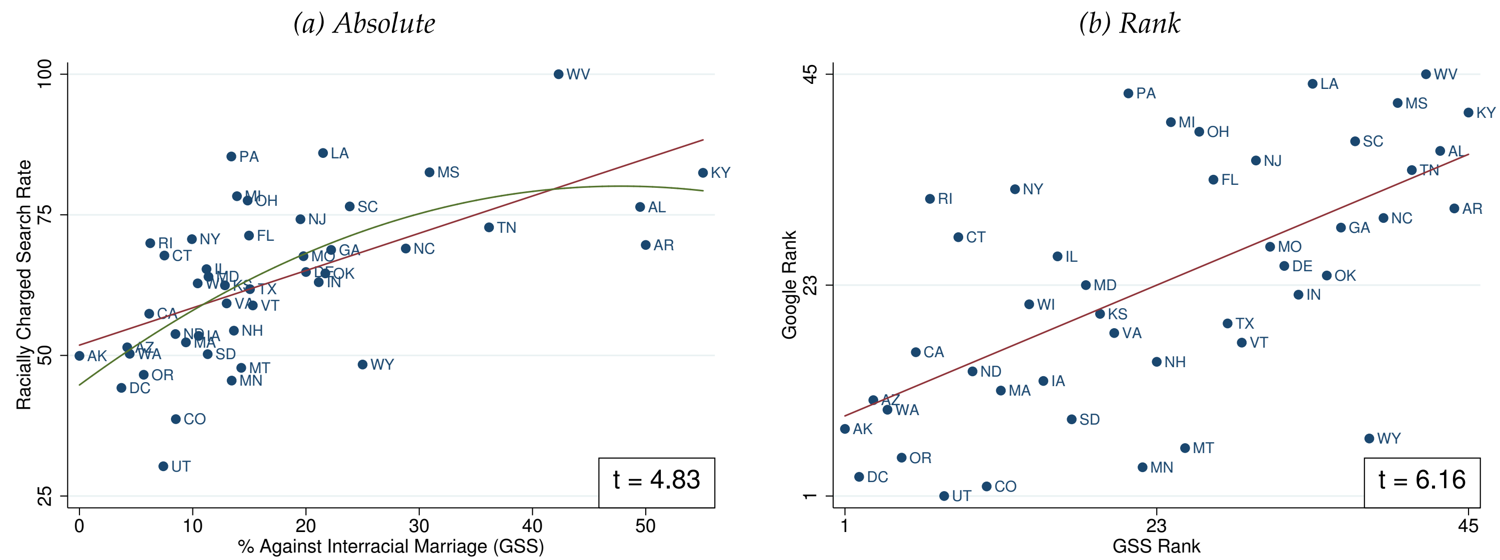

Notes: The x-axis in panels (a) is the measure of racial attitudes used in Mas and Moretti (2009): percent of whites, from 1990-2004, supporting a law banning interracial marriage. The x-axis in panel (b) is the rank of the 45 states for this measure, with higher numbers corresponding to higher measures of racial prejudice. Thus, the value 45 in Panel (b) means that state (Kentucky) had the highest percentage of whites telling the GSS they supported a law banning interracial marriage. The y-axis for panel (a) uses the unrounded number in Table 1.3 for the 45 states for which GSS data are available; The-y axis panel (b) is the rank of racially charged search for these 45 states, with higher numbers corresponding to higher racially charged search rates.

significantly correlation between Kerry 2004 vote share and racially charged search rate, at either the state or media market level.[15] One potential reason for this discrepancy is that racial prejudice is more socially unacceptable among Democrats. Thus, underreporting of prejudice in surveys will be more severe in areas with more Democrats. And surveys, such as the GSS, will falsely find a negative correlation between percent Democrat and racial prejudice.

Demographics and Use by African-Americans

Table 1.4 shows the demographic predictors of racially charged search rate at the media market level. The demographic factor correlating strongest with racially charged search rate is the percentage of the population with a bachelor's degree. A 10 percentage point increase in college graduates is correlated with almost a one standard deviation decrease in racially charged search rate. Younger and more Hispanic areas are less likely to search the term.

There is a small positive correlation between racially charged search rate and percent black. Readers may be concerned that this is due to African-Americans searching the term, limiting the value of the proxy. This is unlikely to be a major factor: the common term used in African-American culture is "nigga(s)," which Google considers a separate search from the term ending in "er." (Rahman, 2011).[16] Table 1.5 shows the top searches for "nigga(s)." In contrast to the top searches for the term ending in "er," the top searches for "nigga(s)" are references to rap songs. Table 1.5 also shows that, even among the five percent of searches that include the epithet ending in "er" and also include the word "lyrics," the 'top searches' are for racially charged country music songs.

The positive correlation between racially charged search rate and percent black is better explained by racial threat, the theory that the presence of an out-group can threaten an in-group and create racial animosity (Key Jr., 1949; Glaser, 1994; Glaser and Gilens, 1997). Racial threat predicts a quadratic relationship between the percentage of the population that is black and racial animus (Blalock, 1967; Taylor, 1998; Huffman and Cohen, 2004; Enos, 2010). Zero African-Americans

[15]The lack of a relationship holds controlling for percent black, as well.

[16]Rap songs including the version ending in 'a' are roughly 45 times as common as rap songs including the version ending in 'er.' – Author's calculations based on searches at http://www.rapartists.com/lyrics/.

Table 1.4: *Predictors of an Area's Racially Charged Search Rate*

	Dependent Variable: Racially Charged Search Rate			
	(1)	(2)	(3)	(4)
Percent Age 65 or Older	6.884*	3.341	6.492*	3.757
	(3.650)	(3.447)	(3.668)	(3.495)
Percent w/ Bachelor's Degree	-9.309***	-8.532***	-10.104***	-9.459***
	(2.105)	(2.147)	(2.004)	(2.129)
Percent Hispanic	-2.620***	-2.298***	-2.659***	-2.297***
	(0.462)	(0.554)	(0.454)	(0.486)
Percent Black	2.556***	0.283	11.245***	6.734**
	(0.826)	(1.268)	(2.158)	(3.172)
(Percent Black)-squared			-24.731***	-16.517***
			(5.613)	(6.070)
Observations	196	196	196	196
R-squared	0.36	0.49	0.41	0.50
Census Div. FE		X		X

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: Standard errors, clustered at the state level, are in parentheses. Racially Charged Search Rate is as defined in Equation 1.1, obtained by the algorithm described in Appendix IV, normalized to its z-score. The demographic variables are individuals in the group divided by total individuals; thus a one-unit change represents a change from 0 to 100 percent. The demographics variables are from the American Community Survey '05-'09. All county-level variables are aggregated to the media market level using Gentzkow and Shapiro (2008).

Table 1.5: *Music, "nigger," and "nigga," 2004-2007*

Rank	*Top searches for 'nigger lyrics'*	*Top searches for 'nigga(s)'*
1.	nigger song	nigga lyrics
2.	nigger song lyrics	my nigga
3.	nigger jokes	niggas lyrics
4.	white nigger	hood nigga
5.	nigger hatin me	my niggas
6.	white nigger lyrics	lyrics hood nigga
7.	johnny rebel lyrics	nigga stole
8.	johnny rebel	nigga stole my
9.	david allen coe	my nigga lyrics
10.	lyrics alabama nigger	nigga what

Notes: The second column shows the 'top searches' reported for searches including both "nigger" and "lyrics." The third column shows the 'top searches' reported for searches including either "nigga" or "niggas." The method for calculating 'top searches' is discussed in Table 1.2. Also noted there, depending on the particular draw, the ranks and terms might differ somewhat.

means race is not salient and racial animus may not form. Near 100 percent African-American communities have few white people; white individuals with racial animus are unlikely to choose such a community. Columns (3) and (4) of Table 1.4 offer support for this theory. Indeed, the preferred fit between racially charged search rate and percent black is quadratic. The numbers imply that racial animus is highest when African-Americans make up between 20 and 30 percent of the population. Three of the ten media markets with the highest racially charged search rate – Hattiesburg-Laurel, Biloxi-Gulfport, and Florence-Myrtle Beach – are between 20 and 30 percent black. Therefore, the relationship between racially charged search rate and percent black is consistent with racially charged search being a good proxy for racial animus.

III The Effects of Racial Animus on a Black Presidential Candidate

Section II argues that the frequency with which an area's Google searches include the word "nigger(s)" – a word, overall, used about as frequently in searches as terms such as "Daily Show" and "Lakers," with most of them returning derogatory material about African-Americans – give a strong proxy for an area's racial animus. This section uses the proxy to test the effects of racial animus on an election with a black candidate. The section focuses on the significance and robustness of the results. I hold off until Section IV in fully interpreting the magnitude of the effects.

III.A The Effects of Racial Animus on Black Vote Share

To test the effects of racial animus on a black candidate's vote share, I compare the proxy to the difference between an area's support for Barack Obama in 2008 and John Kerry in 2004. I show later that the estimated effects on Obama in 2012 were almost identical to the estimated effects on Obama in 2008.

Define $\%Obama2008_j$ as the percent of total two-party votes received by Obama in 2008 and $\%Kerry2004_j$ as the percent of total two-party votes received by Kerry in 2004. In other words, $\%Obama2008_j$ is an area's total votes for Obama divided by its total votes for Obama or John McCain. $\%Kerry2004_j$ is an area's total votes for Kerry divided by its total votes for Kerry or George W. Bush. Then $(\%Obama2008 - \%Kerry2004)_j$ is meant to capture an area's relative preference for a black compared to a white candidate.

The idea is that the different races of the Democratic candidates was a major difference between the 2004 and 2008 presidential races. The 2004 and 2008 presidential elections were similar in terms of perceived candidate ideology. In 2004, about 44 percent of Americans viewed John Kerry as liberal or extremely liberal. In 2008, about 43 percent viewed Barack Obama as liberal or extremely liberal.[17] There were slightly larger differences in perceived ideology of the Republican candidates. Roughly 59 percent viewed George W. Bush as conservative or very conservative in 2004; 46 percent viewed John McCain as conservative or very conservative in 2008. Neither Kerry nor Obama came from a Southern state, important as Southern states have been shown to prefer Southern Democratic candidates (Campbell, 1992). One major difference between the 2004 and 2008 elections was the popularity of the incumbent Republican president. In 2004, George W. Bush ran as a fairly popular incumbent. In 2008, no incumbent was on the ballot, and the Republican president had an historically low approval rating. We would expect a countrywide positive shock to Obama relative to Kerry.[18]

---

[17] Calculations on perceived ideology are author's calculations using ANES data.

[18] Bush's approval rating from October 17-20, 2008 was the lowest for any president in the history of the NBC News-Wall

Before adding a full set of controls, I plot the correlation between $\textit{Racially Charged Search Rate}_j$ and $(\%Obama2008 - \%Kerry2004)_j$. Figure 1.4, Panel (a), shows the relationship at the media market level.[19] Likely due to the different election conditions in 2004 and 2008, Obama does indeed perform better than Kerry country-wide. (See Table 1.6 for a set of summary statistics, including Obama and Kerry support.) However, Obama loses votes in media markets with higher racially charged search rates. The relationship is highly statistically significant ($t = -7.36$), with the Google proxy explaining a substantial percentage of the variation in change in Democratic presidential support ($R^2 = 0.24$).

Table 1.6: *Summary Statistics*

	mean	sd	min	max
Racially Charged Search Rate	39.78	9.21	16.62	100.00
%Kerry2004	48.83	9.57	19.89	70.06
%Obama2008	53.76	10.18	22.16	75.05
%Obama2012	52.04	11.03	19.66	76.87
%Obama2008 - %Kerry2004	4.93	3.18	-10.98	18.60
%Obama2012 - %Obama2008	-1.72	2.03	-10.61	5.27
%HouseDems2008 - %HouseDems2004	7.26	8.74	-39.16	72.59
ln(Turnout2008) - ln(Turnout2004)	0.07	0.06	-0.10	0.25

Notes: All summary statistics are reported for the 196 media markets for which data on Racially Charged Search Rate and voting data are available. All summary statistics reported are weighted by 2004 two-party turnout, the weighting used in Tables 1.7 and 1.10. Racially Charged Search Rate is as defined in Equation 1.1, obtained by the algorithm described in Appendix IV, normalized to its z-score. All candidate variables are that candidate's percentage points of two-party votes in a given year. Turnout is total two-party presidential votes in a given year. All political variables were downloaded at the county level and aggregated to the media market level using Gentzkow and Shapiro (2008).

Street Journal tracking poll (Hart/McInturff, 2012). He was nearly twice as popular in the run-up to the 2004 election as in the run-up to the 2008 election (Gallup, 2012). Modern political elections are considered, in large part, a referendum on the current administration, even if the incumbent candidate is not running; Obama consistently attempted to tie McCain to the unpopular Bush (Jacobson, 2009).

[19] There are 210 media markets in the United States. Ten of the smallest media markets do not have large enough search volume for "weather" and thus are not included. Two additional small media markets (Juneau and Twin Falls) search "weather" much more frequently than other media markets. Since they often score 100 on both "weather" and "weather" or the racial epithet, I cannot pick up their racial animus from the algorithm. Alaska changed its vote reporting boundaries from 2004 to 2008. I was unable to match the media market data with the boundaries for Alaskan media markets. I do not include data from Alaska. Overall, the 196 media markets included represent 99.3 percent of voters in the 2004 election. All of the summary statistics in Table 1.6 are virtually identical to summary statistics over the entire population of the United States.

Figure 1.4: *Racially Charged Search Rate and Black Candidate Support*

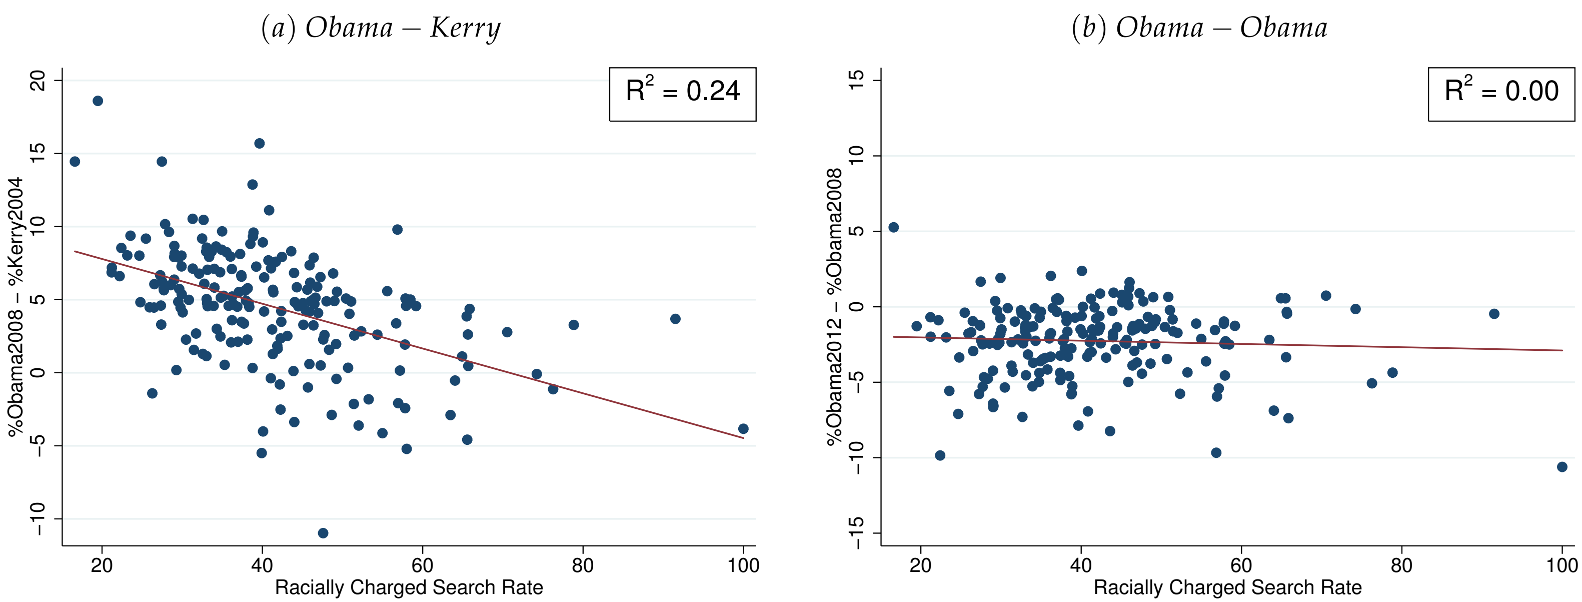

Notes: The x-axis in both panels is a media market's Racially Charged Search Rate, as defined in Equation 1.1, obtained by the algorithm described in Appendix IV. The y-axis in Panel (a) is Kerry's 2004 percentage points of the two-party vote subtracted from Obama's 2008 percentage points of the two-party vote. The y-axis in Panel (b) is Obama's 2008 percentage points of the two-party vote subtracted from Obama's 2012 percentage points of the two-party vote.

One non-racial explanation for the correlation between $Racially\ Charged\ Search\ Rate_j$ and $(\%Obama2008 - \%Kerry2004)_j$ might be that areas with high racially charged search rates were trending Republican, from 2004 to 2008, for reasons other than the race of the candidates. Data using other measures of changing liberalism offer evidence against this interpretation.

Panel (a) of Figure 1.5 shows no relationship between states' racially charged search and changes in states' liberalism, from 2004 to 2008, as measured by Berry et al. (1998). Figure 1.5, panel (b), shows a small, and not significant, negative correlation between media markets' racially charged search and change in Democratic support in House races from 2004 to 2008. (In results shown later, I find that racial animus affected turnout, likely explaining the small relationship with House voting.) Using exit poll data in 2004 and 2008, there is no relationship between racially charged search rate and change in black self-reported support for Obama relative to Kerry ($R^2 = 0.00$); the relationship is driven entirely by white voters ($R^2 = 0.28$).

Furthermore, if the correlation were due to changing partisan preferences correlated with racially charged search rate, other Democratic presidential candidates would have been equally punished in areas with high racially charged search rates around this time period. However, I examine data from SurveyUSA, first used by Donovan (2010), on hypothetical presidential match-ups. I can test whether, matched up against the same Republican candidate, Obama does worse than other Democratic candidates, among white voters, in areas with higher racially charged search. In February 2008, hypothetical match-ups were performed between Hillary Clinton and McCain and Obama and McCain in 50 states. Among white voters, Obama receives significantly smaller vote shares than Clinton in states with higher racially charged search rate ($t = -9.05$; $R^2 = 0.49$). In late September and early October 2007, in 17 states, hypothetical match-ups were performed between John Edwards and three Republican candidates and Obama and the same three Republican candidates. Among white voters, for all three match-ups, Obama receives significantly smaller vote shares than Edwards in states with higher racially charged search rate (Fred Thompson: $t = -3.49, R^2 = 0.45$; Rudy Guiliani: $t = -2.20, R^2 = 0.24$; Mitt Romney: $t = -3.48, R^2 = 0.45$).

Reported voting data are never ideal. However, the results of the alternate match-ups, combined with the race-specific exit polls results, combined with the House voting results, strongly suggest

Figure 1.5: *Change in Liberalism (2004-2008) and Racially Charged Search Rate*

(a) Citizen Ideology *(b) House Voting*

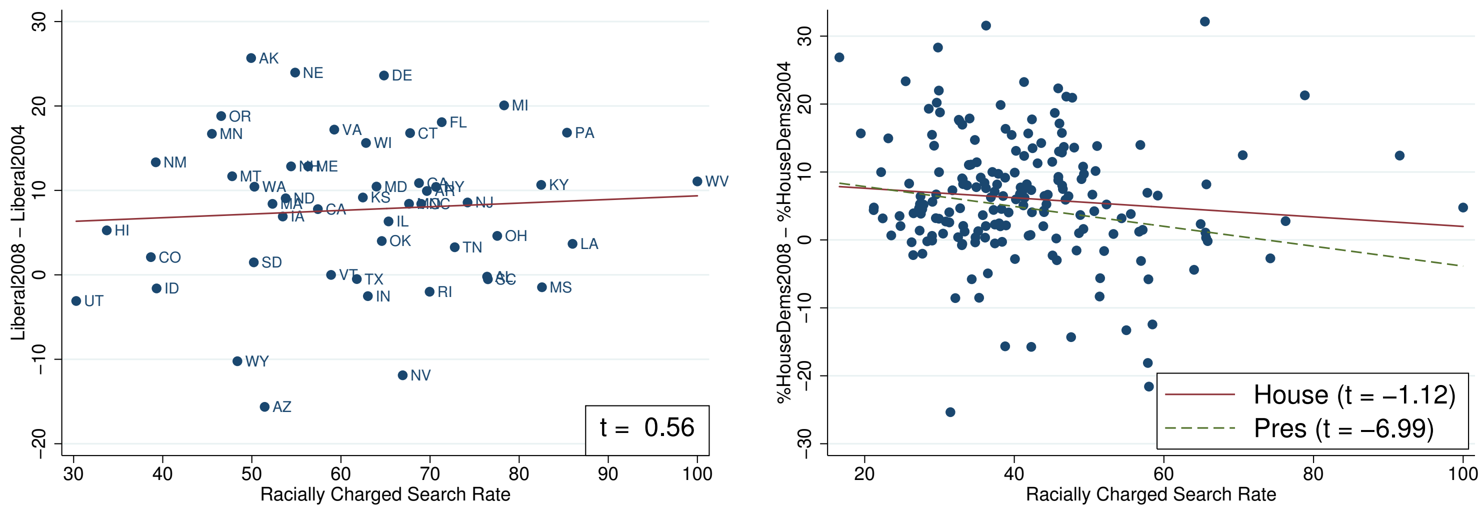

Notes: The x-axis in panel (a) is the unrounded value from Table 1.3. The x axis in panel (b) is the number used in Figure 1.4. The y-axis in panel (a) measures change in liberalism, from 2004 to 2008, according to the "revised 1960-2008 citizen ideology series." This commonly used proxy is described in Berry et al. (1998). To construct the y-axis variable for panel (b), counties are dropped if either major party received 0 votes in House elections in 2004 and/or 2008. Vote totals are aggregated for remaining counties to the media market level, using Gentzkow and Shapiro (2008), and the difference in the two-party share for Democrats is calculated. The dotted line shows the best linear fit of points (not shown) between the difference in Obama 2008 and Kerry 2004 vote shares and racially charged search rate, over the same, limited, 'unopposed' sample. The relationship between racially charged search rate and changing Democratic House support using all races, not just races fielding candidates from both major parties, is not significant, either. However, standard errors are much larger.

that decreased support for Obama in areas with high racially charged search rate is caused by white voters supporting Obama less than they would a white Democrat.

I now return to administrative vote data at the media market level and examine the relationship more systematically using econometric analysis. I add a number of controls for other potential factors influencing voting. I do not find evidence for an omitted variable driving the negative correlation between a media market's racially charged search rate and its preference for Obama compared to Kerry. The empirical specification is

$$(\%Obama2008 - \%Kerry2004)_j = \beta_0 + \beta_1 \cdot Racially\ Charged\ Search\ Rate_j + X_j\phi^1 + \mu_j \tag{1.2}$$

where X_j are area-level controls that might otherwise influence change in support for the Democratic presidential candidate from 2004 to 2008; β_0 is a country-wide shock to Democratic popularity in 2008; and μ_j is noise.

Racially Charged Search Rate$_j$ is as described in Equation 1.1, normalized to its z-score. Thus, the coefficient β_1 measures the effect of a one standard deviation increase in *Racially Charged Search Rate*$_j$ on Obama's vote share. All regressions predicting voting behavior, unless otherwise noted, are weighted by 2004 total two-party votes. All standard errors are clustered at the state level.[20]

The results are shown in Table 1.7. All columns include two controls known to consistently influence Presidential vote choice (Campbell, 1992). I include *Home State*$_j$, a variable that takes the value 1 for states Illinois and Texas; -1 for states Massachusetts and Arizona; 0 otherwise.[21] I also include proxies for economic performance in the run-up to both the 2004 and 2008 elections: the unemployment rates in 2003, 2004, 2007, and 2008.

Column (1), including just the standard set of controls, shows that a one standard deviation

[20] For media markets that overlap, I code the media market as being in the state in which the largest number of people live.

[21] Since I run the regressions at the media market level and some media markets overlap states, I aggregate *Home State*$_j$ from the county level, weighting by 2004 turnout. For the Chicago media market, as an example, *Home State* $= 0.92$, as some counties in the media market are in Indiana.

Table 1.7: *The Effect of Racial Animus on Black Candidate Vote Share*

	Dependent Variable: %Obama2008 - %Kerry2004					
	(1)	(2)	(3)	(4)	(5)	(6)
Racially Charged Search Rate	-1.490***	-1.486***	-1.341***	-2.124***	-2.002***	-1.776***
	(0.305)	(0.258)	(0.260)	(0.435)	(0.259)	(0.304)
Home State	2.616***	4.234***	3.556***	2.481***	4.070***	3.636***
	(0.804)	(1.118)	(1.107)	(0.854)	(1.141)	(0.996)
Observations	196	196	196	196	196	196
R-squared	0.26	0.51	0.62	0.30	0.52	0.62
Standard Controls	X	X	X	X	X	X
Census Div. FE		X	X		X	X
Demographic Controls			X			X
Google Controls				X	X	X

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: Standard errors, clustered at the state level, are in parentheses. OLS regressions are weighted by total two-party presidential votes in the 2004 election. Racially Charged Search Rate is as defined in Equation 1.1, obtained by the algorithm described in Appendix IV, normalized to its z-score. Home State takes the value 1 for Illinois and Texas; -1 for Massachusetts and Arizona; 0 otherwise. Standard controls are Home State and unemployment rates in years 2003, 2004, 2007, and 2008 (from Local Area Unemployment Statistics). Demographic controls are percent African-American, percent Hispanic, percent with bachelor's degree, percent 18-34, percent 65+, and percent veteran (from American Community Survey '05-'09); change from 2000 to 2010 in percent African-American and percent Hispanic (from the Census); and gun magazine subscriptions per capita (from Duggan (2001)). All county-level variables are aggregated to the media market level using Gentzkow and Shapiro (2008). Google controls are normalized search volume for "African-American(s);" "nigga(s);" and "fuck."

increase in a media market's racially charged search rate is associated with 1.5 percentage points fewer Obama votes. Column (2) adds controls for nine Census divisions. Any omitted variable is likely to be correlated with Census division. Thus, if omitted variable bias were driving the results, the coefficient should drop substantially upon adding these controls. The coefficient, instead, remains the same. Column (3) adds a set of demographic controls: percent Hispanic; black; with Bachelor's degree; aged 18-34; 65 or older; veteran; and gun magazine subscriber; as well as changes in percent black and percent Hispanic. Since there is some measurement error in the Google-based proxy of racial animus, one would expect the coefficient to move towards zero as these controls are added. It does. However, the change is not particularly large (less than a 10 percent decline in magnitude) considering the number of controls. The stability of the coefficient to a rich set of observable variables offers strong evidence for a causal interpretation (Altonji et al., 2005).

Adding Google Controls to Reduce Measurement Error

There is not a one-to-one correspondence between an individual's propensity to type the racial epithet into Google and his or her racial animus. Individuals may type the epithet for a variety of reasons other than animus. Individuals harboring racial animus may express it in different ways – either on different search engines or offline.

Any motivations of searches of the word unrelated to animus that do not differ at the area level will not create any bias in the area-level proxy. However, alternative motivations that differ at the area level will lead to measurement error in the area-level proxy. Classical area-level measurement error will cause attenuation bias in the estimates in Columns (1)-(3) of Table 1.7. In Columns (4)-(6), I reproduce the results from Columns (1)-(3) but add controls for an area's search rates for other words correlated with the search term unlikely to express racial animus. This should reduce measurement error in the proxy.

Row (8) of Table 1.2 shows that some searchers are looking for information on the word. I add a control for "African American(s)" search rate to proxy an area's interest in information related to African-Americans. Since a small percentage of searches for the word ending in "er" are looking for particular cultural references, I add a control for "nigga(s)" search rate. Finally, as some areas may be more prone to use profane language on Google, I add a control for an area's search rate for profane language.[22] Columns (4)-(6) show that the coefficient is more negative in each specification when adding the Google controls.

The Cost of Racial Animus on an Incumbent Black Presidential Candidate: Evidence from 2012

Previously, it was found that racially charged search rate significantly predicts Barack Obama's 2008 vote share, controlling for John Kerry's 2004 control. The robustness of the result is evidence for a causal effect of racial animus on Obama.

Was there a similar effect of racial animus on Obama in his second run for president, in 2012? Figure 1.4, Panel (b), shows graphical evidence that the answer is yes. It compares an area's racially charged search rate to the change in Obama's two-party vote share, from 2008 to 2012. If racial animus played a bigger (smaller) role in 2012 than in 2008, we would expect the relationship to be negative (positive). Instead, racially charged search rate shows no correlation with the change in Obama's 2008 and 2012 vote shares. This suggests race played a similar role in 2008 and 2012.

Note, too, that the result in Panel (b), the null relationship between racially charged search rate and change in Obama support, from 2008 to 2012, further supports the causal explanation of Panel (a), the negative relationship between racially charged search rate and change in Kerry 2004 to Obama 2008 support. In particular, the null relationship argues against two alternative explanations. If the negative correlation between racially charged search rate and change in Democratic support from 2004 to 2008 were picking up a trend away from Democratic support in places with high racially charged search rates, one would expect this trend to continue and there to again be a negative correlation in Panel (b). Another, non-causal explanation for the result in Panel (a) is that, by chance, racially charged search rate correlated with random noise in 2008 vote shares. Bias towards finding, and reporting, significant results led to this relationship being found. If this were the case, there should be regression to the mean and a positive correlation in Panel (b). The lack of a significant relationship, instead, adds additional evidence that the correlation in Panel (a) is due to areas with high racially charged search rate punishing Obama.

[22] Following my general strategy of selecting the most salient word if possible, I use the word "fuck."

Table 1.8 examines 2012 data more systematically. Panel (a) reproduces the six regression results from Table 1.7, presenting the identical coefficient on racially charged search rate as shown in the corresponding column in Table 1.7. Panels (b) and (c) of Table 1.8 introduce different dependent variables. In Panel (b), the dependent variable is $\%Obama2012 - \%Obama2008$. This, thus, expands the exercise performed in Figure 1.4, Panel (b). In Panel (c) of Table 1.8, the dependent variable is $\%Obama2012 - \%Kerry2004$. Comparing coefficients in Panel (c) and Panel (a), thus, can be thought of as comparing the size of the effect of racial prejudice in 2008 and 2012.

The regressions in Panel (b) and Panel (c) use the same demographic and Google controls as in Panel (a). However, I use different standard controls to reflect the different election conditions. The standard controls for Panel (b) are: a dummy variable Home State that takes the value 1 for Arizona and -1 for Massachusetts; and the unemployment rates in 2007, 2008, and 2011. In Panel (c), the standard controls are a dummy variable Home State that takes the value 1 for Illinois and Texas and -2 for Massachusetts; and the unemployment rates in 2003, 2004, and 2011.

Panel (b) of Table 1.8 shows that, upon adding the controls, there still is not a significant relationship between racially charged search rate and change in Obama support, from 2008 to 2012. This confirms the robustness of the null result of Figure 1.4, Panel (b). The null result in Panel (b) suggests that racial prejudice played a similar role in 2008 and 2012. Indeed, the coefficients in Panel (c) are roughly similar to the corresponding coefficients in Panel (a).

To summarize, racially charged search rate is a similar predictor of Obama's performance in both 2008 and 2012. In addition, the flat relationship between racially charged search rate and change in Democratic support, from 2008 to 2012, further supports a causal interpretation of the negative relationship between racially charged search rate and change in Democratic support, from 2004 to 2008.

Robustness Checks

Table 1.9 presents a number of robustness checks. Obama received about 20 percentage points more of the two-party vote share in Hawaii than Kerry did. Obama was born in Hawaii. Excluding Hawaii, though, changes the coefficient towards zero by less than 5 percent. The coefficient is of a

Table 1.8: *The Effect of Racial Animus: 2008 Compared to 2012*

	(a) Dependent Variable: %Obama2008 - %Kerry2004					
	(1)	(2)	(3)	(4)	(5)	(6)
Racially Charged Search Rate	−1.490***	−1.486***	−1.341***	−2.124***	−2.002***	−1.776***
	(0.305)	(0.258)	(0.260)	(0.435)	(0.259)	(0.304)
	(b) Dependent Variable: %Obama2012 - %Obama2008					
	(1)	(2)	(3)	(4)	(5)	(6)
Racially Charged Search Rate	0.096	−0.146	−0.027	−0.401	−0.283	0.048
	(0.276)	(0.287)	(0.284)	(0.285)	(0.311)	(0.333)
	(c) Dependent Variable: %Obama2012 - %Kerry2004					
	(1)	(2)	(3)	(4)	(5)	(6)
Racially Charged Search Rate	−1.423***	−1.896***	−1.377***	−2.551***	−2.427***	−1.706***
	(0.467)	(0.425)	(0.284)	(0.577)	(0.469)	(0.457)
Observations	196	196	196	196	196	196
Standard Controls	X	X	X	X	X	X
Census Div. FE		X	X		X	X
Demographic Controls			X			X
Google Controls				X	X	X

Notes: Panel (a) reproduces the six coefficients on Racially Charged Search Rate corresponding to the six coefficients on Racially Charged Search Rate in Table 1.7. Panel (b) presents the coefficients on Racially Charged Search Rate for the same regressions as those used in Panel (a), with a different dependent variable and changed standard controls to adjust for different election conditions. The dependent variable is Obama's two-party vote share in 2012 minus Obama's two-party vote share in 2008. Standard controls are Home State, which takes the value -1 for Massachusetts; 1 for Arizona; 0 otherwise and the unemployment rates in 2011, 2007, and 2008 (from Local Area Unemployment Statistics). Google and Demographics controls are identical to those used in Panel (a) and are listed in Table 1.7. For Panel (c), the dependent variable is Obama's two-party vote share in 2012 minus Obama's two-party vote share in 2004. Standard controls are Home State, which takes the value -2 for Massachusetts; 1 for Illinois; 1 for Texas; 0 otherwise; and unemployment rates in 2011, 2003, and 2004 (from Local Area Unemployment Statistics). Google and Demographics controls are identical to those used in Panel (a) and are described in Table 1.7. Standard errors, clustered at the state level, are in parentheses. OLS regressions are weighted by total two-party presidential votes in the 2004 election.

similar magnitude including changes in House Democratic support from 2004 to 2008 and swing state status.[23]

The main specification requires a somewhat restrictive relationship between Obama and Kerry's vote share. This, though, is not driving the result. The results are of similar magnitudes instead using $\%Obama_j$ as the dependent variable and including $\%Kerry2004_j$ as an independent variable. And they are of similar magnitudes using $\%Obama_j$ as the dependent variable and including a 4th-order polynomial for $\%Kerry_j$ as independent variables. Including this polynomial allows for liberal areas to differ from conservative areas in their relative support for Obama and Kerry. The fact that the coefficient on racially charged search rate is unchanged (perhaps not surprising since racially charged search rate is not significantly correlated with liberalness and voters perceived the candidates as having similar ideologies) offers additional evidence that racial attitudes, not ideology, explains the results. The coefficients are also very similar including trends in presidential Democratic support.

III.B The Effects of Racial Animus on Turnout in a Biracial Election

The robust cost of racial animus on Obama's vote share is the main result of the paper. I can also use the area-level proxy for racial animus to test the effects of racial attitudes on turnout. This both helps us understand the mechanism through which racial prejudice hurt Obama and will also prove useful in interpreting the size of the effects, which I do in the next section.

The effect of racial animus on turnout is theoretically ambiguous. The effect of racial prejudice on Obama's vote share could be driven by any of three reasons, each with different implications for turnout: Individuals who would have voted for a white Democrat instead stayed home (decreasing turnout); individuals who would have voted for a white Democrat instead voted for the Republican (not affecting turnout); individuals who would have stayed home instead voted for the Republican (increasing turnout).

I first use the area-level proxy of racial animus to test the average effect of prejudice on turnout. I regress

[23] I do not include these controls in the main specifications as they could be affected by Obama support and thus not exogenous.

Table 1.9: *Robustness Checks*

Specification	*2008 Coefficient*	*2012 Coefficient*
Baseline (All Controls; Table 1.8, Column (6))	−1.776	−1.706
	(0.304)	(0.304)
Exclude Hawaii	−1.553	−1.463
	(0.230)	(0.411)
Add Control for Change in House Voting	−1.699	−1.610
	(0.284)	(0.452)
Add Control for Swing State	−1.779	−1.647
	(0.317)	(0.442)
Use %Obama as Dependent Variable and Include Control for %Kerry2004	−1.682	−1.661
	(0.285)	(0.460)
Use %Obama as Dependent Variable and Include 4th-Order Polynomial %Kerry2004	−1.648	−1.628
	(0.293)	(0.478)
Add Control for %Kerry2004-%Gore2000	−1.775	−1.694
	(0.312)	(0.439)
Add Controls for %Kerry2004-%Gore2000 and %Gore2000-%Clinton1996	−1.731	−1.642
	(0.329)	(0.453)
Use %Obama as Dependent Variable and Include %Kerry2004, %Gore2000, %Clinton1996	−1.577	−1.547
	(0.326)	(0.459)

Notes: Standard errors, clustered at the state level, are in parentheses. Results in this table are variations on Column (6), Panels (a) and (c), reported in Table 1.8. Swing State status are Battleground States, as defined by *The Washington Post*, available at `http://www.washingtonpost.com/wp-dyn/content/graphic/2008/06/08/GR2008060800566.html`.

$$(ln(Turnout2008) - ln(Turnout2004))_j = \delta_0 + \delta_1 \cdot Racially\ Charged\ Search\ Rate_j + Z_j\phi^2 + \psi_j \quad (1.3)$$

where $(ln(Turnout2008) - ln(Turnout2004))_j$ is the change in the natural log of the total Democratic and Republican votes from 2004 to 2008; Z_j is a set of controls for other factors that might have changed turnout and $Racially\ Charged\ Search\ Rate_j$ is as described in Equation 1.1, normalized to its z-score.

The results are shown in Columns (1) through (3) of Table 1.10. In all specifications, I include percent black and change in the natural log of an area's population from 2000 to 2010. Column (2) adds Census fixed effects. Column (3) adds the same demographic controls used in the vote share regressions in Table 1.7. In none of the specifications is there a significant relationship between racially charged search and turnout. I can always reject that a one standard deviation increase in racially charged search rate – which lowers Obama's vote share by 1.5 to 2 percentage points – changes turnout by 1 percent in either direction.[24]

The null effect of racial attitudes on turnout is consistent with animus not affecting any individuals' decision to turnout (but convincing many who would have supported a Democrat to instead vote for McCain). It is also consistent with racial prejudice driving an equal number of individuals who would have voted for a white Democrat to stay home as it convinced individuals who would have stayed home to vote for McCain.

To better distinguish these two stories, I add to the independent variables in Columns (1) to (3) of Table 1.10 the interaction between an area's percent Democrats and racially charged search rate. If racial attitudes affect some individuals' decisions of whether or not to vote, I expect the following: it should increase turnout when there are few Democrats in an area. (There are few Democrats available to stay home due to racial prejudice.) The effect of racial prejudice on turnout should be

[24]Washington (2006) finds a 2-3 percentage point increase in turnout in biracial Senate, House, and Gubernatorial elections. Perhaps these results can be reconciled as follows: Obama won a close primary. An average black general election candidate would be expected to have won his or her primary by a larger margin than Obama won his by. We would thus expect that the average black candidate would have faced lower racial animus in his or her primary than Obama did in a country-wide Democratic primary. Thus, racial animus among Democrats is lower for the average black candidate in Washington's (2006) sample than for the country as a whole. Thus, relatively few voters would stay home in the general election rather than support the black candidate in the average election in Washington's (2006) sample.

Table 1.10: *Change in Turnout (2004-2008) and Racially Charged Search Rate*

	Dependent Variable: ln(Turnout2008)-ln(Turnout2004)					
	(1)	(2)	(3)	(4)	(5)	(6)
Racially Charged Search Rate	-0.001	-0.001	0.004	0.025**	0.032*	0.033*
	(0.005)	(0.005)	(0.005)	(0.013)	(0.017)	(0.017)
Racially Charged Search Rate · %Kerry2004				-0.056**	-0.071*	-0.064*
				(0.028)	(0.039)	(0.039)
Observations	196	196	196	196	196	196
R-squared	0.67	0.73	0.80	0.67	0.74	0.80
Census Div. FE		X	X		X	X
Demographic Controls			X			X

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: Standard errors, clustered at the state level, are in parentheses. OLS regressions are weighted by total two-party presidential votes in the 2004 election. Dependent variable in all specifications is the natural log of two-party presidential votes in 2008 minus the natural log of two-party presidential votes in 2004. Racially Charged Search Rate is as defined in Equation 1.1, obtained by the algorithm described in Appendix IV, normalized to its z-score. All regressions include change in log population from 2000 to 2010 (from the Census); percent African-American (from American Community Survey '05-'09); and Kerry's share of the two-party vote. Columns (3) and (6) add percent African-American, percent Hispanic, percent with bachelor's degree, percent 18-34, percent 65+, and percent veteran (from American Community Survey '05-'09); change from 2000 to 2010 in percent African-American and percent Hispanic (from the Census); and gun magazine subscriptions per capita (from Duggan (2001)). All county-level variables are aggregated to the media market level using Gentzkow and Shapiro (2008).

decreasing as the percentage of the population that supports Democrats increases.

More formally, the regression is:

$$(ln(Turnout2008) - ln(Turnout2004))_j = \alpha_0 + \alpha_1 \cdot \%Kerry2004_j + \alpha_2 \cdot Racially\ Charged\ Search\ Rate_j$$
$$+\alpha_3 \cdot Racially\ Charged\ Search\ Rate_j \times \%Kerry2004_j + Z_j\phi^3 + \epsilon_j \quad (1.4)$$

where $Kerry2004_j$ is used to proxy an area's percent Democrats.

If racial animus affected Obama vote shares, in part, through changes in turnout, I expect $\alpha_2 > 0$ and $\alpha_3 < 0$.

The coefficients on α_2 and α_3 are shown in Columns (4)-(6) of Table 1.10. In all three specifications, corresponding to the same specifications in Columns (1)-(3), $\alpha_2 > 0$ and $\alpha_3 < 0$. In areas that supported Kerry in 2004, an increase in racial animus decreased 2008 turnout. In areas that supported Bush in 2004, an increase in racial animus increased 2008 turnout. The coefficients tend to be marginally significant, and the standard errors are always too large to say anything precise.

In results not shown, I reproduce the results replacing $ln(Turnout2012) - ln(Turnout2004)$ as the dependent variable. County-level population data near or after the 2012 election are not as-of-yet available, complicating interpretation. However, preliminary results are similar, with no relationship between racially charged search rate and turnout, on average, but a positive (negative) relationship in highly Republican (Democratic) areas.

In sum, the evidence on the effects of racial animus on turnout is as follows: Some Democrats stayed home rather than vote for Obama due to his race; a similar number of individuals who would not have otherwise voted turned out for the Republican due to Obama's race. There is not enough statistical power, though, to determine this number.

IV Interpretation

Section III compares Google racially charged search rate to changing voting patterns from the 2004 all-white presidential election to the 2008 and 2012 biracial presidential elections and finds that racial animus played a significant role in the 2008 and 2012 elections. Section III.A shows the main result of this paper: racially charged search rate is a robust negative predictor of Obama's vote share. Section III.B shows that higher racially charged search rate predicts increased turnout in Republican parts of the country; decreased turnout in Democratic parts of the country; and, on average, no change in turnout. This section aims to give some intuition to the magnitude of the effects of racial attitudes on presidential voting.

How many additional votes would Obama have received if the whole country had the racial attitudes of the most tolerant areas? Media markets' mean racially charged search rate is 2.34 standard deviations higher than the minimum racially charged search rate. Table 1.11 shows the estimated vote shares from different specifications, assuming that no votes were lost in the media market with the lowest racially charged search rate. In 2008, the estimated loss ranges from 3.1 percentage points to 5.0 percentage points.[25] The specification including the full set of controls – Google controls, demographics controls, and Census Division fixed effects, gives a point estimate of 4.2 percentage points. In 2012, the estimated loss ranges from 3.2 percentage points to 6.0 percentage points. The specification that includes the full set of controls yields a point estimate of 4.0 percentage points.

The effects of racial animus on a black compared to a white Democratic candidate can be compared to voters' well-established comparative preference for a home state compared to a non-home-state candidate.[26] Studies show, on average, voters will reward a candidate from their own home-state with about four percentage points of the two-party vote (Lewis-Beck and Rice, 1983;

[25]In estimates using the Google controls, multiplying the coefficient by 2.34 yields an approximation of the true effect. This would be biased upwards if measurement error substantially lowered the measured minimum racial animus. I do not find this is the case. I calculate a new measure of racial animus as the difference in racially charged search relative to predictions from all the controls in Column (4) of Table 1.7. This still leaves Loredo, TX as having the minimum value. Regressing the dependent variable – the difference between Obama and Kerry support – on this measure of racial animus and multiplying the coefficient on the regression by the difference between the mean and the minimum of the measure always yields roughly the same result.

[26]I interpret the results in this paper as the effects of racial animus. An alternative explanation is that this reflects racial attitudes more broadly, with perhaps the Google search proxy correlating with other types of prejudice, such as implicit prejudice. My interpretation is based on: how common the searches are; the clear interpretation of searches as animus; the fact that it is not clear how correlated an area's implicit prejudice and animus are; and some research using individual data that do not find implicit prejudice an important factor when controlling for admitted explicit prejudice (Compare, for example, Piston (2010) to Pasek et al. (2010) and see Kinder and Ryan (2012)). When area-level averages for implicit prejudice are available, this interpretation can be further tested.

Table 1.11: *Country-Wide Effect: Google Compared to Other Measures*

Source	*Obs*	*Measure*	*Controls*	*2008 Cost*	*2012 Cost*
Google	196 Media Markets	Racially Charged Search Rate, '04-'07	Standard	3.5	3.3
				(0.7)	(1.1)
			Standard+Census	3.5	4.4
				(0.6)	(1.0)
			Standard+Census+Demogs	3.1	3.2
				(0.6)	(0.9)
			Standard+Google	5.0	6.0
				(1.0)	(1.3)
			Standard+Google+Census	4.7	5.7
				(0.6)	(1.1)
			Standard+Google+Census +Demogs	4.2	4.0
				(0.7)	(1.1)
GSS	45 States	% Against Interracial Marriage, '90-'04	Standard	2.0	2.3
				(0.6)	(0.6)
			Standard+Census	0.6	2.1
				(1.3)	(1.3)
		Average Prejudice, '72-'04	Standard	2.8	3.0
				(1.1)	(1.0)
			Standard+Census	0.5	2.0
				(1.6)	(1.9)
ANES	Individual	Explicit Prejudice	Piston (2010)	2.3	
				(1.0)	
APYN		Explicit+Implicit Prejudice	Pasek et al. (2010)	2.7	
CCES		Racial Salience	Schaffner (2011)	2.0	

Notes: This table compares the results obtained using the Google data to those using the same specification but measures from the GSS and the estimate obtained by other scholars using individual proxies for racial attitudes and individual reported votes. For all regressions used to calculate the estimated percentage points using Google or GSS, the regressions are weighted by total two-party presidential votes in 2004. The point estimate is then the country-wide effect of moving from the area with the lowest value. Controls are those reported in Table 1.8. The first GSS measure is from Mas and Moretti (2009). The second GSS measure is from Charles and Guryan (2008). Piston (2010) finds that overall prejudice cost Obama 2.66 percent of the white vote. Assuming whites accounted for 87% of the electorate yields the number of 2.3. For the GSS regressions, robust standard errors are shown. For the Google regressions, standard errors clustered at the state-level are shown.

Mixon and Tyrone, 2004). This is roughly consistent with the home-state advantage found in the regressions in Table 1.7. Racial animus gave Obama's opponent the equivalent of a home-state advantage country-wide.

While racial animus obviously did not cost Mr. Obama the 2008 or 2012 election, examining more elections shows that effects of the magnitude found are often decisive. A two percentage point vote loss would have switched the popular vote winner in 30 percent of post-War presidential elections. A four percentage point loss would have changed more than 50 percent of such elections.

IV.A Comparison with Other Methodologies

The effect of racial prejudice found by the methodology of this paper can also be compared to estimates obtained using different data sources and methodology. I find that the effects using Google data are larger than effects found using other methodologies. The specification used in this paper is slightly different from the one used in Mas and Moretti (2009). Mas and Moretti (2009) predict a county's Democratic vote share in 2004 and 2008 House and presidential elections from a set of dummy variables (Year=2008; Election Type=presidential; Election Type=presidential & Year=2008) and an interaction between a state's GSS racial attitudes and the dummy variables. This specification makes it difficult to pick up the effects of racial attitudes on voting for Obama since House elections are high-variance (sometimes, one of the two major parties does not field a candidate, dramatically shifting the Democratic proportion of vote share). A large swing in House voting can falsely suggest a large trend in Democratic voting.[27]

Nonetheless, I do confirm that effects using the GSS measures and the specification of this paper yields a smaller effect and are less robust. Table 1.11 compares the estimates obtained using the Google measure and the specification of this paper to estimates using GSS measures and the

[27] For example, in Mas and Moretti's (2009) Figure 4, the authors compare the difference between the change in Obama and Kerry's vote shares and the change in House voting to their measure of racial prejudice. The difficulty with this comparison is that House elections in which one party does not field a candidate will create enormous noise in the voting metric, swamping any other changes. In 2004 in Vermont, Bernie Sanders won as a highly popular left-wing independent. In 2008 in Vermont, Democrat Peter Welch won with no Republican challenger. Thus, there was a huge gain in Vermont Democratic House support from 2004 to 2008. And the difference between the change in Democratic presidential support and change in Democratic House support, from 2004 to 2008 in Vermont, is -70 percentage points. Adding this kind of noise to the Obama and Kerry difference, and having only 45 state-level GSS observations, it is unlikely that, even if the GSS measure of racial attitudes did predict opposition to Obama, this methodology could pick it up.

specification of this paper. Using either the measure from Mas and Moretti (2009) or Charles and Guryan (2008) always yields smaller estimates of the country-wide effect. The effect picked up using the GSS data is largely due to a few Southern states which measure high on racial prejudice and also voted for Obama significantly less than they voted for Kerry. In contrast to regressions using the Google measure, where the effect is robust to including Census division fixed effects, regressions using the GSS measures tend to lose significance when including the Census division fixed effects.[28] Furthermore, I find that the preferred fit with the GSS measures is quadratic. The fit suggests no effect in just about all parts of the country but an effect in a few southern states. The GSS is ineffective at capturing racial prejudice in all but a few Southern states. Google is also advantageous relative to the GSS in testing for causality: observations from large samples from 196 media markets allows for a rich set of controls and robustness checks, as shown in Tables 1.7, 1.8 and 1.9; this is not possible with 45 state-level observations.

The final row of Table 1.11 includes the estimates from Piston (2010), Schaffner (2011), and Pasek et al. (2010).[29] Each uses individual data and obtains a smaller preferred point estimate. This suggests individual surveys underestimate the true effect of racial attitudes.

There are additional advantages to the empirical specification of this paper relative to studies using individual-level surveys in testing for causality, besides the likely improved measure of racial animus. Individual survey studies rely exclusively on self-reported voting; vote misreporting may be a substantial issue with survey data (Atkeson, 1999; Wright, 1993; Ansolabehere and Hersh, 2011). Further, their measures of racial attitudes are taken from near the election. They thus could potentially pick up reverse causation. Finally, studies testing the effects of racial attitudes on political attitudes have been criticized for omitted variable bias from unmeasured conservative

[28]Highton (2011) located an alternative data source for racial attitudes from The Pew Research Center Values Study. Pew has asked for 20 years individuals whether they approve of blacks dating whites. Aggregating 20 years of data among whites, Highton (2011) constructs a measure available for 51 states and tests the effects of racial animus on voting in the Obama election. While standard errors are still large and the point estimate is always smaller than using Google data, the Pew data source does lead to more robust estimates than the GSS data source, in part due to the six additional observations.

[29]A recent paper by Kam and Kinder (2012) finds that ethnocentrism was a factor against Obama. Tesler and Sears (2010a) also finds an important role of anti-Muslim sentiment in evaluating Obama. Using Google data (such as searches for "Obama Muslim" or "Obama birth certificate") to further investigate this phenomenon is a promising area for future research.

ideology (Schuman, 2000; Feldman and Huddy, 2005; Huddy and Feldman, 2009). This is both because the measures of prejudice, such as whether African-Americans should overcome prejudice "without any special favors," might be connected to conservative ideology and self-reported vote choices in previous elections are even more unreliable than self-reported vote choices in current elections. Thus, individual-level, non-panel studies can only control for self-reported ideology and political beliefs. The empirical specification of this paper, using the unambiguous measure of racial animus and, most importantly, controlling for administrative vote data from a similar election four years earlier, does not seem open to this critique.

IV.B Pro-Black Effect

I find that, relative to the attitudes of the most racially tolerant area, racial animus cost Obama between 3 to 5 percentage points of the national popular vote. Obama, though, also gained some votes due to his race. Was this factor comparatively large?

A ballpark estimate from increased support from African-Americans can be obtained from exit poll data. In 2004, 60.0 percent of African-Americans reported turning out, 89.0 percent of whom reported voting for John Kerry. In 2008, 64.7 percent of African-Americans reported turning out, 96.2 percent of whom reported supporting Barack Obama. Assuming these estimates are correct and, with a white Democrat, black support would have been the same as in 2004, increased African-American support added about 1.2 percentage points to Obama's national popular vote total in 2008.[30] Reported turnout data are not yet available for 2012, though exit polls suggest African-Americans turned out at similar rates in 2012 as they did in 2008. The pro-black effect was limited by African-Americans constituting only 12.6 percent of Americans and overwhelmingly supporting any Democratic candidate.

A variety of evidence suggest that few white voters swung, in the general election, for Obama due to his race. Only one percent of whites said that race made them much more likely to support Obama

[30] Assume 65 percent of whites turned out in 2008 and 47.6 percent of white voters supported Obama. If African-Americans had voted as they did in 2004, Obama would have instead received $\frac{0.126\times0.6\times0.89+0.874\times0.65\times0.476}{0.126\times0.65+0.874\times0.65} = 52.5$ percent of the two-party vote. This is likely an upper-bound, as any Democrat likely would have seen some improvement in black support due to Bush's high disapproval rating among African-Americans.

in 2008 (Fretland, 2008). In exit polls, 3.4 percent of whites did report both voting for Obama and that race was an important factor in their decision. But the overwhelming majority of these voters were liberal, repeat voters likely to have voted for a comparable white Democratic presidential candidate.[31] Furthermore, Piston (2010) finds no statistically significant relationship, among white voters, between pro-black sentiment and Obama support, when controlling for ideology. Although social scientists strongly suspect that individuals may underreport racial animus, there is little reason to suspect underreporting of pro-black sentiment. Finally, in unreported results, I add an area's search rate for "civil rights" to the regressions in Table 1.7. The coefficient on *Racially Charged Search Rate* is never meaningfully changed, and the coefficient on *Civil Rights Search Rate* is never statistically significant.

IV.C Estimated Cost of Race Compared to Actual Performance

This paper suggests a far larger vote loss from racial animus than vote gains from race. This means that Obama would have gotten significantly more votes if race were not a consideration. Is this plausible? Forecasting how many votes a president should receive, based on economic and political fundamentals, lead to a large variance of estimates. In addition, these forecasts tend not to include candidate charisma, or candidate quality more generally, which may be important (Levitt, 1994; Benjamin and Shapiro, 2009). And such forecasts do not adjust for changing composition of the electorate (Judis and Teixeira, 2004). The highly Democratic Hispanic population has grown rapidly, consistently rising from 2 percent of the electorate in 1992 to 10 percent in 2008. This makes every modern election cycle meaningfully more favorable towards Democrats than the previous one. In 2012, had the racial composition of the electorate been the same as it was in 2008, Obama would have lost both Ohio and Florida.

Of the nine 2008 forecasts in Campbell (2008), three predicted that the Democratic presidential candidate candidate would perform at least two percentage points better than Obama did (Lewis-Beck and Tien, 2008; Lockerbie, 2008; Holbrook, 2008). Of the nine 2012 forecasts in Campbell (2012),

[31] Among the 3.4 percent, 87 percent both reported voting for the Democratic candidate in the House race and disapproving of Bush. Among this subset, only 25 percent reported voting for the first time. And, among such first-time voters, 60 percent were 18-24, possibly ineligible to vote in any prior elections.

only Lockerbie (2012) showed a substantial Obama underperformance (1.8 percentage points).

Jackman and Vavreck (2011), using polling data with hypothetical 2008 match-ups, find an "average" white Democrat would have received about 3 percentage points more votes than Obama did. Table 1.6 shows that House Democratic candidates received a 2.3 percentage point larger gain in 2008 relative to 2004 than Obama received relative to Kerry; the results in Section III.B suggest the House Democratic swing would have been even larger absent turnout effects due to Obama's race.

IV.D White Voters Swung by Racial Animus

As another way of giving intuition for the magnitude of the effect, I combine the vote share results in Section III.A with the turnout results in Section III.B. I can then estimate the percent of white voters who would have voted for a white Democrat in 2008 but did not support a black one.

The percent motivated by animus is the number of votes lost due to animus divided by the total number of whites who would have supported a Democrat absent prejudice. Section III.B finds that turnout was unaffected, on average, by prejudice. Thus, the denominator (the percent of whites who would have supported a Democrat, absent prejudice) is the number of whites who supported Obama plus the number of votes lost due to prejudice. Exit polls suggest 41.7 percent of 2008 voters and 38.1 percent of 2012 voters were white Obama supporters. The percent motivated by animus is estimated between $\frac{3.1}{44.8} = 6.9$ and $\frac{5}{46.7} = 10.7$ percent in the 2008 election and between $\frac{3.2}{41.3} = 7.7$ and $\frac{6.0}{44.1} = 13.6$ percent in the 2012 election. Regressions using the full set of controls imply that, among whites who would have otherwise supported a white Democratic presidential candidate, 9.1 percent in 2008 and 9.5 percent in 2012 did not support a black Democratic presidential candidate.

How do these numbers compare to what whites tell surveys? Among whites who told the GSS in 2008 and 2010 that they voted for Kerry in 2004, 2.6 percent said they would not vote for a black president. Three percent of whites told Gallup Obama's race made them much less likely to support him (Fretland, 2008). Approximately 4.8 percent of whites told exit pollsters they voted for McCain and race was an important factor in their vote. Evidence strongly suggests that many whites voted against Obama due to his race but did not admit that to surveys. The numbers can also be compared to other self-reported racial attitudes. In 2002, the last year the question was asked by the GSS, 11.9

percent of white Democrats admitted that they favored a law banning interracial marriage.

For additional intuition on the size of the effect, the numbers can be compared to persuasion rates as calculated by media scholars. Gerber et al. (2009) find that *The Washington Post* persuades 20 percent of readers to vote for a Democrat. Gentzkow et al. (2011) report that, historically, partisan newspapers persuaded fewer than 3.4 percent of readers. DellaVigna and Kaplan (2007) find that Fox News persuades 11.6 percent of viewers to vote Republican. Thus, the proportion of white Democrats who will not vote for a black Democratic Presidential candidate is roughly equivalent to the proportion of Democrats who can be persuaded by Fox News to not vote for a white Democratic Presidential candidate.

V Conclusion

Whether many white Americans will not vote for a black presidential candidate is perhaps the most famous problem complicated by social desirability bias. Scholars have long doubted the accuracy of survey results on this sensitive question. I argue that Google search query data offer clear evidence that continuing racial animus in the United States cost a black candidate substantial votes.

There are many important questions on sensitive topics that may similarly be helped by Google data. In a study of measurement error in surveys, Bound et al. (2001) include the following sensitive behaviors as difficult to measure for surveyors due to social censoring: "the use of pugnacious terms with respect to racial or ethnic groups;" voting; use of illicit drugs; sexual practices; income; and embarrassing health conditions. Words related to all these topics are searched often on Google.

Chapter 2

Who Will Vote? Ask Google

I Introduction

People misreport likelihood of voting to polls. Recent research finds that 67 % of individuals who will not vote tell pollsters that they are almost certain to vote (Rogers and Aida, 2012). Recent research also casts doubt on the reliability of pollsters' tools to screen "likely voters." Even the first screen used by polls, registration status, is subject to large deception. More than 60 % of non-registered voters falsely claim that they are registered (Ansolabehere and Hersh, 2011). And reported voting intention has been shown to have little predictive power, controlling for previous voting behavior (Rogers and Aida, 2012). Vote misreporting is correlated with other variables in important ways that may bias analysis (Ansolabehere and Hersh, 2011; Rogers and Aida, 2012; Vavreck, 2007).

This paper uses Google search data to predict area-level turnout. Previous research argues that Google data can measure socially sensitive behaviors (Stephens-Davidowitz, 2012). And the theory that Google can capture changes in turnout intention over time is compelling: the marginal voter, the individual who only votes in certain elections, is likely to need information prior to voting. Thus, he or she might search for reminders, Googling "how to vote" or "where to vote." The fact that habitual voters are unlikely to need to make these searches does not limit this exercise; these voters' intentions are easy to forecast.

I show that change in search rates for "vote" or "voting" in the October prior to an election, compared to October four years earlier, explains 20 to 40 percent of variation in changes in turnout rates compared to four years earlier for both the 2008 and 2010 elections. The predictive power is little affected by controlling for changes in registration rates, early voting rates, or a state's having a Senate race, three other sources of information available prior to an election that might be used to predict turnout. Prior to the 2012 election, I made out-of-sample predictions for turnout in the 2012 elections. The predictions were always significant and fairly powerful. There were, though, additional lessons that might improve predictions in the future, such as the importance of controlling for trends.

Can more be done with this data than projecting state-wide turnout? Might this data help political forecasters predict election results? And might this help political scientists better understand the causes of voting decisions? I examine two political forecasting applications and mention areas for possible future research.

Forecasters might use Google search data to help choose demographics weightings. In 2008, African-Americans turned out at elevated levels to support the first major-party general election black candidate. In 2012, there was great uncertainty as to whether the black share of the electorate would remain at its historic 2008 levels. Since African-Americans overwhelmingly support the Democratic candidate, a one percentage point increase in the black share of the electorate leads to nearly a one percentage point increase in the Democratic candidate's two-party vote share. The Google data would have correctly predicted a large increase in the black share of the electorate in 2008. The Google data showed large increases in the rates of voting-related searches in black communities prior to the 2008 compared to the 2004 election. And the Google data would have correctly predicted that the black share of the 2012 electorate would be similar to its historic 2008 levels. Many polls that underestimated Obama's support falsely assumed that the black share of the electorate would decline in 2012. Prior to the 2012 election, I predicted the composition of the 2008 electorate based on Google search data. I found that only one major demographic was a robust predictor of changing vote search rates, from October 2008 to October 2012: the share of the population that is Mormon. I thus predicted that the Mormon share of the electorate would increase. Other population groups would remain similar to their 2008 shares, with some adjustments for population increases. (For example, the Hispanic share of the electorate would increase, even though Hispanic turnout rates would not change.) The predictions were largely correct. This methodology would seem particularly promising for measuring a large change in turnout rates for a group whose population differs substantially in different parts of the country. Black turnout rates in 2008 and Mormon turnout rates in 2012 fit this criterion.

Forecasters might adjust their likely voter screen in a state based on Google search activity in that state. If a state is likely to have high turnout, pollsters can let more voters through the screen. Since most marginal voters are Democrats, an increase in the pool of likely voters will likely lead to a higher predicted vote total for the Democratic candidate (Gomez et al., 2008). An increase in Google-related voting activity did, on average, predict better Democratic performance than the polls predicted in the 2010 Midterm elections. However, it did not for the 2008 or 2012 presidential elections. This strategy, thus, appears to only work in Midterm elections.

The data source might prove useful in explaining past behavior. Two possibilities seem particularly promising. First, political scientists might study the determinants of actual voting, controlling for early intent to vote. Some variables that might be studied include negative advertising and voting laws. Second, the data source might be used as a high-frequency measure of intent to vote and can help measure the immediate effects of interventions.

This paper builds on a nascent literature, started by Varian and Choi (2010) and extended by Ginsberg et al. (2009) and Askitas and Zimmermann (2009) in using Google data to forecast the future. It follows Stephens-Davidowitz (2012) in using Google data to proxy a socially sensitive attitude.

II Google Search Proxy

To predict change in area-level turnout, I examine changes in an area's percentage of searches that include "vote" or "voting" in the October prior to an election.

In particular, I am interested in $\Delta ln(Google\ Vote\ Search) =$

$$\ln\left(\frac{Oct.\ Searches\ w/\ "Vote/Voting"_{i,t}}{Oct.\ Searches_{i,t}}\right) - \ln\left(\frac{Oct.\ Searches\ w/\ "Vote/Voting"_{i,t-4}}{Oct.\ Searches_{i,t-4}}\right)$$

Following Stephens-Davidowitz (2012), I choose the most salient words to constrain data-mining. I choose October data so that predictions can be meaningfully delivered prior to an election. One might suspect that likely voters would wait until the last minute to find information on voting. Indeed, search volume for "vote" and "voting" do increase significantly the day of and the day prior to the election. However, perhaps surprisingly, I did not find that including this November data improved the model. In results not shown, I find that the model is a bit worse including a longer range of data (August and September). It is a bit worse shortening the prediction by just including one week in October. While it is impossible to test this with data available, the evidence is consistent with likely voters searching for voting information many times in October and the first days of November.

The measure selected is an "omnibus" measure which captures many different ways that Google searching may predict turnout.

Figure 2.1 shows the words most often used in combination with the word "vote" or "voting" in October 2012. Notably, 25 percent of searches also include the word "early." An obvious concern is that changes in searches that include "vote" or "voting," from October 2008 to October 2012, picks up consistent voters switching to early voting. One possibility is taking out searches that include the word "early." However, I instead use all these searches because of the predictive power these searches have shown, presented shortly, the importance of capturing new voters who happen to choose early voting and would need to Google information, and the fact that I found little evidence for substantially different results excluding the term.

In addition, 20 percent of searches include the word "register" or "registration." I include these in the measure for many of the same reasons I include searches for early voting.

The data are available at the state, media market, and city level from Google Trends. However, if absolute search volume is not high enough, Google Trends does not report data. Stephens-Davidowitz (2012) develops a methodology for obtaining this data. In principle, any data can be obtained with this – or a fairly similar – approach. However, the lower the absolute search volume, the more downloads necessary to obtain the data. I have weighed importance of the data and difficulty of obtaining it in choosing the geographic level of analysis.

Figure 2.1: *Searches for "Vote" or "Voting," October 2012*

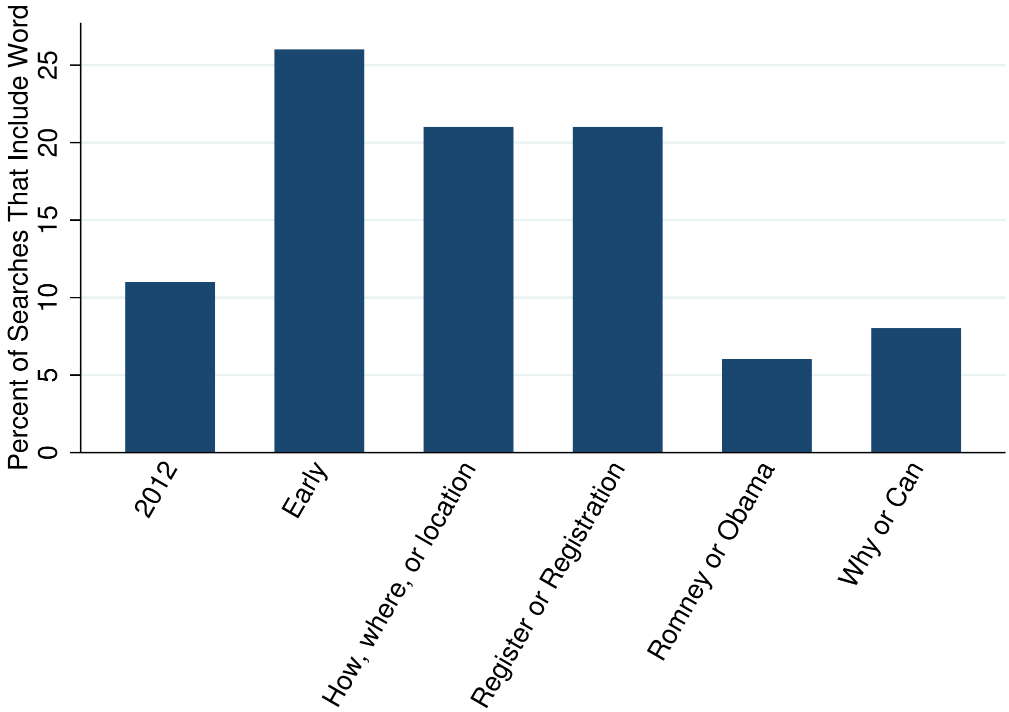

Notes This shows the percentage of searches that include "vote" or "voting" in October 2012 that also included the word(s) shown. This was done by dividing the rate of searches that include either "vote" or "voting" and the word(s) shown by the rate of searches that include "vote" or "voting." All Google data are downloaded from Google Trends.

There is one important issue with the data. Four states – Virginia, Vermont, California, and Delaware – and three media markets – San Francisco, Washington (DC), and Burlington, Vermont – consistently show up as outliers when measuring changes in search volume through time. While I am not entirely sure the reason for this, it seems to be related to incorrect geographic proxies. Google improved its geographic marks on January 2011, and notably these areas – and no others – saw sharp changes in search rates for most words from December 2010 to January 2011.

I do not include these four states and three media markets in the analysis and recommend omitting them from any analysis measuring changes through time prior to January 2011.

III Using Google to Predict Turnout

The previous section develops a proxy for changes in interest in voting information on Google between election cycles. I now compare this to changes in turnout rates. I measure turnout rate as total votes as a percentage of the voting eligible population. I download this data from Dr. Michael McDonald at `http://elections.gmu.edu`.

III.A 2008 and 2010 Elections

Figure 2.2 shows the results for the 2008 presidential election and 2010 midterm election. Panel (a) compares changes in Google voting interest to changes in turnout rates, from 2004-2008. Panel (b) compares changes in Google voting interest to changes in turnout rates, from 2006-2010. The relationships are both highly statistically significant. If Google search activity for voting was higher in October 2008 (2010), than would be expected from October 2004 (2006), turnout was higher in 2008 (2010) than in 2004 (2006).

The relationship is stronger for the midterm elections than the presidential elections.

Table 2.1 examines the relationship between Google searches and turnout for presidential elections more systematically. Since the search proxy includes searches such as "register to vote" or "early voting," one might suspect that Google data is picking up other factors, and alternative publicly available data may be used instead to predict turnout. However, I show that the relationship between Google voting-related searches and turnout is little affected by including changes in registration rates and early voting rates.

In results not shown, I find that Google searches go beyond other factors that might predict midterm turnout. These include changes in registration rates and states that had a Senate race in one year considered but not the other.

III.B 2012 Out-of-Sample Election

Section III.A found that October search volume for "vote" or "voting," in October, compared to four years earlier, was a strong predictor of state-level turnout in both the 2008 presidential and 2010 midterm elections.

Figure 2.2: *Change in Turnout Rates and Change in October Google Voting-Related Search Rate*

(a) 2004 - 2008 *(b) 2006 - 2010*

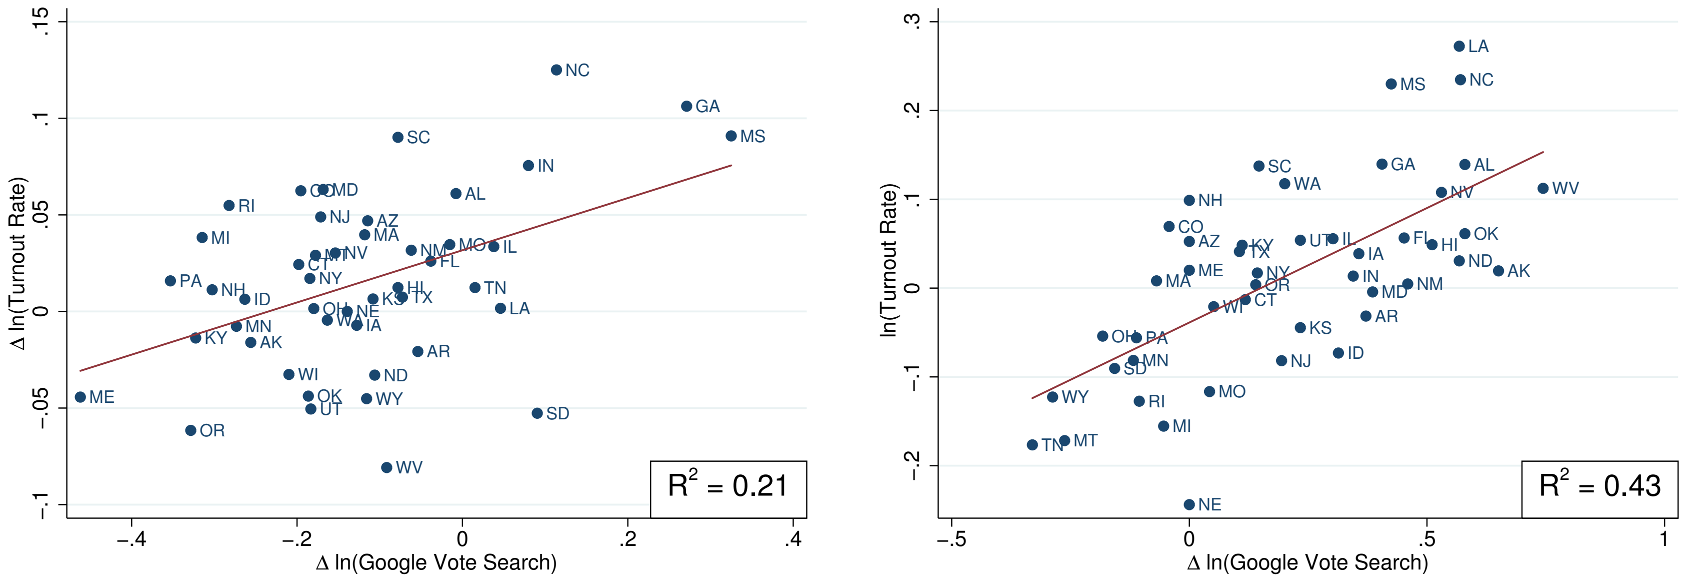

Notes: On the x-axis is the change in the natural log of October search rates for "vote" or "voting." On the y-axis is the change in the natural log of the turnout rate. All Google data are downloaded from Google Trends. Turnout Rate is total presidential votes divided by total eligible voters. These data were downloaded from The United States Elections Project, at `http://elections.gmu.edu/bio.html`.

Table 2.1: *Turnout, 2004-2008*

	(1)	(2)	(3)	(4)
		Δ ln(Turnout Rate)		
Δ ln(Google Vote Search)	0.135***	0.134***	0.151***	0.144***
	(0.038)	(0.037)	(0.039)	(0.031)
Δ ln(Registration Rate)		0.162		0.312***
		(0.109)		(0.090)
Δ ln(Early Vote Rate)		-0.004		-0.004
		(0.015)		(0.011)
Adjusted R-squared	0.19	0.19	0.31	0.41
Observations	46	46	46	46
Weighting	Unweighted	Unweighted	Weighted	Weighted

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: The dependent variable is the change in the natural log of the turnout rate from 2004 to 2008. Turnout Rate is total presidential votes divided by total eligible voters. These data were downloaded from The United States Elections Project, at `http://elections.gmu.edu/bio.html`. $\Delta ln(GoogleVoteSearchRate)$ is the difference between the natural log of Google search rates that include "vote" or "voting," from October 2004 to October 2008. All Google data are downloaded from Google Trends. $\Delta ln(RegistrationRate)$ is change in the change in natural log of registered voters on election day per eligible voters, from 2004 to 2008. $\Delta ln(EarlyVoteRate)$ is change in early votes as a proportion of eligible voting population, from 2004 to 2008. These data were also downloaded from the United States Elections Project, at `http://elections.gmu.edu/bio.html`. California, Vermont, Virginia, Delaware, and Washington D.C. are not included for reasons discussed in the text. Weighted regressions are weighted by 2000 population size, according to the Census.

Since this analysis was all done in the run-up to the 2012 election, the 2012 election presented an excellent opportunity to give an out-of-sample test to this methodology.

Prior to the 2012 election, I calculated the change in Google search rates for "vote" or "voting," from October 2008 to October 2012. I then predicted change in state-level turnout rates based on this data. I made all predictions public on my website the morning of November 6, 2012.

Overall, there was a decrease in country-wide search volumes for "vote" or "voting." I assumed that this was not a meaningful change, as I have often found changes at the national level that did not correlate with national-level changes. I instead assumed that there would be no overall change in turnout but that states in which decreased search rates were more pronounced would see decreased turnout rates. States in which decreased search rates were least pronounced would see increased turnout rates. In fact, the decreased search volume at the country-level did prove meaningful. Overall turnout rates were lower in 2012 than in 2008.

Column (1) of Table 2.2 shows the predicted changes in turnout based on $\Delta ln(Google\ Vote\ Search)_{2012,2008}$ I use the coefficient from the midterm elections, and, as mentioned, assumed that overall turnout would be unchanged.

Column (3) shows the actual change in turnout rates, by state. Figure 2.3 compares the two. The relationship is statistically significant, though somewhat weak, in unweighted regressions. (Panel (a)). Panel (b) of Figure 2.3 weights regressions by 2000 population. The relationship is much stronger. Similar to the 2008 election, smaller states had significantly larger errors in 2012.

Table 2.3 further explores the relationship between a state's predicted change in turnout, from 2008 to 2012, and actual change in turnout, from 2008 to 2012. One interesting finding is that errors in predicting turnout, using Google, were highly serially correlated. States in which Google searches would have under-estimated turnout in both 2008 and 2010 also under-estimated turnout in 2012. This suggests that there are trends in voting-related searches that are not meaningful and should be controlled for in future predictions. This has the potential to greatly improve the methodology.

I also add to the regressions dummy variables that take the value 1 for any state that had at least one fatality caused by Hurricane Sandy (Connecticut, Maryland, New Jersey, New York, North Carolina, and West Virginia). This variable is always negative but does not meaningfully affect the coefficient on the Google-predicted turnout.

As I only have 46 observations, I am limited in how many demographics variables I add. But adding percent black, percent Hispanic, and percent College Grad does not meaningfully change the coefficient on the Google-predicted turnout.

In sum, the out-of-sample predictions were pretty good. The predicted changes in turnout, from 2008 to 2012, correlated with actual changes in turnout, from 2008 to 2012. However, the analysis suggests two changes in future predictions. First, changes in national patterns may have more meaning than I realized. And data should be de-trended. The Google analysis led to consistent under or overestimates in certain states. This would not happen with a proper de-trending.

Table 2.2: *State Turnout Predictions for 2012 Election*

	Predicted 2012 Turnout **FILLED IN PRIOR TO 2012 ELECTION**		*Actual 2012 Turnout*	
	(1)	(2)	(3)	(4)
State	Δ *ln(Turnout)*	*Turnout Rate*	Δ *ln(Turnout)*	*Turnout Rate*
NV	.06	60.2	0	57.1
UT	.05	58.8	-.01	55.4
MD	.05	70.3	-.01	66.2
LA	.05	64.1	-.01	60.4
ND	.04	65.3	-.03	60.6
AR	.03	53.8	-.04	50.5
SD	.03	66.6	-.09	59.4
WY	.03	65	-.06	58.9
AK	.03	69.7	-.14	58.9
MT	.02	67.9	-.06	62.6
FL	.02	67.5	-.04	63.5
NH	.02	73	-.02	70.1
TN	.01	57.6	-.09	52.2
NM	.01	61.6	-.11	54.7
OH	.01	67.6	-.03	64.6
ME	.01	71.6	-.04	68.1
WV	.01	50.5	-.07	46.3
ID	0	63.9	-.06	59.6
MN	0	78.1	-.03	75.7
OR	0	67.4	-.07	63.2
IA	0	69.2	.01	69.9
OK	0	55.6	-.13	49.2
NC	0	65.8	-.01	64.6
HI	-.01	48.5	-.1	44.2
SC	-.01	57.4	-.02	56.6
KY	-.01	57.4	-.05	55.3
IN	-.01	58.6	-.07	55.1
AL	-.01	60.4	-.05	58
WI	-.01	71.7	0	72.5
TX	-.01	53.8	-.08	49.7

	Predicted 2012 Turnout **FILLED IN PRIOR TO 2012 ELECTION**		*Actual 2012 Turnout*	
	(1)	(2)	(3)	(4)
State	*Δ ln(Turnout)*	*Turnout Rate*	*Δ ln(Turnout)*	*Turnout Rate*
MI	-.01	68.3	-.07	64.7
PA	-.02	62.2	-.07	59.4
MA	-.02	65.6	-.01	66.3
WA	-.02	65.1	-.04	64.1
KS	-.02	61.1	-.08	57
NE	-.02	61.5	-.05	60.1
CO	-.02	69.7	-.01	70.3
RI	-.02	60.7	-.06	58
GA	-.03	60.8	-.07	58.4
NJ	-.03	65	-.08	61.9
IL	-.03	61.6	-.08	58.9
CT	-.04	63.9	-.09	60.9
MO	-.04	64.9	-.08	62.5
AZ	-.06	53.1	-.07	52.9
NY	-.06	55.6	-.1	53.2
MS	-.08	55.9	-.02	59.7

Notes: This Table compares state-level turnout predictions, based on Google search data, made prior to the 2012 election, to actual state-level turnout rates. Column (1) shows predicted change in the natural log of the turnout rate, from 2008 to 2012. Column (2) shows the predicted 2012 turnout rate. Column (3) shows actual change in the natural log of the turnout rate, from 2008 to 2012. Column (4) shows actual 2012 turnout rate. Turnout Rate is total presidential votes divided by total eligible voters. These data were downloaded from The United States Elections Project, at `http://elections.gmu.edu/bio.html`. California, Vermont, Virginia, Delaware, and Washington D.C. are not included for reasons discussed in the text.

Figure 2.3: *Change in Turnout Rates and Change in October Google Voting-Related Search Rate, 2008-2012*

(a) Unweighted *(b) Weighted*

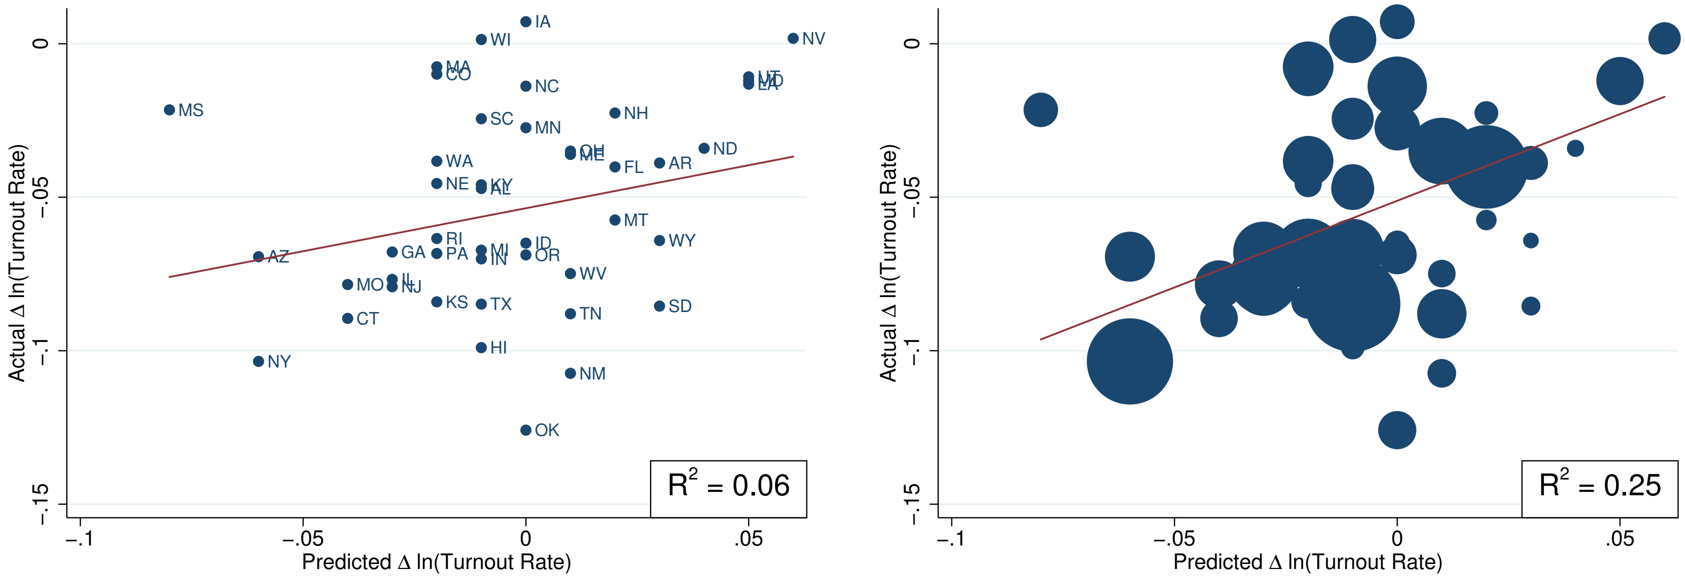

Notes: On the x-axis is the change in the natural log of October search rates for "vote" or "voting." On the y-axis is the change in the natural log of the turnout rate. All Google data are downloaded from Google Trends. Turnout Rate is total presidential votes divided by total eligible voters. These data were downloaded from The United States Elections Project, at `http://elections.gmu.edu/bio.html`. Weighted regressions are weighted by 2000 population, from the Census.

Table 2.3: *Actual Turnout and Predicted Turnout, 2012*

	Δ ln(Turnout Rate)					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ ln(Google Vote Search)	0.293*	0.482***	0.470***	0.571***	0.688***	0.694***
	(0.173)	(0.161)	(0.162)	(0.149)	(0.135)	(0.128)
Residual, 2008		0.205	0.227		0.263**	0.232*
		(0.122)	(0.146)		(0.118)	(0.120)
Residual, 2010		0.205***	0.205***		0.182***	0.217***
		(0.056)	(0.060)		(0.052)	(0.051)
Affected by Hurricane Sandy		-0.003	-0.005		-0.004	-0.006
		(0.013)	(0.015)		(0.010)	(0.010)
Percent Black			-0.000			-0.001
			(0.001)			(0.001)
Percent Hispanic			-0.001*			-0.001***
			(0.001)			(0.000)
Percent College Grad			0.000			0.001
			(0.001)			(0.001)
Adjusted R-squared	0.04	0.29	0.29	0.23	0.44	0.53
Observations	46	46	46	46	46	46
Weighting	Unweighted	Unweighted	Unweighted	Weighted	Weighted	Weighted

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: The dependent variable is the change in the natural log of the turnout rate from 2008 to 2012. Turnout Rate is total presidential votes divided by total eligible voters. These data were downloaded from The United States Elections Project, at `http://elections.gmu.edu/bio.html`. Predicted Δ ln(Turnout Rate) is the Google-based predictions, in Column (3) of Table 2.2. All Google data are downloaded from Google Trends. Residual, 2008 is the residual from a regression of $\Delta ln(TurnoutRate), 2004-2008$ on $\Delta ln(GoogleVoteSearch), 2004-2008$. Residual, 2010 is the residual from a regression of $\Delta ln(TurnoutRate), 2006-2010$ on $\Delta ln(GoogleVoteSearch), 2006-2010)$. Affected by Hurricane Sandy is a dummy variable that takes the value 1 for states Connecticut, Maryland, New Jersey, New York, North Carolina, and West Virginia. The demographics variables are from the 2000 Census. Weighted regressions are weighted by 2000 population, also from the Census. California, Vermont, Virginia, Delaware, and Washington D.C. are not included for reasons discussed in the text.

IV Using Google Data to Estimate Election Demographics

Section III shows that a state's Google searches prior to an election can predict an area's turnout.

Can this information be used to improve election projections?

In this section, I discuss how this information might be used to predict the composition of the electorate.

IV.A 2008 Election

According to exit polls, roughly 88 % of African-American voters supported Democrat John Kerry in 2004. Roughly 95 % of African-American voters supported Obama in 2004.

Roughly 77 % of white evangelical voters supported George W. Bush in 2004. Roughly 74 % of white evangelical voters supported John McCain in 2008.

In addition, Hispanic, Jewish, and young voters lean Democratic. Non-Hispanic white, Mormon, and older voters lean Republican.

Since certain groups are so likely to support a particular party, composition of the electorate is crucial in determining the final vote totals. A one percentage point increase in black turnout will be expected to lead to a nearly one percentage point increase in the Democratic candidate's vote share.

By comparing areas' voting interest, as proxied by Google, to areas' demographics, we might predict whether any demographic can be expected to turnout in unusually high numbers.

This exercise can best be done using a smaller level of aggregation of the data than the state and thus allowing for bigger differences in demographics. I thus do this analysis using data from media markets in the United States.

Consider this exercise for the 2008 presidential election.

Column (1) of Table 2.4 shows the correlation between the percent of the population that is black and change in October voting-related search rates from 2004 to 2008.

This relationship is highly statistically significant and robust. Column (2) adds a number of other demographics variables. The other variable that is consistently statistically significant is the Hispanic population. The Google data consistently predicted increases about 1/3 to 1/2 as

Table 2.4: *Change in Google Search Rate and Demographics, 2004-2008*

	Δ ln(Google Voting Search)			
	(1)	(2)	(3)	(4)
Percent Black	0.722***	0.864***	0.615**	1.028***
	(0.230)	(0.189)	(0.274)	(0.266)
Percent Hispanic		0.301***		0.533***
		(0.108)		(0.147)
Percent 18-34		0.156		0.249
		(0.863)		(0.854)
Percent 65-Plus		-0.085		0.281
		(0.873)		(0.981)
Percent College Grad		-0.488		-0.230
		(0.469)		(0.495)
Percent Kerry		-0.576		-0.865**
		(0.437)		(0.402)
Adjusted R-squared	0.09	0.21	0.16	0.25
Observations	176	176	176	176
Division Fixed Effects	No	No	Yes	Yes

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: The dependent variable is the change in the natural log of Google search rates for the words "vote" or "voting," from October 2004 to October 2008. Regressions are for 176 media markets for which data was available. Percent Kerry is Kerry's share of the two-party vote. The demographics variables are from the 2000 Census. All regressions are weighted by 2000 population, also from the Census.

large among the Hispanic population as among the African-American population. Column (3) reproduces the results of Column (1) but adds Census division fixed effects. Column (4) reproduces the results of Column (2) with Census division fixed effects.

Choosing different variables, in regressions not reported, does lead to some instability in some of the coefficients on the age, college, and percent Kerry variables. However, the percent black and percent Hispanic proved to be remarkably robust.

IV.B 2012 Out-of-Sample Election

Prior to the election, I predicted the composition of the electorate using Google search data.

Table 2.5 shows results analyzing changes in Google search activity, from October 2008 to October

2012. There is a remarkably strong, robust, and stable relationship between changes in Google voting-related search rates and population size. About 50 percent of the variation can be explained by the log of the population rate. And I was unsuccessful in finding any correlate of population size that explained away this result. (There was no similar relationship between population size and changes in Google search rates, from 2004 to 2008, and the results in 2.4) are little affected by adding a control for the log of population.

I was unable to figure out the reason for this relationship. It does not seem to have to do with rural vs. urban areas, which might be affected by internet penetration rates. Nor does it seem to be a universal fact of Google data that population rates correlate with changed search rates over this time period. I assumed that population size would not meaningfully impact the voting composition of the electorate but I did control for it in all my regressions.

Column (1) of Table 2.5, while including the population size control, adds the variable percent black. The relationship is not statistically significant. If we interpret these as actually capturing the likelihood of blacks' voting, as it seemed to in 2008, we can easily reject a drop in the black voting population as large as the gains in black voting population in 2012. Column (2) of Table 2.5 adds the same controls used in Column (2) of Table 2.4. None of the coefficients are statistically significant. at the 5 % level. Column (4) and (5) add division fixed effects to the regressions of Columns (1) and (2), respectively. Again, none of the demographics coefficients are statistically significant at the 10 % level.

Columns (3) and (6) add to the regressions of Columns (2) and (5) religion variables from the 2000 Congregations and Membership Study. Some hypothesized that conservative evangelicals would be less motivated to turn out due to Mitt Romney's Mormon faith. There is no statistically significant relationship between evangelical adherents per capita and changes in Google searches. Intriguingly, of all the religion variables, the only statistically significant coefficient is on Percent Other Religion. This consists mostly of Mormons. There is strong and robust evidence that Mormons will turn out at historic levels to support Romney. The coefficient on Percent Other Religion in Table 2.5 is similar in size and robustness to the coefficient on Percent Black in Table 2.4.

Table 2.6 combines the results of Table 2.5 with changing demographics in the country to make

Table 2.5: *Change in Google Search Rate and Demographics, 2008-2012*

	Δ ln(Google Voting Search)					
	(1)	(2)	(3)	(4)	(5)	(6)
ln(Total Population)	-0.152***	-0.124***	-0.148***	-0.160***	-0.128***	-0.160***
	(0.013)	(0.017)	(0.024)	(0.017)	(0.023)	(0.027)
Percent Black	-0.127	-0.108	-0.121	-0.184	-0.098	-0.197
	(0.155)	(0.163)	(0.183)	(0.198)	(0.215)	(0.211)
Percent Hispanic		0.029	0.032		-0.093	-0.168
		(0.082)	(0.117)		(0.099)	(0.147)
Percent 18-34		0.796	0.055		0.686	-0.352
		(0.836)	(0.669)		(0.903)	(0.745)
Percent 65-Plus		1.219*	0.658		1.230	0.290
		(0.707)	(0.971)		(0.870)	(0.993)
Percent College Grad		-0.554*	-0.672**		-0.565	-0.712*
		(0.281)	(0.328)		(0.359)	(0.371)
Percent Kerry		-0.009	0.061		0.016	0.107
		(0.129)	(0.172)		(0.154)	(0.135)
Percent Evangelical			0.121			0.203
			(0.205)			(0.265)
Percent Mainline Prot			-0.007			0.362
			(0.273)			(0.361)
Percent Catholic			0.063			0.224
			(0.201)			(0.193)
Percent Jewish			0.197			0.841
			(0.918)			(1.389)
Percent Islam			2.550			3.491
			(4.450)			(5.519)
Percent Other Religion			0.480***			0.684***
			(0.129)			(0.240)
Adjusted R-squared	0.67	0.69	0.70	0.69	0.70	0.72
Observations	194	194	194	194	194	194
Division Fixed Effects	No	No	No	Yes	Yes	Yes

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: The dependent variable is the change in the natural log of Google search rates for the words "vote" or "voting," from October 2008 to October 2012. Regressions are for 194 media markets for which data were available. The demographics variables in Columns (2) and (5) are the same ones used in Table 2.4. Religion variables are from 2000 Congregations and Membership Study. All regressions are weighted by 2000 population, from the Census.

Table 2.6: *Predicted 2012 Demographics*

Demographic	*2004 Share* (1)	*2008 Share* (2)	*Predicted 2012 Share* **FILLED IN PRIOR TO 2012 ELECTION** (3)	*Actual 2012 Share* (4)
Black	11	13	13	13
Hispanic	8	9	10	10
White evangelical	23	26	26	26
Other religion	7	6	7	7
Age 18-29	17	18	18	19
Age 65+	16	16	16	16

Notes: Columns (1), (2), and (4) are demographics estimates from CNN's analysis of the 2004, 2008, and 2012 exit polls. Column (3) are predictions made prior to the 2012 election, based on the regressions in Table 2.5.

out-of-sample predictions for the composition of the 2012 electorate. All predictions were made public, as shown, on my website the morning of the election.

The only robust predictor of changing Google search volumes was Other Religion. Based on the magnitude, I predicted that 'other' religion would rise from 6 to 7 percent of the voting population. I predicted that all other voting rates would stay the same. However, since the Hispanic share of the population grew substantially over the previous four years, I assumed that Hispanics would make up 10 percent of the electorate, compared to the 9 percent Hispanic share of the 2008 electorate.

Contrary to many, I did not foresee decreased share of black voters. As shown in Table 2.6, parts of the country with the highest black populations were Googling for voting information in 2012 at the elevated rates of 2008. And I predicted that white evangelicals would make up 26 percent of the electorate, just as they did in the 2008 electorate. I did not foresee major changes in the age composition of the electorate.

The predictions performed well. Hispanics, due to population growth, indeed made up 10 percent of the electorate. Blacks again made up 13 percent of the electorate. The white evangelical share of the electorate did not decrease despite Romney's presence on the ticket. Percent declaring 'other religion' was indeed elevated, and this was almost certainly due to increased Mormon turnout. Utah was, as predicted in Table 2.2, one of the states that ranked highest in the difference between 2012 and 2008 turnout rates.

The only number that was off was the share of the electorate age 18-29. I guessed that it would remain 18 percent of the electorate, wheres young voters rose to 19 percent of the electorate. My methodology may struggle to predict changes in voting rates by age, since media markets do not differ that much in age demographics.

V Using Google Data to Predict Democratic Overperformance

The simplest use of the Google data is just adjusting predictions based on Google-predicted state-level turnout. Pollsters might weaken or strengthen their likely voter cutoff number based on Google search activity.

High turnout is usually a good sign for Democrats (Gomez et al., 2008). Table 2.7 examine the relationship between changes in turnout, compared to four years earlier, and Democratic performance as compared to polls in the 2008 and 2012 presidential elections and 2010 Senate elections. As my measure of predicted performance based on the polls, I use the final predictions of Nate Silver.

V.A 2008 and 2012 Presidential Elections

Columns (1) and (3) of Table 2.7 show that the change in turnout rates, compared to four years earlier, did correlate with Obama's over performance in both the 2008 and 2012 elections. However, of the two presidential elections, the relationship is only statistically significant with the 2008 election. Interestingly, there is positive serial correlation in Obama's over performing the polls. If Obama did better than Nate Silver's final predictions in a state in 2008, he was also expected to do better in 2012.

Columns (2) and (4) of Table 2.7 show that changing Google search rates for voting, would not have successfully predicted Obama performance, relative to Nate Silver's predictions, in either the 2008 or 2012 presidential elections.

However, Columns (3) and (4) show that changes in October Google voting-related searches, while they did predict changes in turnout, did not predict Obama over performance.

V.B 2010 Midterm Election

Column (5) of Table 2.7 shows that change in turnout was a strong predictor of Democratic performance, compared to the polls, in 2010 Senate elections. If turnout was higher in 2010, relative to 2006, the Democratic Senate candidate could be expected to perform better than Nate Silver's final projections. Interestingly, once again Obama 2008 over-performance was positively correlated with Democratic over-performance in a different year. If Obama did better than Nate Silver predicted in 2008, Democratic Senate candidates would be expected to do better than Nate Silver predicted in 2010.

Column (6) shows that, for the 2010 Midterm election, elevated Google search rates were highly statistically significant predictors of Democratic performance compared to Nate Silver's final

Table 2.7: *Democratic Overperformance and Change in Turnout*

	Actual Democrat-Predicted Democrat					
	2008		2012		2010	
	(1)	(2)	(3)	(4)	(5)	(6)
Δ ln(Turnout Rate)	0.138**		0.018		0.099**	
	(0.056)		(0.099)		(0.047)	
Δ ln(Google Vote Search)		-0.014		-0.028		0.053***
		(0.015)		(0.021)		(0.018)
Actual Democrat-Predicted Democrat, 2008			0.310*	0.280*	0.918***	0.996***
			(0.157)	(0.163)	(0.275)	(0.204)
Adjusted R-squared	0.05	-0.01	0.09	0.13	0.46	0.54
Observations	46	46	46	46	31	31

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: The dependent variable is the difference between the Democrat's share of the two-party vote and Nate Silver's final prediction for the Democrat's share of the two-party vote. In 2008 and 2012, this is Obama's vote share. In 2010, this is the Senate Democratic candidate's vote share. Δ ln(Turnout Rate) is the change in turnout rate compared to four years earlier. Δ ln(Google Vote Search) is the change in Google search rates for "vote" or "voting," in October, compared to the October four years earlier. Actual Democrat-Predicted Democrat, 2008 is the difference between Obama's 2008 vote share and Obama's 2008 vote share, as predicted by Nate Silver. Thus, it is also the dependent variable in Columns (1) and (2). Turnout Rate is total presidential votes divided by total eligible voters. These data were downloaded from The United States Elections Project, at `http://elections.gmu.edu/bio.html`. All Google data are downloaded from Google Trends. Nate Silver's predictions come from `http://trumantolong.blogspot.com/2008/11/grading-nate-silver-and-538.html`, `http://elections.nytimes.com/2010/forecasts/senate`, and `http://fivethirtyeight.blogs.nytimes.com/`. California, Vermont, Virginia, Delaware, and Washington D.C. are not included for reasons discussed in the text.

predictions.

Note, too, that the changing predictions are quantitatively significant. Figure 2.2 Panel (b), shows that $\Delta ln(TurnoutRate)$ varies between about 0.25 and -0.25. Column (5) of Table 2.7 shows that changes in turnout of this magnitude would mean an increase in Democratic performance of about 2.5 percentage points or decrease in Democratic performance of about 2.5 percentage points, relative to Nate Silver's predictions. Column (6) shows that Google could have predicted much of this over or under performance.

Thus, Google search rates would not have predicted state-level over or under-performance in the 2008 or 2012 presidential elections. But it would have predicted state-level over or under-performance in the 2010 midterm Senate elections. The huge variance in turnout in Midterm elections may explain the difference. Google data for voting information in October should certainly be considered as a supplementary source in the 2014 midterm elections.

VI Conclusion

I show that search volume for "vote" or "voting," in October, strongly correlates with final turnout, in United States elections. I show two ways that this can help political forecasters.

The data might also prove useful in understanding the determinants of voting decisions. High-frequency search data can be used to study the effectiveness of various interventions in increasing voting intentions. In addition, comparing voting intention – as proxies on Google – to actual voting behavior can be used to better studying how last-minute laws or advertising campaigns influence turnout.

Chapter 3

Unreported Victims of an Economic Downturn

I Introduction

How does an economic downturn affect child maltreatment?[1]

We might expect child maltreatment would rise when times are tough. Indeed, sociologists and public health scholars hypothesize that unemployment – and the resulting stress, anger, and low self-esteem – are major risk factors in child maltreatment (Stith et al., 2009; Dooley et al., 1994; Baum et al., 1986; Linn et al., 1985).

Previous research, however, has generally not found a robust positive correlation between community-level unemployment and reported child maltreatment rates (Lindo et al., 2013). Previous research, in fact, has found some evidence for the opposite relationship. Controlling for state and year fixed effects, Paxson and Waldfogel (1999) find that poor economic conditions are associated with fewer reported victims of maltreatment.[2]

[1]In this paper, when I use the phrase "child maltreatment," I mean child abuse and neglect, as defined by the Federal Government, in CAPTA. This includes one of the following two behaviors: 1) Any recent act or failure to act on the part of a parent or caretaker which results in death, serious physical or emotional harm, sexual abuse or exploitation; or 2) An act or failure to act which presents an imminent risk of serious harm.

[2]Lindo et al. (2013), in addition to an extensive literature review, also present new evidence, using county-level data from California. They find that the male unemployment rate is positively correlated with child maltreatment rates; the

This paper begins by extending this puzzle using data from the recent Great Recession. I show that states that were most affected by the recession saw the largest decreases in referral rates for maltreatment. This relationship survives a fairly large set of controls.

How can we reconcile the strong theoretical arguments that poor economic conditions should increase child maltreatment with the area-level, time-series evidence to the contrary?

The paper next suggests and finds evidence for a reconciliation: a recession decreases the reporting rates of child maltreatment. Evidence suggests that the majority of maltreatment cases are not reported to authorities (Finkelhor et al., 2005; Hussey et al., 2006; Sedlak et al., 2010). During the Great Recession, state and local budgets were slashed in hard-hit areas. Budget cuts to agencies likely to deliver and receive reports might lower reporting rates. Individuals have reported long waits at child maltreatment hotlines; many hang up before getting through (Cardona, 2011; Valdez, 2012; Eckholm, 2009).

To test this hypothesis, I use two alternative proxies for area-level maltreatment rates less likely to be biased by reporting rates: rates of child mortality from neglect and the fraction of Google searches that include the phrase "child abuse" or "child neglect." Child mortalities must be reported, and using an extreme form of an incident with mandatory reporting is a common strategy among economists studying crime (Levitt, 1998; Aizer, 2010). The motivation for the Google proxy is that it can capture community-level suspicion of child maltreatment, including many cases that are never actually reported.

Both proxies comparatively *increased* in hard-hit areas. Each obviously is an imperfect proxy for overall maltreatment, but the likely sources of error are very different. The fact that both showed a comparative rise in recession-hit areas is evidence that the recession caused an *increase* in actual maltreatment.

I find additional evidence for this explanation: First, recession-hit areas saw an increase in the percent of referred cases that were substantiated. This would make sense if an increase in the cost of reporting cases has a bigger effect on the cases less likely to be substantiated. Second, I study high-

female unemployment rate is negatively correlated with child maltreatment rates; and the overall unemployment rate is not correlated with child maltreatment rates. Bitler and Zavodny (2004), Bitler and Zavodny (2002) and Seiglie (2004) find small, and generally insignificant, relationships between economic conditions and reported child maltreatment rates.

frequency, national data for Google search queries likely made by older victims of maltreatment. These searches include "my dad hit me." Such searches rise when weekly unemployment claims are higher, lending further support to the hypothesis that unemployment increases actual maltreatment and any alternative relationship must be due to changes in reporting rates. Third, controlling for the Google proxy for unemployment, areas that spend little money on public services have significantly lower referral rates.

The results imply that a one percentage point increase in the unemployment rate increases actual maltreatment by roughly 2.5 percent but decreases referrals by roughly 3.3 percent. In other words, a one percentage point increase in the unemployment rate decreases the percentage of actual maltreatment cases that are referred by 5.8 percent.

This paper shows that great caution should be used in interpreting results relying on reported rates of maltreatment as proxies for actual maltreatment. If the independent variable affects reporting rates, false conclusions might be drawn about the effects of that variable on actual incidents. In addition, the paper suggests an effect of budget cuts during economic downturns. And the paper leads to an unfortunate conclusion: just when children are most in need of assistance, they are least likely to get it.

The paper builds on earlier work suggesting Google data can be useful in social science research (Varian and Choi, 2010; Ginsberg et al., 2009; Scheitle, 2011; Askitas and Zimmermann, 2009). It follows Stephens-Davidowitz (2012) in using Google data to measure a variable complicated by social desirability bias. It is the first paper to suggest using Google to study crime. Reporting issues complicate analysis of both child maltreatment and crime more generally (Levitt, 1998). Google search data provide an additional data source. In Section III.B and the Conclusion, I discuss the potential of – and some reasons for caution with – using Google data to proxy crime.

It is important to note that the welfare effects of being referred for child maltreatment remain uncertain and limit fully understanding the welfare impacts of these results (Doyle Jr., 2007). Both estimating these impacts, and designing foster care systems that are unambiguously better than abusive households, remain crucial.

II Child Maltreatment in the United States

This section briefly reviews some facts about child maltreatment in the United States.

In 2010, more than 3 million (or roughly 1 in 25) children are reported victims of maltreatment. (See Table 3.1.)

Table 3.2 shows the sources for reported maltreatment cases in 2010. A slight majority (58.6%) came from professionals. The most common professional sources were education personnel (16.4%) legal and law enforcement personnel (16.7%), and social services personnel (11.5%). Almost all the nonprofessional sources (27.7%) were anonymous sources (9.0%), parents (6.8%), other relatives (7.0%), and friends and neighbors (4.4%). Very few reports came from the alleged victim (0.4%) or the alleged perpetrator ($\approx$ 0.0%). The rest of the reports (13.7%) came from other or unknown sources.

Maltreatment cases proceed in three stages.

1. An individual refers a case to child protective services.
2. Child protective services investigates the case.
3. Child protective services substantiates the case.

In 2010, more than 50 percent of referred cases were investigated. Roughly 24 percent of investigated cases were substantiated. Some scholars argue that most unsubstantiated cases should actually classify as maltreatment and even that alleged perpetrators of unsubstantiated cases are just as dangerous as alleged perpetrators of substantiated cases (Kohl et al., 2009; Drake, 1996; Hussey et al., 2005).

Scholars suspect that child maltreatment is significantly underreported to authorities (Finkelhor et al., 2005; Hussey et al., 2006). The Fourth National Incidence Study of Child Abuse and Neglect estimated that only 32 percent of children that experienced observable harm from maltreatment were investigated by authorities: most cases were not reported to authorities (Sedlak et al., 2010).

The welfare effects of child maltreatment, evidence suggests, are substantial. Compared to demographically similar adults who were not victims, grown-up victims of child maltreatment

Table 3.1: *Child Maltreatment*

Maltreatment	*Annual Incidents*
Google Searches for "child abuse" or "child neglect"	$\approx$ 8.4 million
Referrals to agencies	$\approx$ 3.3 million
Responses by agencies	$\approx$ 2 million
Substantiated incidents	436,321
Child mortalities from neglect	1,560

Notes: According to Google AdWords, on 3/22/12, there were an average of 673,000 monthly searches in the United States, on desktops and laptops, including the phrase "child abuse." There were 27,100 including "child neglect." Multiplying by 12 and adding yields the estimate. Other estimates are from Child Maltreatment 2010.

Table 3.2: *Sources of Maltreatment Reports, 2010*

Report Source	*Percent*
PROFESSIONAL	**58.6**
Child Daycare Providers	0.7
Educational Personnel	16.4
Foster Care Providers	0.5
Legal and Law Enforcement Personnel	16.7
Medical Personnel	8.2
Mental Health Personnel	4.6
Social Services Personnel	11.5
NONPROFESSIONAL	**27.7**
Alleged Perpetrators	0.0
Alleged Victims	0.4
Anonymous Sources	9.0
Friends and Neighbors	4.4
Other Relatives	7.0
Parents	6.8
OTHER AND UNKNOWN	**13.7**
Other	7.9
Unknown	5.8

Notes: Source: *Child Maltreatment 2010 (2011)*.

have higher probability of mental illness (Brown et al., 1999; Mathews et al., 2008), are more likely to engage in criminal behavior (Widom, 1989; Lansford et al., 2007; Currie and Tekin, 2012), and earn substantially less (Currie and Spatz Widom, 2010).

III The Great Recession and Child Maltreatment

III.A The Effects of the Great Recession on Reported Child Maltreatment Rates

I now examine the effects of the recent recession on reported rates of child maltreatment. To do so, I utilize the fact that different parts of the United States were differently affected by the recession. And such differences were largely for idiosyncratic reasons that one would not expect to otherwise be correlated with changes in child maltreatment rates or child maltreatment reporting rates. Following Stevenson and Wolfers (2011) I average two years of data (2006 and 2007) as pre-economic crisis and two years of data (2009 and 2010) as post-economic crisis. I then measure the change in unemployment rates over this time period as a proxy for exposure to the recent crisis. I confirm that all the results in this paper are little affected by using alternative measures of economic performance, such as GDP, or slightly different time periods that similarly capture recession exposure.

In this section, I measure reported cases using both referral rates per child and response/investigation rates per child. I will discuss substantiated cases in Section III.D. I study referral rates and response rates now to focus more strongly on reporting pressures and because some scholars argue, due to flaws in the substantiation process, that this is the better measure of maltreatment incidence (Kohl et al., 2009; Drake, 1996; Hussey et al., 2005). Note that many other scholars study only substantiated cases. I hope to clarify the differences between previous results and my results more clearly in Section III.D.

There are two factors that influence the referral and investigation rates of maltreatment (reported maltreatment). First is the actual maltreatment and second is the proportion of maltreated cases that are reported.

In other words,

$$reported\ maltreatment_{i,t} \equiv reporting_{i,t} + maltreatment_{i,t} \tag{3.1}$$

The economic crisis may have changed the actual maltreatment rate. In other words,

$$\Delta maltreatment_i = \beta_0 + \beta_1 \Delta Unemp_i + \beta_3 X_i + \epsilon_i \tag{3.2}$$

It might also have changed the reporting rate of maltreatment. In other words,

$$\Delta reporting_i = \alpha_0 + \alpha_1 \Delta Unemp_i + \alpha_3 X_i + z_i \tag{3.3}$$

Overall, I can test the effects of the economic crisis on the rate of reported maltreatment cases. The empirical specification is:

$$\Delta reported\ maltreatment_i = \hat{\beta}_0 + \hat{\beta}_1 \Delta Unemp_i + \hat{\beta}_3 X_i + w_i \tag{3.4}$$

where $\hat{\beta}_1 = \beta_1 + \alpha_1$

The results, without any controls, are presented in Figure 3.1. Both referred cases and investigated cases per capita declined in states comparatively affected by the recession. The result is statistically significant for both outcomes.

These correlations, of course, do not mean that the economic crisis caused a decrease in reported child maltreatment cases. One alternative possibility is that demographics changes were correlated with economic conditions in such a way as to affect reported maltreatment rates. This is theoretically unlikely, as demographics would move too slowly to significantly affect reported maltreatment rates. Columns (2) and (4) of Table 3.3 shows that the results are little affected by inclusion of changes in the African-American, Hispanic, and very young populations. The change in the percent of the population that is African-American is tiny in all states except District of Columbia. Thus, the coefficient on this variable, which is often statistically significant, can be explained by large changes in District of Columbia. I confirm that results are similar excluding this variable or excluding District of Columbia from the analysis.

Another possibility is that exposure to the economic crisis was correlated with an alternative factor that was also correlated with changes in child maltreatment rates. If there were such a factor, the coefficient should drop upon adding additional controls (Altonji et al., 2005). Columns (3) and (6) of Table 3.3 add controls for percent Hispanic, percent African-American, and percent college graduate. These data average data from 2006 and 2007 using the American Community Survey. The coefficient is little affected by adding these controls.

In sum, there is a negative relationship between exposure to the economic crisis and reported

Figure 3.1: *Severity of Recession and Change in Reported Maltreatment*

(a) Δ ln(Referral Rate) *(b) Δ ln(Response Rate)*

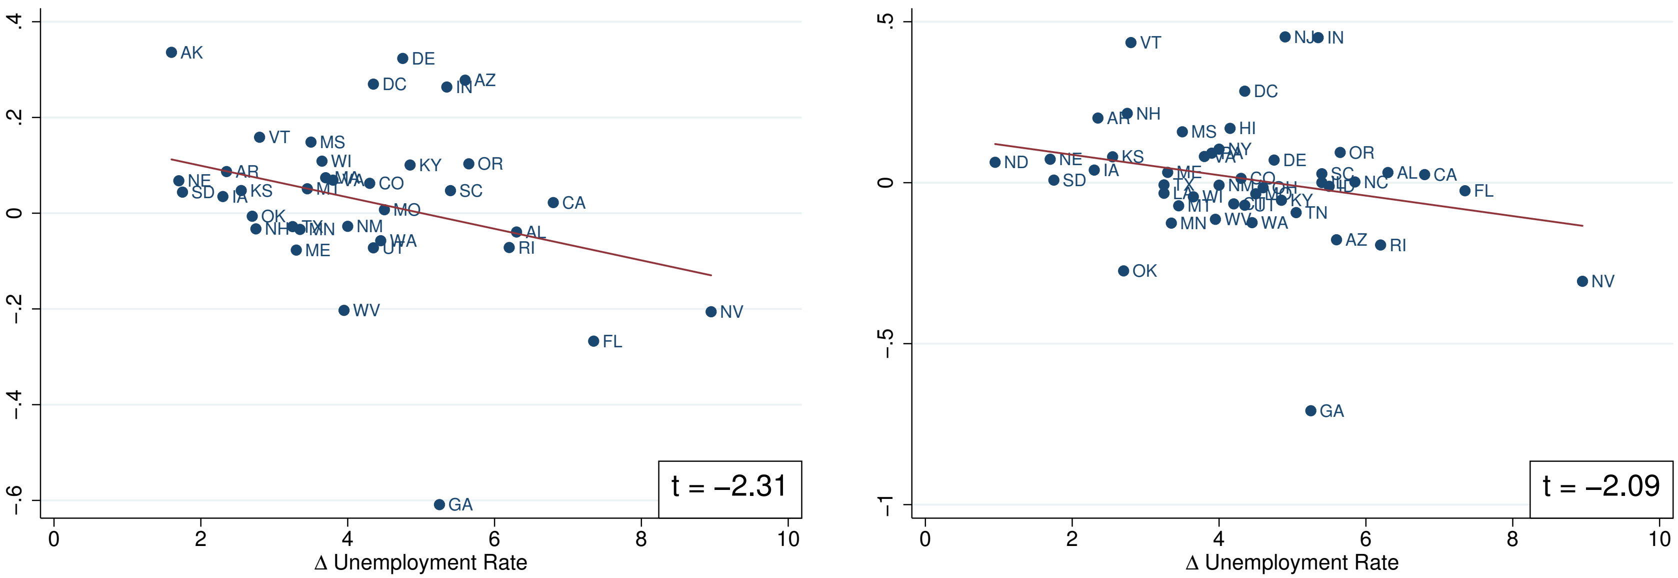

Notes: This figure shows the relationship between change in unemployment and change in two proxies for child maltreatment. Changes for all variables are the difference between the 2009-2010 average value and the 2006-2007 average value. The referral rate is referrals per child, from *Child Maltreatment*. The response rate is responses per child, from *Child Maltreatment*.

Table 3.3: *Reported Cases of Child Maltreatment and Severity of Recession*

	Δ ln(Referral Rate)			Δ ln(Response Rate)		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Unemployment Rate	-0.033**	-0.022*	-0.034**	-0.032**	-0.033*	-0.027*
	(0.014)	(0.012)	(0.014)	(0.015)	(0.017)	(0.015)
Δ % Hispanic		-0.022			-0.041	
		(0.054)			(0.071)	
Δ % Black		-0.076**			-0.107**	
		(0.030)			(0.040)	
Δ % Age 0-4		0.097			-0.148	
		(0.247)			(0.223)	
% Hispanic			0.000			-0.002
			(0.002)			(0.002)
% Black			0.001			0.000
			(0.003)			(0.003)
% College			0.004			0.009*
			(0.003)			(0.005)
Adjusted R-squared	0.07	0.08	0.01	0.04	0.09	0.05
Observations	36	36	36	46	46	46

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: Robust standard errors are in parentheses. Δ variables represent the average, for 2010 and 2009, minus the average, for 2006 and 2007. Referral Rate and Response Rate are from the Children's Bureau Child Maltreatment Annual Reports. Unemployment Rate is from the Bureau of Labor Statistics. Demographics variables are from the Current Population Survey.

cases of maltreatment. And the near-randomness of the crisis and the robustness of the regression results suggest a causal role. There are, by definition, two potential reasons for this: First, the recession caused a drop in actual maltreatment ($\beta_1 < 0$). Second, the recession caused a drop in the reporting rate of maltreatment ($\alpha_1 < 0$). Distinguishing these stories requires alternative proxies for maltreatment that are not likely to be affected by the reporting rates of maltreatment. I now suggest such proxies.

III.B Alternative Proxies for Maltreatment

This section discusses alternative proxies for maltreatment less likely to be affected by reporting bias. The first is child mortality from neglect. Since authorities must report any child mortality and such events must be investigated by child protective services agencies, this is not likely to face reporting bias. The limitation, of course, is that it is an extreme outcome. This means both that it is rare, creating noise in the proxy, and that it may systematically differ from less extreme forms of maltreatment.

Its rareness makes using the natural log of the measure impractical. Large swings from states with few mortalities will create large standard errors. I thus use the change in the fatality rate, which is simply fatalities divided by child population.

In addition, two states note changes in the coding of child fatalities during the time period studied, Mississippi and California. I do not include these states in any regressions with child fatality data, though the main results of this paper are similar including them.

Using Google to Proxy Maltreatment

I supplement the proxy with a new source for crime data: Google search queries. There are two reasons Google search data may have meaningful information on an area's child maltreatment rate. First, victims old enough to use the search engine may look for information. Second, individuals who suspect that a friend, neighbor, or acquaintance is a maltreatment victim may look for information. Many of these suspected cases likely will never be reported to child protective services. Victims rarely report cases themselves. And a large percentage of individuals who suspect maltreatment do not go through with reporting the case. For this to contain information besides official data of referrals, it must be that there are individuals who suspect maltreatment but do not go through with reporting maltreatment. A recent survey found that 28 percent of physicians admitted to suspecting maltreatment but not reporting it, despite being mandated to report it (Gunn et al., 2005).

As a constraint on data-mining, following Stephens-Davidowitz (2012), I use the most salient words. The baseline Google proxy is as follows:

$$Google\ Maltreatment_{i,t} \equiv \frac{[searches\ w/\ (\text{"}child\ abuse\text{"} \mid \text{"}child\ neglect\text{"})]_{i,t}}{[total\ searches]_{i,t}} \tag{3.5}$$

Figure 3.2 shows the returns for "child abuse." The top returns fit with the idea that a large number of searchers are suspecting child maltreatment. Table 3.4 shows the 'top searches' for the proxy. All of them return results similar to those shown in Figure 3.2.

Table 3.4: *Top Searches for Google Proxy*

Top searches for 'child abuse+child neglect'
child abuse neglect
child abuse statistics
child abuse prevention
about child abuse
child abuse report
child abuse reporting
child services
child sexual abuse
sexual abuse
child abuse services

Notes: These show the 'top searches' for "child abuse+child neglect," 2004-present. It is downloaded on 11/27/2012. Results would be similar regardless of time period selected. Depending on the draw, the 'top searches' might be slightly different. Top searches, according to Google, 'are related to the term,' as determined 'by examining searches that have been conducted by a large group of users preceding the search term you've entered, as well as after,' as well as by automatic categorization.

Prior to the crisis, the proxy positively correlates with an area's child fatality rate, though this relationship is not statistically significant. (Figure 3.3, Panel (a)). The proxy has a statistically significant correlation with an area's referral rate. (Figure 3.3 Panel (b)).

Figure 3.2: *Search for "child abuse"*

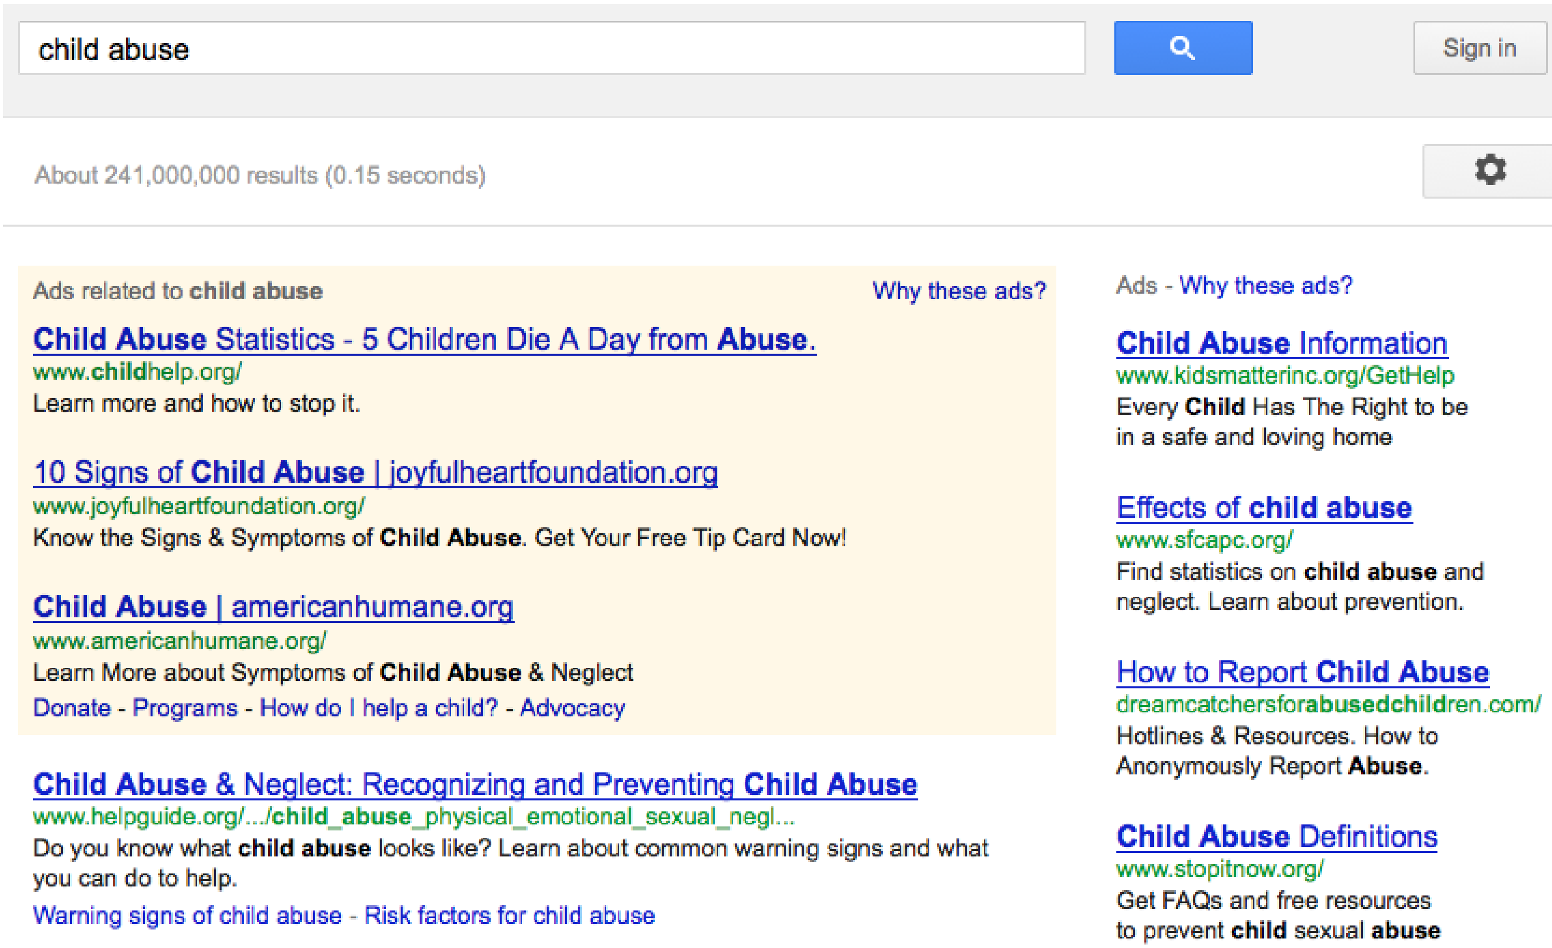

Notes: This shows the returns for a search of "child abuse" on 9/3/2012.

Figure 3.3: *Google Maltreatment and Other Proxies, Pre-Recession*

(a) Fatality Rate *(b) ln(Referral Rate)*

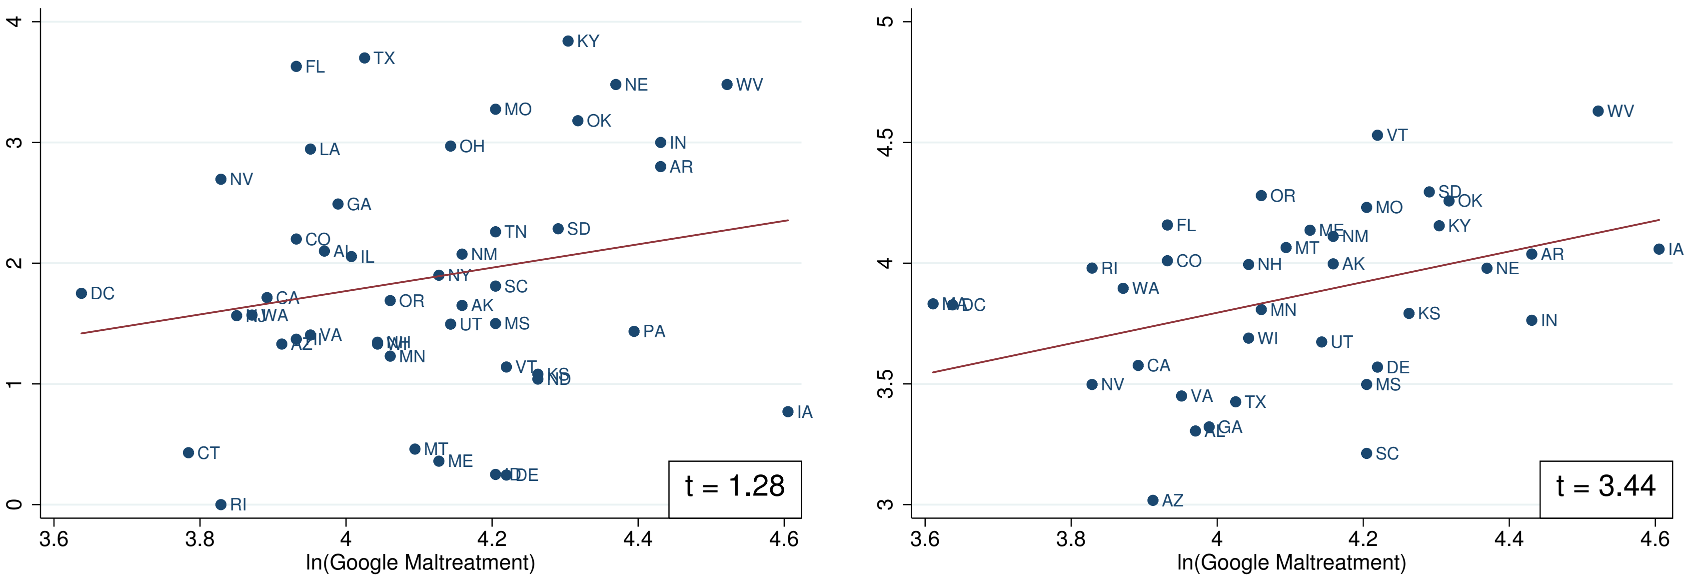

Notes: All variables are averaged for 2006 and 2007. Fatality rate is fatalities per 100,000 children, from *Child Maltreatment*. Referral rate is referrals per child, from *Child Maltreatment*. Google Maltreatment is the percent of Google that searches include "child abuse" or "child neglect," from Google Trends.

In January 2011, Google updated its algorithm to improve precision of geographical estimates. Four states – Vermont, Virginia, California, and Delaware – experience patterns of "child abuse" searching that seem to be errors. In particular, they see large increases prior to January 2011 and then dramatic declines from December 2010 to January 2011 when the geographic codes changed. (They do not see such declines in any other January.) I have found that these states are consistently large, and unexplained, outliers when measuring changes through time for most words. The main analysis would be similar with including these states. However, I do not include these four states in any Google-related analysis.

The majority of searches in the previous Google proxy are, evidence suggests, by individuals suspecting maltreatment. Google would seem to offer another way to measure maltreatment: searches by victims after an incident. Of course, very young victims of maltreatment are not going to use Google. But older victims may very well look for help or advice on Google.

Figure 3.4 shows a return for the search for "dad hit me." The evidence strongly suggests that individuals making searches are young teenagers who have recently been hit by their fathers.

The second Google maltreatment proxy (Google Maltreatment Victim) combines a few of these types of searches that return similar websites and, evidence suggests, are predominantly by recent maltreatment victims.

$$\begin{gathered} Google\ Maltreatment\ (Victim)_t \\ \equiv \\ \frac{[searches\ w/\ (''dad\ hit(s)\ me''|\ ''dad\ beat(s)\ me''\ |\ ''mom\ hit(s)\ me''\ |\ ''mom\ beat(s)\ me'')]_t}{[total\ searches]_t} \end{gathered} \tag{3.6}$$

Caution Using Google Search Data

The Google data might greatly assist researchers trying to understand the causes of maltreatment (as well, of course, as crime more generally). However, there are many reasons the Google proxies might change that do not indicate changes in maltreatment rates. I mention three reasons now; these possibilities motivate many of the data and regression choices in the next section and must always be considered by researchers using the data source for this, or similar, reasons.

Figure 3.4: *Search for "dad hit me"*

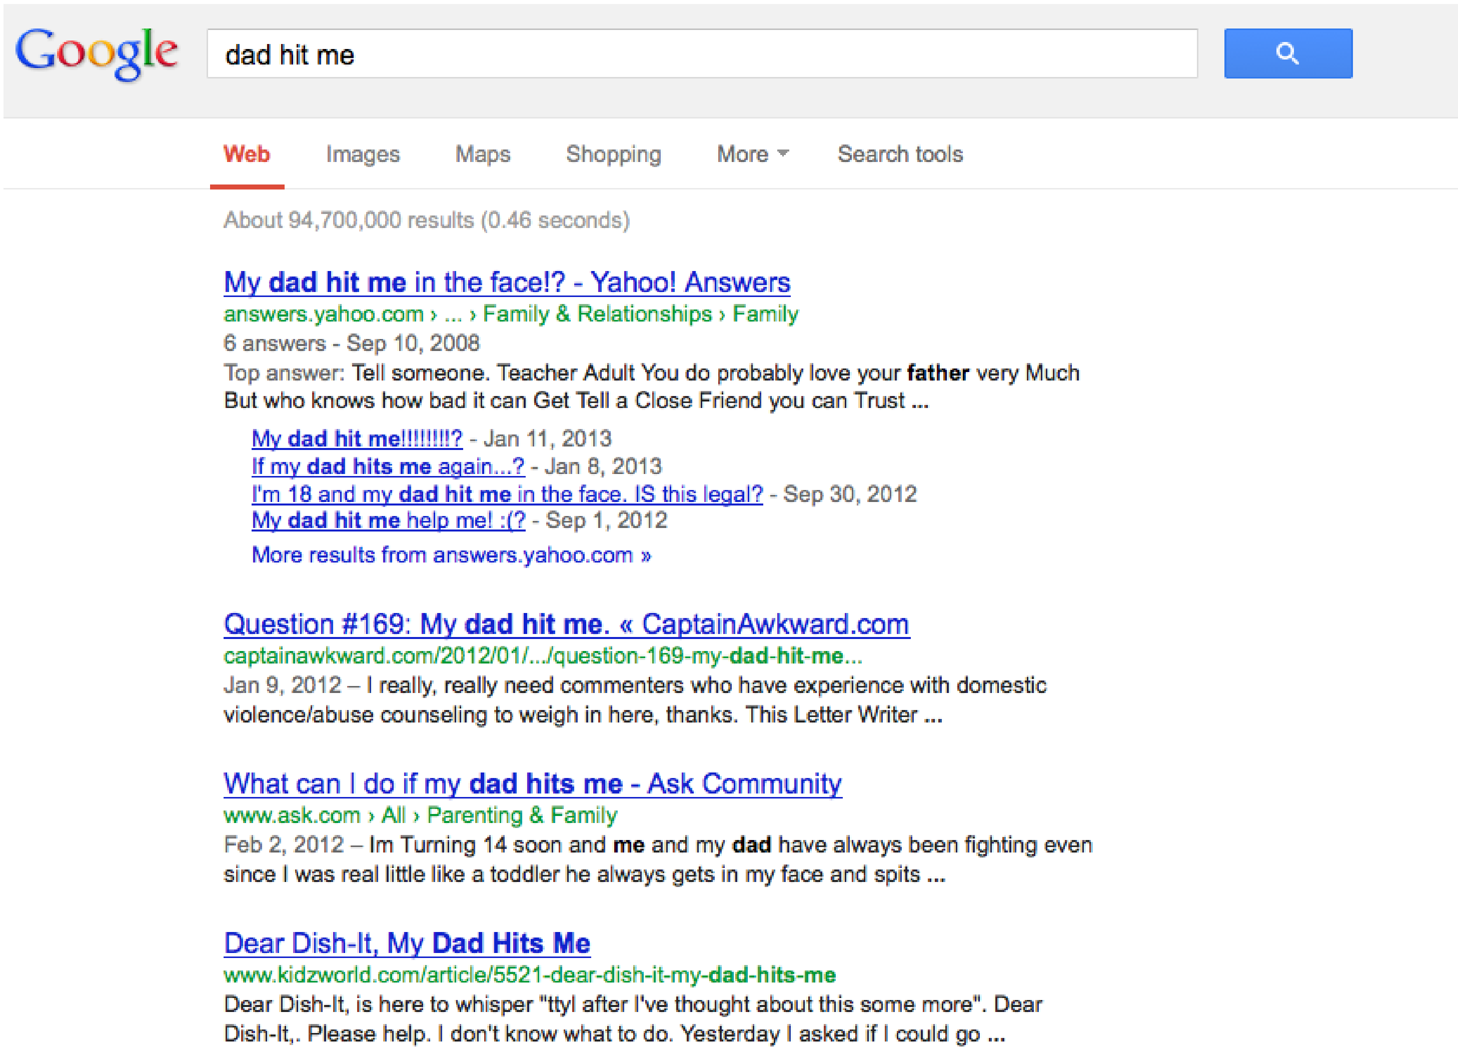

Notes: This shows the returns for a search of "dad hit me" on 1/11/2013.

First, search rates for "child abuse" or "child neglect" can pick up individuals searching for news stories or research on maltreatment, rather than suspecting an incident. An increase in these searches due to more news coverage should not be interpreted as an increase in actual incidents. In the next section, I suggest some controls for this hypothesis. Even so, these tests might not always be conclusive. And researchers must also use intuition in interpreting the evidence.

Second, Google has gained users through time. And the composition of searchers has changed. Long-term trends in search rates may be due to changing composition of searchers. For example, the percent of Google searches that include "science" has consistently dropped through time in the United States. This does not mean, though, that interest in "science" has dropped through time. Rather, early Google users were likely among the demographics most interested in "science." Over time, they have contributed a smaller share of total Google searches. Note that this problem is likely substantially smaller when comparing changes across different areas. Many of the changing demographics are likely to show up, similarly, across different areas. For example, if people uninterested in science were late to use Google, we would expect this pattern would, to some degree, hold in different areas. Thus, all areas would see, on average, decreases in search rates for "science." However, the magnitude of the decrease might depend, in large part, on changes in actual science interest. Note, also, that this problem is likely substantially smaller when comparing high-frequency changes. The composition of searches changes slowly. A big increase in a search on a particular week or day is unlikely due to changing composition of searchers. Having spent many years studying Google data, I have found that comparing the size of changes in different areas and studying high-frequency changes usually lead to meaningful conclusions, whereas studying long-term national trends very often do not. This motivates many of the econometric choices used in this paper; that said, demographics changes driving results are still possible, and I try to test for them as best I can.

Third, the second Google proxy (*Google Maltreatment* (*Victim*), includes very rare searches. Particularly in the early years of Google data, the data are very noisy. Using different combinations of searches can lead to different results. And there is not an obvious constraint on data-mining. Due to this, I limit the use of this data to a high-frequency study of national search data, where each of the different searches tends to follow the same pattern and similar results obtain using different choices of searches. This proxy, though, may have broader use in future studies, relying on more recent data, since the data have become less noisy as the Google user population has grown.

III.C The Effects of the Great Recession on Actual Child Maltreatment Rates

The goal now is to use the alternative proxies for maltreatment to test whether the recession decreased actual maltreatment, or whether the decrease in reported maltreatment is instead due to decreased maltreatment rates.

Assume that both $fatalities_{i,t}$ and $google\ maltreatment_{i,t}$ are noisy proxies of maltreatment.

$$fatalities_{i,t} = \alpha_0 + \alpha_1 maltreatment_{i,t} + \gamma_i + \psi_t + \epsilon_{i,t} \tag{3.7}$$

$$google\ maltreatment_{i,t} = \lambda_0 + \lambda_1 maltreatment_{i,t} + \zeta_i + \eta_t + \mu_{i,t} \tag{3.8}$$

Now I test

$$\Delta fatalities_i = \grave{\beta}_0 + \grave{\beta}_1 \Delta Unemp_i + \grave{\beta}_3 X_i + z_i \tag{3.9}$$

and

$$\Delta google\ maltreatment_i = \check{\beta}_0 + \check{\beta}_1 \Delta Unemp_i + \check{\beta}_3 X_i + q_i \tag{3.10}$$

The results are shown in Figure 3.5 and Table 3.5. A comparatively large exposure to the economic crisis is associated with comparative *increases* in child maltreatment, using both proxies.

There are two explanations for the different results in Sections III.A and III.C. Either $\alpha_1 < 0$ (i.e. the reporting rates of maltreatment declined due to the recession) or $\beta_1 < 0$, $Cov(\Delta Unemp_i, z_i) > 0$, and $Cov(Unemp, q_i) > 0$. In other words, the recession lowered overall maltreatment cases while increasing both child fatality rates and Google searches for "child abuse" or "child neglect."

There is little evidence for the second possibility, though. Overall, there is little reason to suspect that error in the child fatality proxy is positively correlated with error in the Google proxy ($Cov(z_i, q_i) > 0$). The reasons for high fatality rates, controlling for total child maltreatment

Figure 3.5: *Severity of Recession and Change in Actual Maltreatment*

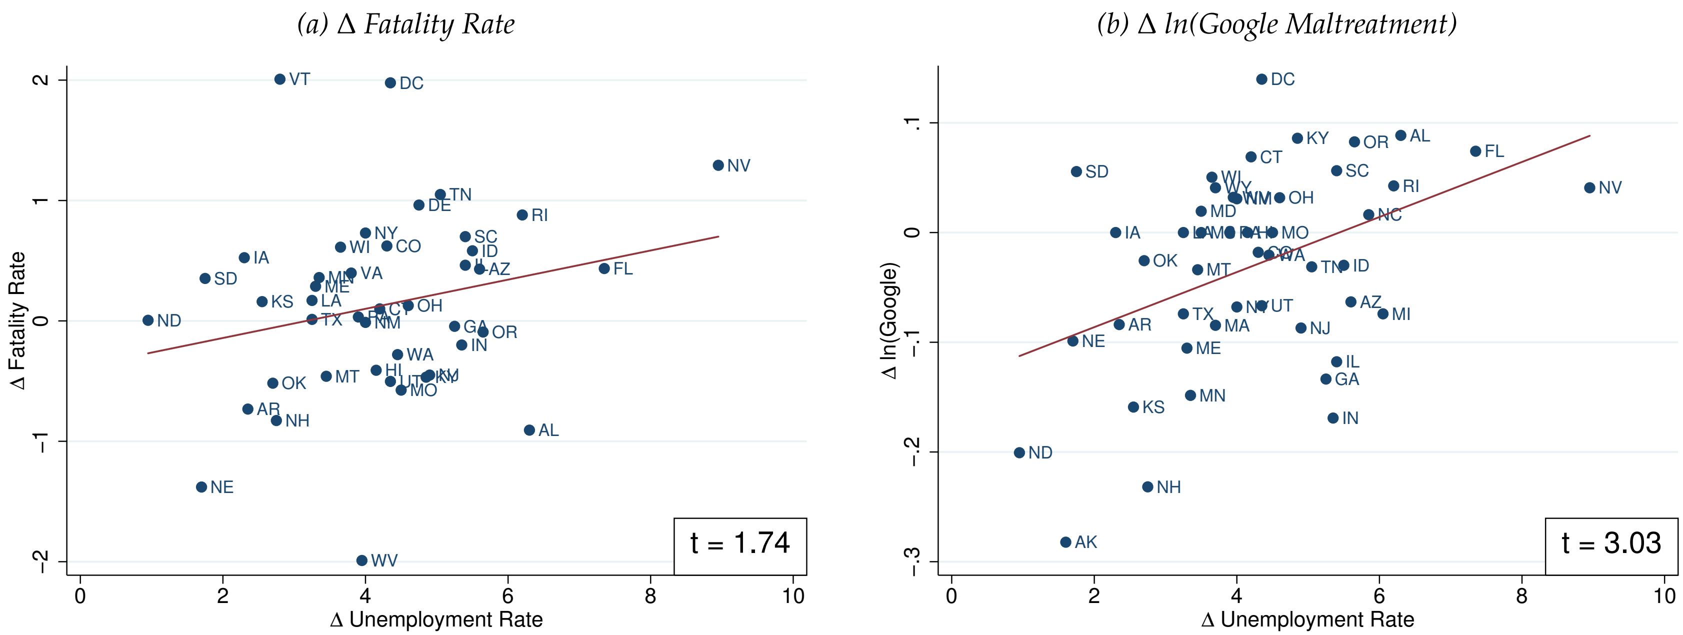

Notes: Δ variables represent the average, for 2010 and 2009, minus the average, for 2006 and 2007. Unemployment Rate is from the Bureau of Labor Statistics. Fatality Rate is from the Children's Bureau Child Maltreatment Annual Reports. Google Maltreatment is as defined in Equation 3.5, from Google Trends.

Table 3.5: *Child Maltreatment and Severity of Recession*

	Δ (Fatality Rate)			Δ ln(Google Maltreatment)		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Unemployment Rate	0.121*	0.110	0.111	0.025***	0.035***	0.022**
	(0.070)	(0.104)	(0.075)	(0.008)	(0.008)	(0.008)
Δ % Hispanic		0.034			0.030	
		(0.399)			(0.027)	
Δ % Black		-0.401***			-0.039***	
		(0.145)			(0.014)	
Δ % Age 0-4		-0.256			0.187**	
		(0.921)			(0.078)	
% Hispanic			0.003			0.000
			(0.008)			(0.001)
% Black			0.010			0.002*
			(0.009)			(0.001)
% College			0.057**			-0.001
			(0.023)			(0.002)
Adjusted R-squared	0.03	0.06	0.20	0.17	0.26	0.16
Observations	43	43	43	47	47	47

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: Robust standard errors are in parentheses. Δ variables represent the average, for 2010 and 2009, minus the average, for 2006 and 2007. Fatality Rate is from the Children's Bureau Child Maltreatment Annual Reports. Unemployment Rate is from the Bureau of Labor Statistics. Demographics variables are from the Current Population Survey.

incidents, and high Google search rates, controlling for total child maltreatment incidents, are likely very different.

In addition, I did not find evidence for alternative explanations for either $Cov(\Delta Unemp_i, z_i) > 0$ and $Cov(Unemp, q_i) > 0$.

One important issue in interpreting the Google data is Google does not report absolute search volumes, only searches normalized as in Equation 3.5. Since the major analysis of interest is how maltreatment is affected by an economic crisis, the proxy could be problematic if the denominator – total Google searches – is affected by economic conditions. It is not obvious which way such bias would work – whether an economic crisis leads to more or less total searches. As evidence against a large change in total Google searches caused by the economic crisis, I find that comparative exposure to the crisis is not correlated with changes in normalized search volume for common words, including "weather," "the," and "a." In results not shown, I divide the Google proxy by normalized search volume for "weather." With this normalization, the proxy is the ratio of child abuse related searches compared to searches that include "weather" instead of compared to total Google searches. The results are little affected by this normalization.

A source of error in the Google proxy is undoubtedly interest in maltreatment, independent of suspected cases. Perhaps in bad economic times individuals are more curious in this topic. If this were the case, then one would expect more media reports on the topic to satiate the demand. However, there is no statistically significant correlation between exposure to recession and change in percentage of stories that include the words "child abuse" or "child neglect." In addition, this would not explain the effects on child mortality.

In sum, I do not find evidence for either $Cov(\Delta Unemp_i, z_i) > 0$ or $Cov(Unemp, q_i) > 0$. And it is unlikely for there to be some missing factor that positively correlates with *both* q_i and z_i. The child maltreatment proxies are very different, with very different sources of error; yet, they yield similar results.

Figure 3.6 shows the relationships between changes in the various child maltreatment proxies over the time periods used. Panel (a) shows a small, though not statistically significant, positive correlation between changes in fatality rates and changes in Google maltreatment searches. The lack of statistical significance is likely explained, in large part, due to the noise in the fatality rate measure. Panel (b) of Figure 3.6 shows no relationship between changes in Google searches suspecting maltreatment and referral rates. This is in strong contrast to the results prior to the recession, shown in Figure 3.3. Prior to the recession, Google search rates suspecting maltreatment were significantly positively related to maltreatment referral rates. The lack of relationship between the changes in these variables, during this time of economic hardship, supports the interpretation that the change in referral rates was largely unrelated to actual changes in maltreatment.

III.D Additional Evidence

Composition of Reported Cases

Section III.A finds that the recent economic downturn led to a significant decrease in reports and investigations for child maltreatment. Section III.C finds that two proxies unlikely to be affected by reporting rates point to the opposite story: an economic downturn increases actual maltreatment.

The suggested reconciliation is that an economic downturn decreases the reporting rates of maltreatment.

Some additional evidence from official data can better test this hypothesis.

We might expect that, if a recession increased the cost of reporting – perhaps by increasing the wait to report – reporting pressures would be most sensitive among individuals least sure about their case.

Section III.A used data on reported cases of maltreatment. However, we can also use data on the percent of reported cases that are substantiated by authorities. While there is certainly error in this rate, we would expect that true cases of maltreatment are more likely to be substantiated. Thus, if the cost of reporting went up in the recession and reporting pressures were most sensitive to individuals least sure about their case, we would expect that the cases that were actually reported in recession-harmed areas were more likely to be substantiated.

There is indeed evidence for this, as shown in Figure 3.7, panel (a), and Table 3.6, Columns (1) through (3). The dependent variable is the percent of investigated cases that are substantiated. The greater the recession, the higher probability of investigated cases being substantiated. This suggests

Figure 3.6: *Changes in Google Maltreatment and Changes in Other Proxies*

(a) Δ Fatality Rate *(b) Δ ln(Referral Rate)*

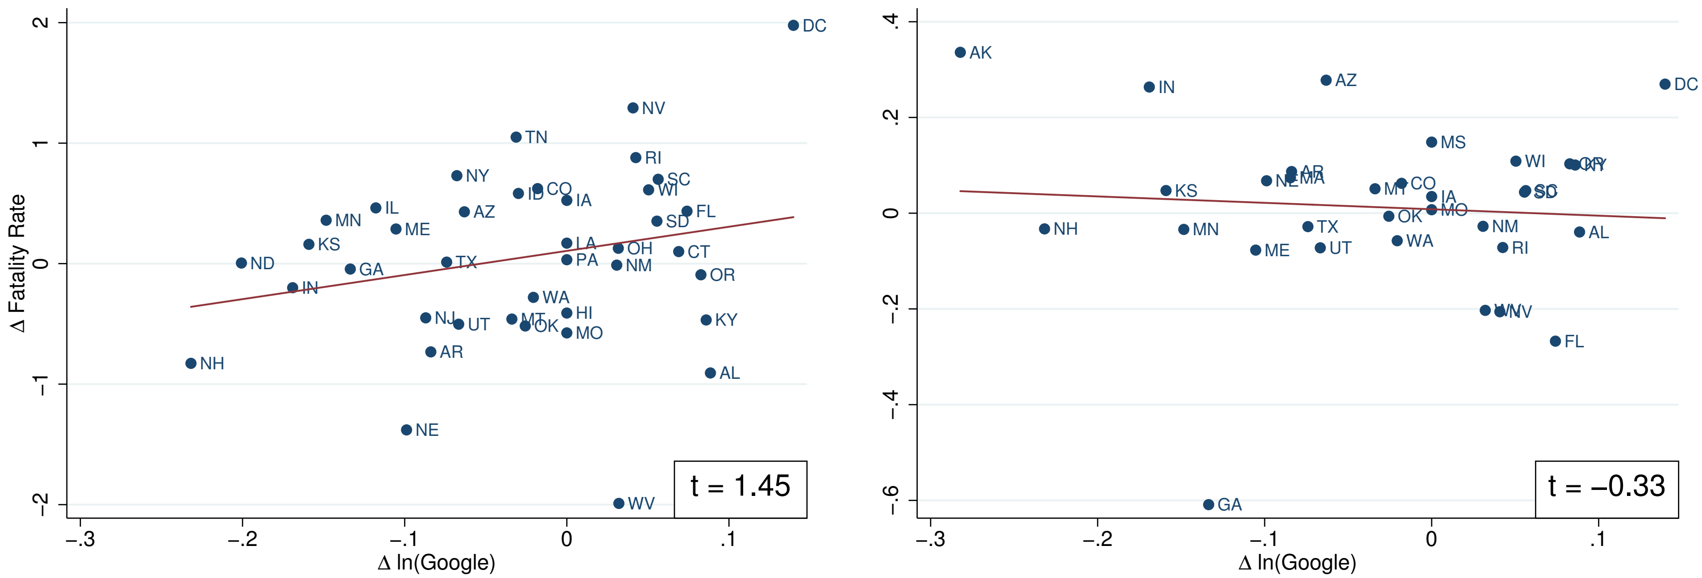

Notes: Robust standard errors are in parentheses. Δ variables represent the average, for 2010 and 2009, minus the average, for 2006 and 2007. Fatality Rate, Referral Rate, and Response Rate are from the Children's Bureau Child Maltreatment Annual Reports. Google Maltreatment is as defined in Equation 3.5, from Google Trends.

that the recession decreased the rates of reporting cases more for cases that were less likely to be substantiated.

Figure 3.7, panel (b), and Table 3.6, Column (4) through (6) show the overall effect of the recession on substantiated cases. The dependent variable is substantiated cases per child. There is not a statistically significant relationship between substantiated cases per child and the economic downturn. In other words, the recession caused a significant decrease in referred and investigated cases of maltreatment. However, since the unreported cases were less likely to be substantiated, it did not cause as large a decrease in substantiated cases of maltreatment.

Even if the substantiation process were perfect, comparing the results on substantiated cases to the results on Google searches and mortality still suggests that the recession decreased the percentage of actual maltreatment cases that were substantiated. The effect of the recession on substantiated cases is slightly negative, while the effects of the recession on child mortality and Google fatalities is always positive.

In sum, the substantiated rate lends further support to the economic downturn decreasing reported rates of maltreatment. The substantiated rate fails to accurately pick up the positive relationship between economic distress and increased maltreatment.

High-Frequency Google and Unemployment Claims Data

Section III.C uses Google data to argue that bad economic conditions during the Great Recession increased child maltreatment. I show that states most affected by the downturn saw the largest increases in searches suspecting child maltreatment.

However, this methodology appears to not fully take advantage of the Google data in understanding the relationship between unemployment and child maltreatment. The Google data can be obtained over a high frequency. They deliver potentially the only meaningful weekly measure of maltreatment incidents.

This allows for an additional test of the effects of unemployment on maltreatment. We can compare the maltreatment proxy to weekly unemployment claims data.

Figure 3.8 shows the second Google maltreatment proxy, *Google Maltreatment* (*Victim*), through

Figure 3.7: *Severity of Recession and Change in Substantiated Cases*

(a) Δ % Substantiated *(b) Δ ln(Substantiated Rate)*

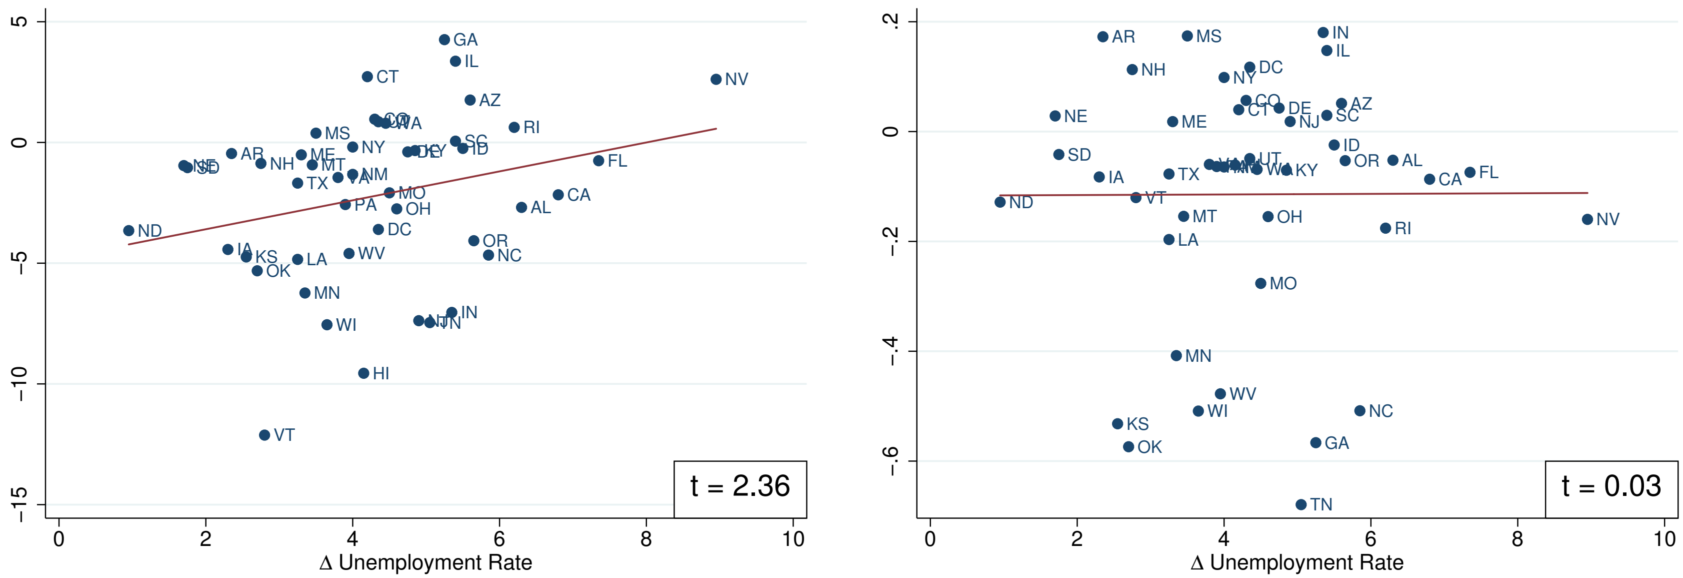

Notes: Δ variables represent the average, for 2010 and 2009, minus the average, for 2006 and 2007. Unemployment Rate is from the Bureau of Labor Statistics. Percent Substantiated and Substantiated Rate are from the Children's Bureau Child Maltreatment Annual Reports.

Table 3.6: *Substantiated Cases and Severity of Recession*

	Δ % Substantiate			Δ ln(Substantiated Rate)		
	(1)	(2)	(3)	(4)	(5)	(6)
Δ Unemployment Rate	0.599**	0.139	0.336	0.001	-0.023	-0.010
	(0.254)	(0.317)	(0.254)	(0.017)	(0.022)	(0.021)
Δ % Hispanic		1.824			0.019	
		(1.362)			(0.085)	
Δ % Black		0.488			-0.084***	
		(0.403)			(0.028)	
Δ % Age 0-4		-3.381			-0.392*	
		(3.963)			(0.207)	
% Hispanic			0.084*			0.004*
			(0.043)			(0.002)
% Black			0.035			0.002
			(0.045)			(0.003)
% College			-0.038			0.008
			(0.100)			(0.006)
Adjusted R-squared	0.05	0.06	0.04	-0.02	0.00	-0.01
Observations	46	46	46	46	46	46

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: Robust standard errors are in parentheses. Δ variables represent the average, for 2010 and 2009, minus the average, for 2006 and 2007. Percent Substantiated and Substantiated Rate are from the Children's Bureau Child Maltreatment Annual Reports. Unemployment Rate is from the Bureau of Labor Statistics. Demographics variables are from the Current Population Survey.

time, for the United States. On the same graph are total weekly unemployment claims in the United States. Both measured are in natural logs and normalized to lie between 0 and 1, for comparison.

Visually inspecting the graph seems to show a relationship. Both unemployment claims and Google searches by maltreatment victims were decreasing in the beginning of this time period. However, they both change direction at roughly the same time.

Table III.E tries to test for a relationship at a high-frequency level. The dependent variable in each regression is the change, compared to the previous week, in $ln(Google\ Maltreatment\ (Victim))$, in the United States. The independent variable is the change, compared to the previous week, in the natural log of total unemployment claims in the United States.

All regressions also include a once-lagged value of both the dependent and independent variable. Column (1) shows a positive correlation. Column (2) and (3) show that the relationship increases upon increasing more lags of the dependent and independent variables.

Columns (4), (5), and (6) show that the relationship stays roughly the same, or is slightly larger, including fixed effects for month of year, week of year, and year. These variables are calculated using the Sunday of a given week.

The evidence of Table III.E does suggest that the more people collecting unemployment claims in a given week, the more victims of child maltreatment.

The coefficients mean that a 100% increase in unemployment claims (doubling) unemployment is associated with between a 114% and 181% increase in child maltreatment searches, which I assume is a proxy for victims.

How do these compare to the results found earlier, studying the shock from the Great Recession? I estimated earlier that a doubling of the unemployment rate during the Great Recession increased maltreatment incidents by 10% to 24% percent during the Great Recession. Thus, the estimates from the high-frequency analysis using the victim-specific Google proxy are substantially higher. One possibility is that the effects of unemployment are more concentrated around changes in unemployment. Children might be particularly vulnerable when individuals have just lost their job and particularly safe when an individual has just obtained a new job. This explanation would also explain why the coefficient seems to rise including more lagged variables in the analysis.

Figure 3.8: *Unemployment Claims and Google Maltreatment (Victim), Weekly in United States*

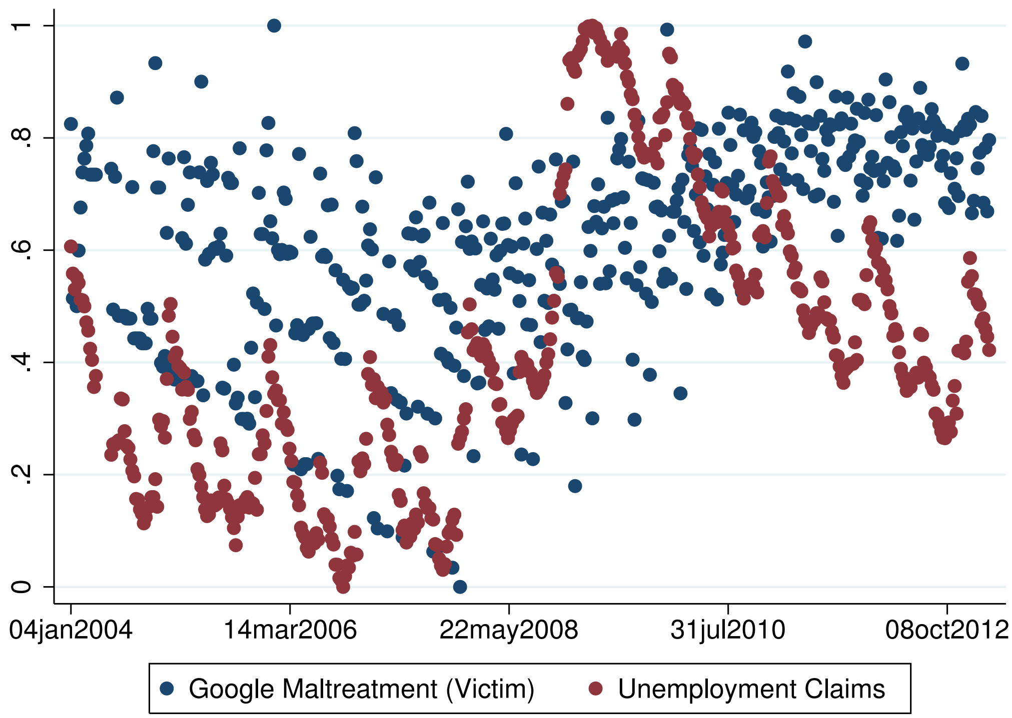

Notes: Data are weekly, for the United States, beginning in 2004 through the week of March 10, 2013. The Google weekly data is from Sunday through Saturday. Unemployment claims data are from Monday through Friday. Google Maltreatment (Victims) is as defined in Equation 3.6. It is the percent of Google searches that include "dad hit(s) me", "dad beat(s) me", "mom hit(s) me", or "mom beat(s) me." Unemployment Claims are total unemployment claims – continuing plus initial – for the United States, downloaded at FRED. Both variables are shown, after taking the natural log, and scaling to lie between 0 and 1.

Table 3.7: *Weekly Google Maltreatment (Victim) and Unemployment Claims*

	Δ ln(Google Maltreatment (Victim))					
	(1)	(2)	(3)	(4)	(5)	(6)
Δ ln(Unemployment Claims)	1.135*	1.280**	1.421**	1.392	1.672*	1.809**
	(0.610)	(0.552)	(0.558)	(1.023)	(0.952)	(0.915)
Adjusted R-squared	0.24	0.32	0.34	0.26	0.33	0.36
Observations	469	465	461	469	465	461
Lags	1	3	5	1	3	5
Year FE	No	No	No	Yes	Yes	Yes
Month FE	No	No	No	Yes	Yes	Yes
Week FE	No	No	No	Yes	Yes	Yes

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: Data are weekly, for the United States, beginning in 2004 through the week of March 10, 2013. The Google weekly data is from Sunday through Saturday. Unemployment claims data are from Monday through Friday. Δ variables represent changes compared to the previous week. Google Maltreatment (Victims) is as defined in Equation 3.6. It is the percent of Google searches that include "dad hit(s) me", "dad beat(s) me", "mom hit(s) me", or "mom beat(s) me." Unemployment Claims are total unemployment claims – continuing plus initial – for the United States, downloaded at FRED. Lags represent number of lagged variables included in the regressions. Lagged values of both the independent and dependent variable are included. Year, month, and week-of-the-year fixed effects are all based on the Sunday of a given week.

III.E The Effects of State and Local Spending on the Reporting Rate of Maltreatment

A reason for the decreased reporting rates of maltreatment may be decreases in resources for government agencies brought about by the recession. Figure 3.9 shows, not surprisingly, that areas most exposed to the recent economic downturn saw relative cuts in government spending. The relationship was particularly strong in education.

Unfortunately, it was difficult to test the effects of government cuts on reporting rates using the same empirical strategy earlier. This is due to the combination of the limited number of observations and high collinearity between exposure to the downturn and budget cuts. Including both change in unemployment rate and change in spending, the coefficients on both depend upon specification choice and do not prove robust.

Table 3.8 offers some suggestive evidence, from before the recession, that both of the channels discussed influence reporting rates. As shown in Table 3.2 and discussed earlier, Roughly 16 percent of total reports of maltreatment come from educators. Columns (1) compares the percent of total maltreatment reports that originate from educators to a state's spending per pupil on education. The more resources devoted to education, the higher the fraction of maltreatment reports that originate from educators. Column (2) shows that the effect is, if anything, slightly larger upon adding controls, including for overall state and local spending per capita. In other words, spending more money on educators leads to more referrals from educators.[3] Thus, we might expect that cutting funds from educators – as was done in areas most affected by the recession – would lead to fewer referrals.

Column (3) compares a state's referral rate to its public spending per capita, controlling for its Google maltreatment proxy. Public spending per capita, here, is a rough proxy for resources to a number of organizations related to children. Controlling for the Google searches – a proxy for actual maltreatment – the more a state spends on public welfare, the more referrals. The implied elasticity is large, with each additional 1 % of spending on public welfare leading to .35% additional referrals. Columns (4) shows that the effect is higher with a broad set of controls, including overall

[3] Of course, this does not test for any general equilibrium effect. It is not known whether the additional referrals from educators were for cases that would have been referred by others, had they not been referred by educators.

Figure 3.9: *Severity of Recession and Budget Cuts*

(a) Δ ln(Government Spending per Capita) *(b) Δ ln(Education Spending per Pupil)*

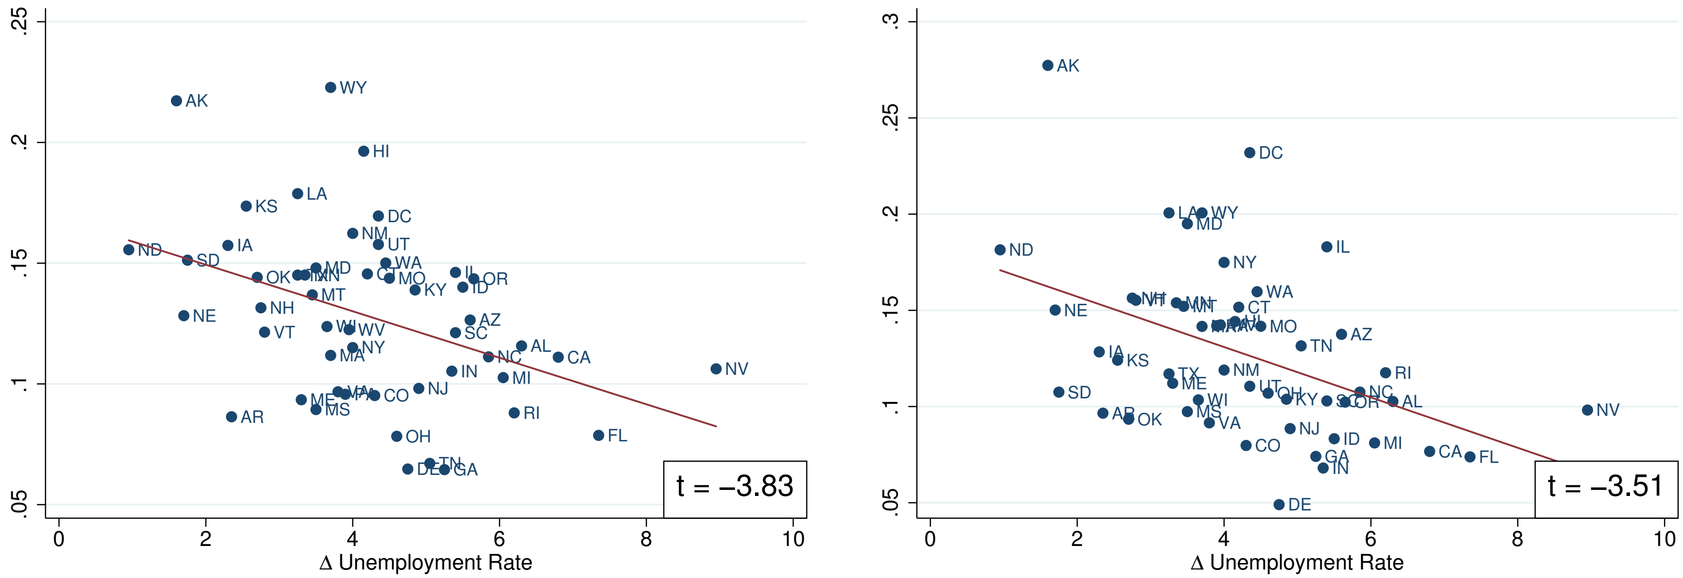

Notes: Changes for all variables are the difference between the 2009-2010 average value and the 2006-2007 average value. Government Spending Per Capita and Education Spending per Pupil are from Census Survey of State and Local Governments.

Table 3.8: *Government Spending and Reports of Maltreatment, Pre-Recession*

	Pct Referrals from Educators		ln(Referral Rate)	
	(1)	(2)	(3)	(4)
ln(Education Spending Per Pupil)	0.061**	0.096***		
	(0.026)	(0.034)		
Expenditure Per Capita		-0.068*		-0.150
		(0.036)		(0.272)
% Hispanic		0.000		-0.005
		(0.000)		(0.005)
% Black		-0.000		-0.015***
		(0.000)		(0.004)
% College		0.001		0.014
		(0.001)		(0.010)
ln(Google Maltreatment)			0.689***	0.641**
			(0.221)	(0.273)
ln(Public Spending Per Capita)			0.370**	0.447**
			(0.142)	(0.194)
Adjusted R-squared	0.10	0.09	0.19	0.39
Observations	47	47	36	36

* $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$

Notes: All variables are averages for 2006 and 2007. Percent referrals from educators is from the National Data Archive on Child Abuse and Neglect. Education spending per capita and expenditure per capita are from U.S. Census Bureau State and Local Government Finances. Child welfare spending is from the Casey Child Welfare Financing Survey.

government spending, suggesting that omitted variables may bias the effect towards zero.

The estimated reductions in reporting rates from the recession are greater than would be predicted by the cross-sectional correlations between spending and reporting rates.

IV Conclusion

This paper tests what happens in a large recession to child maltreatment and child maltreatment reporting. The evidence suggests that actual child maltreatment goes up; reports of maltreatment go sharply down; and overall assistance to maltreated children (substantiated cases) stays about the same.

The paper also suggests that official data on reported maltreatment can lead to misleading conclusions on the effects of the recent recession on actual maltreatment.

Some of the evidence in this paper uses Google searches to proxy child maltreatment. Google data should be considered by scholars studying this very important topic. Using Google searches to proxy domestic violence and crime more generally also seems promising.

When using Google search data to study crime, scholars should test for alternative explanations for changing search trends, such as changing media attention. And Google data can most fruitfully be used when combined with other data sources. If both Google searches and an extreme, always reported outcome show similar trends, this is more convincing evidence than either data point alone can provide.

References

Aizer, Anna, "The Gender Wage Gap and Domestic Violence," *American Economic Review*, 2010, *100* (4), 1847–59.

Alesina, Alberto and Eliana La Ferrara, "Who Trusts Others?," *Journal of Public Economics*, 2002, *85* (2), 207 – 234.

____ **and Howard Rosenthal**, *Partisan Politics, Divided Government, and the Economy*, New York: Cambridge University Press, 1995.

____ **, Ed Glaeser, and Bruce Sacerdote**, "Why Doesn't the US Have a European-Style Welfare System?," *Brookings Papers on Economic Activity*, August 2001, *2001* (2), 187–254.

Allport, Gordon Willard, *The Nature of Prejudice*, New York: Addison-Wesley, 1979.

Altonji, Joseph G., Todd E. Elder, and Christopher R. Taber, "Selection on Observed and Unobserved Variables: Assessing the Effectiveness of Catholic Schools," *Journal of Political Economy*, 2005, *113* (1), 151 – 184.

Ansolabehere, Stephen and Eitan Hersh, "Pants on Fire: Misreporting, Sample Selection, and Participation," *Working Paper*, 2011.

Askitas, Nikolaos and Klaus F. Zimmermann, "Google Econometrics and Unemployment Forecasting," *Applied Economics Quarterly*, 2009, *55* (2), 107 – 120.

Atkeson, Lonna Rae, "'Sure, I Voted for the Winner!' Overreport of the Primary Vote for the Party Nominee in the National Election Studies," *Political Behavior*, 1999, *21* (3), 19.

Baum, Andrew, Raymond Fleming, and Diane M. Reddy, "Unemployment Stress: Loss of Control, Reactance and Learned Helplessness," *Social Science & Medicine*, 1986, *22* (5), 509–516.

Benjamin, Daniel J. and Jesse M. Shapiro, "Thin-Slice Forecasts of Gubernatorial Elections," *The Review of Economics and Statistics*, August 2009, *91* (33), 523–536.

Berggren, Niclas, Henrik Jordahl, and Panu Poutvaara, "The Looks of a Winner: Beauty and Electoral Success," *Journal of Public Economics*, 2010, *94* (1-2), 8–15.

Berinsky, Adam J., "The Two Faces of Public Opinion," *American Journal of Political Science*, 1999, *43* (5), 1209 – 1230.

____ , "Political Context and the Survey Response: The Dynamics of Racial Policy Opinion," *The Journal of Politics*, July 2002, *64* (02).

Berry, William D., Evan J. Ringquist, Richard C. Fording, and Russell L. Hanson, "Measuring Citizen and Government Ideology in the American States, 1960-93," *American Journal of Political Science*, 1998, *42* (1), 327–348.

Bertrand, Marianne and Sendhil Mullainathan, "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination," *American Economic Review*, 2004, *94* (4), 991 – 1013.

Bitler, Marianne P. and Madeline Zavodny, "Child Abuse and Abortion Availability," *American Economic Review*, 2002, *92* (2), 363–367.

____ **and** ____ , "Child Maltreatment, Abortion Availability, and Economic Conditions," *Review of Economics of the Household*, June 2004, *2* (2), 119–141.

Blalock, Hubert, *Toward a Theory of Minority-Group Relations*, New York: Wiley, 1967.

Bound, John, Charles Brown, and Nancy Mathiowetz, "Measurement Error in Survey Data," *Handbook of Econometrics*, 2001, *5*, 3705 – 3843.

Brown, Jocelyn, Patricia Cohen, Jeffrey G. Johnson, and Elizabeth M. Smailes, "Childhood abuse and neglect : Specificity of effects on adolescent and young adult depression and suicidality," *Journal of the American Academy of Child and Adolescent Psychiatry*, 1999, *38* (12), 1490–1496.

Burns, Enid, "U.S. Search Engine Rankings, April 2007," *searchenginewatch.com*, May 2007.

Campbell, James E., "Forecasting the Presidential Vote in the States," *American Journal of Political Science*, 1992, *36* (2), 386–407.

____ , "Editor's Introduction: Forecasting the 2008 National Elections," *PS: Political Science & Politics*, October 2008, *41* (04), 679–682.

____ , "Forecasting the 2012 American National Elections," *PS: Political Science & Politics*, October 2012, *45* (04), 610–613.

Card, David, Alexandre Mas, and Jesse Rothstein, "Tipping and the Dynamics of Segregation," *The Quarterly Journal of Economics*, 2008, *123* (1), 177 – 218.

Cardona, Claire, "Poor Economy Leaves More Children at Risk," December 2011.

Charles, Kerwin Kofi and Jonathan Guryan, "Prejudice and Wages: An Empirical Assessment of Becker's The Economics of Discrimination," *Journal of Political Economy*, 2008, *116* (5), 773–809.

____ **and** ____ , "Studying Discrimination: Fundamental Challenges and Recent Progress," 2011.

Child Maltreatment 2010, U.S. Department of Health and Human Services, Administration for Children and Families, Adminostration on Children, Youth, and Families, Children's Bureau, 2011.

Conti, Gregory and Edward Sobiesk, "An Honest Man Has Nothing to Fear," in "Proceedings of the 3rd Symposium on Usable Privacy and Security - SOUPS '07" ACM Press New York, New York, USA July 2007, p. 112.

CPS, "Computer and Internet Use in the United States: October 2007," 2007.

Currie, Janet and Cathy Spatz Widom, "Long-Term Consequences of Child Abuse and Neglect on Adult Economic Well-Being," *Child Maltreatment*, April 2010, *15* (2), 111–120.

____ **and Erdal Tekin**, "Understanding the Cycle: Childhood Maltreatment and Future Crime," *J. Human Resources*, April 2012, *47* (2), 509–549.

Cutler, David M., Edward L. Glaeser, and Jacob L. Vigdor, "The Rise and Decline of the American Ghetto," *Journal of Political Economy*, 1999, *107* (3), 455 – 506.

DellaVigna, Stefano, "The Obama Effect on Economic Outcomes: Evidence from Event Studies," 2010.

____ **and Ethan Kaplan**, "The Fox News Effect: Media Bias and Voting," *The Quarterly Journal of Economics*, August 2007, *122* (3), 1187–1234.

Donovan, Todd, "Obama and the White Vote," *Political Research Quarterly*, August 2010, *63* (4), 863–874.

Dooley, David, Ralph Catalano, and Georjeanna Wilson, "Depression and unemployment: Panel findings from the Epidemiologic Catchment Area study," *American Journal of Community Psychology*, December 1994, *22* (6), 745–765.

Doyle Jr., Joseph J., "Child Protection and Child Outcomes: Measuring the Effects of Foster Care," *The American Economic Review*, 2007, *97* (5), 28.

Drake, Brett, "Unraveling "Unsubstantiated"," *Child Maltreatment*, August 1996, *1* (3), 261–271.

Duggan, Mark, "More Guns, More Crime," *Journal of Political Economy*, 2001, *109* (5), 1086 – 1114.

Eckholm, Erik, "States Slashing Social Programs for Vulnerable," April 2009.

Enos, Ryan D., "The Persistence of Racial Threat: Evidence from the 2008 Election," *American Political Science Association Annual Meeting*, 2010.

Feldman, Stanley and Leonie Huddy, "Racial Resentment and White Opposition to Race-Conscious Programs: Principles or Prejudice?," *American Journal of Political Science*, 2005, *49* (1), 16.

Finkelhor, David, Richard Ormrod, Heather Turner, and Sherry L. Hamby, "The victimization of children and youth: a comprehensive, national survey.," *Child Maltreatment*, February 2005, *10* (1), 5–25.

Fretland, Katie, "Gallup: Race Not Important to Voters," *The Chicago Tribune's The Swamp*, June 2008.

Gallup, "Presidential Approval Ratings: George W. Bush," 2012.

Gentzkow, Matthew and Jesse Shapiro, *Introduction of Television to the United States Media Market, 1946-1960; Dataset 22720*, Ann Arbor, MI: ICPSR, 2008.

——, ——, and Michael Sinkinson**, "The Effect of Newspaper Entry and Exit on Electoral Politics," *American Economic Review*, 2011, *101* (7), 2980–3018.

Gerber, Alan S., Dean Karlan, and Daniel Bergan, "Does the Media Matter? A Field Experiment Measuring the Effect of Newspapers on Voting Behavior and Political Opinions," *American Economic Journal: Applied Economics*, 2009, *1* (2), 18.

Gilens, Martin, Paul M. Sniderman, and James H. Kuklinski, "Affirmative Action and the Politics of Realignment," *British Journal of Political Science*, January 1998, *28* (01), 159–183.

Ginsberg, Jeremy, Matthew H. Mohebbi, Rajan S. Patel, Lynnette Brammer, Mark S. Smolinski, and Larry Brilliant, "Detecting Influenza Epidemics Using Search Engine Query Data.," *Nature*, February 2009, *457* (7232), 1012–4.

Glaser, James, "Back to the Black Belt: Racial Environment and White Racial Attitudes in the South," *The Journal of Politics*, 1994, *56*, 21 – 41.

—— **and Martin Gilens**, "Interregional Migration and Political Resocialization: A Study of Racial Attitudes Under Pressure," *Public Opinion Quarterly*, 1997, *61* (1), 72–86.

Gomez, Brad T., Thomas G. Hansford, and George A. Krause, "The Republicans Should Pray for Rain: Weather, Turnout, and Voting in U.S. Presidential Elections," *The Journal of Politics*, July 2008, *69* (03).

Gunn, Veronica L., Gerald B. Hickson, and William O. Cooper, "Factors Affecting Pediatricians' Reporting of Suspected Child Maltreatment," *Ambulatory Pediatrics*, 2005, *5* (2), 96–101.

Hart/McInturff, "Study 12062," *NBC News/Wall Street Journal*, January 2012.

Highton, Benjamin, "Prejudice Rivals Partisanship and Ideology When Explaining the 2008 Presidential Vote Across the States," *PS: Political Science & Politics*, July 2011, *44* (03), 530–535.

Holbrook, Thomas, "Incumbency, National Conditions, and the 2008 Presidential Election," *PS: Political Science & Politics*, 2008, *41* (04), 709 – 712.

Hopkins, Heather, "Yahoo! Draws Younger Audience; Google Users Big Spenders Online," *Hitwise Intelligence*, February 2008.

Huddy, Leonie and Stanley Feldman, "On Assessing the Political Effects of Racial Prejudice," *Annual Review of Political Science*, June 2009, *12* (1), 423–447.

Huffman, Matt L. and Philip N. Cohen, "Racial Wage Inequality: Job Segregation and Devaluation Across U.S. Labor Markets," *American Journal of Sociology*, 2004, *109* (4), 902–936.

Hussey, Jon M., Jane Marie Marshall, Diana J. English, Elizabeth Dawes Knight, Anna S. Lau, Howard Dubowitz, and Jonathan B. Kotch, "Defining Maltreatment According to Substantiation: Distinction without a Difference?," *Child Abuse & Neglect*, 2005, *29* (5), 479–492.

——**, Jen Chang, and Jonathan B. Kotch**, "Child Maltreatment in the United States: Prevalence, Risk Factors, and Adolescent Health Consequences," *Pediatrics*, September 2006, *118* (3), 933–42.

Jackman, Simon and Lynn Vavreck, "How Does Obama Match-Up? Counterfactuals & the Role of Obama's Race in 2008," 2011.

Jacobson, Gary C., "The 2008 Presidential and Congressional Elections: Anti-Bush Referendum and Prospects for the Democratic Majority," *Political Science Quarterly*, 2009, *124* (1), 30.

Judis, John B. and Ruy Teixeira, *The Emerging Democratic Majority*, Scribner, 2004.

Kam, Cindy D. and Donald R. Kinder, "Ethnocentrism as a Short-Term Force in the 2008 American Presidential Election," *American Journal of Political Science*, February 2012, *56* (2), 326–340.

Kennedy, Randall, *Nigger: The Strange Career of a Troublesome Word*, New York: Vintage Books, 2003.

Key Jr., Valdimer O., *Southern Politics in State and Nation*, New York: A.A. Knopf, 1949.

Kinder, Donald R. and Allison Dale-Riddle, *The End of Race?*, Yale University Press, 2012.

____ **and Timothy J. Ryan**, "Prejudice and Politics Re-Examined: The Political Significance of Implicit Racial Bias," *APSA 2012 Annual Meeting Paper*, July 2012.

Knowles, John, Nicola Persico, and Petra Todd, "Racial Bias in Motor Vehicle Searches: Theory and Evidence," *Journal of Political Economy*, 2001, *109* (1), 203 – 232.

Kohl, Patricia L, Melissa Jonson-Reid, and Brett Drake, "Time to Leave Substantiation Behind: Rindings from a National Probability Study," *Child Maltreatment*, February 2009, *14* (1), 17–26.

Kreuter, Frauke, Stanley Presser, and Roger Tourangeau, "Social Desirability Bias in CATI, IVR, and Web Surveys: The Effects of Mode and Question Sensitivity," *Public Opinion Quarterly*, January 2009, *72* (5), 847–865.

Kuklinski, James H., Michael D. Cobb, and Martin Gilens, "Racial Attitudes and the 'New South'," *The Journal of Politics*, May 1997, *59* (02), 323–349.

Lansford, Jennifer E, Shari Miller-Johnson, Lisa J Berlin, Kenneth A Dodge, John E Bates, and Gregory S Pettit, "Early Physical Abuse and Later Violent Delinquency: A Prospective Longitudinal Study," *Child Maltreatment*, August 2007, *12* (3), 233–45.

Levitt, Steven D., "Using Repeat Challengers to Estimate the Effect of Campaign Spending on Election Outcomes in the U.S. House," *Journal of Political Economy*, 1994, *102* (4), 777 – 98.

____ , "Relationship Between Crime Reporting and Police: Implications for the Use of Uniform Crime Reports," *Journal of Quantitative Criminology*, 1998, *14* (1), 61–81.

Lewis-Beck, Michael S. and Charles Tien, "The Job of President and the Jobs Model Forecast: Obama for '08?," *PS: Political Science & Politics*, October 2008, *41* (04), 687–690.

____ **and Tom W. Rice**, "Localism in Presidential Elections: The Home State Advantage," *American Journal of Political Science*, 1983, *27* (3), 548–556.

____ **, Charles Tien, and Richard Nadeau**, "Obama's Missed Landslide: A Racial Cost?," *PS: Political Science & Politics*, January 2010, *43* (01), 69–76.

Lindo, Jason M., Jessamyn Schaller, and Benjamin Hansen, "Economic Downturns and Child Abuse," *Working Paper*, 2013.

Linn, Margaret W., Richard Sandifer, and Shayna Stein, "Effects of Unemployment on Mental and Physical Health," *American Journal of Public Health*, May 1985, *75* (5), 502–506.

List, John A., "The Nature and Extent of Discrimination in the Marketplace: Evidence from the Field," *The Quarterly Journal of Economics*, 2004, *119* (1), 49 – 89.

Lockerbie, Brad, "Election Forecasting: The Future of the Presidency and the House," *PS: Political Science & Politics*, 2008, *41* (04), 713 – 716.

____, "Economic Expectations and Election Outcomes: The Presidency and the House in 2012," *PS: Political Science & Politics*, October 2012, *45* (04), 644–647.

Mas, Alexandre and Enrico Moretti, "Racial Bias in the 2008 Presidential Election," *American Economic Review*, 2009, *99* (2), 323 – 29.

Mathews, Carol A, Nirmaljit Kaur, and Murray B Stein, "Childhood Trauma and Obsessive-Compulsive Symptoms.," *Depression and Anxiety*, January 2008, *25* (9), 742–51.

Mixon, Franklin and J. Matthew Tyrone, "The 'Home Grown' Presidency: Empirical Evidence on Localism in Presidential Voting, 1972-2000," *Applied Economics*, 2004, *36* (16), 1745–1749.

Parsons, Christopher A., Johan Sulaeman, Michael C. Yates, and Daniel Hamermesh, "Strike Three: Discrimination, Incentives, and Evaluation," *American Economic Review*, 2011, *101* (4), 1410–35.

Pasek, Josh, Alexander Tahk, Yphtach Lelkes, Jon A. Krosnick, B. Keith Payne, Omair Akhtar, and T. Tompson, "Determinants of Turnout and Candidate Choice in the 2008 U.S. Presidential Election: Illuminating the Impact of Racial Prejudice and Other Considerations," *Public Opinion Quarterly*, January 2010, *73* (5), 943–994.

Paxson, Christina and Jane Waldfogel, "Parental Resources and Child Abuse and Neglect," *American Economic Review*, 1999, *89* (2), 239–244.

Piston, Spencer, "How Explicit Racial Prejudice Hurt Obama in the 2008 Election," *Political Behavior*, 2010, *32* (4), 431–451.

Price, Joseph and Justin Wolfers, "Racial Discrimination Among NBA Referees," *The Quarterly Journal of Economics*, 2010, *125* (4), 1859–1887.

Rahman, Jacquelyn, "The N Word: Its History and Use in the African American Community," *Journal of English Linguistics*, July 2011.

Rogers, Todd and Masa Aida, "Why Bother Asking? The Limited Value of Self-Reported Vote Intention," *HKS Faculty Research Working Paper Series*, 2012.

Saiz, Albert and Uri Simonsohn, "Downloading Wisdom from Online Crowds," October 2008.

Schaffner, Brian F., "Racial Salience and the Obama Vote," 2011.

Scheitle, Christopher P., "Google's Insights for Search: A Note Evaluating the Use of Search Engine Data in Social Research," *Social Science Quarterly*, 2011, *92* (1), 285 – 295.

Schuman, Howard, "The Perils of Correlation, The Lure of Labels, and the Beauty of Negative Results," in David O. Sears, Jim Sidanius, and Lawrence Bobo, eds., *Racialized Politics: The Debate About Racism in America*, University of Chicago Press, 2000, pp. 302–316.

Sedlak, Andrea J., Jane Mettenburg, Monica Basena, Ian Peta, Karla McPherson, Angela Greene, and Spencer Li, "Fourth National Incidence Study of Child Abuse and Neglect (NIS-4), Technical Report," *US Department of Health and Human Services*, 2010.

Seifter, Ari, Alison Schwarzwalder, Kate Geis, and John Aucott, "The Utility of 'Google Trends' for Epidemiological Research: Lyme Disease as an Example," *Geospatial Health*, May 2010, *4* (2), 135–137.

Seiglie, Carlos, "Understanding Child Outcomes: An Application to Child Abuse and Neglect," *Review of Economics of the Household*, 2004, *2* (2), 143–160.

Stephens-Davidowitz, Seth, "The Cost of Racial Animus on a Black Presidential Candidate: Using Google Search Data to Find What Surveys Miss," *Working Paper*, 2012.

Stevenson, Betsey and Justin Wolfers, "Trust in Public Institutions over the Business Cycle," *American Economic Review*, 2011, *101* (3), 281–87.

Stith, Sandra M., Ting Liu, Christopher Davies, Esther L. Boykin, Meagen C. Alder, Jennifer M. Harris, Anurag Som, Mary McPherson, and J.E.M.E.G. Dees, "Risk Factors in Child Maltreatment: A Meta-Analytic Review of the Literature," *Aggression and Violent Behavior*, 2009, *14* (1), 13–19.

Taylor, Marylee C., "How White Attitudes Vary with the Racial Composition of Local Populations: Numbers Count," *American Sociological Review*, 1998, *63* (4), 512–535.

Tesler, Michael, "The Spillover of Racialization into Health Care: How President Obama Polarized Public Opinion by Racial Attitudes and Race," *American Journal of Political Science*, 2012, *56* (3), 15.

____ **and David O. Sears**, *Obama's Race: The 2008 Election and the Dream of a Post-Racial America*, Vol. 2010, Chicago: University of Chicago Press, 2010.

____ **and** ____ , "President Obama and the Growing Polarization of Partisan Attachments by Racial Attitudes and Race," *APSA 2010 Annual Meeting Paper*, July 2010.

Tourangeau, Roger and Yan Ting, "Sensitive Questions in Surveys," *Psychological Bulletin*, 2007, *133* (5), 859–883.

Valdez, Linda, "Stopping Child Abuse: It Begins With You," June 2012.

Valentino, Nicholas A. and Ted Brader, "The Sword's Other Edge: Perceptions of Discrimination and Racial Policy Opinion after Obama," *Public Opinion Quarterly*, May 2011, *75* (2), 201–226.

Varian, Hal R. and Hyunyoung Choi, "Predicting the Present with Google Trends," *SSRN Electronic Journal*, August 2010.

Vavreck, Lynn, "The Exaggerated Effects of Advertising on Turnout: The Dangers of Self-Reports," *Quarterly Journal of Political Science*, 2007, *2* (4), 325–343.

Washington, Ebonya, "How Black Candidates Affect Voter Turnout," *Quarterly Journal of Economics*, August 2006, *121* (3), 973–998.

Widom, Cathy Spatz, "Child Abuse, Neglect, and Violent Criminal Behavior," *Criminology*, 1989, *27*.

Wolfers, Justin, "Are Voters Rational? Evidence from Gubernatorial Elections," Technical Report March 2002.

Wright, Gerald C., "Errors in Measuring Vote Choice in the National Election Studies, 1952-88," *American Journal of Political Science*, 1993, *37* (1), 291–316.

Appendix

Algorithm to Determine Search Volume at Media Market Level

Google Trends does not report data if the absolute level of search is below an unreported threshold. The threshold is clearly high, such that only the most common searches are available at the media market level. And search volume for "nigger(s)" is only available for a small number of media markets. To get around this, I use the following strategy: Get search volume for the word "weather." (This gets above the threshold for roughly 200 of 210 media markets, since "weather" is a fairly common term used on Google.) Get search volume for "weather+nigger(s)," searches that include *either* "weather" or "nigger(s)." (This, by definition, gets over the threshold for the same 200 media markets, since it captures a larger number of searches. Subtracting search volume from "weather" from search volume from "weather+nigger(s)" will give approximate search volume for "nigger(s)" for the 200 media markets. Complications arise from rounding, normalizing, and sampling.

Here is the algorithm:

Take a sample s in Google.
Let X be a set of possible searches. Denote $X_{j,s}$ as the value that Google Trends gives.
This is $X_{j,s} = x_{j,s}/x_{max,s}$ where $x_{j,s}$ is the fraction of Google searches in area j in sample s that are in X. (See Equation 1.1).
Take two words N and W. And let $C = N \cup W$ and $B = N \cap W$. Then $n_{j,s} = c_{j,s} - w_{j,s} + b_{j,s}$. Denoting x_j as the expected value of x in area j, then $n_j = c_j - w_j + b_j$ Assume we have an area for which, for $x \in \{c, w, n, b\}$, $x_{j,s}$ is independent of $x_{max,s}$. Then $X_j = x_j/x_{max}$. Then

$$N_j = \frac{c_{max}}{n_{max}} C_j - \frac{w_{max}}{n_{max}} W_j + \frac{b_{max}}{n_{max}} B_j \tag{3.11}$$

Assume B_j is negligible, a reasonable assumption for words used in this paper. The issue is that N_j, the word of interest, is only reported for about 30 media markets, whereas C_j and W_j are reported for about 200 media markets. Since N_j depends linearly on W_j and C_j I can find $\frac{c_{max}}{n_{max}}$ and $\frac{w_{max}}{n_{max}}$ using data for any media market that reports all 3 values. I can then use these numbers to find N_j for all 200 that report W_j and C_j. If C_j, W_j, and N_j were reported with no error for media markets, I could find exact numbers. Even with 5,000 downloads, I do not get perfect estimates of C_j, W_j, and N_j. I thus back out the coefficients by regressing the averages for 30 media markets that have all data available. The R^2 on this regression is 0.86, meaning there is minor remaining error. After 5,000 downloads, regressing halves of the samples suggest this strategy has captured about 80 percent of the variation in the actual number. To deal with the minor remaining error, I use the first half sample estimate as an instrument for the second half sample when racially charged search is an independent variable in regressions.

Algorithm in Practice:
1. Download 5,000 samples for "weather," from 2004-2007.
2. Download 5,000 samples for "nigger+niggers," from 2004-2007. (A " + " signifies an "or.")
3. Download 5,000 samples for "nigger+niggers+weather," from 2004-2007.
4. Eliminate any media market that ever scores 0 or 100 for "weather." (A 0 means absolute search volume is too small. A 100 means it scores the maximum.)
(12 of the smallest media markets in the country are eliminated, 10 that never show up and 2 that compete for the top "weather" search spot.)
5. Calculate a media market's average score for "weather," "nigger+niggers," and "nigger+niggers+weather."
6. Regress "nigger+niggers" average score on "weather" average score and "weather+nigger+niggers" average score for the markets that never score a 0 or 100 on "nigger+niggers."
7. Use coefficients from regression to back out "nigger+niggers" for remaining markets, using their average search volume for "weather" and "nigger+niggers+weather."