

Measuring Implicit Racial Priming with Automated Text Analysis

Michael Davidson and Jason Wu*

University of California, San Diego

April 1, 2015

Abstract

Existing theories of racial priming argue that political campaigns use visual appeals, racially coded language, stereotypes, and references to racialized public policies in the presence of favorable audience characteristics to tap into racist sentiments among the broader public. Although this work has sought to measure the potential for racist appeals empirically, it requires the analyst to specify the expected patterns of implicit racial priming *ex ante*, even though identifications of implicit racial primes are often contested. In this paper we outline the potential that text analysis holds in improving measurements of polarized language, which can serve as a more systematic measure of implicit priming. We test this approach on a corpus of campaign ads from congressional, gubernatorial, and presidential elections.

*We thank Marisa Abrajano, Zoli Hajnal and Molly Roberts for comments on this project. Errors are our own.

1 Introduction

Although the Willie Horton ad is now seen as a key example of a campaign ad that implicitly appeals to racial animosity, one of Tali Mendelberg's contributions in *The Race Card* is to point out that this consensus did not begin to form until Jesse Jackson charged the Bush campaign with racist campaigning late in the 1988 campaign (Mendelberg 2001). Since explicit appeals to racial prejudice violate American norms of racial equality, efforts to prime racial prejudice often attempt to do so indirectly, and are not understood as racial appeals until they are identified as such by elites or the media.

Implicit racial priming uses images that draw upon the crime narrative scripts that are common on TV news, while explicit racial appeals use racialized language (Mendelberg 2001). However, not every visual treatment of a racial minority is, according to Mendelberg, an implicit racial appeal, and many ambiguous codewords or subjects exist which often serve as the basis of implicit appeals. In a society that has developed norms of racial equality, explicit racial appeals have become rare in campaign advertising, while implicit appeals predominate because they allow their targets to minimize their own sense of cognitive dissonance.

Despite their importance, however, implicit appeals remain difficult to systematically identify because they rely on their ambiguity for their effectiveness (Hutchings & Jardina 2009). The existing literature on racial priming in electoral campaigns has either used case studies to identify examples of racial priming (Mendelberg 2001, Stephens 2013), or experimental treatments constructed by the researcher, which mimic TV newscasts, campaign ads, or get-out-the-vote appeals, to study the ef-

fects of these appeals (Mendelberg 2001, Valentino, Hutchings & White 2002, Huber & Lapinski 2006). These studies have started a debate about whether implicit primes are more effective (Mendelberg 2001, Mendelberg 2008), or whether explicit primes are equally effective (Huber & Lapinski 2006, Huber & Lapinski 2008). They have also begun to shed light on the heterogenous effects of explicit appeals on women (Hutchings, Walton & Benjamin 2010) and the less-educated (Huber & Lapinski 2006).

This paper focuses on three separate questions that have been difficult to answer systematically because of the dearth of good measures of implicit racial priming. First, who makes implicit racial appeals? Second, what topics are used to make implicit racial appeals? And finally, how are these appeals made? More specifically, how does the language in campaign ads used to make an implicit racial appeal diverge from the language used in other ads on the same topic? We use new unsupervised topic modeling methods in conjunction with data on the racial imagery of campaign ads to construct a new measure of implicit racial priming, and to address these questions.

2 Theory

Implicit appeals to racial prejudice are important to understand because they comprise one of the more troubling examples of the potential for priming effects to shape public opinion. Since most citizens hold weak or low-quality opinions about political issues (Converse 1964, Zaller 1992), efforts by political elites or the media to use

primes to structure public debates can be very influential (Iyengar & Kinder 2010). Priming effects can on occasion prove to be difficult to identify, but systematically isolating instances of implicit racial priming is a more challenging task for empirical analysts because the valence of implicit primes is, to a greater extent, different for each observer, depending on their prior beliefs (McIlwain & Caliendo 2011).

This implies that individual cases of implicit primes assume an "I know it when I see it" character that make systematic study difficult. The subjectivity of this construct limits the scope of the content analysis that Mendelberg uses in her empirical analysis; she addresses this problem by focusing on a case analysis of the Willie Horton ad. Her approach may lack generalizability or miss cases of implicit political appeals that fail to attract an explicit denunciation from an opposing party. Other attempts to measure racist appeals to prejudice held by the majority towards minorities, and to distinguish them from racial appeals which do not appeal to anti-minority sentiments, such as McIlwain and Caliendo's Index of Racist Potential (IRP), select the relevant terms of their content analysis *ex ante* and as such are also relatively blunt instruments (McIlwain & Caliendo 2011).

Hypotheses

One of the expectations of existing theory is that racial appeals are more likely to be used in districts with a large proportion of white voters. In McIlwain and Caliendo's IRP, the demographics of the district are by far the best predictors of the presence or absence of potentially racist appeals (McIlwain & Caliendo 2011). This leads us to our first hypothesis, which governs our expectations about where implicit

racial primes are used.

Hypothesis 1: Campaigns in districts with more African-American, Hispanic, and Asian Voters will produce fewer implicit racial primes.

A second set of theoretical expectations concern the topics where we expect to find implicit racial primes. Of particular interest are the discussions in the advertisements about public policies that are commonly understood to have a racial dimension, such as affirmative action, crime, drugs, and welfare. The Willie Horton ad is just one of many instances of racially inflected discussion about crime, which continue to represent one means of tapping into racial sentiments (Peffley, Hurwitz & Sniderman 1997, Peffley & Hurwitz 2002). Discussions of welfare (Gilens 1999) and immigration (Valentino, Hutchings & White 2002, Brader, Valentino & Suhay 2008) are also often seen through the prism of race. In our paper we will look to see whether these topics are discussed more frequently when a minority appears in the imagery of the ad, and we also explore the extent to which the vocabulary of these ads is affected by the presence of a visible minority in the ad. We use this test because the presence of minorities in the imagery of an ad is not, by itself, sufficient to generate an implicit racial appeal (McIlwain & Caliendo 2011). What constitutes an implicit racial prime is instead the combination of implicit racial imagery and an ambiguous but suggestive narrative (Valentino, Hutchings & White 2002).

Hypothesis 2: Advertisements that focus on law and order issues, such as crime and immigration, or distributional issues, such as welfare, will contain more implicit racial primes.

On the other hand, McIlwain and Caliendo find that issues by themselves are not

strongly associated with the IRP. They find that out of 29 issues, only welfare has a modest positive correlation with their index of racist potential, and that even the association with welfare is rather weak.

Hypothesis 2A: The topic of the advertisement is not associated with the presence or absence of implicit racial primes.

Our last set of hypotheses concern the language used to make implicit racial primes. Experimental treatments have found that codewords, such as "inner city" have generated significant priming effects (Hurwitz & Peffley 2005), but studies of actual advertisements have found that recognizable codewords rarely make it into political ads (McIlwain & Caliendo 2011).

Hypothesis 3: Implicit Racial Primes use recognizable codewords, such as "inner city"

Hypothesis 3A: Implicit Racial Primes do not use recognizable codewords

3 Empirical Strategy

In our paper, we use an unsupervised learning algorithm to discover patterns of implicit racial campaigns. As Mendelberg predicts, explicit references to race are rare in this corpus, which implies that a dictionary method will be limited in its utility because of the limited amount of variation in the dataset. A supervised method is also problematic because it relies on the analyst to code a training set (Hopkins & King 2010); as a result, this coding process runs into the same problems of identifying implicit campaigns that prevent existing studies from generating systematic measures

of racial priming. This suggests that an unsupervised topic model, which allows the algorithm to detect topics within our corpus, will be the most useful because it will allow the analyst to learn from the patterns in the data in an unbiased and systematic fashion. Interpretation by the analyst of course will remain crucial because results from this approach will be sensitive to decisions about cleaning the data and the number of topics that the analyst stipulates.

Data

To test our hypotheses, we collected the text for all of the unique ads aired in the 2000 and 2008 elections in races for Governor, the House, the Senate and the Presidency and combined these texts with a variety of covariates that describe the context behind the ad.

We extracted the texts of these advertisements by scraping the storyboards of the campaign aids provided by the Campaign Media Analysis Group (CMAG) to the Wisconsin Advertising Project ¹. CMAG collects data on the context behind each airing of an individual ad on the four major national networks, CBS, NBC, ABC, and Fox, as well as the 25 national cable networks. We used the Text Extraction Toolkit (TET)² as well as optical character recognition to extract the relevant texts.

3

The Wisconsin Advertising Project dataset uses the storyboards provided by

¹<http://wiscadproject.wisc.edu/>

²See <http://www.pdfli.com/products/tet> for more information.

³For the 2008 corpus, no additional OCR was necessary because of the format of the data, but we needed to conduct OCR to extract texts from 2000. This may introduce more errors into our measurements from 2000, which should bias against our ability to find a result.

CMAG to code a host of relevant covariates (Ridout, Franz & Goldstein 2008), which we integrate into the analysis in the next section. Of special interest to this project, the data includes coding for what type of entity purchased each ad (i.e. candidate, party, interest group, coordinated), what program hosted the ad, whether the ad mentioned the candidate and the opponent, and an estimated cost of the ad. Additionally, each advertisement is coded for tone— that is, for whether it is an attack, contrast, or promote ad. Other information includes a summary of the content of the ad- i.e. whether the ad focuses on personal characteristics or policy issues. A summary of some of the key covariates is included in Table 1.


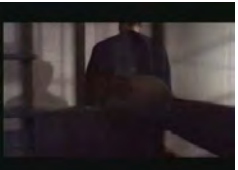



We supplemented the covariates collected by the Wisconsin Advertising Project by asking American users of Mechanical Turk to code the race of the people who appear in the imagery of each advertisement. The presence or absence of a minority in the imagery of each advertisement is one of the key covariates we include in our topic model.

Variable	Levels	n	%	\sum %
Office	Governor	2009	27.3	27.3
	House	2285	31.0	58.3
	President	979	13.3	71.5
	Senate	2097	28.4	100.0
	all	7370	100.0	
Year	2000	2709	27.6	27.6
	2002	4159	42.3	69.9
	2008	2954	30.1	100.0
	all	9822	100.0	
Sponsor: Candidate	0	3032	30.9	30.9
	1	6790	69.1	100.0
	all	9822	100.0	
Sponsor: Party	0	8126	82.7	82.7
	1	1696	17.3	100.0
	all	9822	100.0	
Sponsor: Interest Group	0	9046	92.1	92.1
	1	775	7.9	100.0
	all	9822	100.0	

Table 1: Summary Statistics. Note that each unique advertisement is only counted once, regardless of how many times and how many markets it was aired in.

GOV/NC SMITH ILLEGAL IMMIGRATION

Brand: POL-GOVERNOR (B333)
Parent: POLITICAL ADV
Aired: 03/04/2008 - 03/04/2008
Creative Id: 6174964

 <p>[Announcer]: Republican candidate for governor Fred Smith understands that no one should be rewarded</p>	 <p>for coming to this country illegally. That's why he supports giving law</p>	 <p>enforcement the tools they need to do their job. [Smith]: "I would try to make sure that</p>
 <p>every sheriff's department has the resources to participate</p>	 <p>in the 287G program so that they can catch,</p>	 <p>contain and deport illegal immigrants who are violating our laws."</p>
 <p>[Announcer]: Fred Smith, working to uphold the rule of law. Fred Smith for governor, 2008.</p>	 <p>[PFB]: FRED SMITH, GOVERNOR 2008</p>	

AdDetector 
www.PoliticsOnTV.com

Copyright 2008 TNS Media
 Intelligence/CMAG

 media intelligence / cmag
 1-866-559-CMAG

Figure 1: Example Storyboard from 2008 North Carolina Governor's Race

Model

To analyze the data we use the Structural Topic Model, which is based on the Latent Dirichlet Allocation. LDA is a computer-assisted text analysis algorithm that infers latent topics from a corpus of words (Blei, Ng & Jordan 2003). This model is unsupervised, allowing for the automated discovery of such themes without relying on many a priori assumptions. Topics in LDA are represented as distributions over a vocabulary of words that have an interpretable theme. The analyst first chooses the number of topics for the model. Then the model creates a distribution over the topics, chooses a topic assignment for each word, and then chooses a word from the distribution over the vocabulary of the corpus (Blei 2012). After examining the cohesiveness of the topics that are created in this fashion, the analyst can then refine the model by changing the number of topics selected to achieve greater semantic cohesiveness.

Formally, in LDA, the documents are modeled as random mixtures over latent topics (meaning that documents are classified as belonging to several topics), where each topic is characterized by a distribution over words (Blei, Ng & Jordan 2003). The definition of LDA that follows comes from the original (Blei, Ng & Jordan 2003) paper.

Let a document be defined as a sequence of N words, $\mathbf{w} = (w_1, w_2, \dots, w_N)$, where w_n is the n th word in the sequence. We also define a corpus as a collection of M documents denoted by $D = (\mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_M)$. The generative process for a document w contains the following steps:

1. Choose the number of words in each document $N \sim Poisson(\xi)$

2. Choose a distribution over topics $\theta \sim \text{Dirichlet}(\alpha)$
3. For each of the N words w_n :
 - a) Choose a topic $z_n \sim \text{Multinomial}(\theta)$
 - b) Choose a word w_n from $p(w_n|z_n, \beta)$, a multinomial probability conditioned on the topic z_n .

The Structural Topic Model (STM) is based on LDA. Like LDA, this model is unsupervised, and relies upon the analyst to stipulate a total number of topics in the documents. The attraction of using the STM is that it has the added advantage of allowing the analyst to incorporate covariates, such as time, region, or author to organize the distributions in our topics (Roberts, Stewart, Tingley, Lucas, Leder Luis, Gadarian, Albertson & Rand 2014). The result is better topics and an ability to estimate the effect of these covariates on a topic's prevalence. The advantage of using the Wisconsin Ads Corpus is that many covariates of interest have already been coded, including the party of the candidate and the source of the ad (i.e. an interest group, a party, a candidate, or a coordinated effort).

In our core results, which are presented in the next section, we use the presence or absence of a minority in the imagery of the ad as our content covariate. First, we will explore how this covariate affects the prevalence of a topic. In other words, is a topic discussed more or less when featuring a minority? This seeks to answer the second of our three questions, which looks at what topics are used to make racial appeals. Importantly, we include variables like the number of people in an ad, the

state of the ad sponsor and the candidate’s incumbency status as controls.⁴

Next we explore how the language used in an ad varies within a given topic when minorities are featured. The presence or absence of a minority becomes the axis along which we examine variation in language use.

4 Results

Our model results generate topics that show face validity. In figure 2, a fifty topic STM shows a collection of topics which possess clear themes (e.g. topic 13: taxes, hikes, cut). Our model results, presented in the next three subsections, also allow us to isolate the effect of our key independent variable, the presence of a minority in the imagery of the advertisement, on the specific language used within a topic.

4.1 Who uses Racial Priming?

First, we explore who uses racial priming. What are the characteristics of the race, the constituency and the candidate that lead to the use of racial priming? To answer these questions, we explore how characteristics correspond with the appearance of minorities in ads, controlling for the number of people shown. Rather than display all 50 topics, we focus here on the topics featured in figure 3.

We use the topics produced by the STM as controls in a regression of minority appearance on candidate, district and race characteristics. The data suggest that

⁴These are referred to as prevalence covariates in the STM model. Our prevalence covariates include a host of other variables, namely dummy variables for the identity of the sponsor of the ad (candidate, party, interest group, coordinated), dummies for the type of race (house, senate, governor, president), as well as dummies for the party of the sponsoring entity (republican, democrat).

Top Topics




Figure 2: Estimated Topic Proportions for 50 topic STM with 2000 and 2008 House, Senate, Governor, and Presidential Advertising Data.

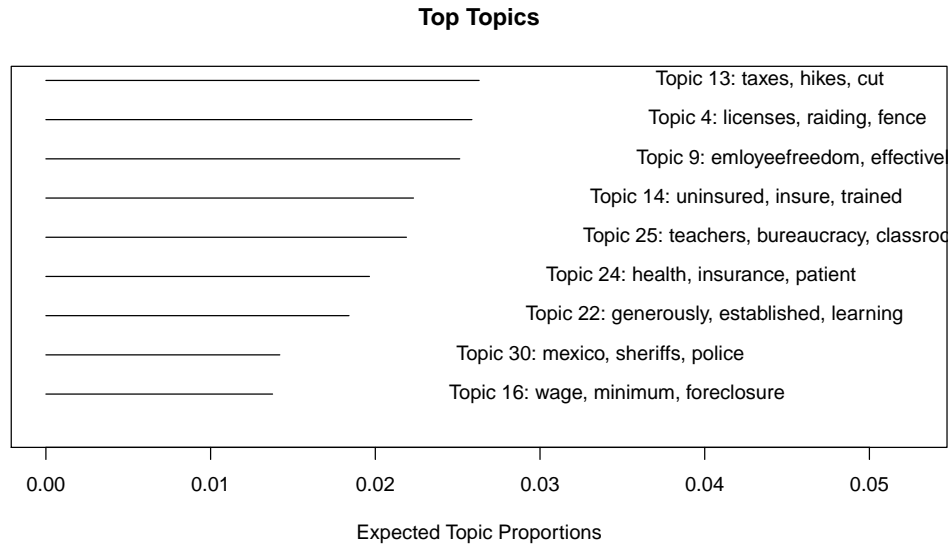


Figure 3: Estimated Topic Proportions for topics of emphasis.

all three play an important role in predicting the use of a minority ads. Republican candidates and incumbents who air ads are less likely than their democratic and challenger counter-parts to include minorities. Of the 4 different offices included in this data set, Gubernatorial races are the most likely include minorities. The closeness of the race has no discernable correlation with minorities being shown in the ad.

Constituency demographics also play an important role in the choice of whether or not to include minorities in the ads. Larger shares of Blacks, Hispanics and Asian all correlate with more minority-inclusive ads. This allows us to conditionally reject Hypothesis 1, which expects districts with large minority populations to see fewer implicit racial primes. Instead we find that race is more likely to be visually primed

in districts with large minority populations. On its own, this result is not a direct test of McIlwain and Caliendo’s finding that districts with fewer minorities see more racist appeals because we are unable to distinguish racist from racial appeals without further information, but it does cast doubt on the notion that racial appeals are less common in minority-heavy areas.

It is important to stress that this model includes 50 flexible controls (not displayed) for the topic of the advertisements. The topics coded by the Wisconsin Advertising Project are much more rigid and limited, and consequently, likely to introduce bias.

4.2 What Topics Feature Racial Priming?

Figure 4 illustrates the conditional correlation between minorities appearing in ads and topic prevalence. In other words, an if an ad features a minority it is more or less likely to contain language indicative of a given topic.

Interestingly, the results lends little support for Hypothesis 2, which expected implicit racial priming to coincide with topics such as immigration, or distributional issues like healthcare. Ads on taxes and unions are actually less likely to feature minorities. However, we can also reject Hypothesis 2A, which expected topics to have little relationship with the presence of implicit visual appeals. Topics related to social welfare show a positive correlation with minority appearance, as illustrated by the two topics on education and child health insurance. What we may be picking up here is a positive valence to the implicit primes in these ads, which seeks to signal inclusiveness.

Table 2: Logit of inclusion of minorities in ads.

	<i>Dependent variable:</i>	
	Minority	
Republican		-0.693*** (0.173)
Margin of Victory		-0.002 (0.009)
Winner		0.233 (0.182)
MV * Winner		-0.002 (0.010)
% Black Constituency		2.592*** (0.515)
% Hispanic Constituency		2.604*** (0.699)
% Asian Constituency		5.573* (2.994)
House Race		-1.319*** (0.315)
Presidential Race		-0.396 (1.320)
Senate Race		-1.427*** (0.308)
Incumbent		-0.303* (0.157)
State FE		Yes
Topic Proportions [1-50]		Yes
Constant		18.721*** (2.116)
Observations	17	4,459
Log Likelihood		-1,410.570
Akaike Inf. Crit.		3,035.139
<i>Note:</i>		*p<0.1; **p<0.05; ***p<0.01

4.3 How are Racial Appeals Made?

Our last set of hypotheses concerned the language that would accompany implicit racial primes in our ads. We find no support for Hypothesis 3, which expected the presence of a visible minority in the ads to be associated with plausible codewords, such as "inner-city," or even more ambiguous codewords like "liberal." What we find instead is that the words which were more likely to appear alongside minorities in advertisements about issues like immigration are somewhat implausible candidates like "federal," "mexico," and "never," while the words that are most likely to accompany depictions of minorities in education ads are innocuous words like "teacher." We find support instead of Hypothesis 3A, which argues that implicit primes are unlikely to be accompanied by intuitive codewords. More study is needed, however, to parse the narratives that appear when minorities appear in the ads, since the Structural Topic Model and other topic models do not take word order or sentence structure into account when they assess how language is used differently in the presence of different covariates.

5 Conclusion

This paper introduces a new measure of racial priming by combining an analysis of the visuals of American campaign ads with an analysis of their language, and one of its contributions is that it uses this measure to conduct a systematic test of several hypotheses that have been raised in this literature. Its preliminary finding is that systematic implicit racial priming does not appear to be taking place in American

campaigns along the lines that previous scholars have suggested. The presence of minorities in a district leads to more racial appeals, not fewer; the presence of visible minorities in an ad is not associated with a higher likelihood to discuss racially charged issues like immigration, and racial codewords do not appear to be endemic to the advertisements that were run at the congressional, presidential, and gubernatorial level in 2000 or 2008. Future iterations of this project will extend this analysis by refining the model to capture information about the race of the incumbents and challengers, and explore the implications that candidate identity holds for the language of the advertisements in each race.

Effect of Minority in Ad by Topic

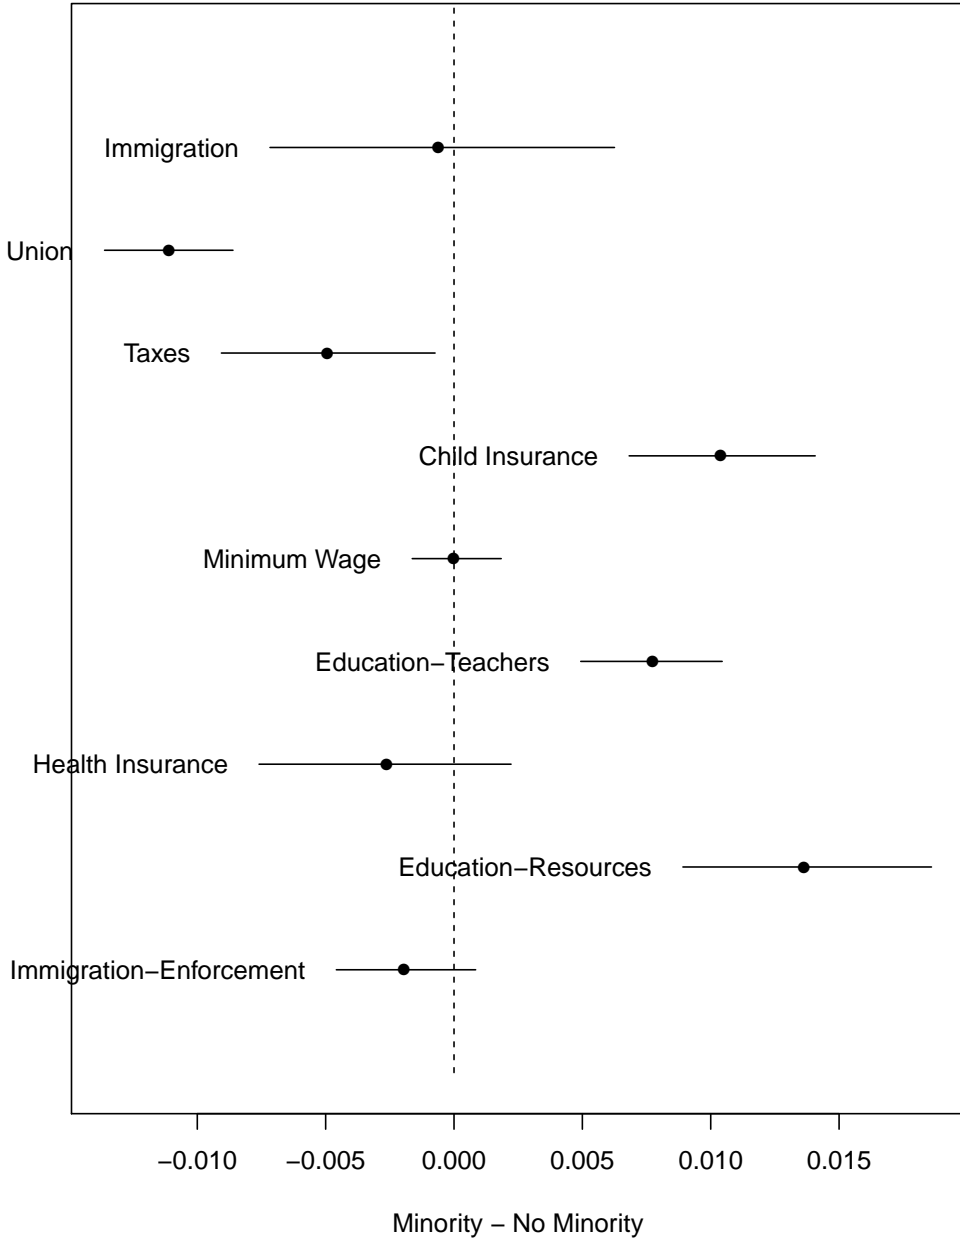


Figure 4: The correlational effect of the presence of a minority on the topic proportion of ads. Control include the party, incumbency status, margin of victory, number of people appearing in the ads, and district demographics.



Figure 5: Effect of a visible minority on the distribution of the vocabulary from nine topics from the STM run on then 2000 and 2008 combined advertisement corpus

References

- Blei, David, Andrew Ng & Michael Jordan. 2003. “Latent dirichlet allocation.” *Journal of Machine Learning and Research* 3:993–1022.
- Blei, David M. 2012. “Probabilistic topic models.” *Communications of the ACM* 55(4):77–84.
- Brader, Ted, Nicholas A Valentino & Elizabeth Suhay. 2008. “What Triggers Public Opposition to Immigration? Anxiety, Group Cues, and Immigration Threat.” *American Journal of Political Science* 52(4):959–978.
- Converse, Philip E. 1964. The Nature of Belief Systems in Mass Publics. In *Ideology and Discontent*, ed. David E Apter. New York: The Free Press of Glencoe pp. 206–261.
- Gilens, Martin. 1999. *Why Americans Hate Welfare*. Race, Media, and the Politics of Antipoverty Policy Chicago: University of Chicago Press.
- Hopkins, Daniel J & Gary King. 2010. “A Method of Automated Nonparametric Content Analysis for Social Science.” *American Journal of Political Science* 54(1):229–247.
- Huber, Gregory A & John S Lapinski. 2006. “The “Race Card” Revisited: Assessing Racial Priming in Policy Contests.” *American Journal of Political Science* 103(04):645–668.

- Huber, Gregory A & John S Lapinski. 2008. "Testing the Implicit-Explicit Model of Racialized Political Communication." *Perspectives on Politics* 6(01):125–134.
- Hurwitz, Jon & Mark Peffley. 2005. "Playing the Race Card in the Post–Willie Horton Era The Impact of Racialized Code Words on Support for Punitive Crime Policy." *Public Opinion Quarterly* 69(1):99–112.
- Hutchings, Vincent L & Ashley E Jardina. 2009. "Experiments on Racial Priming in Political Campaigns." *Annual review of political science* 12(1):397–402.
- Hutchings, Vincent L, Hanes Walton & Andrea Benjamin. 2010. "The Impact of Explicit Racial Cues on Gender Differences in Support for Confederate Symbols and Partisanship." *The Journal of Politics* 72(04):1175–1188.
- Iyengar, Shanto & Donald R Kinder. 2010. *News That Matters*. Television and American Opinion, Updated Edition University of Chicago Press.
- McIlwain, Charlton & Stephen M Caliendo. 2011. *Race Appeal*. How Candidates Invoke Race in U.S. Political Campaigns Philadelphia, PA: Temple University Press.
- Mendelberg, Tali. 2001. *The Race Card*. Campaign Strategy, Implicit Messages, and the Norm of Equality Princeton: Princeton University Press.
- Mendelberg, Tali. 2008. "Racial Priming Revived." *Perspectives on Politics* 6(01):109–123.
- Peffley, Mark & Jon Hurwitz. 2002. "The Racial Components of "Race-Neutral" Crime Policy Attitudes." *Political Psychology* 23(1):59–75.

- Peffley, Mark, Jon Hurwitz & Paul M Sniderman. 1997. "Racial Stereotypes and Whites' Political Views of Blacks in the Context of Welfare and Crime." *American Journal of Political Science* 41(1):30.
- Ridout, Travis N, Michael Franz & Kenneth Goldstein. 2008. "Measuring Exposure to Campaign Advertising." University of Wisconsin-Madison Advertising Project.
- Roberts, Margaret E, Brandon M Stewart, Dustin Tingley, Christopher Lucas, Jetson Leder Luis, Shana Kushner Gadarian, Bethany Albertson & David G Rand. 2014. "Structural Topic Models for Open-Ended Survey Responses." *American Journal of Political Science* 58(4):1064–1082.
- Stephens, Lafleur Nadiyah. 2013. The Effectiveness of Implicit and Explicit Racial Appeals in a "Post-Racial" America. PhD thesis.
- Valentino, Nicholas A, Vincent L Hutchings & Ismail K White. 2002. "Cues that Matter: How Political Ads Prime Racial Attitudes During Campaigns." *American Political Science Review* 96(01).
- Zaller, John R. 1992. *The Nature and Origins of Mass Opinion*. New York: Cambridge University Press.