Elevated Threat-Levels and Decreased Expectations: How Democracy Handles Terrorist Threats *

Tabitha Bonilla † Justin Grimmer ‡

May 20, 2013

Abstract

A persistent concern in democracies is that terror threats make the public willing to restrict freedoms for increased safety. But a large literature has struggled to determine how terrorist threats affect the public's policy preferences. To more credibly estimate the effects of terror threats, we exploit elevations of the U.S. government's color coded alert system. Using this design, a statistical model for texts and a new collection of news stories, we show that media outlets allocate substantially more attention to terrorism after an alert. The alerts have, however, only a limited effect on the public. The terror alerts raise the public's perceived likelihood of a terror attack, but opinion about President Bush's job performance, preferences for foreign intervention, or willingness to restrict civil liberties changes little in response to the alerts. Rather, the only consistent result is decreased economic expectations—consistent with the strong economic downturn after the 9/11 attacks and the types of stories published after the terror alerts are elevated. Terror alerts, then, did not exercise direct influence on the public's policy preferences. Instead the alerts changed the topic of conversation.

*For helpful discussions we thank Barry Burden, Shanto Iyengar, Clayton Nall, Gary Segura, and Sean Westwood

†Ph.D. Candidate, Department of Political Science, Stanford University; Encina Hall West 616 Serra St., Stanford, CA, 94305

‡Assistant Professor, Department of Political Science, Stanford University; Encina Hall West 616 Serra St., Stanford, CA, 94305

1 Introduction

The catastrophic terrorist attacks of September 11, 2001 caused dramatic shifts in U.S public policy and drastic shifts in public opinion. Shortly after the attacks a large war effort began in Afghanistan, the Patriot Act expanded the government's surveillance powers, and the public almost uniformly approved of President Bush's job performance. The shifts in policy and opinion also lead to concerns that Americans would hastily restrict civil liberties for policies that created a perception of safety. Russ Feingold (D-WI), the lone senator to oppose the Patriot Act, argued that the expansion of federal powers was evidence that "we must...redouble our vigilance to preserve our values and the basic rights that make us who we are" (Feingold, 2001). More conspiratorial concerns were also expressed—that politicians would exploit fears about terrorism to gain electoral support (Nagourney and Halbfinger, 2004).

With deep worries about the viability of democratic governance social scientists have undertaken wide ranging investigations into how the threat of terrorism affects the American public, with conflicting conclusions. Studies sometimes show a relationship between perceptions of threat and shift in policy preferences—though usually the effects of threat are found only among particularly responsive groups (Davis and Silver, 2004; Huddy et al., 2005; Merolla and Zechmeister, 2009; Hetherington and Suhay, 2011; Mahlhotra and Popp, 2012). The mixed evidence is due, in part, to the diverse research designs social scientists have employed. Some studies use observational survey evidence to examine relationships between perceptions of threat and policy preferences (Davis and Silver, 2004; Willer, 2004; Huddy et al., 2005; Kam and Kinder, 2007; Hetherington and Suhay, 2011). The survey evidence, however, is often criticized as deeply confounded—with respondents who are already sympathetic to restrictive policies more likely to feel threatened by terrorism (Mahlhotra and Popp, 2012; Kam and Kinder, 2007). To limit the confounding, other scholars use ex-

perimental manipulations (Merolla and Zechmeister, 2009; Gadarian, 2010; Mahlhotra and Popp, 2012). But the manipulations are also criticized. For example, Hetherington and Suhay (2011) argue that "the relationship between exogenous threats created by the experimenter and real-world threats is too often unclear" (Hetherington and Suhay, 2011, 551). What is needed, then, is a research design that provides the realistic conditions observed when terror threats are high, while also limiting the confounding that occurs in observational studies of threat (Gabel and Scheve, 2007).

In this paper we offer one such design. We exploit elevations of the Department of Homeland Security's (DHS) color-coded terror alert system to examine how the media and public respond to increased threat. While the terror alert system was much maligned (McDermott and Zimbardo, 2007), it is nearly ideal for studying how the public responds to attempts to raise the perceived likelihood of an attack. The seven official alerts were difficult to anticipate and, as we will show, received substantial media attention (Reese, 2003). Because of both trends in public opinion and the potential for political manipulation, however, studying terror alerts in extended time series creates the chance for deep confounding (Willer, 2004). Instead, we examine the period shortly before and after an alert level is raised. By analyzing the time immediately before and after the alert level is raised we limit the potential confounding in our estimates, increasing our certainty that the changes we observe are attributable to the elevated terror alert.

Using this new research design, we show that the terror alerts exerted substantial influence over the media's agenda, but this did not translate into broad effects on the public. We characterize the media's response to alerts using a topic model for texts that allows us to estimate the effect of alerts on the media's attention, (Blei, Ng and Jordan, 2003; Grimmer, 2012; Wallach, 2008) applied to a new collection of contemporaneous newspaper stories and major network nightly news broadcasts. This reveals the dramatic effects of alerts on the media, which substantially increased the attention allocated to terror threats in response to

the alert. The increased attention, however, was fleeting—attention to terrorism decreased sharply the day after a terror alert (Boydstun, 2012).

While the alerts caused a sudden and substantial shift in media coverage, the alerts had a more limited effect on public opinion. To measure the effect of the increased alert level, we use a set of serendipitous surveys—surveys that happen to be in the field when the alert level is raised. Using the elevated alert level as the treatment, this creates a control group—those interviewed before the alert—and a treated group—those interviewed after the alert (Holland, 1986). We show that the serendipitous surveys approximate random assignment. Across a variety of characteristics, we show that the treated and control groups look strikingly similar. Further, the use of serendipitous surveys ensures that we are using an actual intervention—avoiding criticisms about experimental manipulations others have made (Hetherington and Suhay, 2011).

Using this credible research design, we show that terror alerts do increase the perceived likelihood of a terror attack and the salience of terrorism. But this does not systematically shift respondents' evaluations of Bush's job performance, nor their policy preferences. The terror alerts do, however, cause a substantial increase in economic pessimism. This increased pessimism, we show, is deeply connected to how the public forms its economic expectations (De Boef and Kellstedt, 2004).

Our study contributes to the literature on the public policy consequences of terror warnings and the social science literature on how democracies handle terror threats. Much of the literature worries about a public that will be easily fooled by opportunistic politicians. We show, however, that the public's preferences are largely robust. In the time period we analyze—the years immediately following the 9/11 terrorist attacks—boosts in the salience of terrorism are not met with a corresponding shift in the public's preferences towards restrictions of civil liberties or support for foreign interventions. The alerts do change the topic of media coverage, shifting attention towards the alert. And the alerts boost the salience

of terrorism for the public. The result is that alerts appear to perform their intended job: they call the public's attention to a threat, without causing the public to shift their other priorities.[1]

The results of our study also contribute to the analysis of the public policy of terror warnings. While previous studies have demonstrated the substantial psychological costs of terror alerts (McDermott and Zimbardo, 2007), our results document an economic cost not previously considered (De Boef and Kellstedt, 2004)—a cost other studies suggest are reflected in stock trading (Pagano and Strother, 2008). This economic cost illuminates the challenges that government officials face when deciding whether to raise terror alerts. Not only must officials consider the need to raise the public's awareness of potential threats. But elected officials must also consider the potential side-effects of the alerts, which may include increasing economic pessimism in an already weak economy.

Given the importance of terror alerts for our research design, we first provide a brief overview of the Department of Homeland Security's (DHS) alert system and the political controversy surrounding the alerts.

2 Research Design: The Sudden Elevation of Terror Threat Levels

To examine how the media and public respond to terror threats, we exploit elevations of the DHS's terror alert system, which the Bush administration announced on March 12, 2002. The alert system was intended to "recommend protective measures for federal departments and agencies to prevent, prepare for, mitigate against, and respond to, terrorist attacks" (Reese, 2003).

[1]We caution that the finding about the public's non-response to elevations of the alert level likely depends on the context of our study. Our study examines how boosts in the salience of terrorism affect an already anxious public. When there is less anxiety in the public, however, the elevation of the alert levels (or similar elite statements) may have a larger effect on public policy preferences.

The color coded alert system had five threat levels, but only two were used throughout its existence. The alert level was usually kept at "Yellow" or an elevated threat level. But when the administration suspected an increased likelihood of a terror attack—either from intelligence or world events—the Bush administration raised the threat level to"Orange" or a high threat level. We call this elevation a *terror alert*. Table 1 provides the date of each elevation of the threat level and the reason the Bush administration provided for raising the threat level. Elevations of the threat level tended to occur suddenly, with no information to the general public before the elevation occurred. For example, Tom Ridge, Secretary of the Department of Homeland security, held a press conference at 12:30 pm Eastern Standard Time to announce the February 7 terror alert. The timing of the alert is crucial. It provides ample time for nightly newscasts to cover the alert on their newscasts and for newspapers to publish stories about the alert the next day. The result is that respondents to surveys would have been able to quickly receive information about the elevated alert.

Table 1: Elevations of Alert from Yellow to Orange

Number	Elevation Date	Reason
1	September 10, 2002	Anniversary of 9/11 attacks
2	February 7, 2003	Al Qaeda plans to attack "lightly secured targets"
3	March 17, 2003	Start of Iraq war
4	May 20, 2003	Terror attacks in Saudi Arabia and Morocco
5	December 21, 2003	Spike in "threat related intelligence reports"
6	August 1, 2004	Financial sector threat
7	July 7, 2005	Mass transit, London attacks

Elevations of the terror alert level from yellow to orange. Source: http://www.dhs.gov/homeland-security-advisory-system.

While the alert system was intended to replace vague and haphazard warnings that the Bush administration issued initially after the 9/11 attacks (Reese, 2003), the alert system was regularly criticized as ineffective (McDermott and Zimbardo, 2007). There were also allegations that the Bush administration used the alert system for electoral gain. This

concern was expressed regularly during the 2004 presidential campaign. After the alert on August 1, 2004, which followed the Democratic National convention, Howard Dean thought the Bush administration used alerts to affect campaign dialogue. Dean argued "that every time something happens that's not good for President Bush, he plays this trump card, which is terrorism" (Nagourney and Halbfinger, 2004). Tom Ridge, Secretary of DHS under Bush, fueled speculation that the Bush administration used terror alerts for political gain. In his autobiography, Ridge wrote that in response to a tape from Osamba bin Laden days before the election that "[Attorney General John] Ashcroft strongly urged an increase in the threat level, and was supported by [Secretary of Defense Don] Rumsfeld. There was absolutely no support for that position within our department. None. I [Ashcroft] wondered, 'Is this about security or politics?' " (Bloom and Ridge, 2009, 226-227).

Terror alerts, then, were announced unexpectedly and were potentially subject to manipulation. The sudden announcement makes them ideal for studying how sudden changes in perceptions of threat affect both the media and public in American democracy. The timing of alerts and the potential for strategic manipulation imply that a narrow window around alerts is preferable to a long time series. Table 1 shows that most of the alerts are concentrated early in President Bush's first term—a time when terrorism was relatively more salient, the administration was more popular, and the public was more supportive of policies that restricted liberties to increase safety. A time series model that examines the effect of the alerts will likely conflate the background conditions of the time the alerts were raised with the effect of the alerts (Willer, 2004). Similarly, if the Bush administration is strategic in issuing terror alerts, a time series model will likely fail to credibly estimate the effect of the alerts.

Rather than time series with alerts as interventions, we focus on the period immediately before and after the alert is raised (Shadish, Cook and Campbell, 1978). By focusing on the time around an alert, we ensure that neither trends in media coverage and public opinion,

nor the strategic use of the alerts, confound our analysis (Shadish, Cook and Campbell, 1978). This is because the slowly changing trends in all of our variables are essentially identical before and after the alert, ensuring that the time period before the alert provides a reasonable estimate of the counterfactual values for the responses and media coverage after the alert. The result are differences that are attributable to the alerts, or more valid causal inferences (Angrist and Pischke, 2008).

Using this design built around alert elevations, we are ready to begin our inquiry into how terror alerts affect pubic attitudes.

3 How Terror Alerts Affect Media Coverage

We begin our inquiry with how the alerts affect the media's agenda. To do this we gathered a collection of 51,766 newspaper stories and transcripts of stories from major network nightly newscasts (ABC, CBS, NBC), from the two days prior to an alert, the day of the alert, and the two days after. To collect the newspaper stories and transcripts we use the *Lexis-Nexis* online archive, searching for front page and front section stories and nightly news transcripts on the dates surrounding six of the terror alerts. Our collection of newspaper stories includes a wide array of papers from across the country, ensuring that our analysis is not dependent upon national newspapers only.[2]

Before applying our topic model to estimate the effects of the alerts, we first examine differences in the rates words are used before and after the terror alert. While simplistic, the differences in word rate usage foreshadow our findings using our more sophisticated topic model—that alerts cause immediate, though short lived, shifts in media attention and that the alerts distract attention from other salient policy disputes. Specifically, we use the per-document word usage and identify words with the largest rate increase after a terror alert

[2]In both our analysis of text and surveys we avoid the alert issued during the Iraq war. This is because that alert change is confounded by the start of a war, a higher salience event whose effects likely overwhelm those the terror alerts.

(Monroe, Colaresi and Quinn, 2008). The per-document word usage provides an intuitive summary of the words that distinguish documents after the terror alert. This approximates more complicated models used to identify predictive words (Monroe, Colaresi and Quinn, 2008; Taddy, 2013), while maintaining an easy to interpret parameter estimate.

The words with the largest rate increase are presented in Figure 1. The left-hand column contains the words, which are sized proportional to how well they separate stories reported before and after an alert. The middle-column shows the average per-document word rate (vertical axis) for nightly newscasts over the five days included in our window (horizontal axis). The points are averages across the alerts. Because all alerts are issued in time for the news casts, the third day of our window (the day of the alert) is the first post-alert reporting. The right-hand column presents the average per-document word rate for newspapers.

Figure 1 shows that *alert* has the largest rate increase after a terror alert is issued. Before the terror alert, the word *alert* is used only 0.07 times per nightly newscast. But on the day after the terror alert, *alert* is used 0.81 times per story—representing the substantial shift in focus to the alert. Both *threat* and *terror* have similarly large jumps. This increase is much more substantial for nightly newscasts than newspapers. This is due, in part, to the national focus in nightly newscasts that are not found in local papers. It also reflects the larger share of other stories that newspapers can report in their front section, while nightly newscasts have a more limited number of stories they can report each evening. Figure 1 also shows that the media's attention to the alert is fleeting: the burst in attention to terror related words is followed by an equally precipitous decline in attention (Boydstun, 2012). For each of the words, we see that the rate the word is used return to pre-alert levels by the fifth day.

The words with the largest rate decreases reveal that the terror alerts crowd out coverage of policy debates or elections. Some of the words deal with the Iraq war—such as *Iraq*, *Saddam*, and *Powell*. Others reflect the timing of the alerts, including *convention* and *presidential*. Alerts do focus the media's attention on possible threats, but they also appear

Figure 1: Changes in Word Rates Around Terror Alerts

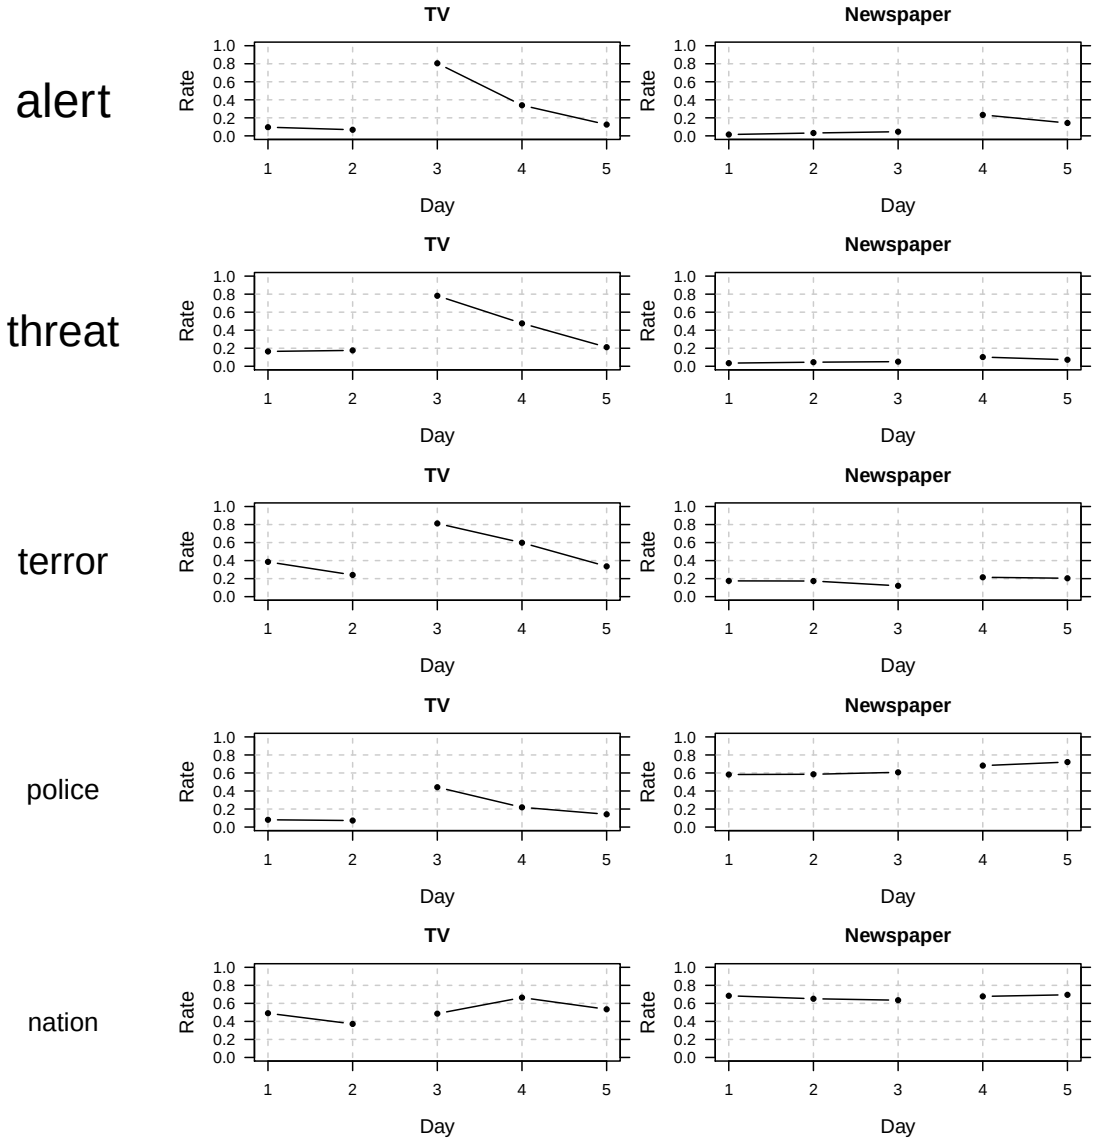

This figure shows that words related to terrorism are used much more often after an alert than before. The left-hand column shows the words with usage rates that change most drastically after an alert. The middle column shows how the word usage rates change for television news, while the right hand column shows the change in usage rate for newspaper stories. This reveals the dramatic shift in attention in television news surrounding an alert.

to crowd other news stories.

3.1 A Topic Model for Measuring Shifts in Attention

Raw changes in word usage rates suggest that alerts shift what the media cover, while crowding out other policy disputes. To better understand how the alerts affect coverage we introduce and then apply our topic model to the collection of news stories. Building on *Latent Dirichlet Allocation* (LDA) (Blei, Ng and Jordan, 2003), our model measures the topics that occur in the newspapers and evening news stories. We will suppose that there are 24 topics that occur in collection of newspaper stories and evening news stories, a number chosen using both substantive and statistical criteria (Quinn et al., 2010; Grimmer and Stewart, 2013).[3]

Our primary quantity of interest is how newspapers and nightly newscasts divide their attention over the topics both before ($t = 0$) and after ($t = 1$) a terror alert. We will call the proportion of stories that newspaper i, for alert a, dedicates to topic k, as $\text{Attention}_{i,a,t,k}$. Collecting across topics we then have our primary quantity of interest,

$$\textbf{Attention}_{i,a,t} = (\text{Attention}_{i,a,t,1}, \text{Attention}_{i,a,t,2}, \ldots, \text{Attention}_{i,a,t,24}) \tag{3.1}$$

Conditional on each newspaper's $\textbf{Attention}_{i,a,t}$ we suppose that the topic of each story j is drawn and then conditional on the topic, the content of the document is drawn, $\boldsymbol{y}_{ijat}$. As Grimmer and Stewart (2013) explain about similar models, our model shares a common structure with LDA. However, we assume that each document has only one topic, a more substantively appropriate assumption for our analysis. While the multiple topics per document is often a useful assumption for analyses (Blei, Ng and Jordan, 2003), when ana-

[3]We chose this number of topics using two different criteria. Statistically, the number of topics presented here aligns with the estimated number of topics from a non-parameteric topic method. (Specifically, a Dirichlet process prior on a mixture of multinomial distributions). While useful, this sort of model has its own built in assumptions that affect the number of topics that are used (Wallach et al., 2010), so we also engage in substantive search, as outlined in (Quinn et al., 2010). We varied the number of topics, before determining that 24 topics provided a reasonable comprise. We should note that our findings are largely robust to varying the number of topics over a reasonable range.

lyzing newspaper stories assuming one topic can be useful: it allows us to classify stories as primarily about terrorism.[4]

The data generating process, which we provide in Equation 3.2, illuminates the shared structure with LDA. To efficiently borrow information across texts from similar time periods, we suppose that each $\mathbf{Attention}_{i,a,t}$ are drawn from an alert and time specific Dirichlet distribution. Then, conditional on $\mathbf{Attention}_{i,a,t}$ we draw each document's topic from a Multinomial distribution. Conditional on the document's topics, we then draw its contents from a topic-specific multinomial distribution, with the rate of words determined by $\mathbf{Word\ Rate}_k$.

$$\begin{aligned}
\mathbf{Word\ Rate}_k &\sim \text{Dirichlet}(\mathbf{1}) \\
\alpha_{t,a,k} &\sim \text{Gamma}(0.25, 1) \\
\mathbf{Attention}_{i,a,t} &\sim \text{Dirichlet}(\boldsymbol{\alpha}_{a,t}) \\
\mathbf{Topic}_{i,j,a,t} &\sim \text{Multinomial}(1, \mathbf{Attention}_{i,a,t}) \\
\mathbf{y}_{i,j,a,t} &\sim \text{Multinomial}(n_{i,j}, \mathbf{Word\ Rate}_k)
\end{aligned} \tag{3.2}$$

To apply the model to the collection of newspaper stories and evening news transcripts, we first represent the texts quantitatively. To do this, we follow a set of steps, including: discarding word order, stemming words, eliminating punctuation, and removing stop (or place holder) words. (See Grimmer and Stewart (2013) for details). We estimate the model using a variational approximation—a fast and deterministic alternative to Gibbs sampling (Jordan et al., 1999; Grimmer, 2011).

Table 2 presents the topics that result from applying the model to the collection of newspaper stories and nightly newscasts. We used two methods to label the topics. Following Quinn et al. (2010), we randomly sampled 20 documents assigned to each category, read the

[4]Often the multitopic assumption is advocated as a way to deal with content unrelated to a story in a newspaper article. Our experience is that this material rarely determines the final classification as has little effect on our topics. Further, our preprocessing steps eliminate much of this additional content, with little to no consequence for our estimated topics.

documents, then produced hand labels. We recorded the hand coded labels in the left-hand column. We also used the mutual information between words and topic labels to automatically generate labels for the topics (Manning et al., 2008), which we record in the middle-column. Finally, we report the proportion of the stories that are assigned to each story in the right-hand column.

Table 2: Topics of News Stories and Evening News Broadcasts

Label	Discriminating Stems	%
Memorial	peopl,time,famili,work,world,live,home,like	11.28
Local Small Business	citi,plan,project,000,develop,counti,properti,million	7.83
Criminal Prosecution	polic,offic,arrest,old,sentenc,prison,prosecutor,told	6.96
Iraq/World	al,attack,iraq,offici,palestinian,bomb,american,peopl	6.86
Local Philanthropy	section,load,date,type,photo,newspap,english,music	6.83
Business	compani,busi,market,million,price,peopl,work,time	6.40
Local Awards	educ,board,teacher,budget,fund,colleg,high,million	4.92
Law and Order	polic,man,car,arrest,offic,vehicl,old,accid	4.64
Medical Law	compani,law,drug,feder,lawsuit,million,worker,medic	4.38
Personal Interest Stories	team,game,player,run,win,sport,season,second	4.04
State/Local Elections	vote,candid,voter,democrat,campaign,race,counti,ballot	3.93
2004 President Campaign	bush,nation,convent,columbia,presid,space,war,democrat	3.79
Local	school,center,group,inform,denver,program,island,week	3.67
National Political News	bush,law,senat,rule,feder,govern,commiss,court	3.55
Local Crime	school,town,island,polic,committe,charg,hous,case	3.28
Weather	hurrican,storm,wind,rain,flood,area,damag,beach	3.22
Iraq War	iraq,bush,unit,war,saddam,nuclear,powel,nation	2.75
Terror Alerts	alert,attack,offici,terrorist,terror,al,homeland,threat	2.68
Medical Research	hospit,drug,studi,research,percent,dr,care,medic	2.45
Local Budgeting	budget,million,citi,increas,fund,revenu,fiscal,cost	2.14
Religion and Politics	church,law,appeal,rule,court,marriag,sentenc,order	1.98
Life Events	funer,church,home,memori,son,daughter,sister	1.67
Obituary	funer,late,servic,burial,sister,wife,born,niec	0.64
Legal Notice	record,hear,island,probat,sold,certifi,clerk,notic	0.09

Table 2 shows that the model was able to identify substantively distinct and interesting topics. The largest topic deals with *memorials*, and in particular memorials to the victims

of the 9/11 attacks. This is appropriate, as the first terror alert occurred on the anniversary of the 9/11 attacks. Other topics reflect the major national policy debates at the time, such as the *Iraq war*, or the local news interest of papers—including *local philanthropy* and *state and local elections*. And critically, the model identifies a *terror alert* topic—a collection of stories that announce that the alert level was raised and explain the evidence used to justify the elevation.

3.2 How Terror Alerts Affect Media Attention

With the credibly labeled topics and estimates of attention for newspapers and nightly newscasts, we now examine how attention shifts in response to the terror alerts. We use the measures of how media outlets divide their attention before and after an alert to estimate the effect of alerts on the media's agenda. Figure 2 presents the average change in attention to topics (columns), for each of the six terror alerts we include. In each panel the dot is the average change and the lines are 95 percent confidence intervals for the change.

Figure 2 shows that terror alerts usually substantially increase the media's attention to terrorism. The largest increase in attention occurred after the December 21 alert—with media outlets allocating 6.3 percentage points more of their space to terror articles than before (95 percent confidence interval, [5.25, 7.44]). A similar shift occurred in four of the five other alerts. After an alert is issued, newspapers from across the country and nightly newscasts dedicate substantially more attention to terrorism.

The exception is the terror alert issued on September 10, 2002, where there was only a 0.6 percentage point increase in attention to domestic terrorism: a small increase with 95 confidence intervals that substantially overlap zero (95 percent confidence interval [-0.6, 1.76]). With the anniversary of the 9/11 attacks approaching, the media already allocated a substantial amount of attention to the threat of terrorism. And as the bottom-left panel of Figure 2 shows, there was a much larger shift in attention towards memorializing the victims

Figure 2: Terror Alerts Shift Attention to Terrorism, Though the Shift Varies Across Alerts

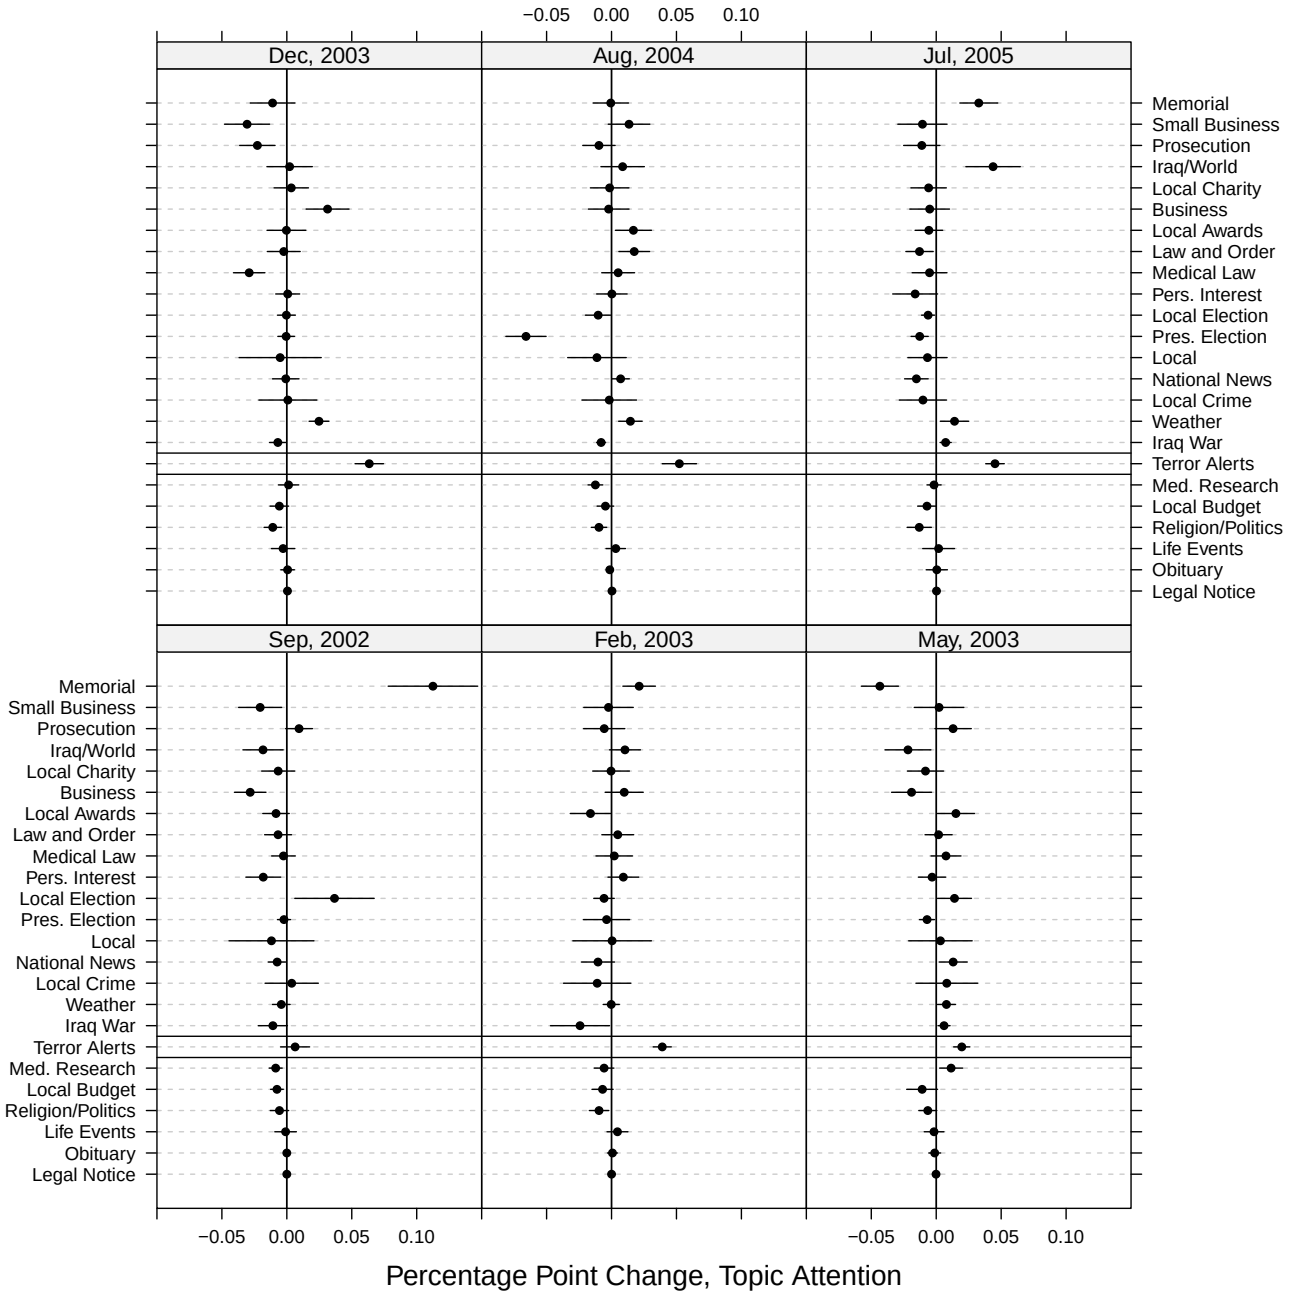

This figure shows the within-outlet change in attention to topics that occurs in response to alerts. This reveals that terror alerts are widely reported, though sometimes other events receive more attention.

of the attacks than publicizing the elevated alert. Together, this ensured that the terror warning that coincided with 9/11 did little to affect the media's agenda.

But when terror alerts do exert influence it is usually short lived. The pace of decline is particularly remarkable for nightly news broadcasts. Averaging across all terror alerts, there is a 27 percentage point increase in the proportion of stories about terrorist threats. But the third day after the terror alert, the attention declines substantially—to only a 7 percentage point increase over the day before the alert was raised. Newspapers follow a similar trajectory. The day after an elevation of the threat level, newspapers allocate nearly three times the space to domestic terrorism (increasing the space allocated 3.3 percentage points). By the next day, however, newspapers are allocating only 1.7 percentage points more than the day before the alert was raised.

Not only is the effect of the terror alerts fleeting, Figure 2 shows that the terror alerts also draw media attention away from other major policy disputes. For example, the terror alert on February 7, 2003, caused a decline of 2.4 percentage points in attention to the Bush administrations push for the Iraq war (95 percent confidence interval, [-4.70, -0.00]), while the terror alert on May 20, 2003, caused a decline in attention to the conduct of the war (2.1 percentage point decrease, 95 percent confidence interval [-3.94, -0.00]). And due to its close proximity to the Democratic National Convention, the August 1 terror alert caused a 6.6 percentage point decline in attention to the 2004 presidential election (95 percent confidence interval, [-8.14, -5.04]).

This section shows that terror alerts affect the media's agenda. Alerts cause an immediate and substantial focus of attention on the increased risk of terrorism and the justification for the alert elevation. Because media space is scarce this diverts attention away from other stories—such as the justification for the Iraq war or the 2004 presidential campaign. The Bush administrations' terror alerts, then, appear to accomplish one of their objectives: to increase the media attention to the potential for a terror attack. This effectiveness also lead some to be concerned about the potential for the alerts' misuse. As politicians and pundits recognized, the alerts were a powerful agenda setting tool for the Bush administration—

providing a tool to focus on the risk of terrorism and creating the potential to manipulate public opinion.

4 Terror Alerts, Serendipitous Surveys, and Economic Expectations

We now show that this potential for manipulation does not manifest: there is only a shift in the public's economic expectations, and not their policy preferences. This narrow effect of the alerts is surprising, but prior literature offers some insights into why the public's economic expectations will be more responsive to terror alerts than their policy preferences. MacKuen, Erikson and Stimson (1992) present evidence that the public forecasts when considering current economic conditions (but see Haller and Norpoth 1994; Norpoth 1996). Because terror alerts will remind the public of the detrimental effects of a terror attack—and because the public's economic expectations are malleable (De Boef and Kellstedt, 2004)—we would expect that alerts will have a negative effect on the perceived economic conditions. In contrast, the public's policy preferences were already divided along partisan lines when the Bush administration began issuing terror alerts (Jacobson, 2007). Given this division, it is difficult to move the sorted partisan positions, even if the alerts raise the salience of terror attacks (Zaller, 1992).

To assess the effects of the terror alerts, we use a set of serendipitous surveys that happened to be in the field when the alert level was raised. This uses the alert elevation to partition the survey into a control group—those interviewed before the alert level was raised—and treatment group—those interviewed after the elevation of the alert level. We used the *Roper Center* archive to identify the surveys, which are contained in Table 3. There are many advantages to serendipitous surveys, but as Table 3 shows we are limited by the surveys that happened to be in the field when the alert level was raised. This implies that

we are limited in the number of surveys available for each alert, the firm which conducted the surveys, and even the questions that are available for each alert. The most drastic consequence is that we fail to find any surveys for one alert: we were unable to identify a serendipitous survey for the terror alert on July 7, 2005, after the London bombings.

Table 3: Surveys Used to Assess the Effect of Terror Alerts

Alert Date	Survey Firm	Pre-Alert Dates	Post-Alert Dates
Sep. 10, 2002	Gallup (Investor)	Sep. 1-9	Sep. 10-15
Feb. 7, 2003	ABC-Washington Post	Feb. 6,	Feb. 7-9
	Gallup (Investor)	Feb. 1-6	Feb. 7-15
May 20, 2003	Gallup (Adult)	May 19	May 20-21
Dec. 21, 2003	Pew	Dec. 19-20	Dec. 21-Jan. 4, 2004
	National Annenberg Election Study	Dec. 16-20	Dec. 21-25
Aug. 1, 2004	National Annenberg Election Study	Jul. 27-31	Aug. 1-5

While restricting our focus to serendipitous surveys may seem overly restrictive, it provides numerous advantages. For example, it ensures that we use a realistic intervention—the government's actual attempt to elevate the threat level. This means that we avoid the experimental treatments that observational scholars critique (Hetherington and Suhay, 2011). The use of an actual intervention also ensures that our analysis includes the broad consequences of terror alerts—such as corresponding shifts in media attention, fluctuations in the stock market, and the spreading of information through social networks.

The surveys are useful because they concentrate attention on when the terror alerts are most likely to have an effect: immediately after the alert is issued. Our analysis of the media's attention shows that the alerts have an immediate and brief effect on the topics of stories that news outlets produce. Given the most likely source of information for the public about the alert are the news media, it follows that the reaction to the alerts would have to be immediate as well. Further, the longer the time window after the alert, the more likely other factors confound estimates of the terror alert's effect (Shadish, Cook and Campbell,

1978).

The serendipitous surveys, then, provide a realistic intervention and focus attention on the time when the effect of the alert is most clearly attributable. The serendipitous surveys also create treatment assignments that approximate the random assignment found in an experiment (Mahlhotra and Popp, 2012). Figure 3 shows this *balance* between the treatment and control group. Each plot presents the test statistics for the difference in means between the respondents interviewed after the alert (the treated group) and the respondents interviewed before the alert (control group). The closer the points lie towards the center line, then the more similar the treated and control group's background characteristics. If the dots fall outside of the black lines, then we can reject the null hypothesis of no difference at the 95% confidence level.

Across the surveys, Figure 3 shows that the treatment and control groups are similar—for most surveys and background characteristics, the treatment and control groups are close to zero. And across the surveys, there are only six instances where we reject the null hypothesis of no difference between the treatment and control groups. But there are some instances where there is imbalance between the treatment and control groups. This can occur because of systematic selection in when respondents are interviewed (Burden, 2000). Respondents who are harder to contact—those who are younger, lower income, and Democrats—are more likely to be interviewed after the terror alert is elevated (though this difference is often small). To mitigate the influence of the moderate remaining confounding, we will estimate the effect of the terror alerts conditional on the respondent's background characteristics (Ho et al., 2007). This provides a robust research design for measuring the influence of terror alerts on the public.

Figure 3: Assessing Balance Across Treatment and Control Groups

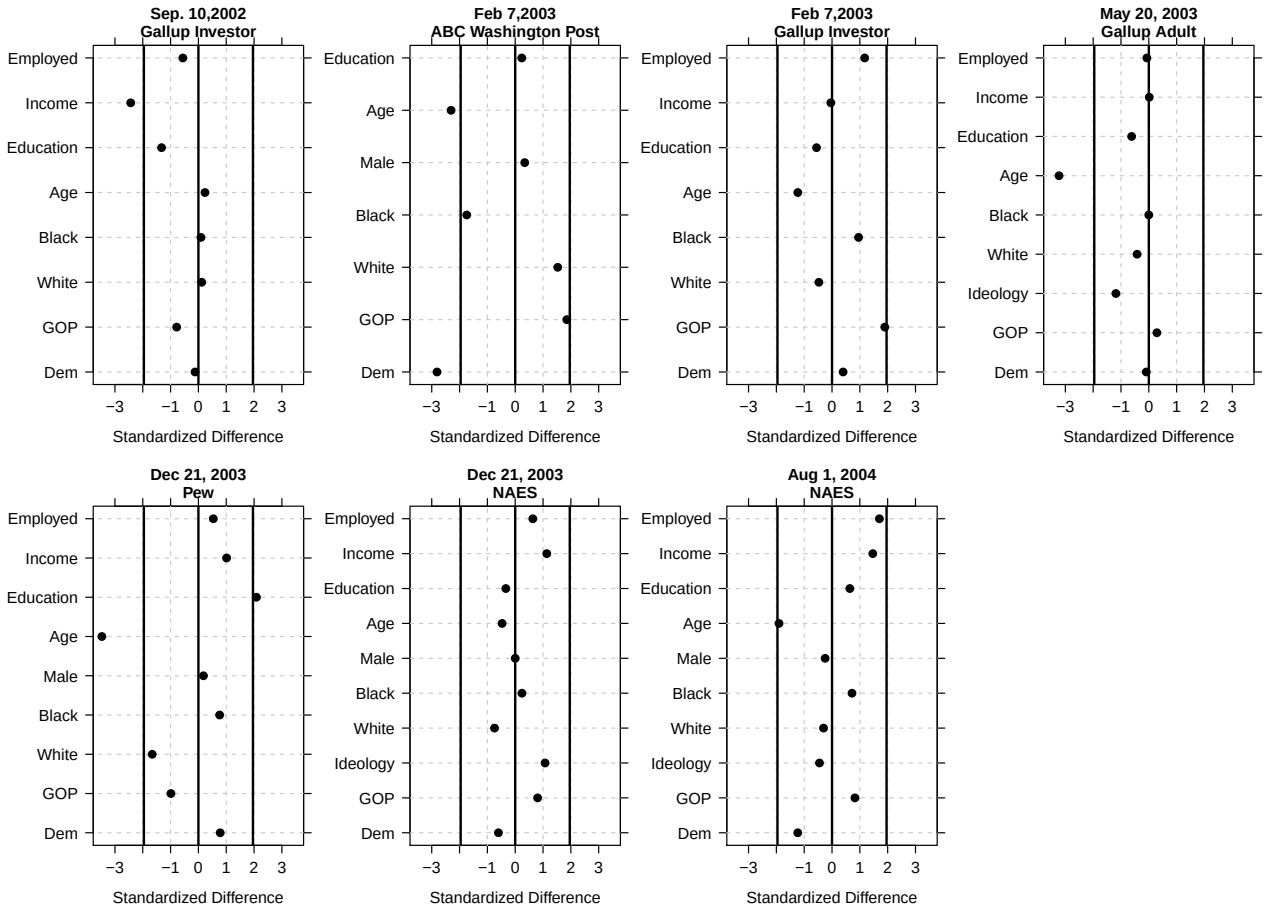

This figure presents the balance across the treatment and control groups, presenting the standardized difference in means for the background characteristics available in the surveys. Across surveys, we find that treatment and control groups appear quite similar.

4.1 Elevated Threats, Decreased Expectations

Serendipitous surveys provide realistic interventions and approximate the conditions of a randomized experiment.[5] Using the surveys, we first assess how the terror alerts affect the

[5]In what follows, we compare many effects. This creates the real chance of false discovery. To mitigate against this possibility before conducting the analysis we identified the questions of interest, based on a set of dependent variables where we expected to observe influence. Further, to mitigate against false discovery we attempt to replicate our major findings across surveys–decreasing the likelihood that a finding is merely the result of random error and multiple comparisons.

perceived likelihood of a terror attack.[6] The left-hand plot in Figure 4 presents the effect of the terror alert on the perceived likelihood of a terror attack. Across the figures, the points are the estimates of the average treatment effect and the thick and thin lines are 80 and 95 percent confidence intervals, respectively.

Figure 4: Terror Alerts Increase the Perceived Threat of Terror Attacks

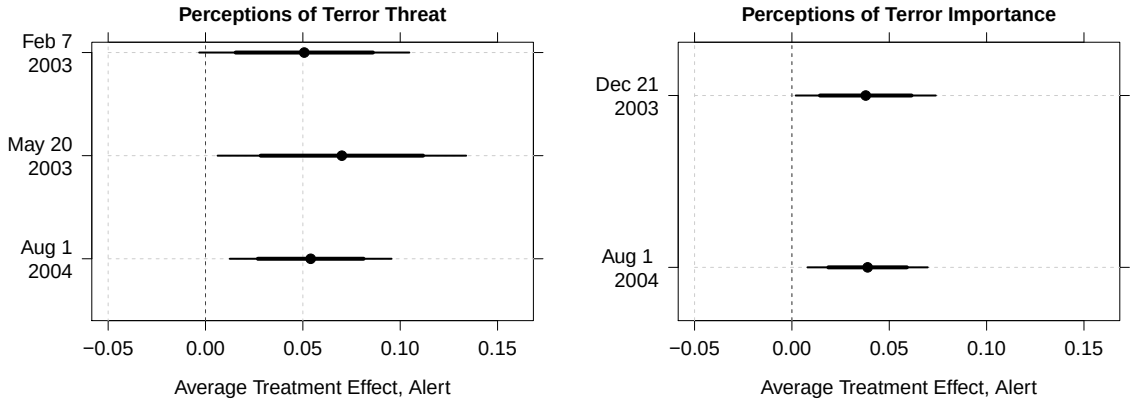

This figure shows that terror alerts increased the threat of an attack (left-hand plot) and the likelihood that respondents identified terrorism as the most important problem (right-hand plot).

Figure 4 shows that the elevation of the terror alert and the shift in media attention caused an increase in the perceived likelihood of a terror attack. For example, after the May 20 terror alert there was a 7 percentage point increase in the proportion of the public who thought that there would be "an act of terror over the next several weeks" (95 percent confidence interval, [1, 14]). A similar increase was found after the February 7 terror alert (5.1 percentage point increase, 95 percent confidence interval [-0, 10]) and the August 1, 2004, alert (5.4 percentage point increase, 95 percent confidence interval [1.25, 10.0]).

The terror alerts also cause respondents to be more likely to identify terrorism as the most

[6]The Gallup Investor's survey during the February 7 terror alert asked respondents if a terror attack was the most likely threat to the investment environment, the Gallup general population survey during the May 20 terror alert asked adults if a terror attack is likely, and the NAES asked respondents during the August 1 terror alert about the likelihood of a terror attack.

important problem facing the country. The right-hand plot in Figure 4 shows that there is a substantively interesting increase in the proportion of the public who identify terrorism as the most important problem. After the December 21 alert there was a 3.8 percentage point increase in the proportion of respondents who identified terrorism as the most important problem facing America (95 percent confidence interval, [0.0, 7.6])—a nearly *50* percent increase in the proportion of respondents who identify terrorism as the most important problem facing America.

Figure 4 shows that the alerts appear to have an effect on the public—both increasing the perceived likelihood of an attack and the importance of terrorism. However, the plots in Figure 5 show that this does not affect how the public evaluates President Bush, nor the public's policy preferences. Consider the left-hand plot, which examines the effect of the terror alert on President Bush's approval rating. Across the four alerts with available data (and in two surveys from the December 21 terror alert), we find that the alerts have little effect on Bush's approval rating: the effects are inconsistent and their confidence intervals have substantial overlap with zero. Further, the lack of an effect on approval is not because we lack sufficient statistical power. To show this, we pooled across surveys in the field during the terror alert. We included a fixed effect for each survey, ensuring that baseline differences in Bush's approval rating would not affect our results. The bottom line in the left-hand plot of Figure 5 is the pooled estimate, revealing that the alert levels cause a small increase in President Bush's approval ratings—1 percentage point—an increase neither substantively nor statistically significant, though it is relatively precisely estimated (95 percent confidence interval [-1.36 , 3.25]).

The center and left-hand plots in Figure 4 reveal the alerts also do little to change policy preferences. The center plot shows how the alerts affect support for the Iraq war.[7] No alert

[7]During the February 7 terror alert, the ABC-Washington post poll asked respondents if they favor military action against Sadaam Hussein in Iraq, during the December 21 terror alert a Pew poll asked respondents if they favor keeping troops in Iraq, and during the alerts on December 21, 2003, and August 1,

Figure 5: Terror Alerts Have Small Effect on GWB Approval, Support for Iraq War, or Support for Civil Liberty Restrictions

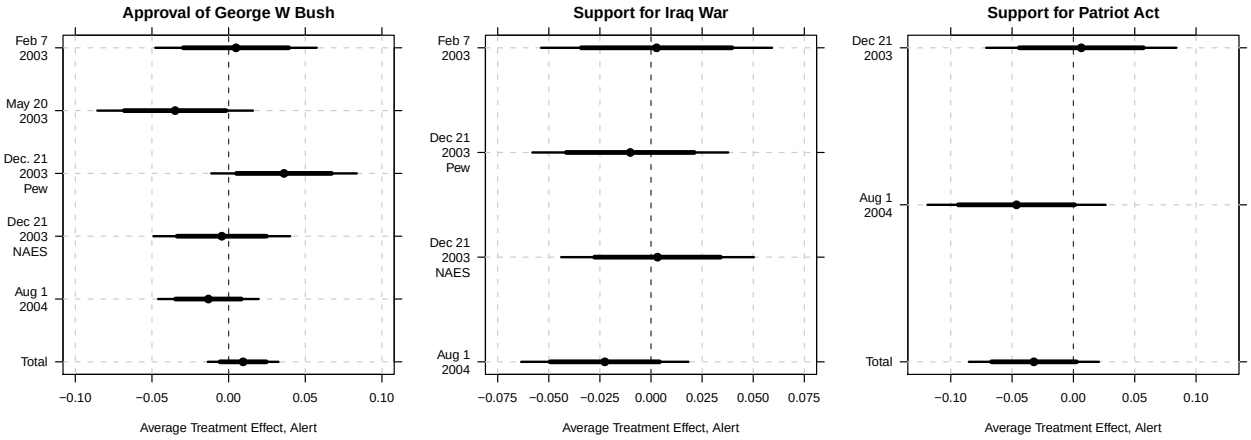

This figure shows that the alerts do little to affect approval of President Bush, support for the Iraq war, nor support for the Patriot Act.

substantially affected the public's preference for war in Iraq after the conflict began, nor did the alert affect the evaluation of the war after the invasion started. For example, the Bush administration was advocating for the Iraq war when the February 7 terror alert was issued. Despite the concerns of pundits, this alert had essentially no effect on the public support for the Iraq invasion—increasing support only 0.3 percentage points—a minute increase with 95 percent confidence intervals that substantially overlap with zero (95 percent confidence interval, [-6, 5]). The terror alerts also did little to affect the public's support for the Patriot Act.[8] During the terror alerts on December 21, 2003, and August 1, 2004, we find little change in support for the Patriot Act after the terror alert is issued. Pooling the observations across the alerts suggests that the alerts *decreased* support for the Patriot Act 3.2 percentage points, but this difference is not statistically significant (95 percent confidence

2004, the National Annenberg Election Study asked respondents if they approve of President Bush's handling of the Iraq war.

[8]During the December 21 and August 1 terror alerts, the NAES asked if the Patriot act is good for the country.

interval, [-8.52, 2.08]).

While terror alerts do not shift policy preferences or evaluations of President Bush, they do increase economic pessimism. The left-hand plot of Figure 6 shows that the terror alerts cause respondents to be more likely to expect worse economic conditions in the future.[9] The May 20 alert caused a substantial increase in economic pessimism—with respondents 10.8 percentage points more likely to expect the economy to be worse in the future (95 percent confidence interval, [5.84, 16.95]). This increase in pessimism is replicated in the December 21 alert, with a 4.6 percentage point increase in the number of respondents expressing pessimism (95 percent confidence interval, [1.0, 8.0]). And while the difference for the August 1 alert is not significant, pooling the surveys together reveals that the alerts cause a 3.3 percentage point increase in the proportion of respondents who believe the economy will be worse in a year (95 percent confidence interval, [1.1, 5.5]).

The right-hand plot shows that the general economic pessimism was paralleled by an increase in investment pessimism, though only for one terror alert.[10] The February 7 alert caused 7.5 percentage point decrease in the proportion of respondents who said that now was a good time to invest (95 percent confidence interval, [-13.82, 1.22]). In the same survey, respondents also expressed a significant decrease in optimism. Pooling across surveys, we find that investors are more pessimistic after terror alerts.

The right-hand plot of Figure 6 shows that the September 10 terror alert does not substantially affect the public's economic pessimism. This, however, may be due to the circumstances of the alert. With the anniversary of the terror attacks already approaching, terrorist threats were already salient for the public. And as we demonstrated in the analysis of news-

[9]During the May 20 alert, Gallup asked respondents if the economic conditions are getting worse or better and during the December 21 and August 1 the National Annenberg Election study asked respondents if the economy would be better or worse a year from now.

[10]During the September 10 and February 7 terror alerts, Gallup asked a sample of investors if it was a good time to invest and if the investors were pessimistic about the stock market. Both surveys were conducted on homes with at least ten-thousand dollars of investment.

Figure 6: Terror Alerts Increase Pessimism about Economic Future

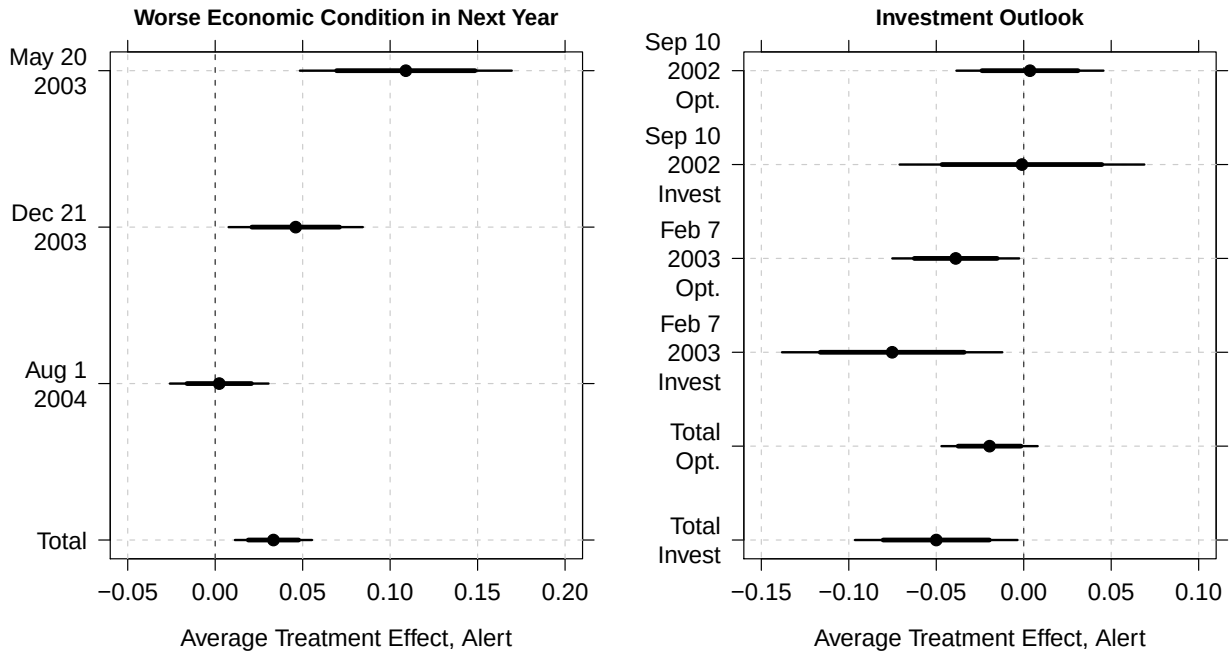

This figure shows that the alerts affect the public more pessimistic about the economic future.

paper stories, the September 10 alert had essentially no effect on the media's agenda. With no influence on the media's agenda and a public already worried about terrorist threats, this terror alert had little opportunity to affect economic pessimism.

4.1.1 Mechanisms: Explaining the Increased Pessimism

Explaining the public's increased economic pessimism in response to terror alert is, of course, difficult to do. Indeed, our current design is insufficient to rigorously isolate the mechanism (Imai et al., 2011). That said, we suggest three different reasons that the terror alerts increased economic pessimism.

Stock Market and Volatility The terror alerts not only affect media coverage—they also tended to exercise influence on the stock market. Pagano and Strother (2008) show that there is an increase in the volatility of important stocks, indicative of an increase in perceived risk among investors. Further, this effect is particularly concentrated among knowledgeable investors—those most likely to be sensitive to increases in the market risks. This suggests that the pessimism expressed in the survey data is also expressed with traders in stock markets.

The Economic Consequences of the 9/11 attacks De Boef and Kellstedt (2004) show that both economic and political events influence consumer confidence and MacKuen, Erikson and Stimson (1992) show that the public is forward thinking in their economic evaluations. In particular, De Boef and Kellstedt (2004) show that sudden political events can cause decreases in consumer confidence. Terror alerts are particularly well suited to cause this decrease in confidence. After the 9/11 attacks the U.S. economy entered a deep recession, a recession the public was still anxious about as the alert levels were raised. Therefore, it may have been that the public associated the increased risk of a terror attack with an increased risk of an economic downturn.

Terror Alerts and the Airline Industry The media may have also helped the public connect the terror alerts to the effect of terrorism on the ailing airline industry. To show this, we applied LDA to all newspaper stories and nightly news transcripts that our model classified as about terrorism after an alert level was raised (Blei, Ng and Jordan, 2003). This reveals the distinctive language associated with each alert—such as the threats to the financial sector after the August 1 alert or the threat to transit after the London bombing. But it also shows that the alerts regularly focus, unsurprisingly, on the detrimental effects of the attacks on the airline industry and the often substantial costs. Across alerts, 8.1 percent of stories were about the airline industry—with a maximum of 18 percent. Raising the salience of the airline industry reinforces the serious consequences of terror attacks on the airline industry and the economy more generally.

5 Conclusion

We have offered a new research design, built around the Department of Homeland Security's terror alert system, to show how American democracy responds to terrorist threats. Using a new collection of over 50,000 newspaper stories and nightly newscasts, we demonstrated that the alerts exert substantial influence on the media's agenda, though only briefly. And then using a set of serendipitous surveys that happened to be in the field when the alerts were issued, we demonstrated that the shift in media attention has muted effects on the public. The terror alerts do increase the perceived likelihood of a terror attack and the salience of terrorism, but does not affect the public's policy preferences. Instead, the terror alerts cause the public to be more pessimistic about the economic future.

Our study has wide ranging implications for the public policy of terror alerts. Recently, scholars have worried about the psychological consequences of warning the public about the possibility of a terror attack (McDermott and Zimbardo, 2007). Our study adds an economic concern to the psychological worries. After the 9/11 terror attacks, the economy recovered

slowly and the public's confidence was easily shaken. While we believe the effect of the alerts are short lived, even brief declines in consumer confidence can have real consequences for stock prices (Pagano and Strother, 2008).

Our study also shows that concerns about manipulation using terror alerts were misplaced. Pundits and political scientists worried that the public would be manipulated to advance political interests of the Bush administration. However, we have demonstrated that the public only became more pessimistic about the economy, an outlook not beneficial to incumbent politicians. Our results show that the media both allocated more attention to terrorism after an alert was issued *and* dedicated less space to other policy disputes. This suggests that Howard Dean was partly right in his critique of the Bush administration. Whether intentional or not, the terror alerts and the robust media response resulted in a change in the topic of conversation.

Finally, our study contributes to a social science debate about how the public handles terror threats. Using our more realistic treatments, we find more muted effects than found in either observational studies subject to confounding or experiments that use strong treatments. This, we believe, shows how effective serendipitous surveys can be for understanding how democracies handle the threat from terrorism.

References

Angrist, Joshua D. and Jörn-Steffen Pischke. 2008. *Mostly Harmless Econometrics: An Empircist's Companion*. Princeton University Press.

Blei, David, Andrew Ng and Michael Jordan. 2003. "Latent Dirichlet Allocation." *Journal of Machine Learning and Research* 3:993–1022.

Bloom, Lary and Tom Ridge. 2009. *The Test of Our Times: America Under Siege And How We Can Be Safe Again*. MacMillan.

Boydstun, Amber. 2012. "Puzzles in the News: How Policy Issues Make the Front Page." University of California, Davis Mimeo.

Burden, Barry. 2000. "Voter Turnout and the National Election Studies." *Political Analysis* 8(4):389–398.

Davis, Darren and Brian Silver. 2004. "Civil Liberties vs. Security: Public Opinion in the Context of the Terrorist Attacks on America." *American Journal of Political Science* 48(1):28–46.

De Boef, Suzanna and Paul M. Kellstedt. 2004. "The Political (and Economic) Origins of Consumer Confidence." *American Journal of Political Science* 48(4):633–649.

Feingold, Russ. 2001. "On Opposing the U.S.A Patriot Act." Speech Given to Associated Press Managing Editors Conference.

Gabel, Mathhew and Kenneth Scheve. 2007. "Estimating the Effect of Elite Communications on Public Opinion." *American Journal of Political Science* 51(4):1013–1028.

Gadarian, Shana Kushner. 2010. "The Politics of Threat: How Terrorism News Shapes Foreign Policy Attitudes." *Journal of Politics* 72(2):469–483.

Grimmer, Justin. 2011. "An Introduction to Bayesian Inference via Variational Approximations." *Political Analysis* 19(1):32–47.

Grimmer, Justin. 2012. "Appropriators Not Position Takers: The Distorting Effects of Electoral Incentives on Congressional Representation." *American Journal of Political Science* . Forthcoming.

Grimmer, Justin and Brandon Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* . Forthcoming.

Haller, H. Brandon and Helmut Norpoth. 1994. "Let the Good Times Roll: The Economic Expectations of US Voters." *American Journal of Political Science* 38(3):625–650.

Hetherington, Marc. J. and Elizabeth Suhay. 2011. "Authoritarianism, Threat, and Americans' Support for the War on Terror." *American Journal of Political Science* 55(3):546–560.

Ho, Dan, Kosuke Imai, Gary King and Elizabeth Stuart. 2007. "Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference." *Political Analysis* 15(3):199–236.

Holland, Paul. 1986. "Statistics and Causal Inference." *Journal of the American Statistical Association* 81(396):945–960.

Huddy, Leonie, Stanley Feldman, Charles Taber and Gallya Lahav. 2005. "Threat, Anxiety, and Support of Antiterrorism Policies." *American Journal of Political Science* 49(3):593–608.

Imai, Kosuke, Luke Keele, Dustin Tingley and Teppei Yamamoto. 2011. "Unpacking the Black Box of Causality: Learning about Causal Mechanisms from Experimental and Observational Studies." *American Political Science Review* 105(4):765–789.

Jacobson, Gary C. 2007. *A Divider, Not a Uniter: George W. Bush and the American People, The 2006 Election and Beyond*. Longman.

Jordan, Michael et al. 1999. "An Introduction to Variational Methods for Graphical Models." *Machine Learning* 37:183–233.

Kam, Cindy D. and Donald R. Kinder. 2007. "Terror and Ethnocentrism: Foundations of American Support for the War on Terrorism." *Journal of Politics* 69(2):320–338.

MacKuen, Michael B., Robert S. Erikson and James A. Stimson. 1992. "Peasants or Bankers? The American Electorate and the US Economy." *American Political Science Review* 86(3):597–611.

Mahlhotra, Neil and Elizabeth Popp. 2012. "Bridging Partisan Divisions Over Antiterrorism Policies: The Role of Threat Perceptions." *Political Research Quarterly* . Forthcoming.

Manning, Christopher et al. 2008. *Introduction to Information Retrieval*. Cambridge University Press.

McDermott, Rose and Philip G. Zimbardo. 2007. The Psychological Consequences of Terrorist Alerts. In *Psychology of Terrorism*, ed. Bruce Bongar et al. Oxford University Press.

Merolla, Jennifer L. and Elizabeth J. Zechmeister. 2009. *Democracy at Risk: How Terrorist Threats Affect the Public*. University of Chicago Press.

Monroe, Burt, Michael Colaresi and Kevin Quinn. 2008. "Fightin' Words: Lexical Feature Selection and Evaluation for Identifying the Content of Political Conflict." *Political Analysis* 16(4):372.

Nagourney, Adam and David M. Halbfinger. 2004. "Race Dogged by Terror Fight." The New York Times.

Norpoth, Helmut. 1996. "Presidents and the Prospective Voter." *Journal of Politics* 58(3):776–792.

Pagano, Michael S. and T. Shawn Strother. 2008. "Risk, Uncertainty, and the Perceived Threat of Terrorist Attacks: The Advisory System that Cried Wolf?" Villanova University School of Business Mimeo.

Quinn, Kevin et al. 2010. "How to Analyze Political Attention with Minimal Assumptions and Costs." *American Journal of Political Science* 54(1).

Reese, Shawn. 2003. "Homeland Security Advisory System: Possible Issues for Congressional Oversight." Congressional Research Service Report. Available at: http://www.au.af.mil/au/awc/awcgate/crs/rl32023.pdf.

Shadish, William R., Thomas D. Cook and Donald T. Campbell. 1978. *Experimental and Quasi-Experimental Designs for Generalized Causal Inference*. Wadsworth.

Taddy, Matt. 2013. "Multinomial Inverse Regression for Text Analysis." *Journal of the American Statistical Association* . Forthcoming.

Wallach, Hanna. 2008. Structured Topic Models for Language PhD thesis University of Cambridge.

Wallach, Hanna, Lee Dicker, Shane Jensen and Katherine Heller. 2010. An Alternative Prior for Nonparametric Bayesian Clustering. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics (AISTATS)*. Vol. 9.

Willer, Robb. 2004. "The Effects of Government-Issued Terror Warnings on Presidential Approval Ratings." *Current Research in Social Psychology* 10(1):1–12.

Zaller, John. 1992. *The Nature and Origin of Mass Opinions*. Cambridge University Press.

Tabitha Bonilla (Ph.D. Candidate, Stanford University) is a graduate student in the Stanford University department of political science. Her research interests include in political campaigns.

Justin Grimmer (Ph.D., Harvard University) is an Assistant Professor of Political Science at Stanford University. His research examines political representation and communication in American politics.