

The Mediating Effect of the Media in the 2004 Presidential Election: The Case of the “Swift Boat” Ads

Scott Ashworth*
Joshua D. Clinton†

May 3, 2011

Abstract

Does media coverage of political advertisements influence voter opinion more than the advertising campaigns themselves? Ads may influence election outcomes either through the exposure resulting from the ad buys made by the political participants or through the indirect exposure that results from the media’s coverage of campaign events and so-called “ad-watches.” Perhaps the most famous example are the advertisements of the 527 organization “Swift Boat Veterans for Truth” which produced ads during the 2004 presidential election. So much so, in fact, that the verb “swift-boating” is now routinely used when discussing campaign strategies. Using both individual-level data as well as data from state-level polls across time we isolate the effect attributable to primary and secondary airings. We show that to the extent that the ads affected the outcome, it was only because of the coverage that was given by cable news to the ads. 2% of the change in national opinion can be attributed to the coverage of cable news, which means that 27 % of the overall move towards Bush over the course of the election is due to the coverage given by cable news.

*Associate Professor, Harris School of Public Policy, University of Chicago. Email: sashwort@uchicago.edu

†Associate Professor, Department of Political Science, Vanderbilt University. Email: josh.clinton@vanderbilt.edu

During the 2004 Presidential campaign, candidate John Kerry's military record was strongly attacked in ads produced by a 527 group called the Swift Boat Veterans for Truth. The press quickly decided that the ads took a major toll on Kerry's campaign. *The Wall Street Journal* wrote:

cable news networks jumped on the story and gave the ads heavy play, even as they probed the veracity of the commercials. Meanwhile, the cheaply-funded effort helped bring down Mr. Kerry's poll numbers (Cummings and Flint, 2004),

while *The New York Times* wrote:

[t]he advertisements are one example of how Republican groups have used incendiary commercials...to propel themselves to the forefront of the presidential race in weeks (Justice and Rutenberg, 2004).

This paper asks if these claims are correct. In particular, we ask: what are the maximum effects attributable to the direct broadcast of these ads? Is there any evidence that these ads would have been effective were it not for the fact that the news media and political pundits focused so heavily on these ads? If the latter, is the effect driven simply by the provision of information about the campaign, or can the effects be tied to more specific facts about media presentation?

The popular wisdom that the Swift-Boat campaign was an important factor in Kerry's defeat exemplifies a more general belief that elections are affected by highly salient media events, with examples ranging from Nixon's uncomfortable appearance in the 1960 debates to LBJ's 1964 "Daisy" ad to the "Willie Horton" ad in 1988. But political scientists tend to be skeptical that such events have sizable effects. The old conventional wisdom is that campaigns have only a "minimal effect" (e.g., ? and ?). Recent work has emphasized that these findings apply to the ability of campaigns to change voters' basic preferences, leaving open the possibility that more subtle mechanisms like agenda setting or priming can affect the way voters weight their basic preferences, and thus their vote choices. Iyengar and Kinder

(???) demonstrate such effects in an elegant series of experiments, but scholarship based on field data is mixed, with a majority of the work finding a small effect (e.g., ?, ?, ?, and ?).

The belief that campaign effects are small is bolstered by the finding the finding that election results are highly predictable ex-ante, without taking campaign events into account (Gelman and King ???). While this work is convincing about the aggregate effect of campaigns, it is harder to interpret in terms of the individual campaign events that journalists make their living discussing. In particular, this literature cannot distinguish between the hypotheses that individual campaign events have small effects and the alternative, that events have large but usually offsetting effects.

Finally, most of this work has difficulty identifying the mechanism through which campaign events matter (to the extent that they do)—is it due to the events themselves or the way that event is covered by commentators and journalists, as suggested by the *Journal* writers? The latter question is particularly important in thinking about the media’s role as “watchdog”—if events are effective only because of the choice made by reporters and journalists rather than the choices made by the candidates themselves, then media “watchdogs” may actually be counterproductive in terms of policing the campaign airwaves and that so doing may actual undermine their goal of keeping the campaigns “honest.”¹ (See Ansolabehere and Iyengar ??? for more discussion of this ripple effect of negative advertising.)

The Swift Boat campaign is a particularly useful case in which to examine media effects. The conventional wisdom is that the campaign had a sizable impact, and that the Kerry campaign had a very difficult time parrying the attacks. Thus a convincing null effect here would be a strong falsification of the large-media-effects hypothesis. In addition, the precise strategy used by the Swift Boat Veterans for Truth created several sources of variation in voter exposure to their claims. The ad buys themselves were small, concentrated in just four states, so there is geographic variation in potential exposure. These small ad buys were

¹We obviously cannot assess the counterfactual of would have happened had the ads not been run at all, or identify the effect of the ads due solely to the discussion of journalists and political commentators. Instead, we can only identify the effect of running the ads on the air versus having them discussed by pundits and journalists.

designed to attract media coverage, and were massively successful in doing so. Thus regular consumers of newspapers and television news were more likely to be exposed to the message than other voters. Finally, the ads and resulting media storm were in August, leaving plenty of time both before and after the ads for a differences-in-differences approach to control for unobserved differences in vote intention confounded with geography of media consumption.

We find evidence for a sizable effect of the campaign, with our preferred estimate attributing over two thirds of Kerry’s decline relative to Bush over the period ??? to ??? to cable news discussions and rebroadcasts of the ads. There is no detectable effect of the actual ad buys themselves—the population targeted by the purchased ad buys did not noticeably change their opinion more or less than other voters. There is also no effect among individuals closely following the campaign through the print media, even though there was substantial coverage of the ads by print journalists and commentators. The fact that the effect shows up for cable viewers but not for paper readers is suggestive of an important difference in media effects by media type.

The paper is organized as follows. Section 1 provides background for the case study, including a summary the ad campaign of the 527 organization *The Swift Boat Veterans for Truth* and its coverage in newspapers and on cable news. Section 2 details the data we use, our identification strategy, and the effect of bias in self-reported exposure measures. Section 3 estimates the effect of direct and indirect exposure and calculates the amount of the increased Bush support that is attributable to primary and secondary exposure to the Swift boat ad campaign. Section 4 concludes.

1 The Case

The Swift Boat Veterans for Truth (SBVT) was a political group composed primarily of veterans of the Navy from the Vietnam war. They organized to oppose the Presidential bid of John Kerry, who was also served in the Navy during the war. As has become typical

since the McCain-Feingold campaign finance bill, the group organized as a 527 group—named for a provision of the tax laws that allows groups independent of political parties or candidates to raise and spend money outside the regulatory orbit of the FEC, so long as they do not explicitly advocate the election or defeat of a specific candidate. Consistent with this limitation, the group did not explicitly advocate for the election of Kerry’s opponent, George W. Bush; instead they criticized both Kerry’s activities while on active duty and his subsequent statements about the war.

The SBVT’s primary intervention into the campaign was a pair of ad buys on broadcast television.² The first ad, called “Any Questions?”, presented testimony by veterans, many of whom claimed to have “served with John Kerry”. Some examples of the claims: “John Kerry has not been honest” and “When the chips were down, you could not count on John Kerry” (Quotations from the transcript available at ???.) The second ad, called “Sellout”, claimed that Kerry lied about being in Cambodia in 1968.

Laboratory evidence suggests that “Any Questions” in particular was an effectively crafted ad. During the campaign, Borick (???) showed 19 advertisements from the 2004 election campaign to groups of voters, eliciting vote intentions both before and after viewing. Before viewing “Any Questions”, 79% of independents said they were either definitely or most likely going to vote for Kerry. After viewing, only 39% answered definitely or most likely. The percentage not sure if they would vote for Kerry increased from 0% to 28%. None of the other ads Borick examined were remotely as effective; the other ads induced an average of less than 2% of voters to move from supporting a candidate to being unsure, with the second largest change being only 9% (Borick, Table 11).

Effective as “Any Questions” was in the laboratory, the ads did not air enough to affect too many voters. Table 1 reports the ads only appeared on broadcast television in four states, with relatively low ad buys. Even so, the ads were targetted for maximum exposure. The *New York Times* reported that the strategy of the advertising campaign, according to

²The group produced four ads in total, but the others were never broadcast on broadcast television and were instead released directly to the internet.

one of its leaders, John O’Neil, was to make “small advertising purchases in states with large populations of veterans, tries to time the commercials to Kerry campaign stops, and has its members appear in those states at the same time” (Justice and Rutenberg).

Ad Title	Air Data	Ad Buy	States
“Any Questions”	August 5, 2004	\$550,000	Ohio, Wisconsin, and West Virginia
“Sellout”	August 24, 2004	\$800,000	Iowa, Ohio, Wisconsin, and cable

A spokesperson for the group said that “the strategy from the get-go” was to make small ad buys to win free time on talk shows. “Everyone was surprised as to how well it worked.” (*Wall St. Journal*, Sept. 17, 2004). Figure 3.2 documents this success. It plots the number of stories and news programs mentioning the phrase for each week starting the week of July 7, 2004 and continuing until the week after Election Day.

Press coverage of the Swift boat ads peaked during the week of August 22nd to August 28th – the week immediately preceding the 2004 Republican convention – with 275 news stories and 190 broadcasts mentioning the ad campaign. Because the peak coverage occurred prior to the Republican National Convention, opinions in the week of the 22nd are only possibly affected by the Swift Boat coverage, whereas opinions in the week of August 29th may be affected by both the coverage and the Republican National Convention.

Although much of the coverage was negative towards the ads, it was roughly coincident with a decline in Kerry’s support. Figure 2 graphs the support for Kerry according to the 2004 Annenberg Rolling Cross-Section. The area shaded in grey are the periods of the most intense media coverage of the Swift Boat campaign.

Noting the dramatic increase in support for Bush immediately following the airing of the ads, many in the media blamed the ads and resulting coverage for Kerry’s decline. For example, in an analysis of Kerry’s decline, the *LA Times* wrote: “Although a solid majority of Americans say they believe Kerry served honorably in Vietnam, the poll showed that the fierce attacks on the senator from a group of Vietnam veterans criticizing both his performance in combat and anti-war protests at home have left some marks... [N]ow

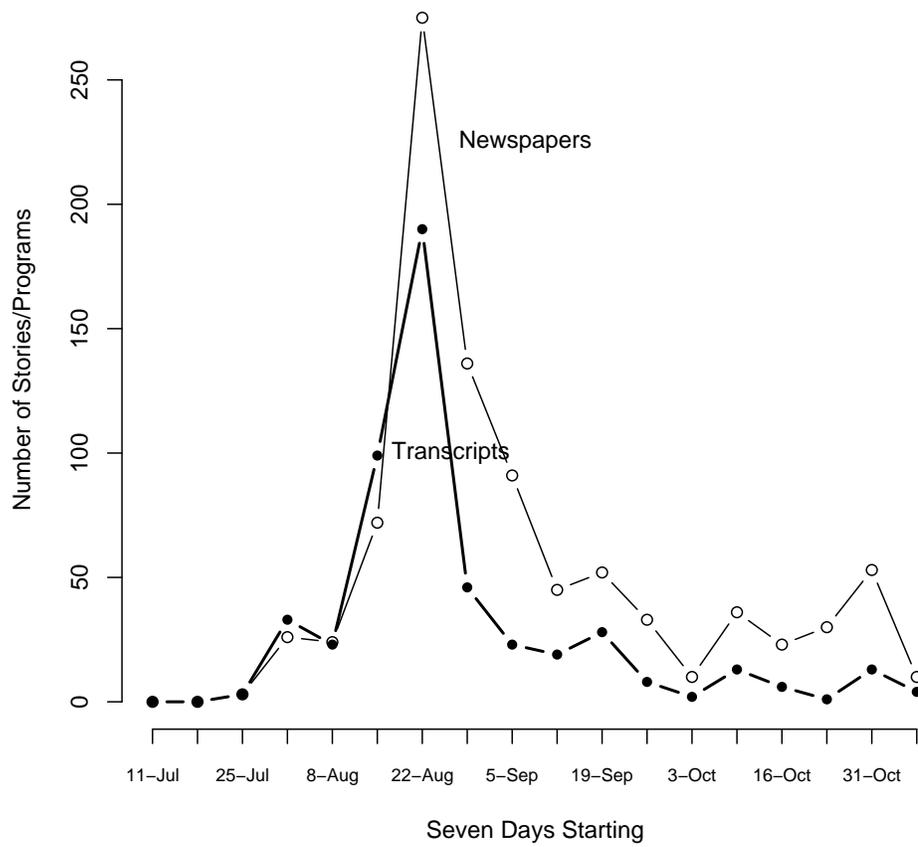


Figure 1: NUMBER OF NEWSPAPER STORIES AND NEWS BROADCASTS MENTIONING “SWIFT BOAT VETERANS FOR TRUTH. Thick line denotes number of news transcripts mentioning the phrase and thin line denotes the number of newspaper stories.”

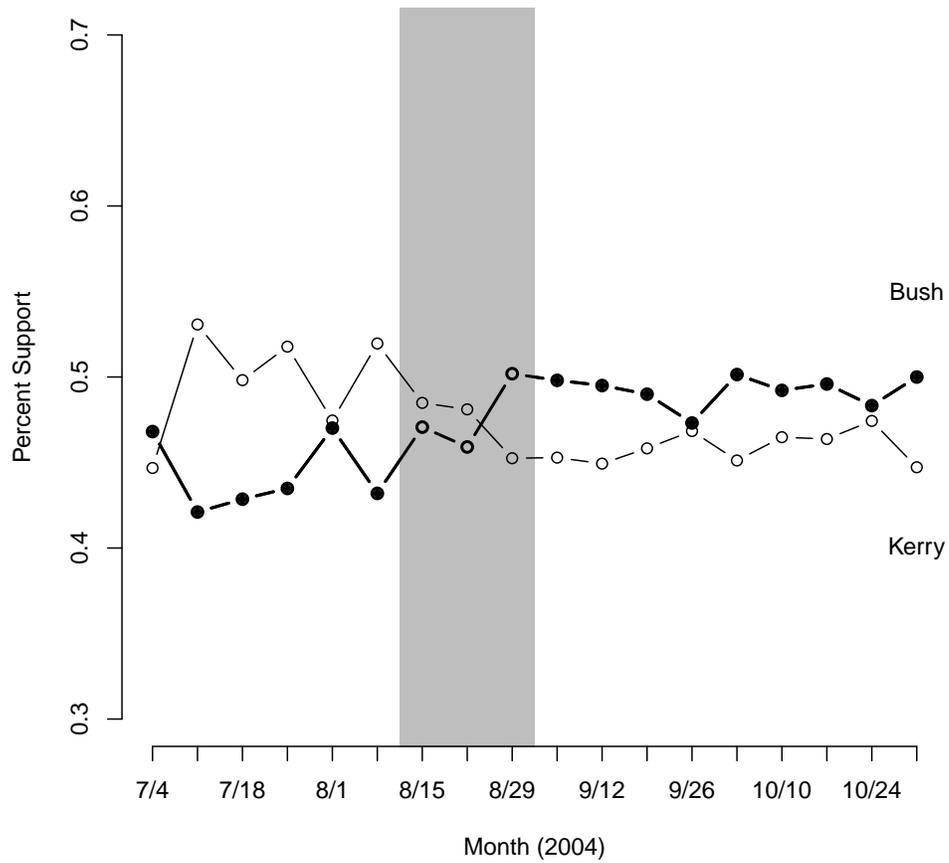


Figure 2: SUPPORT FOR BUSH AND KERRY IN THE ANNENBERG ROLLING CROSS-SECTION BY WEEK

that the SwiftVets have begun their assault on Kerry he seems clueless on how to respond. His entire campaign is struggling just to answer the SwiftVet charges. And while Kerry is distracted with that problem, the president has been using the power of the incumbency to highlight his record.”

2 Identification Strategy

The simplest way to estimate the effect of the Swift Boat ads is to compare the behavior of voters who report seeing the ads (or knowing about them) to voters who do not. It’s well-known that such a comparison is likely to be biased. Voters who end up seeing controversial campaign ads are likely to differ from voters who do not in many ways—being more interested in politics, more likely to live in battleground states, etc. Our strategy uses two features of the data to deal with these issues. First, we are able to divide voters into subgroups based on characteristics that make them more or less likely to be exposed to the Swift Boat campaign: voters who live in the states where the ads aired vs. those who do not, voters who regularly read a newspaper vs. those who do not, and voters who regularly watch cable news vs. those who do not. Second, we use high-frequency data so that we can examine any differences that show up between the media-defined subgroups are focused on the sharp increase in Swift-Boat content in mid August, while controlling for persistent differences between the groups.

2.1 Data

Our main analysis uses data from the Annenberg National Election Study’s rolling cross-section collected during the 2004 campaign.³ To avoid possible question wording effects, we only use responses to the question that is identical before and after the Swift Boat ad

³The panel component did not re-interview respondents on a schedule that is useful for our identification strategy.

campaign in the analysis that follows.⁴ We classify individuals according to whether they reside in a state where a Swift Boat ad was broadcast on network television, and whether or not they self-report watching cable news or read a newspaper at least once a week.⁵ Putting aside the well-known problems associated with self-reported media usage for the moment (Vavreck, Prior), this yields four groups of voters – those who watch cable news and read a newspaper, those who neither watch cable news or read a newspaper, and those that consume just one of the two sources. According to the respondents, 54% read a newspaper and watched cable news at least once in the last week, 23% read a newspaper but did not watch cable news in the last week, 13% watched cable news but did not read a paper in the last week, and 11% admitted that they consumed neither.

The survey also asks about partisanship, gender, age, age², education level, household income, race, gender, self-reported ideology, views on gay marriage, perception of personal economic situation, whether the individual is a born-again christian, and whether anyone who is, or has been, in the military lives in the household.

We also exploit data from every state poll conducted between May 1, 2004 and October 31, 2004. These data are from the collection of polls available on the *National Journal's* online poll archive for the 2004 presidential election. To eliminate potential “house effects”—i.e., the idiosyncratic factors that are associated with the sampling and weighting of the data associated with each polling firm—we collect only those state-level polls for which there is a poll before and after the August broadcast of the Swift Boat Veterans for Truth ad buys.

⁴From 7/6 until 11/1 the question wording included the names of the vice-presidential candidates as well. (Question cRC14: “ If the 2004 presidential election were being held today, would you vote for George W. Bush and Dick Cheney, the Republicans; John Kerry and John Edwards, the Democrats; or, Ralph Nader ad Peter Camejo of the Reform Party?”) From 4/30 until 7/20 the question included only the names and parties of the Presidential candidates. (Question cRC13: “If the 2004 presidential election were being held today, would you vote for George W. Bush, the Republican; John Kerry, the Democrat; or Ralph Nader?”) Between 7/6 and 7/20, the sample was split between each question wording. Conducting the analysis using both questions prior to the ad campaign does not alter the conclusions.

⁵The precise question wording was: “How many days in the past week did you watch a 24-hour cable news channel, such as CNN, FOX News Channel, or MSNBC?” Of the 55,579 valid responses, 33.83% indicate that they did not watch cable news last week. Insofar as the likely response bias is to inflate new watching because of social desirability, this is probably an underestimate of the percentage who were not exposed to cable news.

We only use state-level polls from the same survey organizations using samples of “likely voters.” The definition of “likely voters” is specific to each polling firm, but using the same polling firm and differencing the before and after polls should successfully control for possible definitional differences.

2.2 Identification

Our approach to identification uses subgroups defined by characteristics that should affect their likelihood of being exposed to the Swift-Boat campaign. To see how this works, consider the case of two groups, $g \in \{1, 2\}$, and two time periods, $t \in \{b, a\}$ (*before* and *after*). Write Y_{gt} for the Bush-Kerry margin in group g at time t . Consider a simple linear model for this margin:

$$Y_{gt} = \mu_g + \delta_t + \gamma\rho_{gt},$$

where

- μ_g is a group fixed effect,
- δ_t is a time-period fixed effect,
- ρ_{gt} is the fraction of group g exposed to the Swift-Boat campaign at time t , and
- γ is the conversion rate of the campaign.

Of course, $\rho_{gb} = 0$ for both groups. Thus a differences-in-differences estimand is purged of both the group-specific and time-specific effects:

$$\begin{aligned} DID &\equiv \mathbb{E}[(Y_{1b} - Y_{1a}) - (Y_{2b} - Y_{2a})] \\ &= (\delta_t + \gamma\rho_{1a}) - (\delta_t + \gamma\rho_{2a}) \\ &= \gamma(\rho_{1a} - \rho_{2a}). \end{aligned}$$

The first round of differencing removes any time-invariant confounders. If, for example, one group lives in a more active media market, that might affect the perception of particular candidates, but because that environment is held fixed within each group, differencing identifies the amount of change that can be attributed to the ad buys themselves. The second round of differencing removes the effect of time-specific shocks affecting all voters.

A sample analogue of the DID estimand can be easily estimated via regression. Index individuals by i , with $g(i)$ denoting the group person i is a member of. Define

$$y_{it} = \begin{cases} 1 & \text{if person } i \text{ intends at } t \text{ to vote for Bush} \\ 0 & \text{if person } i \text{ has no vote intention at } t \\ -1 & \text{if person } i \text{ intends at } t \text{ to vote for Kerry.} \end{cases}$$

Then the group g -time t sample mean \bar{y}_{gt} is an unbiased and consistent estimator of Y_{gt} , and the sample analogue of *DID* is numerically identical to the coefficient β_3 in the regression

$$y_{it} = \beta_0 + \beta_1 G_i + \beta_2 T_t + \delta \cdot G_i \cdot T_t + \epsilon_{it}$$

where the regressors are the dummies $G = \mathbb{I}[g = 1]$ and $T = \mathbb{I}[t = \text{after}]$.

This example has been limited to just two groups and two time periods, with no covariates. But it is easy to extend the regression formulation to remove all of those limitations. Our estimating equation does just that. Let I_t^T be a dummy variable for times after T :

$$I_t^T = \mathbb{I}[t \geq T].$$

Also let \mathbf{x}_{it} be a vector of covariates. Then we estimate

$$y_{it} = \beta_0 + \beta_1 G_i + \beta_2 I_t^T + \beta_3^T \cdot G_i \cdot I_t^T + \eta' \mathbf{x}_{it} + \epsilon_{it}.$$

The estimate $\hat{\beta}_3^T$ for T equal to the time of the Swift-Boat campaign is our estimate of the *DID*-effect.

This approach relies on two important assumptions. First is the common-trend assumption: absent the Swift-Boat campaign, the different groups would have followed the same changes in margin over the course of the campaign. (This is the upshot of assuming all the time-specific shocks are common to all groups.) We take a couple of approaches to evaluate this assumption, both of which exploit the high frequency of surveys both before and after the Swift-Boat campaign. First, we use graphs to see if the groups look similar prior to the Swift-Boat campaign. Second, we estimate a rich set of placebos, estimating β_3^T for the “wrong” T . If the common trend hypothesis is violated before the SBVT came to prominence, some of these placebos will show significant effects.⁶

The power of the difference-in-differences design is that it controls for the known-differences in the types of media consumers that exist (Prior XXXX). Intuitively, we are simply comparing how the relative preference for Bush over Kerry changes within different groups (organized by their media consumption) and whether it changes more for some groups than others in ways that are consistent with an effect that can be attributable to the Swift Boat ads. This is problematic if the composition of these groups change over the course of the election in ways that are confounded with the swift boat campaign. For example, if the set of respondents to the weekly cross-section became more Republican, or more likely to consume cable news around the time of the “treatment” (perhaps in response to the treatment) this would violate the common trend. The appendix reveals that this is not the case – the characteristics of the cross-section are unchanged over the course of the campaign. The appendix also reveals that the characteristics of the media consumers does not change over the course of the election – for example, self-reported cable news viewers are not more likely to be Republican after the Swift Boat campaign.

⁶This is related to the randomization-inference approach of Abadie et al REF.

The second important assumption is constancy of the conversion rate across groups. Imagine for example that the campaign had no effect on cable news viewers, but that it moved non-viewers who happened upon it strongly in Kerry’s direction. Also assume that cable viewers were more likely to be exposed to the campaign. Then the *DID* estimate would be positive for Bush, even though no one actually moved in Bush’s favor when exposed to the campaign.

2.3 The Effect of Self-Reported Media Consumption Bias

It is well known that individuals overreport their consumption of the media. It is implausible that 67% actually watched cable news at least one day, and 77% read a newspaper at least once during the week as the self-reported exposure measures indicate. The relevant question for us, however, is how does the bias in self-reported exposure affect our ability to estimate the effects that may be due to the various media outlets?

It can be shown that the well-documented bias in self-reported consumption measures does not affect our ability to estimate the effects of interest. Suppose that the total change that we are trying to explain is given by Δ_{All} and the change in cable news viewers is given by Δ_{Cable} . If the true fraction ρ of the population actually watches cable news, to calculate the percentage of the total change Δ_{All} that is due to cable viewers we need to calculate:

$$\frac{\rho \times \Delta_{Cable}}{\Delta_{All}}.$$

Suppose that we do not know the actual proportion instead use self-reported consumption p to measure ρ and $p > \rho$ due to the tendency of individuals to overreport their actual exposure and we instead calculate $\frac{p \times \Delta_{Cable}}{\Delta_{All}}$. Despite the bias in the exposure estimate p , this is the correct estimate of the effect of cable news exposure because the over-reporting bias cancels out in the numerator. To see this suppose that w denotes the probability that an individual who claims to be a cable news watcher is an actual viewer so that $p \times w = \rho$ gives the true exposure (which may be unknown). This means that the true fraction of cable news viewers is $p \times w$ (rather than just p), but also that effect Δ_{Cable} is incorrect because

some of the estimated effect is due to non-watchers saying they are cable news viewers when they are not. As a consequence, the estimated effect of exposure must also be corrected to account for the over reporting using $\frac{\Delta_{Cable}}{w}$. These corrections exactly cancel and we can calculate the estimated effect *even if the magnitude of the bias w is unknown*.

The implied calculations and cancellations imply that the fraction of the total effect that is attributable to cable news viewers is given by:

$$\frac{p \times (\Delta_{Cable})}{\Delta_{All}} = \frac{p \times w \times \left(\frac{\Delta_{Cable}}{w}\right)}{\Delta_{All}}.$$

The same logic applies to the fact that not every respondent who watches cable news will have been exposed to the coverage of the Swift boat ads on cable news. If s denotes the percentage of cable news watchers who were exposed to the Swift Boat ads during the course of watching cable news, the numerator is $p \times w \times s \times \left(\frac{\Delta_{Cable}}{w \times s}\right)$ which again reduces to just $p \times \Delta_{Cable}$. Using the noisy measure of media consumption therefore provides the correct estimate even though not every self-reported viewer actually watches cable news or watches cable news coverage of the Swift boat ad campaign.

3 Results

There are two ways that ads may reach an audience. There is the direct effect of the ad being broadcast on television and potentially influencing individuals who were watching the programs on which the ads were shown. The second way that ads may be effective is through the indirect effect they have through their coverage on secondary sources. This may occur either by affecting how the campaign is covered by the press (e.g., the way coverage of the campaign is framed and the issues that the media focuses upon), or by giving coverage and exposure to ads to a broader audience than they might otherwise receive.

More concretely, the Swift Boat ad campaign may have affected the media coverage of the campaign by helping to create an emphasis on certain issues more than others. If so,

this would presumably affect both the print and broadcast media alike. An effect unique to the broadcast media – and perhaps especially for cable news networks given the amount of programming they produce – is that the broadcast news may rebroadcast the ads themselves and effectively increase the audience share of the ads beyond the potentially limited audience living in the direct buy media markets. Alternatively, the format in which such discussions take place – visual rather than text – political debate rather than story may produce different effects.

Existing research does not disentangle these two mechanisms. Looking only at the correlations of ad buys and voter behavior does not reveal whether there might be different effects – e.g., does the media mediate or even magnify the effect of ads – because the different ways that citizens may become aware of the ads is unknown.

3.1 Direct Broadcast Effects

We first take up the direct effects of the ads, asking if exposure to them on broadcast television moved voters in the direction of supporting Bush. We have leverage on this question because of the geographic restrictiveness of the ad buy—residents of only three states were potentially exposed to the more effective “Any Questions” ad.

For a first cut at answering this question, we use the state polling data described above. The ads aired in August, so we compare Bush and Kerry support in September to support in July. Figure 3 graphs the estimated change in support for Bush and Kerry for each poll pairing accounting for the margin of each using the bootstrap.⁷

Although it is true that Figure 3 reveals that the support for Bush increased between July and August in every one of the six poll pairs conducted in states where the Swift Boat ads were run, only one is statistically distinguishable from zero at conventional levels. Moreover, although the support for Kerry decreases between July and August in the states where paid

⁷We use the reported margin of error and the sample size, along with the assumption of independent normal errors to generate the implied errors associated with each poll to account for the sampling error when differencing the polls.

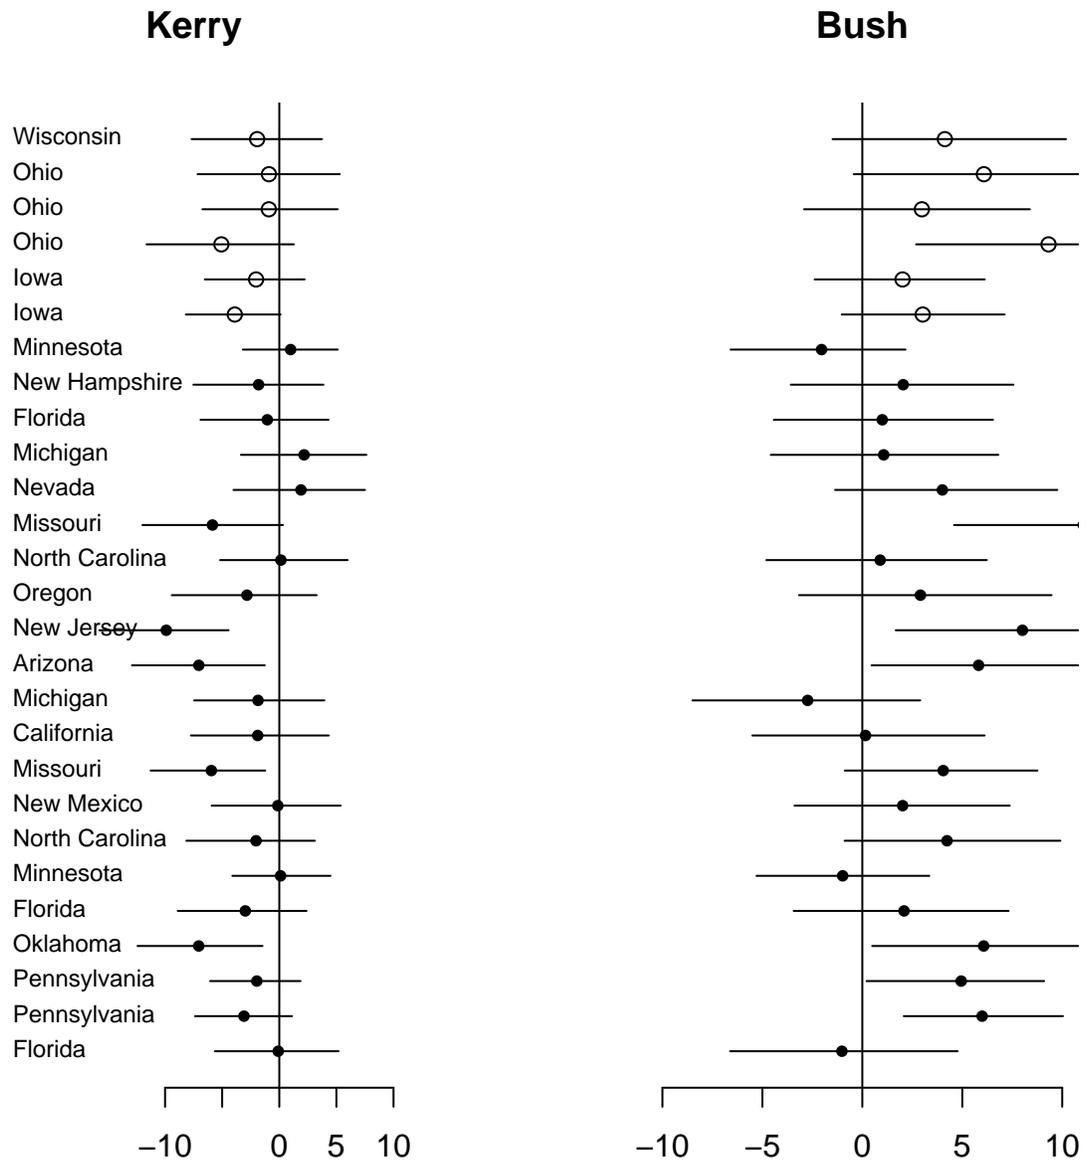


Figure 3: CHANGE IN CANDIDATE SUPPORT BETWEEN JULY AND SEPTEMBER FOR STATE-LEVEL POLL PAIRS. Open (solid) circles denote change in poll pairs in states where Swift Boat ads did (did not) air during August. Lines denote 95% Confidence intervals accounting for margin of error in polls.

ad buys occurred, no change is statistically significant.

More importantly, these changes are similar to those of many of the unexposed states. We calculate the average effect using the 6 state poll pairs in which the ads were run and compare it to the average effect in the 21 poll pairs in which the ads were not run.⁸ The difference in the average support for Bush in states where the ad was and was not directly broadcast is 1.69, with a 95% confidence interval of -0.74 and 4.16. The difference in Kerry's support is -0.09, with a 95 % confidence interval of -2.67 and 2.58. The excess move against Kerry in the exposed states is effectively zero. The estimated excess move for Bush is positive, but we cannot rule out that the campaign had no effect.

The individual-level analysis tells the same story. Using the identification strategy of section 2.2 and the individual level data from the rolling cross-sections conducted by the Annenberg National Election Study over the course of the campaign, we can estimate how the changes in opinion for the various groups of media consumers compare before and after various points in the campaign. The table of regression coefficients are reported in the appendix, but the effect of residing in a state where the ads were aired is plotted in Figure 4

Figure 4 reveals that the lack of effect from the direct ad buys evident in Figure 3 is confirmed using individual level data. There is no effect that can be attributed to the direct exposure to the Swift Boat ad campaign because residents of these states did not experience more of a pro-Bush change in opinion than non-residents following the airing of the ads. To the extent that the ads affected the election, their impact must therefore be through means other than the exposure that the direct ad buys provided.

⁸To do so requires accounting for two types of uncertainty: the margin of error associated with each polling estimate, and uncertainty at the aggregate level that results from averaging the effect over all poll-pairs that are, or are not, in an ad-buy state. We account for these two levels of uncertainty using the bootstrap. That is, we generate bootstrap samples of each poll accounting for the associated margin of error, we calculate the pre-post difference using these bootstrap samples, and then we bootstrap the average difference in the average change in the states who are and are not exposed to the direct ad campaign.

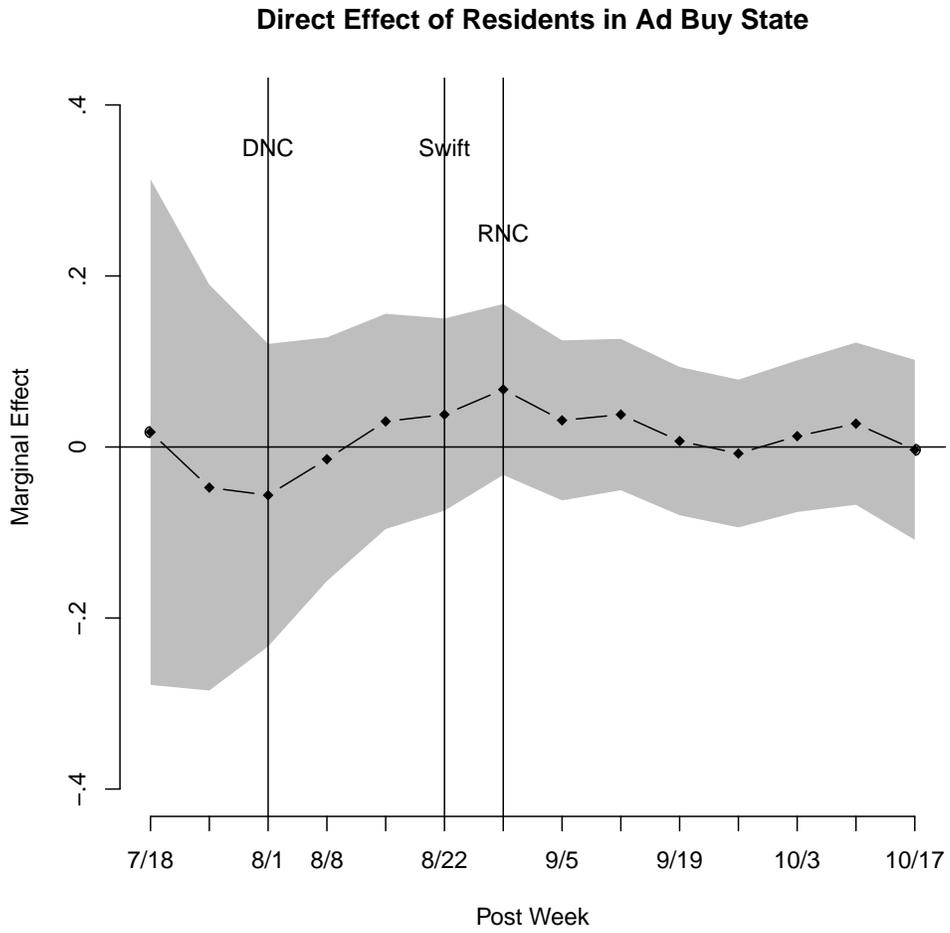


Figure 4: THE DIFFERENCE IN THE DIFFERENCE FOR SUPPORT FOR BUSH BY WEEK FOR RESIDENTS OF AD BUY STATES.

3.2 Media Effects

Next we turn to the effects of the media's coverage of the Swift-Boat campaign. We look at two subgroups of media consumers: newspaper readers and cable news viewers and those that consume both and neither. It's instructive to start with a graphical approach. Figure 5 graphs the average margin for Bush among the sample of respondents who were interviewed during the noted week by the Annenberg National Election Survey. The graph presents how the weekly average margin changes across during the campaign from the seven days starting with July 7, 2004 to the seven days ending on November 7, 2004 for respondents grouped by their self-reported media coverage.

As Figure 3.2 revealed, coverage of the Swift Boat ad campaign by both the print and cable news media peaked during the week of August 22nd. If there are media effects on voter behavior that are driven by media coverage of the Swift Boat ad campaign, we should observe the largest effects following the extensive coverage.

A more formal comparison comes from regressions. If Swift Boat ads have an affect through newspapers or cable news programs, we should see differences in the support for Bush emerge following the advertising campaign once coverage of the ads commenced – especially during the week of August 22nd which was the week of peak media coverage. As section 2.2 reveals, the effect of interest is revealed by comparing the coefficient on the interaction of media consumption and an indicator denoting how post-exposure opinions differ from pre-exposure opinions between consumer and non-consumers. Table 1 presents the estimated effects for four groups: Readers only, Watchers only, Readers and Watchers, and neither Readers nor Watchers to isolate the effect of cable news and newspaper coverage given the fact that some respondents consume both media sources and the effects may differ. For example, the coefficient on $Post \times Cable\ News$ in Table 1 reveals how the average support for Bush relative to Kerry differs in the pre and post period relative to the average support of individuals who consume neither cable news or newspapers at least once a week.

For robustness, we estimate the relationship varying the week used to denote the start

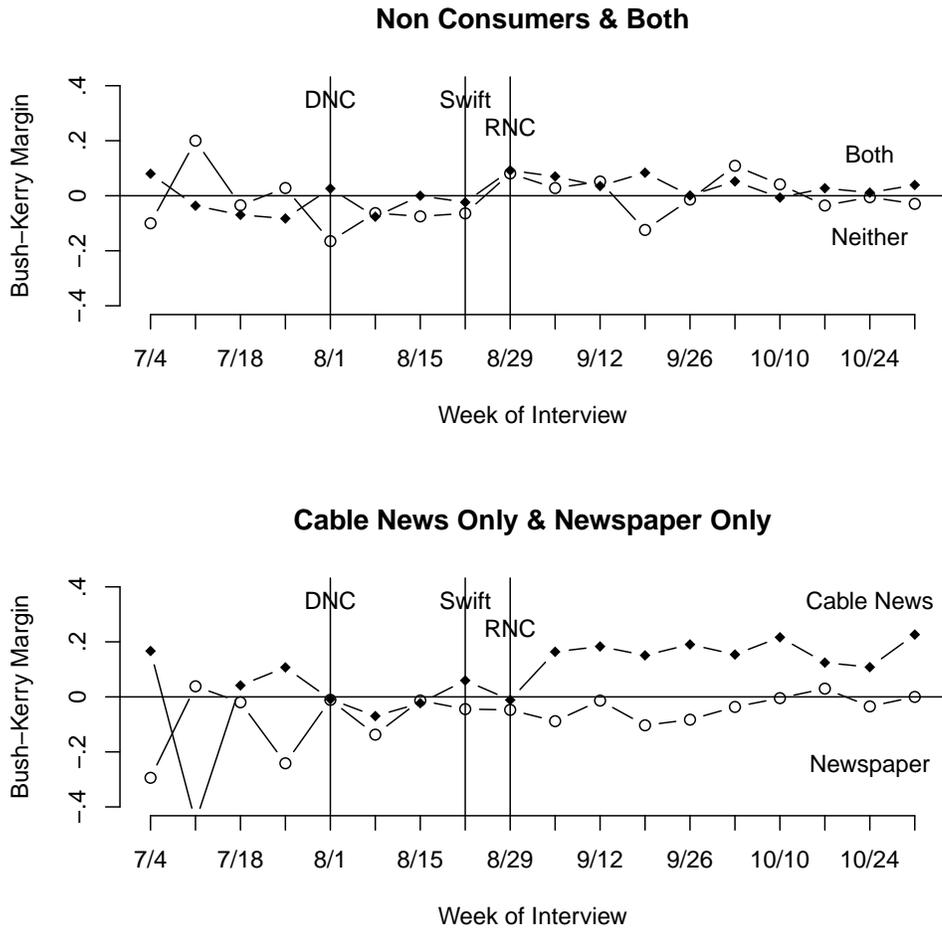


Figure 5: THE AVERAGE MARGIN FOR BUSH BY WEEK BY MEDIA CONSUMER TYPE. The top figure graphs the average margin for consumers of both cable news and newspapers (solid) and consumers of neither. The bottom figure graphs the average margin for consumers of just cable news (solid) and just newspapers (open). The averages are based on rolling cross-sectional surveys conducted in each week (starting with the week of 7/4) surveys from a rolling cross-section.

of the post-exposure effect. If the Swift Boat ads have an effect, the effect should be the largest for the week following the actual airing of the ad. If there is a difference prior to the broadcast of the ads, the effect obviously cannot be attributed to coverage of the Swift Boat ads. Model 1, for example, compares the change in opinion before and after August 8th. This is essentially a placebo test, although it is somewhat confounded by the fact that an active election season means that events or varying significance are occurring nearly every week which could produce sizeable effects.

Model #:	1	2	3	4	5	6	7
Post Week:	8/8	8/15	8/22	8/29	9/5	9/12	9/19
Newspaper (Std. Err)	-.055 (.055)	-.089 (.046)	-.070 (.041)	-.068 (.037)	-.051 (.035)	-.045 (.033)	-.029 (.031)
Cable News	.039 (.069)	-.005 (.056)	-.009 (.050)	.003 (.044)	.012 (.041)	.056 (.039)	.091* (.037)
Newspaper × Cable News	.040 (.090)	.078 (.072)	.067 (.063)	.047 (.056)	.059 (.051)	.033 (.048)	-.010 (.045)
Week	.004* (.002)	.002 (.002)	.001 (.002)	.000 (.002)	.002 (.002)	.005 (.003)	.007* (.003)
Post × Newspaper	.016 (.054)	.056 (.045)	.036 (.041)	.036 (.037)	.015 (.035)	.006 (.035)	-.023 (.034)
Post × Cable News	.091 (.069)	.146* (.057)	.159* (.051)	.155* (.046)	.153* (.044)	.101* (.043)	.054 (.042)
Post × Cable News × Newspaper	-.099 (.092)	-.148* (.075)	-.143* (.067)	-.127* (.060)	-.152* (.057)	-.128* (.055)	-.070 (.054)
Constant	-.055 (.032)	-.024 (.032)	-.019 (.033)	-.007 (.033)	-.031 (.035)	-.060 (.037)	-.086* (.037)
R ²	.474	.525	.539	.554	.548	.520	.504

Table 1: TREATMENT EFFECT FOR CABLE NEWS VIEWERS AND NEWSPAPER READERS: For the 72 means, * indicates two-tailed significance at .05 or higher.

Several results emerge from Table 1. First, there is no effect of the Swift Boat campaign on newspaper viewers. While newspaper viewers are more pro-Kerry than the other groups (as is evident by the negative coefficient on *Newspaper* – they do not change their opinions any more or less than respondents who consume neither cable news or newspapers. For every pre/post comparison, the coefficient of change on *Newspaper* × *Cable News* is nearly zero.

Second, there are sizable effects for viewers of cable news. Inspecting the coefficients on the interaction $Post \times Cable\ News$ reveals effects for pre/post comparisons using the weeks of August 15th, August 22nd, August 29th, September 5th, and September 12th. According to the press coverage graphed in Figure , these are the weeks of the most intense coverage of the Swift Boat ads. Moreover, the largest estimated effect occurs using the break point defined by the week of the most intense media coverage – August 22nd. Relative to the change evident among respondents who did not consume newspapers or cable news at least one time during the week, the average change among respondents who consumed only cable news was 15.9% more pro-Bush.

Third, for individuals who both watched cable news and read the newspaper, there was no net effect of the Swift Boat ads – the coefficient on the triple interaction $Post \times Cable\ News \times Newspaper$ nearly exactly offsets the effect of the cable news exposure (i.e., $Post \times Cable\ News$). This suggests that to the extent that cable news has an effect, either the set of respondents who consume both cable news and newspapers were unaffected by the ad campaign or else the coverage in newspapers were able to counteract the effect.

For clarity, Figure 7 plots the estimated effect for the three groups relative to the baseline change evident in respondents who report consuming neither media source for each of the pre/post weeks used.

For every comparison, the effect of consuming just newspapers is nearly zero. There is an increase around the time of the Democratic National Convention but the effect is statistically indistinguishable from zero. For viewers of cable news, there was a large move in a pro-Bush direction, but only after the Swift Boat campaign. Reassuringly, the largest difference is the week immediately following the second ad and during the period of the most intense media coverage of the Swift Boat ads (August 22nd).⁹

⁹As Figure 7 indicates, interpreting the effect is made more difficult by the fact that the Republican National Convention was also occurring during this week. To determine how much of the effect is due to coverage of the Swift Boat and as opposed to possible convention effects we need to estimate the counterfactual of – what would the effect be if there had not been a convention? While we obviously cannot change history and we cannot compute the effect of the Swift Boat coverage in the absence of a convention, we can compute the effect of conventions in the absence of a Swift Boat campaign. Using the estimated effects for

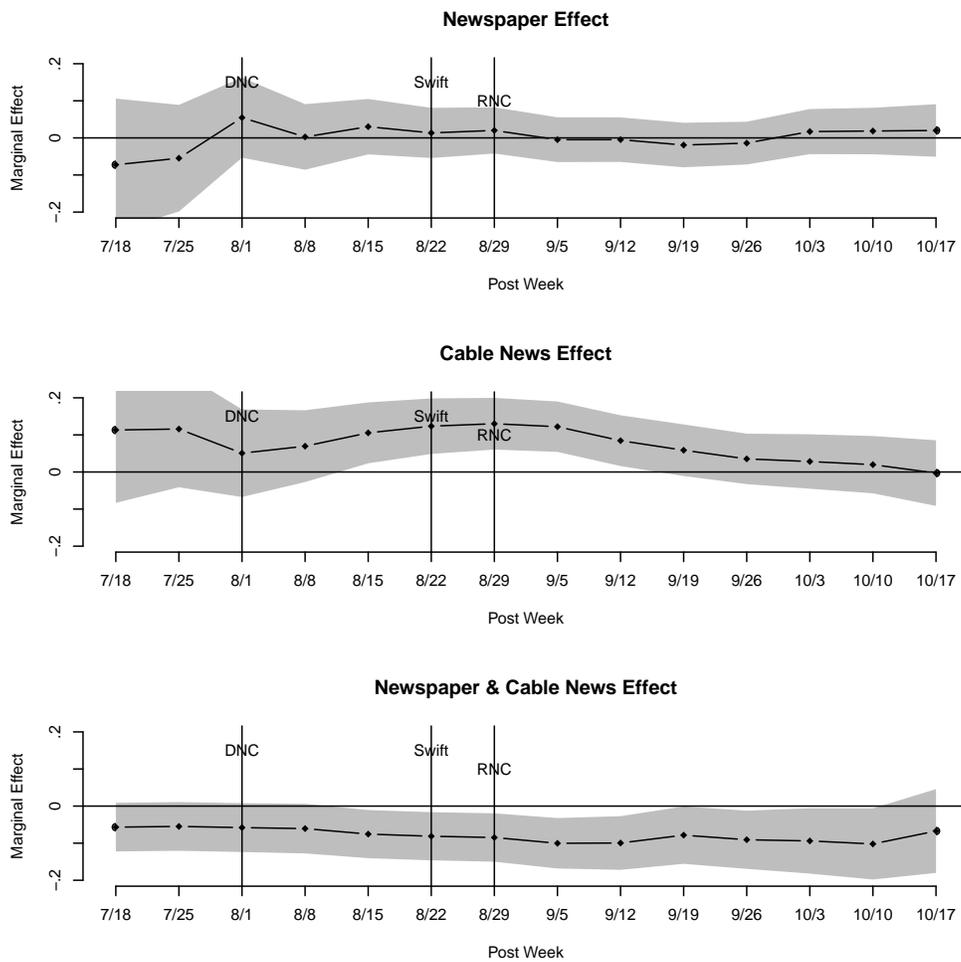


Figure 6: THE EFFECT OF EXPOSURE BY MEDIA CONSUMPTION AND WEEK: The area denotes the 95% confidence interval for the estimated effect.

How should we interpret the effects? The fact that the effects are only evident for those individuals who only watch cable news suggest that the effect is not due to how the media “frames” the campaign. If the Swift boat ads changed how events were interpreted by the media we would expect to find an effect among individuals who read the newspaper because the newspaper coverage would presumably be affected by similar concerns. The fact that only cable news viewers are affected suggest that the effect is peculiar to the cable news medium. Whether this is a difference due to the visual nature of cable news, the fact that information is transmitted by voice rather than text, or the fact that the ads themselves were rebroadcast we cannot say.

Similarly, the effect may be due to other circumstances or events that affect the subset of respondents who only consume cable news. If such individuals are systematically different than the others, or if they engage in activities that others do not we may be misattributing the effect. We observe a difference in these individuals relative to the others which we are assuming is due to the Swift Boat campaign, but it may be due to other things. That said, the appendix reveals that these individuals differ only slightly from the others. The characteristics with the largest non-zero predicted effects for the probability that an individual consumes cable news only include: female (+1.2%), black (-1.6%), born-again christian (+1.4%), and Democrat (-1.2%). In other words, relative to other groups, cable news only viewers are slightly more likely to be female, white, born-again, and non-Democrat (although the Republican indicator is statistically and substantively insignificant).

To examine this possibility, we replicate this analysis by partisanship – i.e., conduct the analyses separately for self-identified Republicans, Democrats, and independents. So doing reveals that the effect evident in Table 7 for cable news consumers is almost entirely attributable to changes in the opinions of self-identified Republicans and Democrats. Al-

the Democratic convention (which occurred prior to the week of August 1st) reveals no pro-Kerry bounce, but the effect may be confounded by the fact that the post period includes the Republican convention and the Swift boat ads. [USE A NARROWER POST WINDOW?] We can also estimate the effect that is due to the RNC convention using the effect of the DNC in the 2000 election. Both represent the second convention and occur at roughly the same point in the election calendar. If we use the difference in the effects as an estimate of the Swift boat ad coverage we get xxxx. The appendix contains the complete set of analyses.

though the estimates are imprecise due to small sample sizes, there is no evidence consistent with the claim that cable news coverage of the Swift boat ads changed the opinion of political independents.

In contrast, there are effects among Republicans and Democrats.

We might also wonder if there are differences by media outlet – given that many have argued that Fox News is appreciable different from the other cable stations (although recently MSNBC has apparently become more ideological in its programming as well (XXXX)) – perhaps the effect we document for cable news is actually a Fox News effect. To explore this possibility we use a similar design to compare the reactions of Fox News watchers to CNN and MSNBC watchers to disentangle the cable news effect by the network that is watched. [[STATS ON VIEWERSHIP]]

The treatment in this analysis is identical to the design employed in the paper except now we are interested in identifying and distinguishing between Fox News viewers and those viewers who watched CNN or MSNBC.

Because we are subsetting the sample even further to distinguish those viewers who report watching Fox News and not CNN or MSNBC as well as those who watch CNN or MSNBC and not Fox News, our estimates are less precise. Nonetheless, Figure 8 reveals that there is no difference in the change in opinion between Fox News viewers and viewers of CNN and MSNBC. The implications of this is that the results we document are not due to the framing of a particular cable station. Instead, the effect is due to characteristics that are common to cable news coverage.

One possibility is that the visual nature that matters – the effect of the ads are talked about by those in the cable news organization and that the effect we document is a consequence of cable news rebroadcasting the ads to a larger audience while trying to set up the story. Alternatively, perhaps the discussion format that cable news used to discuss the Swift Boat ads were sufficiently similar so as to produce similar effects. If both organizations discussed the ad campaign using pundits on each side of the issue, perhaps that format was

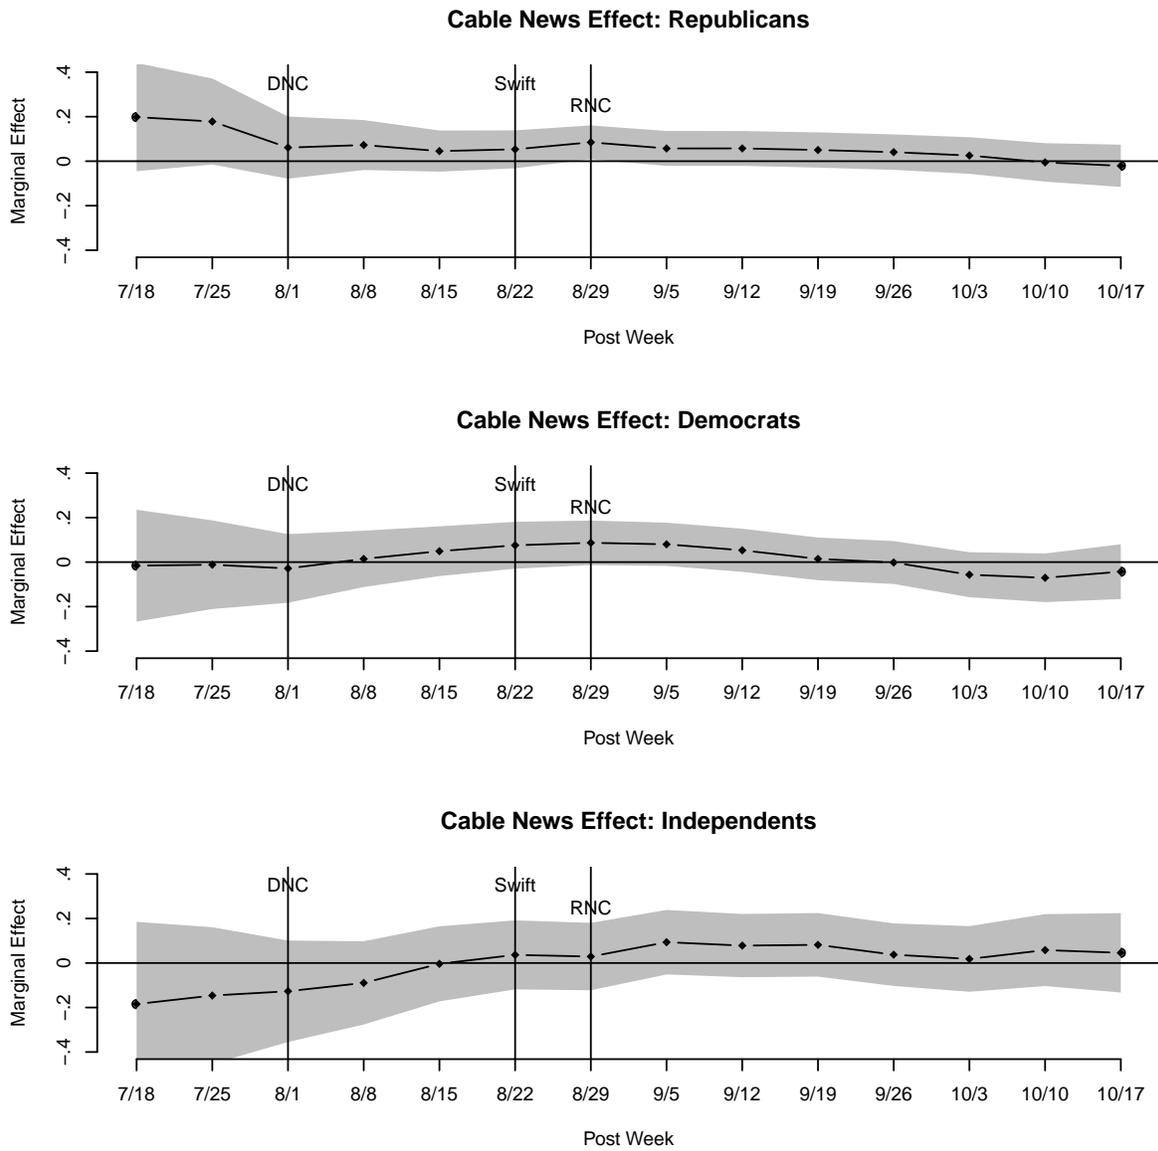


Figure 7: THE EFFECT OF CABLE NEWS EXPOSURE BY PARTISANSHIP AND WEEK: The area denotes the 95% confidence interval for the estimated effect.

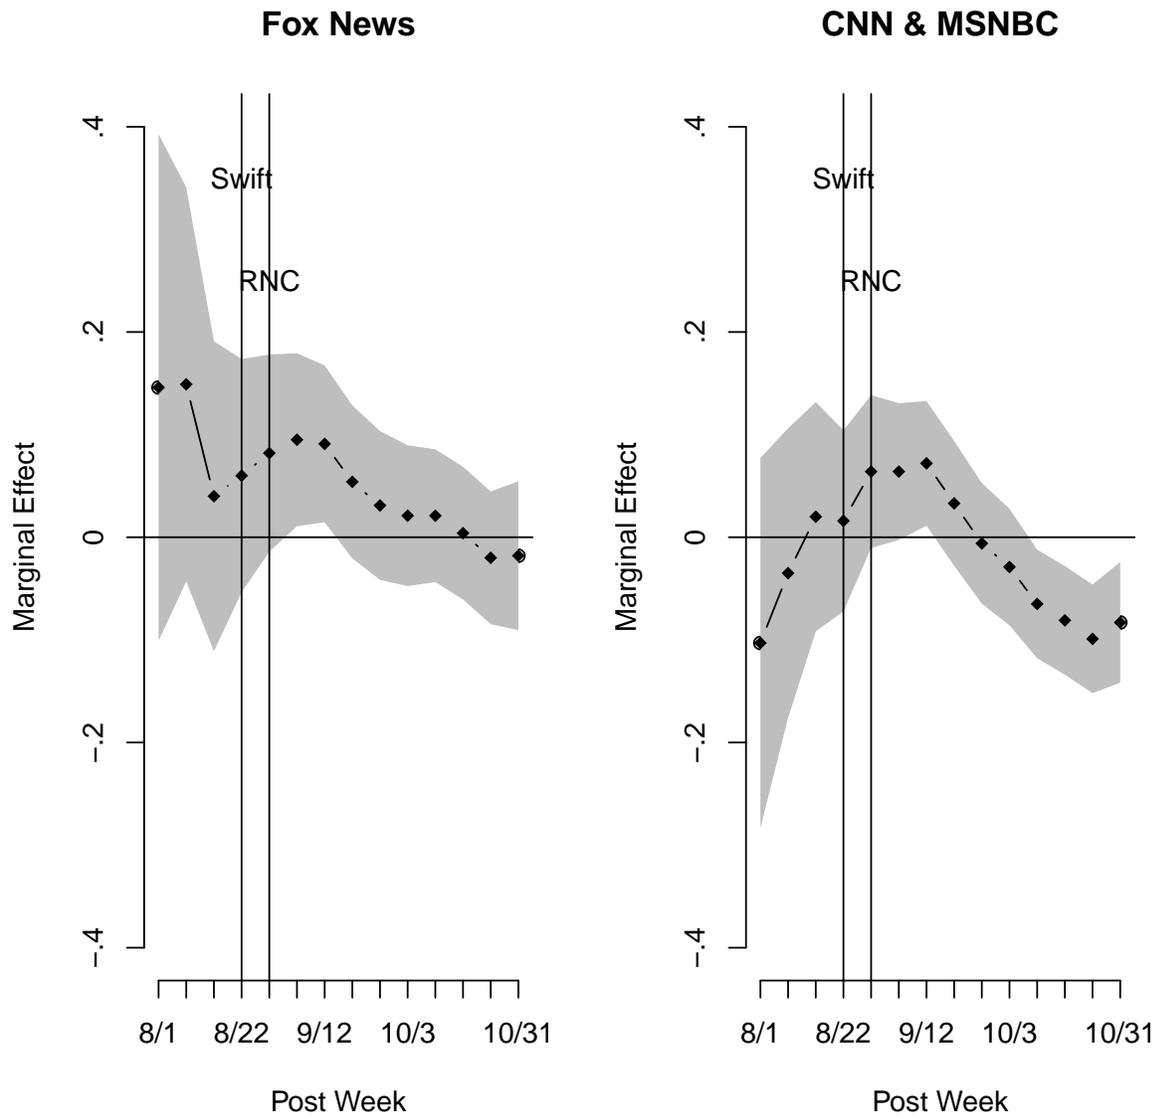


Figure 8: DIFFERENTIAL EFFECTS OF CABLE NEWS STATIONS? Shaded regions denote 95% confidence intervals for the estimated media effect.

sufficient to sway voters in a single direction. Of course, these possibilities are not exclusive – perhaps the ads were effective because CNN, MSNBC, and Fox News ended up rebroadcasting the ads several times to set up discussions and debates about the ads and provided many viewers with the ability to see the ads themselves, and then they were also able to hear and evaluate competing viewpoints on the ads which may have pushed them away from John Kerry. While this is purely speculative, it is an interpretation that is consistent with the effects we document.

3.3 Interpreting the Effects: How much is plausibly due to the Swift Boat ads?

As discussed above, the effects we estimate in prior sections are a combination of the effect of exposure to the Swift-Boat campaign and of the difference in exposure rates between the groups. The overall change in the margin between Bush and Kerry over the period we examine is about a 0.077 change in favor of Bush. The treatment effect we estimate for cable news viewers in the prior section is .159 based on Model 3 in Table 1. What percentage of the change in voter opinion is therefore plausibly attributable to exposure to the Swift Boat ads rather than other media-specific effects given that not every consumer of cable news was exposed to coverage of the Swift Boat ads?

While self-reports of exposure are known to be problematic (?), looking at actual exposure rates is nonetheless instructive. Table 3.3 reports the self-reported percentage of respondents in each group that indicate that they either saw the Swift Boat ads or the heard about it using several questions that were asked of respondents following the flurry of press coverage surrounding the Swift Boat ad campaign.

Using the estimator of section 2.3, based on the results of the models reported in Table 3.3 and the fact that 13% percentage of respondents reportedly consume cable news and not

	Pct. of Sample	Seen Ad	Only Heard of Ad	Seen or Heard	Not Know
Resident in Ad Buy State	9%	50.0%	19.6%	69.6%	30.4%
Resident in non-Ad Buy States	91%	36.6%	26.2%	62.8%	37.2%
Cable news, Newspaper	54%	44.3%	26.8%	71.1%	28.8%
Cable news, No Newspaper	13%	41.1%	20.3%	61.4%	38.6%
No Cable news, Newspaper	23%	26.7%	29.1%	55.8%	44.2%
No Cable news, No Newspaper	11%	23.5 %	19.9%	43.4%	56.6%

Table 2: Self-Reported Exposure to Swift Boat Ads

newspapers, the amount of the effect that is attributable to cable news coverage of the Swift Boat ad campaign is:

$$\frac{.1304 \times .159}{.077} \times 100 \approx 27\%.$$

Put differently the change in the national numbers that is plausibly due to the change in opinion experienced by individuals who only watch cable news is 2.07%, which represents 27% of the total pro-Bush move of 7.7%. This is the effect that can be attributable to the coverage of the Swift boat ad campaign, but recall that the change reflects any differences that affect only those respondents who only watch cable news so the estimate is likely the maximum plausible effect.

4 Discussion

Political science has spent a great deal of attention trying to document whether ads have an effect, and, if so, characterizing how different ads might affect different types of individuals. Where less progress has been made is in considering the various ways that ads may affect opinion and trying to disentangle the effects that are attributable to direct exposure through the ad buys of the campaigns themselves relative to the exposure that ads may be given through coverage by the media. Such “downstream” effects are difficult to assess.

We use perhaps the most famous, and most influential ad that is thought to have “downstream” effects – the Swift Boat ads in the 2004 presidential election – to investigate the

nature and magnitude of the indirect effects of advertising. Using both individual level and state level data, we demonstrate that the effects of the ad campaign was entirely through the secondary exposure provided by media coverage of the ad campaign. Using a difference-in-differences design, we show that there is evidence that the widely credited “Swift Boat” campaign did appear to have an effect, but the effect was likely through the actual re-broadcast of the ads on cable television rather than the effects of media framing. There is no detectable effect of the actual ad buys themselves – the population targeted by the purchased ad buys did not noticeable change their opinion. There is also no effect among individuals closely following the campaign through the print media even though there was substantial coverage of the ads by print journalists and commentators. However, there is an effect among those who follow the campaign on cable news – an outlet which frequently replayed the actual ads themselves in front of a larger audience

We show that cable news coverage of the Swift Boat campaign produced (at most) a 2.7% increase in the support for Bush. To the extent that events are effective only because of the choice made by reporters and journalists rather than the choices made by the candidates themselves, then media “watchdogs” may actually be counterproductive in terms of policing the campaign airwaves and that so doing may actual undermine their goal of keeping the campaigns “honest.”¹⁰

There are three possible explanations for the effect we find. First, the coverage between newspaper and cable news may differ in important ways. If so, the nature of the media may be responsible for the effects. It is unclear whether the relevant differences are differences in content – e.g., the differences between a visual medium and a text medium, the difference between the debate format typically employed in cable news programming versus a narrative story in the newspaper – or differences in the type of information that is talked about the ads, or differences in the ideological leanings of the media outlets.

¹⁰We obviously cannot assess the counterfactual of would have happened had the ads not been run at all, or identify the effect of the ads due solely to the discussion of journalists and political commentators. Instead, we can only identify the effect of running the ads on the air versus having them discussed by pundits and journalists.

Second, the differences in effects we detect may be due to differences in the audiences – respondents who indicate that they watch only cable news may be systematically different from the others in ways that make them especially susceptible to changing their opinion either in response to the Swift boat ads or else around the time of the Swift boat ad coverage for other reasons. While we do not find large differences in terms of measurable differences, there may be other differences that are more difficult to detect.

NOTES:

- Ad buy state-not and cable-not contrasts in seen ad are very similar, but big effect only for cable. Suggests not a simple re-broadcast effect.
- Cable-no and Newspaper-no contrasts in knowledge are very similar, but big effect only for cable. Suggests not simple knowledge effect.
- Live possibilities: coverage differs between newspaper and cable, maybe because of different distributions of slant. (What does literature on slant say?)
- Audiences differ, maybe in terms of partisanship. (Should be able to look at data we already have.)
- TV has aura.