The Influence of Television and Radio Advertising on Candidate Evaluations: Results from a Large Scale Randomized Experiment

Alan Gerber

James G. Gimpel

Donald P. Green

Daron R. Shaw

February 13, 2006

Abstract: This paper reports the results of the first large-scale experiment involving paid political advertising. During the opening months of the 2006 Texas gubernatorial campaign, approximately $2 million of television and radio advertising on behalf of the incumbent candidate were experimentally deployed. In each experimental media market, the launch date and volume of television advertising was randomly assigned. Launch dates of radio advertising were randomly assigned as well. In order to gauge movement in public opinion, a tracking poll conducted brief telephone interviews with approximately 1000 registered voters each day; a total of 30,757 completed interviews over the course of the month-long experiment and brief follow-up one month after the conclusion of the television campaign. Results indicate strong but short-lived effects of television advertising and somewhat equivocal effects of radio advertising. The ephemeral nature of these effects is more consistent with psychological models of priming or emotions than models of Bayesian learning.

Paper prepared for presentation at Stanford University's Methods of Analysis Program in the Social Sciences, February 13, 2007 and the annual meeting of the AAAS, San Francisco, February 16, 2007. The authors are grateful to the Texans for Rick Perry campaign, which made this experiment possible but bears no responsibility for the conclusions drawn in this report.

The persuasiveness of mass media communications has been hotly debated since the dawn of modern social science during the 1930s. Whereas most political observers accorded enormous weight to radio, newsreels, and, later, television communication, early students of political propaganda tended to downplay these effects. The survey findings of Lazarsfeld, Berelson, and Gaudet (1944) coupled with the controlled experiments of Hovland, Lumsdaine, and Sheffield (1949) called into question the persuasiveness of political communication and ushered in the "minimal effects" thesis (Klapper 1960) that would hold sway among academics until the 1980s. Indeed, much of the scholarship of the 1960s and 1970s sought to explain the social psychological mechanisms that account for minimal effects, such as selective exposure, acceptance, and recall (Sears and Whitney 1973), as well as simple inattentiveness to and forgetfulness of political content of television programming (Neumann 1986; Patterson and McClure 1976).

The 1980s and 1990s saw a reversal of the overall tenor of media research, spurred by a series of laboratory experiments involving simulated viewing situations. This wave of experiments indicated that television news shapes what issues the public considers important (Iyengar, Peters, and Kinder 1982), that the tone of televised advertising influences voter cynicism (Ansolabehere and Iyengar 1995), and that the content of news stories shapes public opinion toward race-related policies (Gilliam and Iyengar 2000). More to the point for our purposes, campaign advertisements were found to significantly increase voters' support for the sponsoring candidates (Ansolabehere and Iyengar 1995, p.186; Brader 2005, p.395; Kaid 1997, p.1088; Valentino, Hutchings, and Williams 2004, p.349). As this new wave of experimental findings took root, its results were echoed in observational studies, which increasingly found

evidence of campaign effects in survey data (Shaw 1999; Johnston, Hagen, and Jamieson 2004; Freedman and Goldstein 1999; Freedman, Franz, and Goldstein 2004).

Missing from the empirical literature on campaign effects, however, is experimental deployment of mass media advertising in the context of an actual campaign. Random assignment enables laboratory studies to detect causal effects, but the external validity of these measurements remains uncertain because subjects are not exposed to the advertisements in their home environment and know that their pre-survey, viewing, and post-viewing interview are part of a social science study. Observational studies have the advantage of measuring the behaviors or random samples of voters and sometimes use sophisticated measures of advertising exposure, but they cannot identify the causal effects of advertising without invoking strong assumptions about unobserved factors that might be correlated with advertising and vote choice. To be sure, the literature on media effects has become steadily more sophisticated methodologically and nuanced substantively. For example, a series of studies employing regression discontinuity designs have suggested that although presidential advertisements seem to have minimal effects on turnout (Krasno and Green 2005), they do appear to shape preferences (Arceneaux and Huber 2006). This pattern of findings is confirmed by a small but growing set of field experiments that suggest that televised public service announcements influence voter turnout (Vavreck and Green 2005) and vote choice (Panagopoulos and Green 2006). Although these recent experimental studies represent an advance over laboratory studies insofar as they examine the size and nature of unobtrusive media intervention effects on political outcomes, they remain limited in scope and statistical power. To date, field experiments have focused exclusively on nonpartisan messages, and no field experiment has come close to deploying mass media communications on the scale and intensity of an actual national or statewide campaign.

Our research attempts to fill this gap by evaluating a $2 million television and radio campaign. This three week campaign encompassed 20 TV media markets (designated market areas, hereafter "DMAs") and more than 80 AM and FM radio stations. At its peak, the experimental campaign deployed up to 1000 gross ratings points (GRPs[1]) of advertising per week into selected DMAs, which is comparable to what battleground states experience during the fall of a presidential election year (Johnston et al. 2004).

The unprecedented cooperation of a major statewide election campaign allowed us to blend and expand several attractive features of past research designs. Like the Johnston et al. (2004) and Arceneaux and Huber (2006) studies, we track outcomes using a very large daily tracking poll, which gathered approximately 1,000 interviews per day. Like the Vavreck and Green (2005) and Panagopoulos and Green (2006) experiments, we use random assignment to generate unbiased inferences about the effects of television and radio communications. And like the Freedman et al. (2004) and Johnston et al. (2004) studies, we rely on objective measures of advertising broadcasts rather than on survey reports of whether people recalled seeing or hearing the ads.

This essay is structured as follows. We begin by presenting a model that contrasts competing hypotheses about how voters process campaign information. One hypothesis holds that the ads provide information that voters use to update their prior beliefs about the candidates. An alternative hypothesis models holds that advertising momentarily raises the salience of what Zaller (1992) has termed "considerations," views of the candidates' attributes that contribute to voters' overall evaluations. Whereas a Baysian learning model typically suggests a gradual

[1] One GRP is equal to one percent of the viewing audience; 1000 points is ostensibly the equivalent of everyone seeing an ad 10 times, though we note below that this estimate may cover a variety of scenarios. We follow the convention of using GRPs as the measure of advertising volume (see Ansolabehere et al. 1999, p. 903; Shaw 1999, p. 349; Johnston et al. 2004, p. 70).

process of opinion change in the wake of advertising, the priming model implies evanescent effects. Next, we describe the gubernatorial race that provides the setting for our experiment and the television and radio ads that were the focus of our experimental evaluation. The third section describes the design of the experiment and the survey that we used to assess outcomes. We then present an array of cross-sectional time-series models that parameterize the media effects in different ways.

The statistical analysis reveals strong but short-lived effects of television advertising and somewhat equivocal effects of radio advertising. The ephemeral nature of these effects has a variety of important substantive and methodological implications. Substantively, the results are more consistent with psychological models of priming than models of Bayesian learning. Methodologically, the findings underscore the importance of taking decay into account when measuring media effects within the context of laboratory or survey experiments (Clinton and Lapinski 2004). More broadly, this study shows the feasibility of conducting large-scale field experiments on political advertising and provides the first field experimental evidence of the influence of campaign advertising on public opinion.

Modeling Campaign Effects

The literature on the effects of campaign advertising consists of two interrelated research agendas. The first attempts to gauge the effects of political ads and to examine how these effects vary as a function of the ads' content, tone, or emotive qualities. The second attempts to explain the conditions under which ads influence public opinion. While the first strand of this literature emphasizes the short-term influence that ads can have on viewers, the second strand cautions that

ads may never be viewed by their intended targets and that those who do view them may be inattentive or resistant to their message.

The practical implications of this literature are clear: in order to shift voter preferences, an advertising campaign must be of sufficient scope as to expose large number of voters, and its content must be sufficiently engaging to attract the audience's attention and sufficiently persuasive to win over those who might otherwise resist its blandishments. Whether guided by this research literature or their own intuition, campaigns seem to obey these precepts. They advertise as heavily as their means allow and craft messages that they believe large segments of the electorate will accept.

As we move from the core of this research literature to more specific questions about the dynamics of advertising effects, theoretical and empirical guidance diminishes. With respect to the duration of advertising effects, one finds two competing and somewhat antagonistic perspectives. The first perspective emphasizes the *information content* of advertising. From this perspective, many students of advertising contend that the effects of media exposure decay rapidly over time. Whether the content of the ads concerns goods and services (Vakratsas and Ambler 1999; Zielske and Henry 1980; Burke and Srull 1988) or politics (Neumann 1986), it seems clear the viewers find it difficult to recall the content of what they watched only days or hours earlier. On the other hand, it has been argued that voters process information "on-line," forming new opinions and discarding the information on which those opinions were based. As Lodge, Steenbergen, and Brau (1995) argue, "memory for campaign messages not only fades but fades quickly…the half-life of the message typically being less than a week," yet "recall is not a necessary condition for information to be influential" (pp.315, 317- 318). This interpretation

has gradually won acceptance among scholars, who now believe that messages may have enduring effects even after their content is forgotten.

A competing perspective is that advertisements prime voters to *invoke different evaluative criteria* when assessing the candidates. An advertisement that stresses the threat posed by international terrorism might increase the extent to which voters draw on their foreign policy attitudes when evaluating presidential candidates. Iyengar, Peters, and Kinder (1982) and Miller and Krosnick (1996), for example, show that news stories have the capacity to strengthen the relationship between targeted attitude domains and candidate evaluations. Another argument with similar implications holds that the emotional content of advertisements accentuates their effects (Brader 2005). The priming hypothesis implies that the effects of advertising decay as the cognitive accessibility of the primed considerations fades, which is consistent with the fact that priming effects observed during the lab session or survey interview dissipate when subjects are re-interviewed days or weeks later (Gaines, Kuklinski, and Quirk 2007).

These two competing perspectives – learning and priming – have quite different empirical and theoretical implications that have not, to our knowledge, been integrated within a common theoretical framework. Models of learning based on Bayes' rule gained prominence with Achen's (1992) model of rational learning, which posits that rational voters have prior beliefs about an unobserved quantity, the net utility gain associated with the election of a given candidate, and update these beliefs based on efficient use of new information. Gerber and Green (1998) generalized this model to account for the possibility that the underlying value of a given candidate changes over time. Their learning model implies that the speed of opinion change is a function of signal-to-noise ratio, that is, the rate of true change in the underlying attributes of the candidates as opposed to non-diagnostic variation in rhetoric and events. The two corollaries of

the rational learning model are (1) voters change their opinions slowly unless confronted with an especially telling piece of new information suggesting change in the candidates' underlying attributes and (2) as voters learn, they replace old information with new, a proposition that comports with the "on-line processing" hypothesis of Lodge et al. (1995). Combined, these propositions mean that if an ad contains information that causes voters to abruptly prefer candidate A over candidate B, this new state should persist until new evidence is presented to change this view.

From the standpoint of the rational learning model, a transitory advertising effect represents an anomaly.[2] An abrupt shift toward a candidate followed by a return to the prior states suggests an altogether different psychological process, akin to what Zaller (1992) refers to "sampling considerations." In Zaller's model, considerations are ideas that may or may not come to mind as a person formulates and expresses a political preference. The probability that any given consideration comes to mind is in part a function of the political environment; where political communications make a particular consideration salient, its sampling probability increases.

Melding the considerations model with the Bayesian learning model is a matter of changing the translation of beliefs into evaluations. A Bayesian analysis might imply that the voter should altogether ignore the information content of an advertisement as unreliable noise, while a considerations model might suggest that, at least temporarily, the connection between beliefs and evaluations might change as voters come to apply new evaluative criteria. Similarly, the emotional reaction generated by TV ads (Brader 2005) might be modeled as a momentary

[2] Rapid change is not in itself an anomaly, but to explain both rapid change in the wake of a commercial and rapid reversal of this change, one would have to posit that the commercial contained a good deal of pertinent information and that, subsequently, the outside environment supplied a comparable amount of pertinent information in rebuttal. See the appendix for a formal development of this argument.

shock to evaluations that decays as the individual returns to his or her equilibrium emotional state. We show formally in the appendix how either priming or evoking emotions is sufficient to generate a strong but short-lived change in candidate evaluations.

Description of the Experimental Setting

Our analysis focuses on the 2006 re-election campaign of Texas governor Rick Perry. Aside from its receptiveness to experimental evaluation, the Perry campaign looked much like other big-state re-election campaigns. Perry became the Republican governor of Texas in 2000, succeeding George W. Bush when the latter resigned after winning the presidency. Perry had from 1985 to 1991 served two terms as the Commissioner of Agriculture and in the Texas House of Representatives. With deep Texas roots in a rural part of West Texas, Perry had solid conservative credentials that put him squarely in the middle of the Texas mainstream. In 2002, Perry was reelected governor, winning decisively over a well-financed Democratic opponent by a 58 to 40 percent margin.

The 2006 race presented a potentially more difficult reelection test for several reasons. First, early signs indicated that the public had a far dimmer view of President Bush's performance heading into the 2006 race than they had heading into the 2002 contest. Second, Governor Perry faced some internal strife and factionalism within his own legislative majority, after action on an important education finance package stalled during the 2005 legislative session. During the first half of 2005, Perry's popularity ratings declined, and it appeared that he might face a stiff challenge in the GOP primary from U.S. Senator Kay Bailey Hutchinson. But in July 2005, Hutchinson decided not to challenge Perry, and in September Hurricane Katrina put the Governor's leadership on display through a surge in media coverage. His

approval rating rebounded to its highest level in three years, just above 50%. Even so, a threatened primary challenge from state comptroller Carole Keeton Strayhorn kept the Perry camp on edge though the end of 2005.

That a Republican official would bolt the party to challenge an incumbent Republican governor was unprecedented even for a state as politically colorful as Texas. Although less formidable an opponent than Hutchinson, Strayhorn was a credible challenger with her own distinguished record of public service. Strayhorn was the first female mayor of Austin and the only mayor elected for three successive terms. She was also the first woman elected Texas Railroad Commissioner—a position that has often been considered a stepping-stone to the governorship.

The late filing deadline meant that final word on a Strayhorn primary challenge could wait until the beginning of January, which, with early voting in Texas, meant a potentially short and intense battle for the nomination in early March. In the fall of 2005, Strayhorn aired a series of radio ads that were critical of Governor Perry. In December, she purchased several hundred thousand dollars' worth of television time to air in January, presumably to kick off a nomination challenge; the Perry campaign countered by buying three weeks of advertising that form the basis of this experiment. However, on the day of the deadline, January 2nd, she announced her intention to run as an Independent, which meant that January 2nd marked the beginning of the general election campaign. After the March primary decided the Democratic ticket, the 2006 race took shape as a four-way contest between Perry; his Democratic challenger, former U.S. House member Chris Bell; Strayhorn; and a second Independent, county-Western singer Richard "Kinky" Friedman. In the end, Perry was reelected with 39% of the vote.

As noted below, the media campaign launched by TRP was conceived in December, before the nature of the primary election competition had become clear. The Strayhorn campaign deployed a television and radio ad during the first three weeks of January in selected media markets, but did not advertise during the rest of the primary season. The other candidates, for their part, deployed no media advertising in January. Thus, the start conditions of this study – abrupt shifts in the quantity of advertising and a vacuum or profusion of opposing ads – satisfies what Zaller (1996) regards as the ideal conditions for detecting media effect.

Thus, depending on whether the Strayhorn campaign advertised in a given media market, TRP ads either aired unopposed or aired opposed for a three week period. This nuance in the design presents a valuable opportunity to study the effects of contested (two-sided) and uncontested (one-sided) messages. Political observers often argue that campaigns cannot afford to stand idly by while the opponent takes to the airwaves, on the grounds that ads are especially effective when they go unchallenged. We will assess both the average effect of TRP advertising and the extent to which its effectiveness varies depending on whether the ads air concurrently with Strayhorn ads.

The Experimental Campaign Ads

The Perry Campaign kicked off its campaign with ads that highlighted the governor's accomplishments and charisma. The television ad sought to link positive images of the photogenic governor with voters' pride in the state of Texas. The scenes sweep from Texas landscapes to a schoolroom to a doctor's office, with Governor Perry's voiceover:

> I've never been more proud to call myself a Texan. In Texas we've set the national standard for economic development. We gained 300,000 new jobs. Lawsuit reform is

> bringing better healthcare to millions. We've invested ten billion new dollars in our public schools while improving standards of accountability for student performance. Our people are compassionate. Our vision, bold. Our values, strong. The best is yet to come. I'm proud of Texas. How 'bout you?

Two aspects of this ad are noteworthy. First, it appeals only indirectly to voter ideology, focusing instead on their retrospective performance evaluations and what Stokes (1963) dubbed "valence" issues. Oblique ideological reference is made to "lawsuit reform," "improved standards of accountability for student performance," and strong values, but these conservative themes are balanced by willingness to spend more on education. The emphasis throughout is on accomplishments and the pride that Texans should take in their state and the governor's stewardship. The scenes shift from one iconic Texas image to another, interspersed with images of a handsome suit-clad governor milling with workers and children or poised in front of the State Capitol.

Second, the ad makes no mention of opponents, their platforms or attributes. The script does not even conjure up implicit critiques of the challengers by referring to the governor's "proven leadership" or other personal traits. This tactic reflects the strategic setting at the time during which the ads aired. Carole Keeton Strayhorn had limited name recognition (at least with that particular name—she was known to some Texans by a last name that she dropped after her remarriage), as did the two Democratic contenders for their party's nomination. By "going positive," the TPR campaign sought to deny its opponents the salience that comes with back-and-forth in both the paid media and accompanying news coverage.

The radio ad followed a similar format but with more specific references to accomplishments. The text of the ad can be found in the appendix. The overall theme is that

Texans have ample reason to be proud of Texas given the many ways in which the state has excelled under Governor Perry's leadership.

Although no mention is made of a political campaign – no requests for voters' support, for example – the television and radio ads implicitly anticipate criticisms by other candidates. The Perry campaign did not have a copy of Carole Keeton Strayhorn's ad as it produced its own but expected performance-related criticisms along the lines of Strayhorn's radio ads during the fall.

Like the Perry ad, the Strayhorn televised ad features the candidate as narrator and focuses primarily on performance evaluations:

> Partisan politics is making it tough for Texas. After nine legislative sessions full of name-calling, our school funding is in crisis, property taxes are up, and judges are having to do our governor's job. We all know that unless we set politics aside, we'll never fix what's broke. I'm Carole Keeton Strayhorn. I'm a Republican, but I've decided to put partisan politics aside, to run for governor as an Independent. One tough time might just need one tough grandma to shake Austin up.

In contrast to the Perry ad, which placed him in a variety of Texas settings, the Strayhorn ad features her standing in front of a blank white background. Unlike the Perry ad, hers mentions party directly but parallels his strategy by attempting to stand above politics. It does not "go negative" in the conventional sense of attacking the personal failings of the opponent. Instead, it critiques the partisan deadlock of which the governor is a part. The Perry ads in some sense pre-empt or answer this critique by adducing evidence that conditions in Texas are strong and improving.

Experimental Design

Random assignment is the feature of this study that sets it apart from past studies of large-scale media campaigns. This section explains how the random assignment was performed and the implications for the statistical analysis that follows.

For the broadcast television and radio experiments, the cross-sectional unit of observation is the DMA. Texas comprises 20 such markets of varying geographic and population size. These broadcast media markets also define the boundaries of cable TV systems, which are obligated to carry advertising purchased on broadcast stations. Approximately two-thirds of all households subscribe to cable television.[3]

Of these 20 media markets, the campaign was willing to allow experiments in 18, regarding the other two (Houston and Dallas-Fort. Worth) as too politically important to leave to chance. As it turned out, after Strayhorn pulled out of the Republican primary, the Perry campaign pulled out of (only) these two media markets ahead of schedule in order to conserve resources. The randomized intervention, then, consists of the media time that was purchased prior to Strayhorn's announcement. In the analysis that follows, we demonstrate that the results are not affected by whether one excludes these two nonexperimental markets or treats them as part of the experiment.

In light of the political, demographic, and geographic heterogeneity of the DMAs, we matched them as closely as possible into similar groups and then randomly assigned members of each stratum into an ordering that indicated the start date of the broadcast television campaign. Within each weekly rollout bracket, we randomly assigned the quantity of gross ratings points

[3] The experimental also involved the randomization of cable TV markets in select DMAs. This experiment, however, was underpowered and produced results with large standard errors. Including or excluding this aspect of the design in the overall analysis of TV advertising has no effect on the conclusions, which are driven by the broadcast TV results.

(GRPs) to be purchased: 250, 500, or 1000. The rollout dates were then given to the campaign's television media buyer, who arranged to have broadcast TV ads introduced in the specified DMAs each week.[4]

The experiment did not randomize the stations and programs on which the ads were placed. Broadcast TV ads were purchased in a variety of stations based on the campaign consultants' strategic judgment. Table 1 illustrates the sort of advertising purchases that were made in three DMAs that varied in terms of the quantity of advertising to which they were randomly assigned. The DMA assigned 250 GRPs aired most of the Perry ads on morning shows and news programs. The DMA assigned 500 GRPs bolstered the morning and news programming with additional entertainment programming. At 1,000 GRPs, the quantity of entertainment programs increased further, and the range of shows extended into late night entertainment and news programs.

Once the broadcast TV rollout brackets were determined, members of each bracket were randomly assigned to a period of radio advertisement. The GRP weight of the radio ads was not varied randomly. The purchaser of the radio ads was given discretion about how to place the ads, subject to the constraint that each radio station's signal remain largely confined to a single DMA. The research team obtained data on each station's signal propagation zone and worked with the radio purchaser to generate a list of advertising purchases. In the end, the radio purchases were somewhat more circumscribed geographically than would be usual for a political campaign but still achieved a substantial audience.

The time profiles of the broadcast and radio ads in each DMA are depicted in Figure 1. For ease of comparison, all 20 graphs are depicted using the same axes. On the horizontal axis are days, counting from 0 (January 3rd) to 26 (January 29th). Note that the experimental

[4] In one case, the buyer was unable to place ads totaling 500 GRPs and purchased 475 instead.

treatment did not begin until January 5th; the first two days of the study laid down a pre-treatment baseline. Tracing each of the treatment profiles across time shows the variability in the timing of the media rollout. For example, DMA #16 received a full dose of 1,000 television GRPs from the first week on, while DMAs #4 and #14 did so from week two on, and DMA #19 received a burst of 1000 GRPs only during the third week. Comparison across the vertical axis reveals the considerable variation in advertising intensity. DMA #2, for example, received 250 GRPs of television for three weeks, while DMA #10 received this dosage during only the third week.

Looking at the TV and radio trajectories in conjunction, we see a range of different configurations. In 11 DMAs there were periods during which TV ran in the absence of radio. In 4 DMAs, radio ran in the absence of broadcast TV. In 10 DMAs, the two ran concurrently during at least one week. Also noteworthy are weeks in which nothing aired: 11 DMAs aired no media during week 1; 4 were quiescent during week 2; and 2 DMAs ran no ads during week 3. After week 3, no ads were run in any media markets for the duration of the study.

Figure 2 shows the location and temporal relationship between the Perry and Strayhorn ads. The graph makes clear that the Strayhorn ads bear little relationship to placement and timing of the Perry ads in the 18 experimental markets. Sometimes her ads precede the Perry ads; in other cases, they run concomitantly. Fortunately for our purposes, her ads are localized both geographically and temporally, leaving ample statistical leverage for differentiating the effects of the two campaigns competing communications.

In forming the data for our pooled cross-sectional time-series analysis, the state of Texas has been divided into 20 media markets, as shown in Figure 3. The question of how best to configure the time dimension of the data is answered empirically. The overall pattern of results are little affected by whether we operationalize time in terms of weeks, which was the unit of

experimental assignment, or days. The advantage of coarser time units is the ability to track opinion change more reliably, due to greater numbers of survey interviews. The advantage of finer time units is the ability to examine dynamics and lagged effects.

Survey Design

In light of the fact that the unit of assignment is the DMA, we sought to allocate our survey sample in a way that would make the most of the experiment's power. The basic strategy was to spread the surveys more or less evenly across each of these geographic units. By re-weighting the data according to the probability of selection from the voter file, we can approximate the results that would have obtained through a simple random sample, and our results conform closely to the results from a concurrent poll based on a simple random sample from the voter file.[5]

The survey itself was conducted by a polling firm that specializes in "voter identification" calls – which is to say, brief and inexpensive surveys. One of the practical innovations of this experiment was the use of this type of survey, which cost one-tenth as much as a conventional survey per completed interview, allowing for 1,000 completed interviews per day.[6]

[5] A poll by Baselice & Associates during January 16-19 put Perry ahead with 41%, the Democratic candidate with 14%, Strayhorn with 24%, and Friedman with 8%, with the remainder undecided or mentioning other names. Perry has a 17 percentage-point lead over Strayhorn in our poll and a 19 point lead in the Baselice poll. The main difference between the two polls is that ours has a much higher proportion of undecided voters, perhaps due to the fact that those with lower levels of interest in politics received less weight in the likely voter weighting scheme used by Baselice & Associates.

[6] New sample was released every other day, following the first day's survey. Thus, the survey is a tracking poll, with independent samples every two days. The short duration of callbacks is obviously a potential source of bias, but this bias is constant over time.

After a brief introduction[7], the first two questions assessed the favorability ratings of the two leading candidates, Perry and Strayhorn. (See appendix for text.) The third question asked voters whom they would vote for if the election were held today. The fourth and final question on the survey instrument was rotated among three randomly selected alternatives. The first item in the rotation was designed as a manipulation check to assess whether, as expected, respondents in treatment areas were more likely to recall seeing Perry advertisements than respondents in control areas. One-third of survey respondents received this question. A parallel question asking about exposure to Strayhorn ads was directed to one-third of the sample. The final one-third of the sample was asked about their radio listening habits.

Table 2 describes statewide trends in public opinion over time. Although these trends do not speak to the question of advertising effectiveness, they provide a useful description of how the campaign unfolded in early 2006. The incumbent governor enjoyed a lead, but approximately one-third of the respondents declined to express a vote preference. As we move from the ballot test to candidate evaluations, somewhat stronger over-time patterns are apparent. During January, evaluations of Rick Perry became more polarized, with larger proportions of the sample offering strongly favor or strongly unfavorable evaluations. Less pronounced but still noticeable is the week-by-week increase in the proportion of people who provide an evaluation. Evaluations of Carole Keeton Strayhorn were confined to a much smaller proportion of registered voters, as a high proportion of respondents conceded that they were unfamiliar with her. This proportion declined steadily as time goes on, from 58% in the first week after she declared her candidacy to 49% three weeks later.

[7] This introduction was shortened after the first day of interviews in order to increase the response rate. This change in format prevents direct comparisons between the first day of interviewing and the days that follow. However, by including fixed effects for each day (or campaign period), the models presented below take this change into account.

Gauging the Effects of Broadcast TV and Radio

Our analysis is based on aggregate survey data, where the level of geographic aggregation is the media market and the level of temporal aggregation is either the day or the week. The daily aggregation scheme results in a dataset with 580 observations (N=20, T=29). Because the aggregate units encompass varying numbers of survey observations, the analyses below show how the results change when the data are weighted analytically to reflect the number of individual observations in each aggregate unit or when outcomes are assumed to have heteroskedastic errors due to differential sample sizes.

We model the aggregate survey response as a linear function of the campaign's television and radio advertising along with fixed effects for both time and geography.

$$Y_{it} = \alpha + \beta_1 TV\ GRPs_{it} + \beta_2 Radio\ GRPs_{it} + \gamma_1 Day_{it}^{1} + ... + \gamma_{J-1} Day_{it}^{J-1} + \delta_1 Market_{it}^{1} + ... + \delta_{K-1} Market_{it}^{K-1} + u_{it} \quad (1)$$

In equation (1), "TV GRPs" are gross ratings points (in 1,000s) associated with Perry's TV ads; and "Radio GRPs" are gross ratings points (in 1,000s) associated with Perry's radio ads. The use of fixed effects for geography tracks a given DMA over time as ads are randomly rolled out; in effect, we have 20 distinct experiments in which the media buy at any given time is randomly determined. The use of time fixed effects controls for statewide shocks that result from events or the vagaries of survey administration on any given day.

We consider below two alternative dependent variables (Y_{it}). The first is the difference in evaluations given to Perry and Strayhorn on a scale ranging from -4 (Strayhorn is more highly rated) to +4 (Perry is more highly rated). The second outcome measure is the percent expressing an intention to vote for Perry (without excluding "don't know" responses). The results are

unchanged when the latter dependent variable is calculated only for those respondents who express a candidate preference.

An alternative modeling approach incorporates time-series dynamics. Whereas the previous models assume that media effects exert an instantaneous effect, the inclusion of a lagged dependent variable on the right-hand side turns the process of decay into an empirical question:

$$Y_{it} = \alpha + \rho Y_{i,t-1} + \beta_1 TV\ GRPs_{it} + \beta_2 Radio\ GRPs_{it} + \gamma_1 Day_{it}^{1} + ... + \gamma_{J-1} Day_{it}^{J-1} + \delta_1 Market_{it}^{1} + ... + \delta_{K-1} Market_{it}^{K-1} + u_{it} \quad (2)$$

This specification is a nested alternative to equation (1) in that the two are identical when $\rho = 0$. For $0 < \rho < 1$, larger values of ρ imply greater persistence in the effects of media advertisements, whereas smaller values imply that shocks to candidate preferences die out quickly. We can also explore the time-decay of ads by variants of models (1) and (2) in which lagged media buys are included in the model.

These specifications may be augmented by including additional control variables. Of particular interest are the radio and television gross rating points purchased by the opposing Strayhorn campaign. Because the opponent's points were not randomly assigned, their causal effects are not identified. Including measures of Strayhorn's advertising nevertheless serves two purposes. First, it allows us to estimate the Perry ads' effects while controlling for whatever incidental correlation might exist between the randomly assigned Perry ads and their nonrandomly assigned counterparts. In operational terms, this simply involves adding control variables to the equations above – something that turns out to have no real consequence for the estimated effects of the Perry ads. Second, knowing where the Strayhorn ads were aired allows us to conduct an analysis that assesses whether the TRP ads were especially effective where

Strayhorn ads were absent. A further control variable is the average partisanship of a DMA. Partisanship is imputed based on the voter file, which indicates the number of Republican or Democratic primaries in which a person has voted. This measure is a strong predictor of candidate preference.[8] Since this variable predates the campaign, controlling for it dampens the sampling variability associated with our DMA-level averages of Y_{it}.

Results

This section begins by estimating simple models that ignore time-series dynamics and gradually builds up to more complex dynamic models. To preview the findings, we find little evidence of time-series dynamics. Thus, the simple models we present first tell the same substantive story as the more elaborate models that follow.

Table 3 reports the analysis of comparative candidate evaluations tracked on a daily basis. The six columns of Table 3 show how the estimated effects of television and radio advertising change across alternative modeling approaches. The first two columns report least-squares estimates with fixed effects for media market and day. Without covariates, the estimated effects of 1000 gross ratings points of television and radio advertising are .22 and .45, respectively. To put these estimates in perspective, it is helpful to bear in mind the fact that the campaign purchased up to 1000 GRPs of television advertising and 187 GRPs of radio ads. Thus, the maximum dosage of advertising would have boosted Perry's relative standing by approximately (1.0)(.22)+(0.187)(.45)=.30 scale points. Across the 580 observations, this outcome measure's standard deviation is .30 scale points. Advertising thus appears to have the capacity to induce a substantial shift in the relative evaluations of the candidates. This

[8] Note that, as expected, the media variables have no effect on past primary voting patterns when the latter is treated as the dependent variable. This analysis serves no substantive purpose but confirms the basic soundness of our experiment and analytic approach.

conclusion remains unchanged when one introduces controls for the airing of opposing ads or the partisan composition of the sample.

The next four columns of Table 3 explore the consequences of altering the model, method of estimation, and sample size. Column 3 shows what happens when fixed effects for each day are replaced by a smooth function, in this case a fifth-order polynomial in time. The use of a polynomial economizes on degrees of freedom but fails to account for jumps in the time-series resulting, for example, from differential response rates on weekends. Column 4 applies the multiplicative heteroskedasticity regression model described by Greene (2002, p.232-235), which parameterizes the disturbance variance as a function of the sample size used to measure the outcome in DMA *i* at time *t*. A simpler alternative to this approach is presented in column 5, which simply weights the data by sample size. The final column in Table 3 confirms that these results are robust to the decision to include the two nonexperimental media markets in the analysis. Dropping the Houston and Dallas DMAs increases the standard errors slightly but has little effect on the magnitude of the coefficients or their statistical significance. In all, Table 3 shows that across a variety of model specifications and estimation methods, one obtains highly significant effects of TV advertising and marginally significant effects of radio advertising.

Table 4 repeats this analysis, substituting vote preferences for candidate evaluations. The dependent variable is the percentage of respondents in market *i* at time *t* who intend to vote for Perry. In calculating this percentage, we included all respondents, even those who did not express a vote preference; but the results remain unchanged when these respondents are excluded from the calculations. Again, the six models reveal the effects of TV to be large, statistically significant, and robust. The weakest estimate suggests that 1000 GRPs boosted Perry's vote

share by 5.1 percentage-points. The main contrast between Tables 3 and 4 concerns radio. The effects of radio shown in Table 4 are uniformly weak and insignificant.

The same conclusions obtain when the data are aggregated weekly and modeled using first-differences, where *t* indexes week. Weekly aggregation is necessary to build up sufficient interviews in each DMA so that a first-difference is not overwhelmed by sampling variability. First-differencing means that the GRPs aired in a given week are subtracted from the GRPs aired in the preceding week. As the media campaign is rolled out, this difference is positive; as it continues, the difference is zero; when it concludes, the difference is negative. The same differencing procedure is also applied to the left-hand side of the equation, so that the resulting model (N=20, T=5) is:

$$Y_{it} - Y_{i,t-1} = \alpha + \beta_1(TV\ GRPs_{it} - TV\ GRPs_{i,t-1}) + \beta_2(Radio\ GRPs_{it} - Radio\ GRPs_{i,t-1}) + \gamma_1 Week_{it}^{1} + \ldots + \gamma_{J-1} Week_{it}^{J-1} + u_{it} \quad (3)$$

In effect, this specification assesses the effects of changing media volume on changes in public opinion. The use of fixed effects for DMA is inconsequential here, as first-differences effectively eliminate the means of the cross-sectional units.

While the results concerning radio are insignificant using first-differences, Table 5 again shows that TV exerts a strong and statistically significant effect on both candidate evaluations and vote choice. Figure 4 represents the multiple regression results for TV in graphical form, by showing how the residualized dependent variable (the variance that remains after regressing on all right-hand-side variables other than TV) relates to the residualized independent variable (the variance that remains in first-differenced TV after regressing on all right-hand-side variables). For both candidate evaluations and vote share, the introduction of TV produces a gain, while the withdrawal of TV leads to a loss. The plots show an approximately linear relationship between GRPs and average outcomes.

Dynamic models. Working with daily data, we find lagged dependent variables to be weak predictors, with ρ values hovering close to zero. This pattern is not altogether surprising, since survey error in daily data attenuates the estimate of ρ. What is surprising is the fact that ρ is found to be weak even when we consider the statewide time-series, for which we have large daily samples, even when we apply a Kalman filter, which makes explicit allowance for survey error (Green, Gerber, and DeBoef 1999). Applying the Kalman filter to daily results in an estimated ρ of .13; applying the filter to two-day aggregates, we obtain a ρ of .20. Apparently, the time-series tends to revert to its mean quickly.

Table 6 shows how the estimated effects of TV and radio change when dynamics are introduced. The simplest model includes a lagged dependent variable, as in equation (2). This specification has essentially no effect on the estimates. More interesting are the regressions that include lagged values of the media variables. Although including these lags introduces collinearity, thereby undercutting our ability to estimate the effect of any particular lag with precision, the overall pattern is nonetheless suggestive. Radio's effects appear to peak on the first day during which they air. Television's effects appear to peak with a one-day lag. Neither medium has any detectable effect after two days. (Introducing a spurious a one-day lead, as expected, shows no effect for either medium.) Overall, the analysis of dynamics suggests that the effects of TV, although powerful, are short-lived.

The fact that opinion equilibrates quickly in the wake of a televised ad campaign like the one studied here underscores the potential costliness of waging a battle for votes on TV. The immediate effect of a televised ad campaign makes it a cost-effective option in the short-run. To shift opinion 5% among an electorate of 14 million people means persuading 700,000 voters. To

expose the electorate to 1,000 GRPs of television at a rough average cost of $150 per point amounts to approximately $3 million per week. At less than $5 per vote, this expenditure is a bargain, but the difficulty is that there is little to show for one's money a week later.

Two-Sided versus One-Sided Communication. A recurrent theme in both public opinion research and campaign strategy is the hypothesis that media messages are especially effective when they go unrefuted. Scholars have argued that one-sided communication is unusually effective in swaying public opinion (Zaller 1992; Brody and Shapiro 1991). Practitioners argue that media campaigns are a matter of necessity, lest one's political opponents monopolize the airwaves and deploy their messages with special effectiveness.

The experimental study presented here may be thought of as two sub-experiments, one in areas where Strayhorn's ads aired and another where the Perry campaign's ads ran without any competition. The statistical question is whether we see evidence of stronger media effects in the latter experiment. Conveniently, the Strayhorn campaign deployed its TV ads in under half of the media markets. Based on 348 observations in which there were no Strayhorn ads, we find that Perry's TV ads increased his vote share by 5.4 percentage-points (SE=1.8), as compared to a 4.6 percentage-point gain (SE=2.6) in markets where Strayhorn ran her TV spots (N=232). Radio has no significant effect in either setting. Clearly, the two sets of estimates are not statistically distinguishable from one another, and the estimates grow even closer once one controls for the partisanship of the DMA. We therefore find minimal support for the hypothesis that ads are either more or less effective when running concurrently with opposition ads.

Discussion

The use of random assignment to assess the effects of mass media advertisements represents a major advance. This research method offers the first unbiased estimate of these causal effects in a field setting. Although ours is by no means the first study to find sizeable advertising effects, one should not dismiss it as merely reaffirming the results of laboratory and observational research. The large effects observed in the lab are routinely challenged on the grounds that they fail to tell us how media exposure translates into votes in the context of an actual campaign. The large effects found in observational studies are similar open to the charge that campaigns target their ads strategically, so that even perfect measures of media exposure may produce spurious effects. As Gerber, Green, and Kaplan (2004) point out, the contribution of field experiments may be underestimated because researchers using other research methods report only one measure of uncertainty, sampling variability, when in fact the standard errors of their estimates properly include uncertainty associated with modeling assumptions and with extrapolation from the lab to the outside world. When uncertainty is properly accounted for, random assignment of an actual campaign's mass media campaign provides an unusually convincing demonstration of causality.

From the standpoint of theory, the most provocative finding to emerge from this experiment concerns the rapid decay of advertising effects. In the course of a typical campaign, advertising targets are usually bombarded with messages without interruption; our experimental campaign enabled us to study how opinion equilibrates after the TV messages subside. This unusual design feature sheds light on the psychological mechanisms by which an ad designed to prime positive associations with the governor influences voter preferences. The pattern of abrupt change and equilibration appears to be inconsistent with a model of rational learning. Learning models can explain rapid change in the wake of new information, but they have difficulty

explaining the rapid change and reversion we see in our data. Extending the learning model to incorporate priming or affect seems necessary to account for the observed pattern.

This finding opens up a new research agenda. The theoretical arm of this broader research project involves formally characterizing the properties of different types of ads and the mechanisms through which they influence voter preference. Rational learning models are relatively well-developed formally; more work needs to focus on formalizing propositions about the priming and emotion-evoking effects of advertising. This paper has taken a step in that direction, proposing alternatives to rational learning that might account for the evanescent effects of television ads.

This theoretical enterprise must be complemented by empirical exploration. Progress requires assessment of a range of different ads' effects in different political contexts. Do the dynamics of opinion change look different when, instead of offering evocative imagery, an advertisement campaign reveals new information? Would, for example, a negative campaign ad that leveled a specific charge about an opponent's malfeasance in office produce an enduring shift in opinion? Do the effects of emotion-evoking ads dissipate more rapidly than ads with comparable content that lack emotional content? Are the effects of advertising muted toward the end of a campaign, when voters have acquired more information about the candidates, or do the priming or emotional mechanisms remain undiminished in influence?

As the empirical consequences of different types of ads come to be better understood, scholars will be better positioned to develop arguments about the strategic logic of different types of advertising campaigns. The current study suggests that valence appeals have the capacity to influence large segments of the electorate, but while such ads may succeed in associating a candidate with a popular value in the short run, they face the challenge of

presenting information that will leave a lasting impact on the way that voters think about the candidates. Hard-hitting attack ads (e.g., charges of corruption) also appeal to valence dimensions but may adduce evidence that has a more enduring impact on the way that a candidate is evaluated. From a strategic standpoint, the latter strategy seems riskier, which may explain why the incumbent front-runner studied here, who faced an array of challengers from across the ideological spectrum, chose to open his campaign with a positive broad-based appeal.

It should come as no surprise that a study such as this one generates as many research questions as it answers. This experiment is among the first of its kind in any discipline to estimate the effects of a large-scale media campaign using random assignment. At a minimum, this study must be replicated in other political contexts in order to answer basic questions, such as how much the effectiveness of advertising depends on the tone and content of the ads, proximity to Election Day, and the competitiveness of the race. In a similar vein, the relative effectiveness of radio and television advertising awaits further experimentation, given the limited power of the current study to detect radio's influence.

That said, the current study represents an important advance, both because it shows that large-scale experiments of this sort are possible and because it for the first time provides reliable measurement of causal effects. Although, like any scientific effort, this study requires replication and refinement, the causal estimates it generates force a rethinking of several key topics in the field of political psychology. If one-sided information is especially influential, why were Perry's ads undiminished in effectiveness when presented in markets that aired Strayhorn ads? If campaign ads have powerful but short-lived effects on voter preference, what are we to conclude about the "minimal effects" thesis? It can hardly be said that people are so inattentive or resistant to new information that they fail to change, but on the other hand – consistent with

models of priming or emotion but inconsistent with on-line processing – the effects of these messages appear to fade quickly. The advent of field experimental investigation of campaign advertising's effects, in sum, has the potential to profoundly shape the empirical foundation on which theories of political communication rest.

References

Achen, Christopher H. 1992. Breaking the Iron Triangle: Social Psychology, Demographic Variables and Linear Regression in Voting Research. *Political Behavior* 14 (3): 195-211.

Ansolabehere, S., and S. Iyengar. 1995. Going Negative: How Attack Ads Shrink and Polarize the Electorate. New York: Free Press.

Ansolabehere, Stephen D., Shanto Iyengar, and Adam Simon. 1999. Replicating Experiments Using Aggregate and Survey Data: The Case of Negative Advertising and Turnout. *The American Political Science Review* 93:901-9.

Arceneaux, Kevin, and Gregory Huber. 2006. Uncovering the Persuasive Effects of Presidential Advertising. Unpublished manuscript, Yale University.

Brader, Ted. 2005. Striking a Responsive Chord: How Political Ads Motivate and Persuade Voters by Appealing to Emotions. *American Journal of Political Science* 49(2): 388-405.

Brody, Richard, and Catherine R. Shapiro. 1991. A Reconsideration of the Rally Phenomenon in Public Opinion. In *Political Behavior Annual*, Volume 2, edited by S. Long. Boulder: Westview Press.

Burke, Raymond R., and Thomas K. Srull. 1988. Competitive Interference and Consumer Memory for Advertising. *Journal of Consumer Research* 15(1): 55-68.

Clinton, Joshua D., and John S. Lapinski. 2004. "Targeted" Advertising and Voter Turnout: An Experimental Study of the 2000 Presidential Election. *Journal of Politics* 66 (1), 69–96.

Freedman, Paul, Michael Franz, and Kenneth Goldstein. 2004. Campaign Advertising and Democratic Citizenship. *American Journal of Political Science* 48:723-41.

Freedman, Paul, and Ken Goldstein. 1999. Measuring Media Exposure and the Effects of Negative Campaign Ads. *American Journal of Political Science* 43:1189-1208.

Gaines, Brian J., James H. Kuklinski, and Paul J. Quirk. 2007. The Logic of the Survey Experiment Reexamined. *Political Analysis* 15: 1-20.

Gerber, Alan S., and Donald P. Green. 1998. Rational Learning and Partisan Attitudes. *American Journal of Political Science*. 42: 794-818.

Gerber, Alan S., Donald P. Green, and Edward H. Kaplan. 2004. The Illusion of Learning from Observational Research. In Ian Shapiro, Rogers Smith, and Tarek Massoud, eds., *Problems and Methods in the Study of Politics*. New York: Cambridge University Press, pp. 251-73.

Gilliam, Jr., Franklin D. and Iyengar, Shanto. 2000. Prime Suspects: The Influence of Local Television News on the Viewing Public. *American Journal of Political Science* 44:

560-574.

Green, Donald P., Alan S. Gerber, and Suzanna L. De Boef. 1999. Tracking Opinion over Time: A Method for Reducing Sampling Error. *Public Opinion Quarterly* 63:178-92.

Greene, William H.. 2002. *Econometric Analysis*. New York: Prentice Hall.

Hovland, C., Lumsdaine, A. and Sheffield, F. 1949. *Experiments in Mass Communication: Studies of Social Psychology in World War II*. Ann Arbor, Association for Consumer Research.

Iyengar, Shanto, Mark Peters, and Donald R. Kinder. 1982. Experimental Demonstration of the Not-So-Minimal Consequences of Television News Programs. *American Political Science Review* 76: 848–858.

Johnston, Richard, Michael G. Hagen, and Kathleen Hall Jamieson. 2004. *The 2000 Presidential Elections and the Foundations of Party Politics*. Cambridge, England: Cambridge University Press.

Kaid, Lynda Lee. 1997. Effects of the Television Spots on Images of Dole and Clinton. *American Behavioral Scientist* 40(8): 1085-1094.

Klapper, J. T. 1960. *The Effects of Mass Communications*. New York: Free Press.

Krasno, Jonathan S., and Donald P. Green. 2005. Do Televised Presidential Ads Increase Voter Turnout? Evidence from a Natural Experiment. Unpublished manuscript, Yale University.

Lazarsfeld, Paul F., Bernard Berelson, and H. Gaudet. 1944. *The people's choice: How the voter makes up his mind in a presidential campaign*. New York: Columbia University Press.

Lodge, Milton, Marco Steenbergen, and Shawn Brau. 1995. "The Responsive Voter: Campaign Information and the Dynamics of Candidate Evaluation." *American Political Science Review* 89(2):309-26.

Miller, Joanne M, and Jon A. Krosnick. 1996. News Media Impact on the Ingredients of Presidential Evaluations: A Program of Research on the Priming Hypothesis. In Diana Mutz, et al. (eds.), *Political Persuasion and Attitude Change*. Ann Arbor: University of Michigan Press.

Neumann, Russell. 1986. *The Paradox of Mass Politics: Knowledge and Opinion in the American Electorate*. Cambridge: Harvard University Press.

Panagopoulos, Costas, and Donald P. Green. 2006. The Impact of Radio Advertisements on Voter Turnout and Electoral Competition. Paper presented at the Annual Meeting of the Midwest Political Science Association, Chicago, IL, April 20-23.

Patterson, Thomas E., and Robert D. McClure. 1976. *The Unseeing Eye: The Myth of Television Power in National Elections*. New York: G.P. Putnam's Sons.

Sears, David O., and R. E. Whitney. 1973. *Political Persuasion*. Morristown, NJ: General Learning Press. 1973.

Shaw, Daron R. 1999. The Effect of TV Ads and Candidate Appearances on Statewide Presidential Votes, 1988-96. *The American Political Science Review* 93:345-61.

Shaw, Daron R. and Bartholomew H. Sparrow. 1999. "From the Inner-Ring Out: News Congruence, Cue-Taking, and Campaign Coverage." *Political Research Quarterly.* 52 2:323-352.

Stokes, Donald E.. 1963. "Spatial Models of Party Competition." *American Political Science Review* 57(2):368-77.

Vakratsas, Demetrios, and Tim Ambler. 1999. How Advertising Works: What Do We Really Know? *Journal of Marketing* 63(1): 26-43

Valentino, Nicholas A., Vincent L. Hutchings, and Dmitri Williams. 2004. The Impact of Political Advertising on Knowledge, Internet Information Seeking, and Candidate Preference. Journal of Communication 54(2): 337-354.

Vavreck, Lynn, and Donald P. Green. 2005. Mobilizing Voters Through TV Public Service Announcements: A Field Experiment Unpublished report prepared for the Center for Information and Research on Civic Learning and Engagement.

Zaller, John. 1992. *The Nature and Origins of Mass Opinion*. New York: Cambridge University Press.

Zaller, John. 1996. The Myth of Massive Media Impact Revived: New Support for a Discredited Idea. In Diana Mutz, et al. (eds.), *Political Persuasion and Attitude Change*. Ann Arbor: University of Michigan Press.

Zielske, Hubert, and Walter Henry. 1980. Remembering and Forgetting Television Advertisements. *Journal of Advertising Research* 20(2): 7-13.

Appendix A: Text of Campaign Ads, Survey Instrument

Perry Radio:

I've never been more proud to call myself a Texan. This is Governor Rick Perry. In Texas we've set the national standard for economic development. We've gained 300,000 new jobs. Recently, we're called the best business climate nationwide. Our lawsuit reform's protecting our small businesses and ensuring that decisions about your healthcare are made by you and your doctor. In Texas, we've invested ten billion new dollars in our public schools while improving standards of accountability and student performance. We were the first state to make the college prep coursework the standard curriculum in our schools. And Texas is one of only three states to make significant improvement on reading and math exams. Our people are compassionate. Our vision, bold. Our values, strong. The best is yet to come. I'm proud of Texas. How 'bout you? Political ad paid for by Texans for Rick Perry. For more information, visit proudoftexas.com.

Survey Instrument

Candidate Favorability

I'd like to read you the names of some people who will be running for governor this year. For each one, please tell me whether you have heard of that person and if so, whether you have a positive or a negative impression of that person. If you do not recognize the name, just let me know. Here is the first name...Rick Perry.

[Wait for response, then ask:] Would that be a strong or just somewhat (positive/negative) impression?

...Carole Keeton Strayhorn… Would that be a strong or just somewhat (positive/negative) impression?

Ballot Test

If this year's election for governor of Texas were held today, and you had to make a choice, which of the following candidates would you vote for? Rick Perry, the Republican candidate; Carole Keeton Strayhorn, an Independent candidate; Richard F. "Kinky" Friedman an Independent Candidate; or whichever Democratic candidate wins the nomination?

Campaign Recall

Thinking about Rick Perry for a moment -- Have you seen, read, or heard anything in the last three or four days[9] about Rick Perry on TV or Radio?

[9] After the first day of interviewing, the time frame was changed from "recently" to "in the last three or four days."

Appendix B: Formal Models of Rational Learning, Priming, and Emotional Response

Having sketched out the contours of the Bayesian learning model and its extensions in the text, we now lay out the argument in more formal detail. Following Gerber and Green (1998), voters seek to make inferences about the true attributes of the politician. By attributes we include things like the politician's level of competence, attractiveness, as well as the match between the politician's attributes and the political and policy making challenges he or she faces. Voters begin with some assessment of the politician. This initial belief, which is denoted a_0, has a variance, P_0, which captures the voter's initial uncertainty about the politician. The true attributes of the politician at time t are:

$$\alpha_t = \alpha_{t-1} + \eta_t \ , \qquad (1)$$

where η is an innovation term, independently distributed N(0,q). Since our empirical application covers a relatively short period of time (the interval of three weeks during which the ads were aired and the weeks on either side of this period), it is expected that the politician's true underlying attributes will be nearly fixed, which implies that q will be relatively small. In each period the voter gets some information that can be used to produce an updated assessment of the politician. At time t the voters observe a noisy measure of the true attributes of the politician:

$$y_t = \alpha_t + \varepsilon_t \,, \qquad (2)$$

where ε_t is an error term, independently distributed N(0,h). The voter updates beliefs using the new information, y_t. It can be shown that the assessment of the politician's current attributes that minimizes the mean squared error is:

$$a_t = a_{t-1} + K_t(y_t - a_{t-1}), \qquad (3)$$

Where $K_t = \dfrac{P_{t-1} + q}{P_{t-1} + q + h}$. The variance of the estimator a_t is P_t, where $P_t = hK_t$.

Is Bayesian learning compatible with rapid opinion change and re-equilibration in the wake of an advertisement? Suppose the information environment were modeled as a function of whether the politician is engaging in a media campaign. To account for this possibility, we modify equation (2) to permit the information environment to vary over time:

$$y_t = \alpha_t + Z \ + \varepsilon_t, \qquad (2)'$$

Equation (2') allows the media campaign to shift the information the voter observes by Z and also admits the possibility that the media campaign alters the amount of noise associated with the voter's observation by altering h, the variance of ε.[10]

Using (2') and (3), we can describe how the voter responds to a media campaign. Assume the media campaign begins period t, lasts through period t+1, and then stops. Substituting y_t into (2), the change in voter assessments between period t-1 and period t is:

$$a_t - a_{t-1} = K_t Z + K_t(\alpha_t + \varepsilon_t - a_{t-1}). \qquad (4)$$

Using (3) and (4) we can evaluate whether the basic pattern of media effects observed in our experiment is consistent with the learning model described by the updating equation. Suppose we observed that, following an initial jump, opinion stabilized at a new level during the media campaign and then returned to the original opinion level immediately after the media campaign stopped. Translating these empirical findings into the formulas used in the learning model, the change in opinion during the media campaign (periods t to t+1) is then:

$$a_{t+1} - a_t = K_{t+1}(y_{t+1} - a_t) = 0, \qquad (5)$$

which implies that: $y_{t+1} = a_t$.

[10] We assume that the voters do not adjust their assessments to account for the more positive message flow denoted by Z that the voter observes due to the media campaign; if such adjustment is partial, then Z is what remains after the partial adjustment.

Next, the effect of initiating the media campaign can be described using the model. The change in opinion from the period before the media campaign (t-1) to the first period of the campaign (t) is:

$$y_{t+1} = a_{t-1} + K_t(y_t - a_{t-1}). \quad (6)$$

Thus, y_t can be written as:

$$y_{t+1} = \alpha_t + \eta_t + Z + \varepsilon_{t+1} = y_t + \varepsilon, \quad (7)$$

where ε is a collection of error terms. Combining (5) and (7) yields:

$$y_t = (1 - K_t)a_{t-1} + K_t(y_t) - \varepsilon \quad (8)$$

The assessments can be scaled so that ε is small relative to *a*. The flow of information, y_t , will differ from the assessment made prior to the media campaign, a_{t-1}, since this assessment does include the periods when the information flow included Z. As a result, equation (7) is satisfied only when K_t is approximately equal to 1. By a similar logic, the empirical result that the assessment returns to the pre-media campaign assessment can be shown to imply that $K_{t+2} \approx 1$.

For $K_t = 1$, it must be that h is 0. In words, when h=0 then y_t, the instantaneous flow of information, provides a perfect signal of the true underlying trait, in which case voters rely on this signal to the exclusion of previous information. Not only does this value of h seem implausible, it also produces an empirical anomaly for the learning model. Under the assumption that the innovations in the candidate's underlying traits is very small (q near 0), if h=0 then all y_t should equal y_{t+1}. In other words, all assessments should remain roughly constant. In fact, the assessment moves noticeably during the media campaign. If it is assumed that h=0, then it cannot

also be the case that q is near 0, so h=0 leads to a contradiction between what is observed and the model's assumptions.

In order to address this anomaly, one may extend the model in either of two ways. The first is consistent with the idea that advertising primes certain evaluative criteria. Consider a utility function $U(\gamma a_1, a_0)$, where $\gamma(Z)$ is the weight placed on consideration 1, and Z=1 when the media campaign is on the air and 0 otherwise. Let $\gamma(1) > \gamma(0)$. In this case, the advertising campaign has the effect of switching which assessment, a_1 or a_0, is used when forming an evaluation of the candidate. One can easily extend this line of analysis to include many possible types of evaluative criteria and patterns of decay whereby the cognitive accessibility of a primed criterion dissipates over time.

A second approach is to augment the model so that the voter forms evaluations based on two inputs, one rational and the other emotional. On the rational side, the voter's assessment of the politician's current attributes minimizes the mean squared error, following the logic of (3) above. However, the overall assessment equals:

$$U_t(X) = a_t + \gamma D(Z),$$

where D(Z) is the media campaign's contribution to positive feelings about the candidate. The form of the emotion-related function D(Z) can be fairly general, can include lags, etc. Under conditions where the function D(Z) decays exponentially with time as the voter's emotional reaction returns to its baseline level, an advertising campaign might have little effect on the Bayesian component a_t but still produce a temporary shift in utility as a result of changes in D(Z).

Table 1: Illustration of Programming on which Broadcast TV Ads Aired in Three DMAs, by Gross Ratings Points

250 GRPs	500 GRPs	1000 GRPs
5a Morning News	5a Morning News	5a News
6a News	6a News	6a News
7a Today Show	7a Today Show	7a Today Show
7a Good Morn Amer.	7a Good Morn Amer.	7a Good Morn Amer.
7a Early Show	7a Early Show	7a Early Show
9a Regis	9a Regis	9a Regis
Noon News	10a Price is Right	9a Ellen
5p News	10a The View	10a Price is Right
6p News	11a Young & Restless	11a Young & Restless
630p Wheel of Fort.	Noon News	Noon News
630p Millionaire	3p Dr. Phil	1230 Soaps
10p News	4p Oprah	3p Dr. Phil
Sunday Today	430 Jeopardy	4p Oprah
Sunday AM News	5p News	430 Jeopardy
Sunday 5p News	6p News	5p News
Dateline 6-7p	630p Wheel of Fort.	6p News
60 Minutes	630p Enter. Tonight	630p Wheel of Fort.
Sunday 10p News	630p Millionaire	630p Millionaire
Friday 20/20 9-10p	9p News	9p News
	10p News	10p News
	1030pTonight Show	1030pTonight Show
	Sunday Today	1030p Letterman
	Sun Meet the Press	1030p Nightline
	Sunday AM News	Sat Jeopardy
	Sunday 5p News	Sat 6p News
	60 Minutes	Sat 10p news
	Sunday 10p News	Sunday Today
	Fri 20/20 9-10p	Sunday AM News
	Fri Law & Order 9-10p	Sunday Today
	Sat Cops	Meet the Press
	Sat Am Most Wanted	Face the Nation
	Tues Navy NCIS	This Week
		NFL Playoff Game
		Sunday 5p News
		Dateline6-7p
		60 Minutes6-7p
		Ext. Home Makeover
		Sun Cold Case 8-9p
		Law & Order9-10p
		Sun Movie 8-10p
		Sunday 10p News
		Sunday Sports Ext.
		Tues Navy NCIS
		Weds Law & Order
		Thurs. CSI
		Thurs Primetime Live

Table 2: Survey Results Over Time

(a) Ballot Test Results by Week of Study

Unweighted N	January 5-11	January 12-18	January 19-25	January 26-29	March 5-6
Perry	33.8%	32.8%	32.7%	31.7%	33.7%
Strayhorn	13.3%	14.1%	14.9%	15.3%	11.4%
Friedman	4.1%	4.0%	4.5%	4.7%	5.3%
Democrat	11.5%	12.8%	13.0%	12.2%	12.0%
Other	2.0%	3.1%	3.7%	2.3%	3.2%
Don't Know	35.3%	33.2%	31.3%	33.7%	34.4%
Total %	100.0%	100.0%	100.0%	100.0%	100.0%
N of Cases	7040	7059	7108	4032	2044

(b) Favorability Ratings of Rick Perry

	January 5-11	January 12-18	January 19-25	January 26-29	March 5-6
Strong Positive	30.3%	33.3%	34.9%	35.1%	38.5%
Weak Positive	14.1%	12.9%	13.0%	11.8%	17.1%
No opinion	33.2%	31.8%	29.5%	29.5%	24.1%
Weak Negative	7.1%	6.7%	6.4%	5.7%	6.8%
Strong Negative	15.2%	15.3%	16.2%	18.0%	13.5%
Total %	100.0%	100.0%	100.0%	100.0%	100.0%
N of Cases	7157	7178	7221	4087	2044

(c) Favorability Ratings of Carole Keeton Strayhorn

	January 5-11	January 12-18	January 19-25	January 26-29	March 5-6
Strong Positive	16.2%	20.1%	23.8%	22.9%	21.1%
Weak Positive	11.7%	12.2%	11.9%	10.7%	16.2%
No opinion	57.7%	52.1%	47.8%	48.8%	45.7%
Weak Negative	5.9%	5.7%	5.6%	4.7%	7.1%
Strong Negative	8.6%	10.0%	11.0%	13.0%	9.9%
Total %	100.0%	100.0%	100.0%	100.0%	100.0%
N of Cases	7094	7113	7146	4056	2044

See appendix for question wording. Results here are unweighted for sampling probabilities, but weighted results are similar.

Table 3: Estimates of TV and Radio Advertising's Effects on Candidate Evaluations, Measured Daily

Independent Variables	Model I: No Covariates OLS	Model II: With Covariates OLS	Model III: Polynomial Time Trend OLS	Model IV: Heteroskedastic Errors MLE	Model V: Weights Based on Sample Size WLS	Model VI: Exclude Two Nonexperimental DMAs WLS
TV GRPs (in 1000s)	.218**	.229**	.186**	.177**	.214**	.220**
(Standard Error)	(.061)	(.061)	(.056)	(.054)	(.060)	(.062)
Radio GRPs (in 1000s)	.448*	.449*	.384*	.329	.408*	.382*
(Standard Error)	(.218)	(.217)	(.214)	(.203)	(.212)	(.219)
Polynomial Time Trend	No	No	Yes	Yes	No	No
Fixed Effects for Day	Yes	Yes	No	No	Yes	Yes
Fixed Effects for Media Markets	Yes	Yes	Yes	Yes	Yes	Yes
Controls for Strayhorn TV and Radio	No	Yes	No	Yes	No	Yes
Controls for Voter Partisanship	No	Yes	No	Yes	No	Yes
N	580	580	580	580	580	522

** p < .01. *p < .05, one-sided test. Dependent variable is the difference between evaluations given to Strayhorn and Perry. The outcome measure ranges from -4 to +4, with high scores indicating support for Perry, but the daily means have a standard deviation of 0.30. Partisanship is measured as the number of GOP primaries in which the voter has voted, minus the number of Democratic primaries. Estimating robust cluster standard errors, where clustering occurs at the level of the experimental week within each media market, produces standard errors that are slightly *smaller* than conventional standard errors. Attributing this to sampling variability, we report the larger standard errors generated by OLS. Polynomial time trend is a 5^{th} order polynomial of the number of days since the start of the study.

Table 4: Estimates of TV and Radio Advertising's Effects on Vote Preference, Measured Daily

Independent Variables	Model I: No Covariates OLS	Model II: With Covariates OLS	Model III: Polynomial Time Trend OLS	Model IV: Heteroskedastic Errors MLE	Model V: Weights Based on Sample Size WLS	Model VI: Exclude Two Nonexperimental DMAs WLS
TV GRPs (in 1000s)	5.14**	5.31**	5.17**	5.07**	5.09**	5.05**
(Standard Error)	(1.48)	(1.47)	(1.42)	(1.36)	(1.45)	(1.52)
Radio GRPs (in 1000s)	0.25	0.45	0.19	-0.38	0.45	0.76
(Standard Error)	(5.34)	(5.28)	(5.44)	(5.12)	(5.17)	(5.38)
Polynomial Time Trend	No	No	Yes	Yes	No	No
Fixed Effects for Day	Yes	Yes	No	No	Yes	Yes
Fixed Effects for Media Markets	Yes	Yes	Yes	Yes	Yes	Yes
Controls for Strayhorn TV and Radio	No	Yes	No	Yes	No	Yes
Controls for Voter Partisanship	No	Yes	No	Yes	No	Yes
N	580	580	580	580	580	522

** $p < .01$. Dependent variable is the percent of all respondents (including "don't know" responses) favoring Perry in a trial ballot. Partisanship is measured as the number of GOP primaries in which the voter has voted, minus the number of Democratic primaries. Estimating robust cluster standard errors, where clustering occurs at the level of the experimental week within each media market, produces standard errors that are slightly *smaller* than conventional standard errors. Attributing this to sampling variability, we report the larger standard errors generated by OLS. Polynomial time trend is a 5^{th} order polynomial of the number of days since the start of the study.

Table 5: First-Difference Estimates of TV and Radio Advertising's Effects on Candidate Evaluations and Vote Preference, Measured Weekly

Independent Variables	Candidate Evaluations	Vote Choice
TV GRPs (in 1000s)	.184**	5.49**
(Standard Error)	(.066)	(1.54)
Radio GRPs (in 1000s)	.319	-3.59
(Standard Error)	(.232)	(5.38)
Week 1 Dummy	-.284**	-.08**
Week 2 Dummy	-.187**	-.04**
Week 3 Dummy	-.184**	-.03
Week 4 Dummy	-.023	-.01
N	100	100

** $p < .01$. Dependent variables are defined in Tables 3 and 4, and the model specification is given in equation (3). The N consists of 20 DMAs observed over 6 weekly periods. Differencing eliminates the first week's observations.

Table 6: Dynamic Effects of TV and Radio Advertising, Measured Daily

Independent Variables	Dependent Variable: Candidate Evaluations				Dependent Variable: Vote Preference			
TV (no lag)	.240	.218	.102	.092	5.41	5.14	1.64	1.42
	(.063)	(.061)	(.109)	(.110)	(1.51)	(.1.48)	(2.68)	(2.69)
TV, one day lag			.140	.234			4.21	6.77
			(.110)	(.148)			(2.68)	(3.62)
TV, two day lag				-.010				-2.84
				(.110)				(2.69)
Radio (no lag)	.475	.448	.472	.471	.89	.25	.16	1.04
	(.022)	(.218)	(.386)	(.389)	(5.38)	(5.34)	(9.45)	(9.50)
Radio, one day lag			-.044	.049			-.37	-.63
			(.386)	(.522)			(9.45)	(12.87)
Radio, two day lag				-.010				.66
				(.388)				(9.49)
Lag Outcome	-.077				-.127			
	(.044)				(.044)			

N=580, except for specifications with lagged outcomes, for which N=560.

Table 7: Interaction Between the Effects of Perry Advertising and the Presence of Opposing Ads

	Candidate Evaluations		Vote Choice	
No Covariates	**No Opposition Ads**	**Opposition Ads**	**No Opposition Ads**	**Opposition Ads**
TV GRPs (in 1000s)	.235	.206	5.44	4.59
(Standard Error)	(.074)	(.109)	(1.82)	(2.63)
Radio GRPs (in 1000s)	.387	.034	1.34	-4.21
(Standard Error)	(.281)	(.343)	(6.95)	(8.27)
Controlling for Party and Opposing ads				
TV GRPs (in 1000s)	.248	.254	5.74	5.53
(Standard Error)	(.073)	(.114)	(1.80)	(2.72)
Radio GRPs (in 1000s)	.381	.250	1.10	-4.96
(Standard Error)	(.278)	(.355)	(6.80)	(8.49)
N	348	232	348	232

All models control for day and DMA fixed effects. Data are weighted by survey sample size.

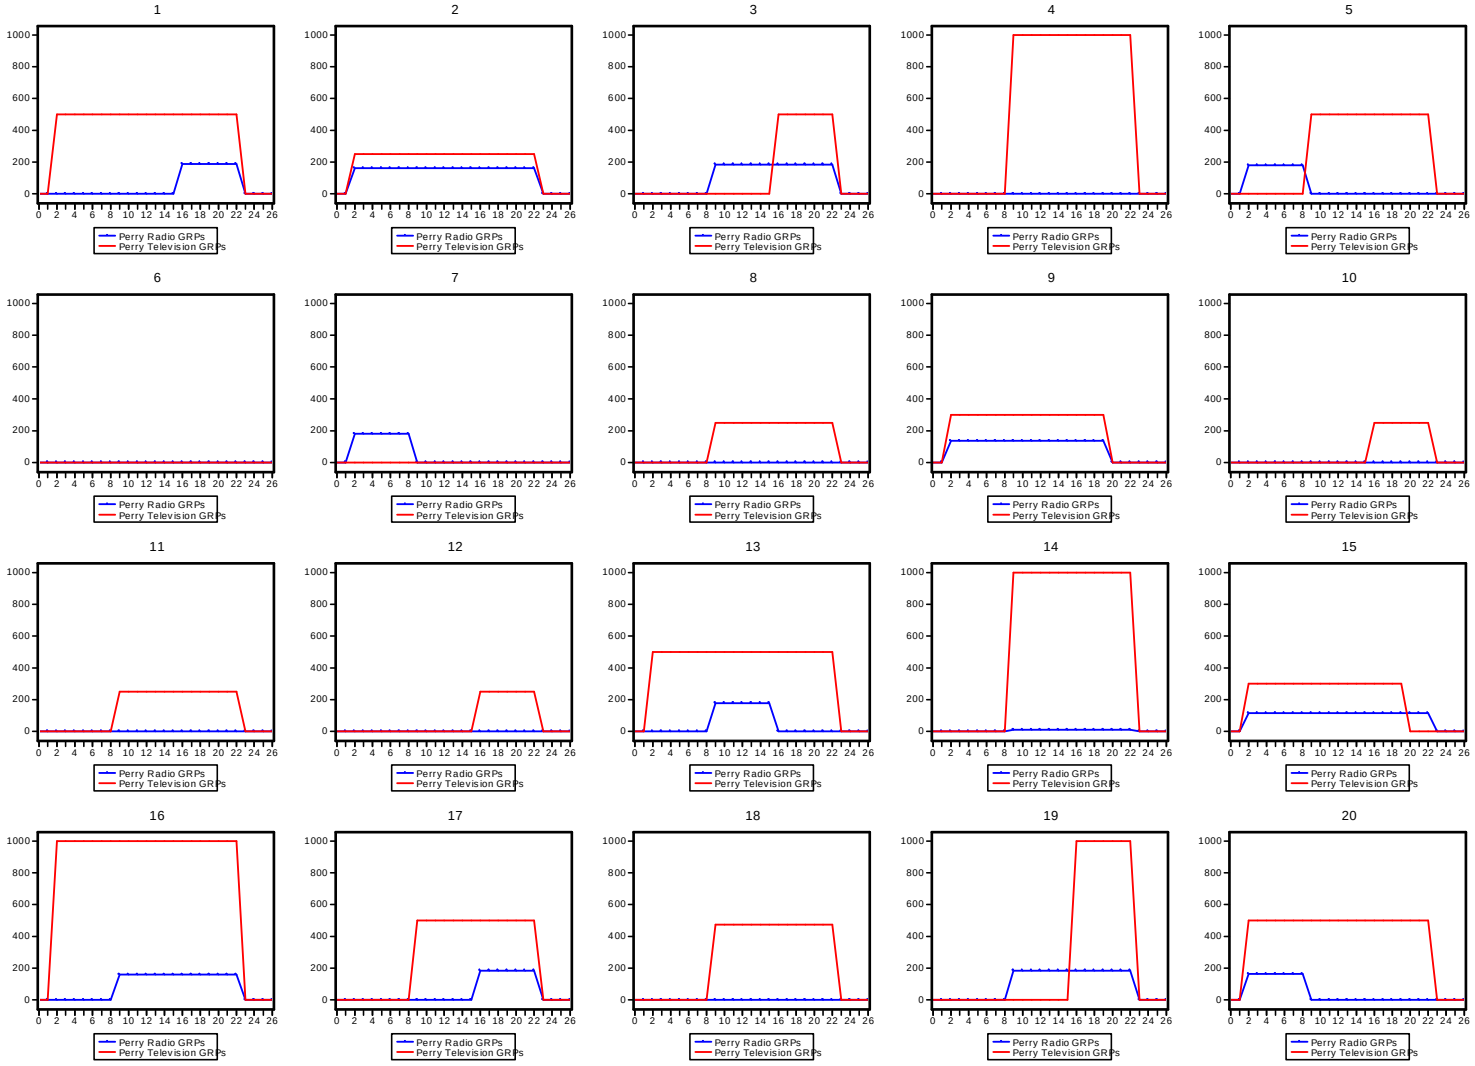

Figure 1: Television and Radio Purchases by Perry Campaign, by Media Market and Day

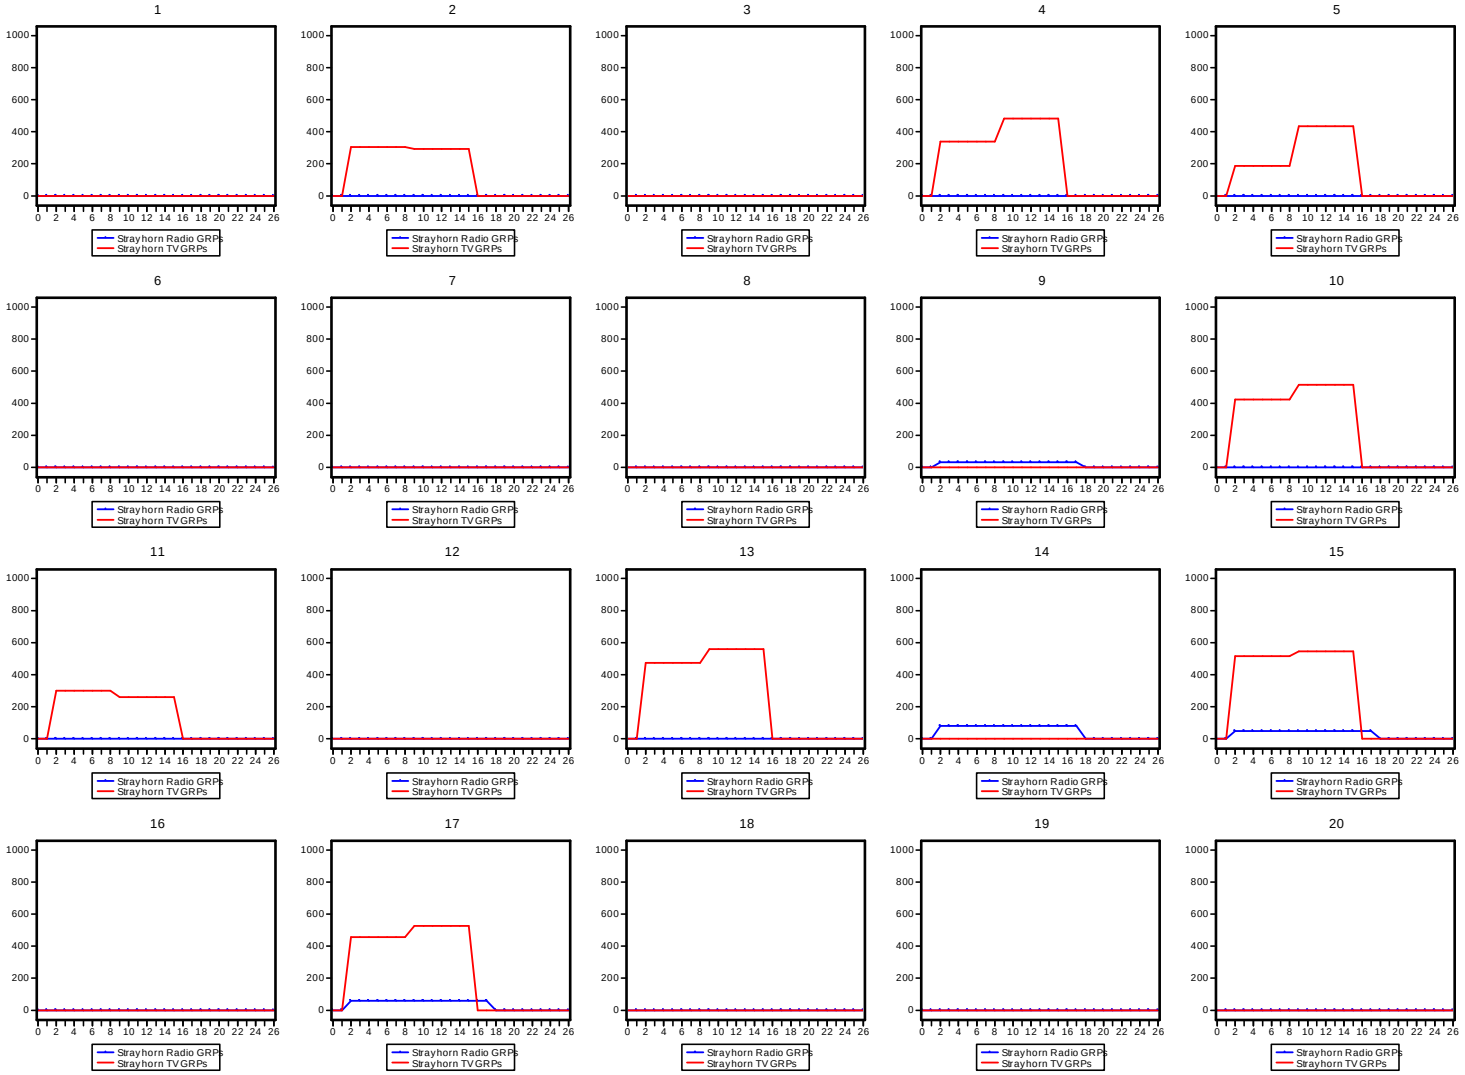

Figure 2: Television and Radio Purchases by Strayhorn Campaign, by Media Market and Day

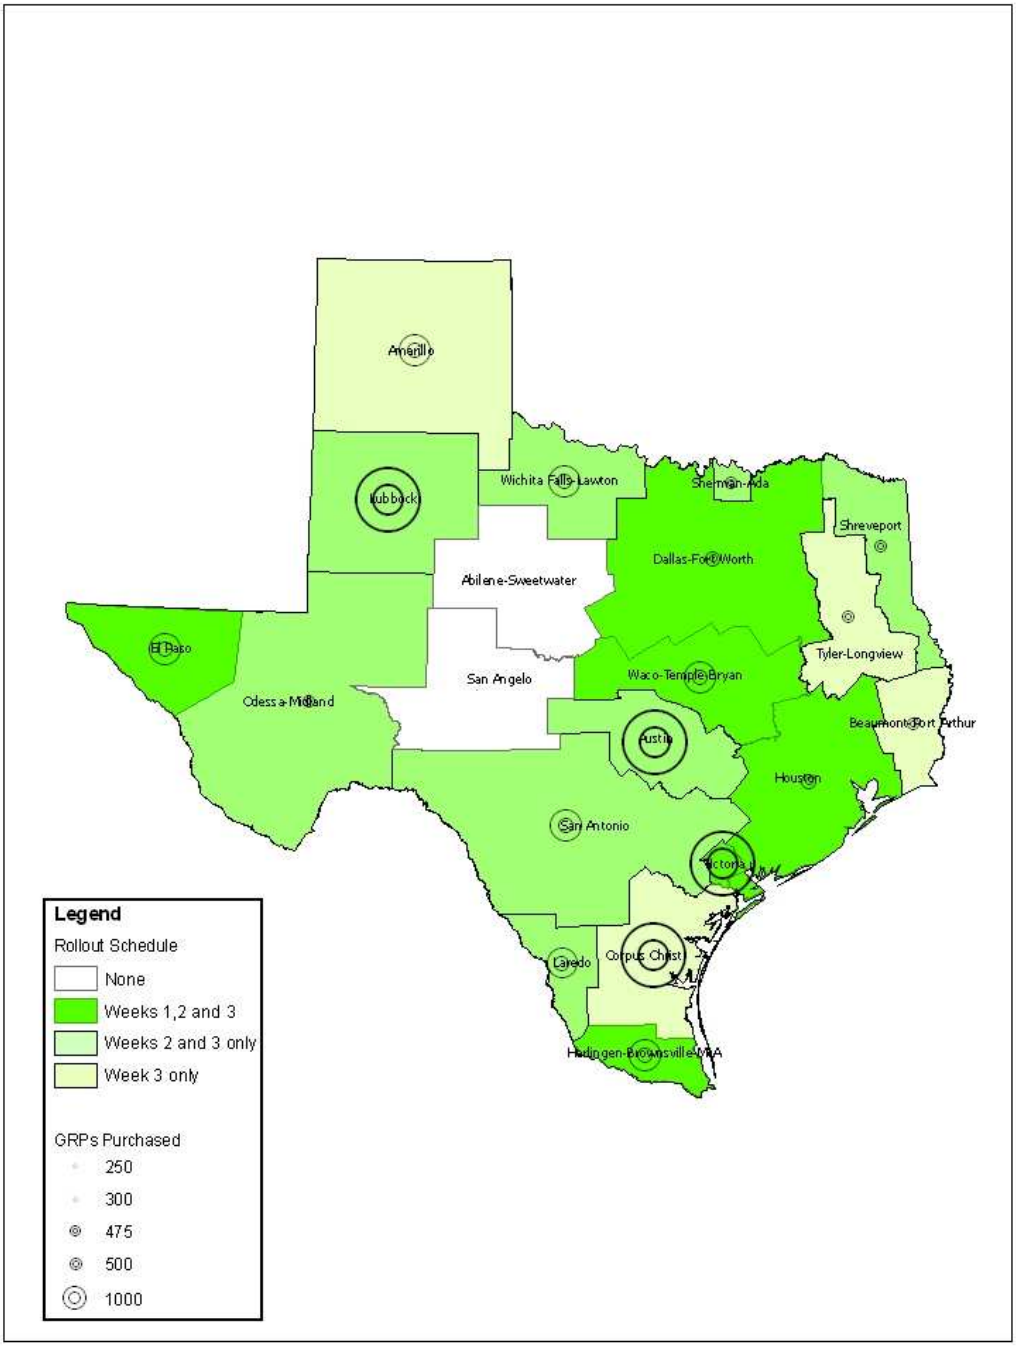

Figure 3: Broadcast Television Advertising Rollout Schedule and GRPs Purchased

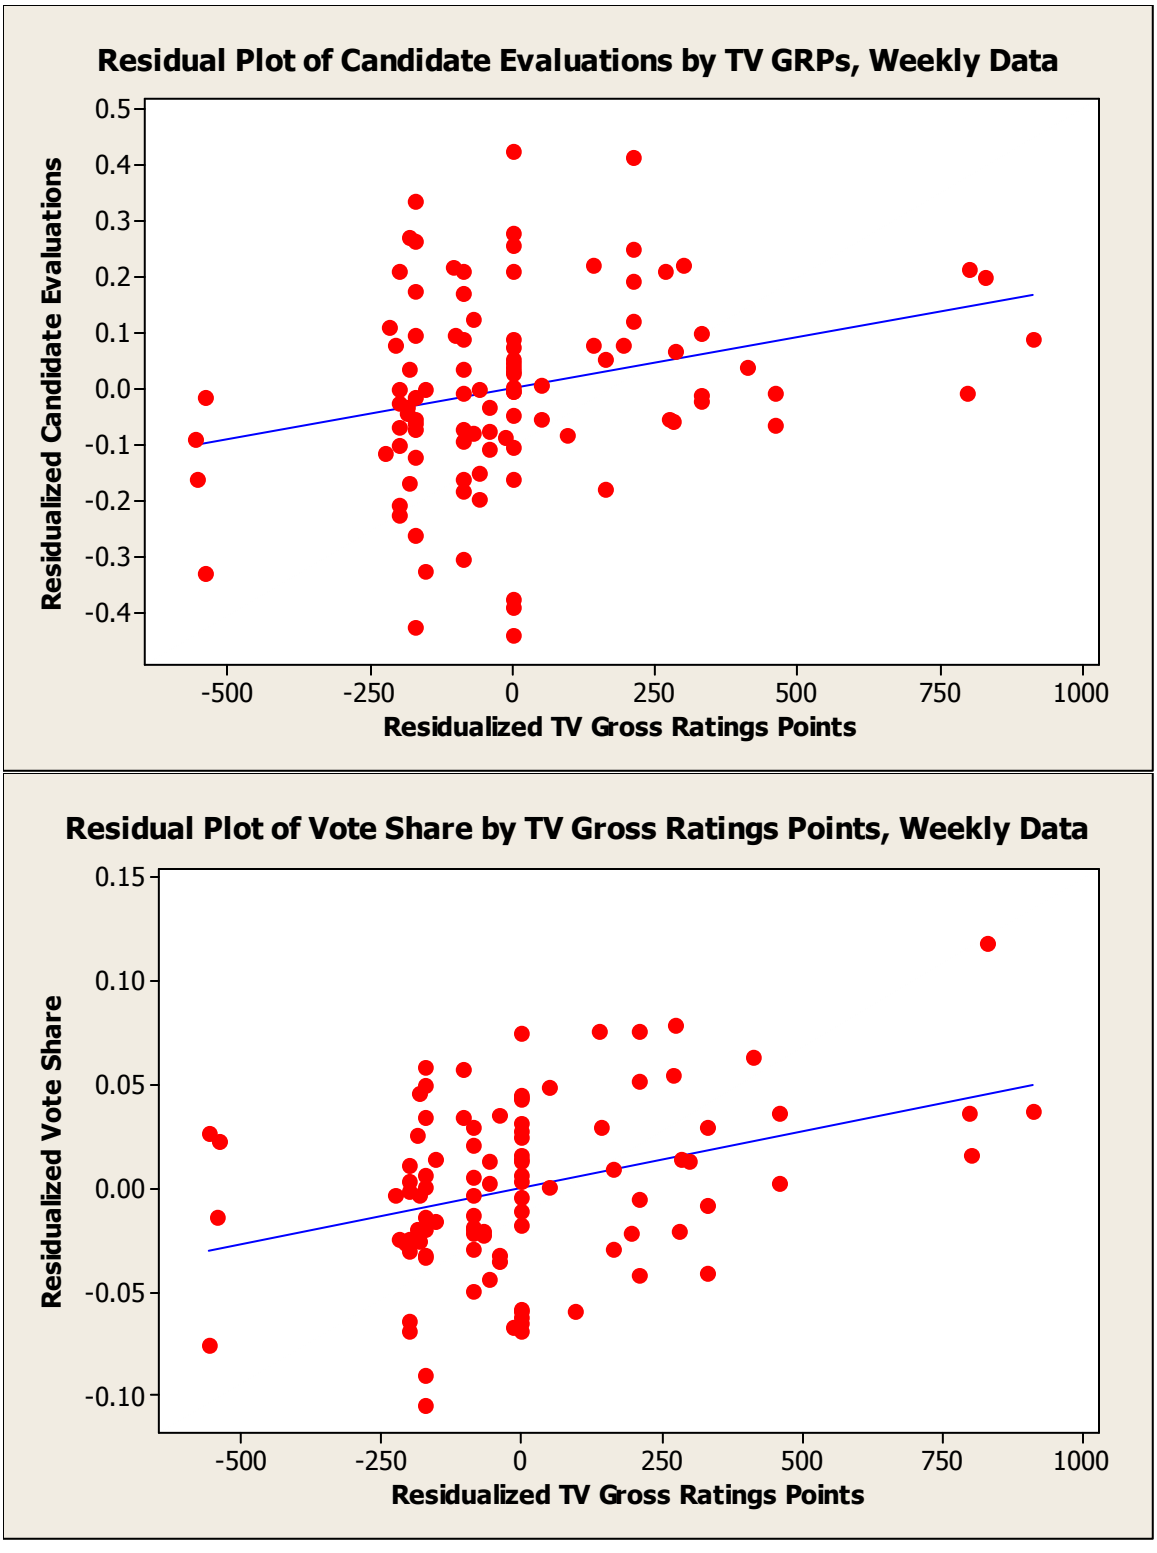

Figure 4: Illustration of the Effects of TV Advertising on Voter Preferences