

Negative Campaigning

Richard R. Lau and Ivy Brown Rovner

Department of Political Science, Rutgers University, New Brunswick, New Jersey 08901;
email: ricklau@rci.rutgers.edu

Annu. Rev. Polit. Sci. 2009.12:285–306

The *Annual Review of Political Science* is online at
polisci.annualreviews.org

This article's doi:
10.1146/annurev.polisci.10.071905.101448

Copyright © 2009 by Annual Reviews.
All rights reserved

1094-2939/09/0615-0285\$20.00

Key Words

attack advertising, demobilization, meta-analysis, negativity

Abstract

The past two decades have seen an explosion of social science research on negative political advertising as the number of political observers complaining about its use—if not negative campaigning itself—has also grown dramatically. This article reviews the literature on negative campaigning—what candidates are most likely to attack their opponent, under what circumstances, and most importantly, to what effect. We also discuss the many serious methodological issues that make studying media effects of any kind so difficult, and make suggestions for “best practices” in conducting media research. Contrary to popular belief, there is little scientific evidence that attacking one’s opponent is a particularly effective campaign technique, or that it has deleterious effects on our system of government. We conclude with a discussion of whether negative political advertising is bad for democracy.

INTRODUCTION

Political candidates invest heavily in strategic campaign communications. Whereas press releases, public appearances, stump speeches, and interviews are geared toward garnering free media coverage, carefully crafted commercial or paid communications represent an industry unto themselves. Substantial percentages of campaign war chests are dedicated to such varied forms of political communication as lawn signs, bumper stickers, direct mail, candidate web sites, and television and radio appeals. Although all of these elements are integral to a successful campaign media strategy, television advertisements remain the most visible, expensive, and presumably cost-effective form of paid political communication.

As one might expect, television advertising is by far the most pervasive communicative technique studied in the political communications literature. In just half a minute, a candidate can deliver a precise message to a broad audience. A single ad can increase a candidate's name recognition, inform voters of salient issues and candidates' policy concerns, or motivate previously reluctant individuals to vote through emotional appeals—if it is repeated often enough. Therein lies the rub: effective political advertisements are expensive to produce, and they become increasingly costly with each ad buy. The more often an ad is shown, the wider its audience, and the more often people are exposed to an ad, the more likely they are to process and remember its message. In theory, there should be a point of decreasing marginal utility at which everyone who is going to see an ad has already seen and processed it, and additional exposures are ignored or, worse, become annoying. How often (if ever) this point is reached in a political ad campaign is not clear. But even the wealthiest and best financed candidates are limited by the size of their war chests, and consequently must decide which messages will be conveyed to the public, what media tactics should be employed, and how often they can afford to present each message. If the electoral district is large enough to include multiple media markets, candidates

also must decide what messages to provide in which media markets.

One of the most important decisions candidates make is whether to run on their own merits—that is, their own policy ideas, past accomplishments, and personal strengths, which most observers would call a “positive” appeal; or if instead their campaign will concentrate on the perceived weaknesses of their opponent's policy proposals, prior policy failures, and/or personal peccadilloes, which in this essay we call a “negative” appeal. Candidates are not restricted to one or the other, of course, and most campaigns employ a combination of both techniques. But in our review, at a very general level, if the appeals a candidate makes are primarily positive, we refer to their campaigns as positive, whereas if the appeals they make are primarily negative, we refer to their campaign as negative.¹ We use the word “tone” to refer generically to the balance of positive and negative appeals in a campaign.

Why a Focus on Negative Political Campaigns?

This essay focuses on negative campaign communications, a widespread phenomenon that has attracted increasing attention from the press and social scientists alike. Although politics in America was never a gentleman's sport, an accelerated proliferation of negative advertising over the past 30 years—or at least a rise in the number of people complaining about it—seems apparent. But that is the most we can say. Before 1996 there are not, in fact, any good data on the actual prevalence of negative campaigning in

¹ Campaigns that employ approximately equal numbers of positive and negative appeals could be called “mixed,” but for simplicity we assume here that all campaigns are either primarily positive or primarily negative. How one could actually get an overall negativism “score” for a campaign that included television ads aired different numbers of times in different media markets, another set of radio ads similarly aired different numbers of times in different markets, newspaper advertisements, billboards, lawn signs, bumper stickers, and personal appearances by the candidate, etc., is a question no one we know of has ever tried to address.

the United States. Geer (2006) reports a steady growth in the percentage of negative appeals in the television advertisements produced by the major party presidential candidates from 1960 through 2004 but has no evidence on how often each ad was shown. Buell & Sigelman (2008) report no growth in the negativism of presidential elections between 1960 and 2004, as judged by the tone of statements from all presidential candidates reported in the *New York Times* between Labor Day and election day. Similarly, Lau & Pomper (2004) report no growth from 1988 to 2002 in the proportion of statements reported in newspapers from major party senatorial candidates (or their spokespersons) criticizing the opponent. But this is secondary evidence, filtered through the lens of the newspaper reporter, and tracks only what the candidates are *saying*. It is only with the advent of “ad detector” technology developed by the Campaign Media Analysis Group, which tracks satellite-based feeds of political advertisements, that we have been able to gather any solid evidence on how often different televised ads are shown in different media markets. As seen in **Figure 1**, in this relatively short timeline, negative political advertising in the United States has not grown at all (at least at the federal level).²

We can be more definitive about increasing *interest* in negative campaigns, however. **Figure 2a** displays the number of stories about negative advertising or negative campaigns appearing in the *New York Times*, the *Washington Post*, the *Christian Science Monitor*, and the Associated Press newswire during presidential election years since 1980. Only 17 stories about either of these topics

appeared in any of these newspaper sources in 1980, but this figure grew to 210 stories in the millennium year before falling back a bit in 2004. **Figure 2b** shows the number of academic articles providing some evidence on the effects (broadly construed) of negative political advertisements or negative campaigns. The first article we could find in the social science literature was published in 1984, and the first study by a political scientist appeared in 1990. At last count (near the end of 2006) there were 110 books, chapters, dissertations, and articles addressing this broad topic empirically, and many more exploring other aspects of negative campaigns. The upshot of the great majority of these articles and news stories is that negative campaigning has become a pervasive and often corrosive aspect of the American political scene—“the electronic equivalent of the plague” (West 1993, p. 51), to quote one of the more colorful characterizations. Negative campaigning and the use of attack ads have been criticized for reducing politics to its least common denominator. Some scholars have gone so far as to attribute diminishing public trust and ever-decreasing turnout rates to the growing use of negative ads.

Why all this attention to campaign negativism if there is no real evidence it is increasing? The obvious answer is that coordinated candidate messages *appear* to grow increasingly negative with each campaign cycle even if they actually are not, making negativism increasingly difficult for any voter to avoid or ignore. This is due primarily to biases in the news media that make stories about negative advertisements much more likely to be written and/or aired than stories about positive ads. For example, West (2005) reports that almost two thirds of the stories about political advertisements broadcast on the CBS evening news between 1972 and 2004 involved a negative advertisement. Unless predicated on false information, positive communications provide little if any entertainment value. After all, how many stories on a candidate’s love of God, country, family, and all-American values can a single media outlet run? By contrast, negative communications

²These data are available from the Wisconsin Advertising Project at the University of Wisconsin. We thank John Sides for providing the numbers in **Figure 1**, and John Geer and Shana Gadarian for data in **Figure 2a**. We may soon have new evidence on how much money presidential candidates from before 1996 spent advertising in the different states, and even possibly how often specific ads aired in different locales, from data currently being gathered by Scott Althaus and Daron Shaw from presidential libraries around the country. Stay tuned. . . .

dangle like red meat in front of journalists hungry for a new angle. It is therefore not surprising that negative ads get covered (and consequently replayed by the media) much more frequently than positive ads—leading to the common impression that negative ads are more pervasive than they are. To the extent that political practitioners recognize these media biases, they certainly take advantage of them, running negative ads within the Washington Beltway media market in order to capture the attention of the political news media.

These same biases contribute to the impression that most negative ads are ad hominem attacks and misleading exaggerations if not out-and-out lies. Some, surely, are, and the media do their best to bring the most outrageous examples to our attention, but it would be a mistake to assume that all or most or even a large minority of negative political advertisements have these characteristics.

Even if it could be definitively shown that negative ads are increasing, such evidence would not explain why candidates attack their opponents. We can assume the answer is that candidates and political practitioners believe attack ads are effective—and there are sound theoretical reasons to think they may be right. Social psychologists have pointed to “negativity”: the tendency for negative information to be more influential than equally extreme or equally likely positive information (Kanouse & Hanson 1971). The “*equally* likely or *equally* extreme” part of the definition is crucial, and is what makes negativity an interesting psychological phenomenon rather than mere common sense. If we describe a candidate as having a pleasant smile but holding foreign policy positions that will lead to nuclear war, of course no one predicts that the combination of these two bits of information will result in neutral evaluations. Kernell (1977) and Lau (1982) have carefully documented negativity effects in candidate evaluation and voting behavior, and Lau (1985) has provided evidence for two different explanations for negativity, one perceptual and one motivational. The perceptual “figure-ground hypothesis” states that negative information

stands out against a generally more positive background; the motivational “cost-orientation hypothesis,” which is presumably part of our genetic makeup, states that most people are inherently more strongly predisposed to avoid costs than they are to seek gains. By this account, negativity explains both the apparent growth of campaign negativism and its alleged corrosive effect on our political system.

In this article, we review the evidence on just how effective negative advertising is in achieving its goals, and just how corrosive those attacks are on our political system. We begin by trying to carefully define what a “negative political campaign” is. Is our very general definition of attacking one’s opponent (rather than talking about oneself) sufficient, or do we need a more precise and nuanced definition? We might believe (echoing Justice Stewart’s famous statement about pornography) that we know it when we see it, but negative campaigning turns out to be a bit more difficult to define precisely than one might think. Having established a workable definition of the phenomenon of interest, we then look at the literature on when candidates choose to “go negative.” There is a fair amount of agreement between the formal and the empirical literature, but we also find some clear theoretical predictions that have been *disconfirmed* by empirical evidence.

Then we turn to arguably the most important question: Do negative campaigns work? Clearly, paid campaign consultants and the reporters who cover campaigns believe they do—otherwise they would not continue to use them. As Laura Mansnerus, writing in the *New York Times* about the highly negative 2005 New Jersey gubernatorial campaign, explained: “The people who produce these ads and the consultants who hire them know that negative campaigning works. These people are paid way too much to be mistaken about whether poison is effective” (Mansnerus 2005). As social scientists, we seek more convincing evidence than this.

There is a long history in political science of doubting that the media affect the political process. But as the “minimal effects” view has been

replaced by the “media matter” view (Iyengar & Kinder 1987), some within the academy have come to believe that negative political campaigns matter so much that they can affect the very nature of our political system, discouraging voters from participating in individual elections and weakening their faith in the political process itself over time. In fact, this is one of the few instances that we can think of where an academic finding has been picked up by the press and become part of the conventional wisdom about political campaigns. The impact of negative campaigns on the American political system is the fourth major topic of our review.

Although we mention some methodological issues as we discuss our first four topics, the penultimate section of the article focuses on a collection of methodological concerns that plague the study of political campaigns (and of media effects more generally). In the concluding section we ask a normative question: Are negative campaigns bad for democracy? Some political scientists argue that, far from demobilizing the electorate, negative political campaigns may even stimulate participation. They believe that negative ads educate voters by providing critical information about a candidate’s policy positions, character, and personal history, and in so doing, allow voters to make more informed political decisions. A review of the literature suggests that categorizing negative campaigns as patently good or bad is at the very least an oversimplification and at most a gross oversight. The question of whether negative campaigns are good for democracy, like the decision to launch a negative attack, is complicated and multilayered. Nonetheless, as negative campaigns continue to have a sizeable presence in our political system, an assessment of their intended and unintended influence on political behavior is warranted.

WHAT IS A NEGATIVE CAMPAIGN?

In order to examine the role of negative campaigning, we need to carefully define it. The literature has emphasized analysis of the use

and effectiveness of negative political *advertisements*, trending toward a simple positive-negative dichotomy in which ads that focus on their sponsors are branded “positive” and those that address the opposition are “negative.” Each advertisement is treated as a single unit of analysis and weighted equally. This dichotomy may lead many people intuitively to categorize positive ads as patently “good” and negative ads as inherently “bad,” but Jamieson et al. (2000) warn that doing so is both incorrect and misleading:

Many scholars mistakenly assume that “attack” is both “negative” and “dirty.” Conflating these terms obscures the important distinction between legitimate and illegitimate attack and minimizes the likelihood that the deceptions found in supposedly “positive” discourse will be probed. (p. 44)

Although the positive-negative dichotomy may be conducive to experimental research—where isolating and manipulating a single variable (e.g., negative communication) in order to measure its effect is the end goal—Jamieson et al. suggest that content analysis that fails to treat “contrast” or “comparative” advertising as a separate and unique category of political discourse not only excludes a key communicative ingredient, but also portrays a simpler landscape than what, in fact, exists. We agree, when it comes to creating typologies of individual political advertisements. This distinction loses most of its power when applied to entire campaigns, however. Campaign communications include television, radio, and print ads, as well as the content of official campaign web sites, speeches, debates, interviews, and direct mail (fund-raising) appeals. Unless a candidate completely ignores his or her opponent (as may occur when a challenger fails to represent a legitimate threat), all campaigns by their very nature are comparative. Not forgetting the tremendous difficulties of measuring and combining the positive or negative tone of every different aspect of a major political campaign, a unidimensional scale (100% positive to

Table 1 Perceived fairness of different types of attacks

Criticism	Percent saying criticism is fair
Talking one way and voting another	80.7
His/her voting record	75.8
His/her business practices	71.0
Taking money from special interest groups	70.7
Taking money from individuals with ethical problems	63.0
Current personal troubles with drugs or alcohol	56.1
Current extramarital affairs	45.1
Political actions of his/her party's leaders	37.1
Past extramarital affairs	27.8
Past personal troubles with drugs or alcohol	25.9
Personal lives of his/her party's leaders	19.1
Behavior of his/her family members	7.7

Source: 1998 Sorenson Institute for Political Leadership Survey of Virginia Voters, as reported by Freedman et al. (1999).

100% negative) seems in theory the best way to think about the negativism of a campaign.

But these are all subtle differences in academic definitions that may be meaningless to the general public. Do laypeople know when they have seen a negative ad? Is there a limit to what they will tolerate in a political campaign? Drawing on two surveys conducted in Virginia, Freedman et al. (1999) maintain that voters are indeed capable of distinguishing between mudslinging and fair play, and that they neither condemn nor condone campaign negativism outright. When asked to rate the fairness of 12 charges that one candidate might make against another, respondents revealed a clear hierarchy of tolerance for attacks (see **Table 1**). Generally speaking, voters believe that attacking a candidate for his actions in office is fair game, whereas bringing up prior personal problems should be out of bounds.

Freedman et al.'s findings also reveal educational and partisan biases. Respondents' thresholds for tolerating different types of attacks increased with education, political information, and Republican party identification. That is, more educated, better informed, and Republican voters were all somewhat less likely to say that any given charge was unfair. Freedman et al. did not, however, find significant variation with respect to either gender or race. A

follow-up study examining the repercussions of counterattacks reveals that voters ascribe similar standards of fairness to rebuttals. The less "fair" the response, the more willing voters are to vote for the initial attacker, regardless of the perceived fairness of his or her original charge.

But do voters' theoretical opinions of campaign negativism translate into the reality of mass political behavior? Not necessarily, according to Sides et al. (2003). They report a study of the 1998 California gubernatorial election, in which Democrat Lieutenant Governor Gray Davis opposed Republican Attorney General Dan Lungren. They link rolling cross-sectional survey data from 2902 respondents interviewed during the last six weeks of the campaign with data on the actual tone of the advertising aired by the two candidates during each day of that period. This design allows Sides et al. to determine how closely citizens' perceptions of a campaign's negativism are tied to the actual nature of that campaign. Not surprisingly, there was a very strong partisan component to these perceptions, with Democrats perceiving Davis's campaign to be much less negative than Lungren's campaign, and Republicans perceiving the opposite pattern. What is much more surprising is that the realities of the campaign—how negative Davis's and Lungren's campaigns actually were—had no effect on perceptions. Indeed, the *only* significant predictor in the equations for both Davis and Lungren was partisanship. Clearly, negativism remains in the eye of the beholder. It seems that few voters during actual political campaigns come anywhere close to the objectivity that a researcher would apply in judging the content or tone of a campaign.

Sigelman & Kugler (2003) agree, suggesting that scholars' systematic examinations and even our definitions of campaign negativism are out of sync with voter perceptions. Our objective, systematic coding of ads implicitly assumes that scientific realities guide perceptions, but that assumption suggests that voters consume campaign information uniformly and as judiciously as scholars conduct their own research. Sigelman & Kugler point out that

whereas scholars seek out the breadth of political communications and afford each statement equal weight, voters do not; instead they seek the information that is most concise and readily available, often at the expense of accuracy, and pay disproportionate attention to that with which they most agree. Given those differences, it is not surprising that Sigelman & Kugler uncover inconsistencies between voter perceptions of negativity and scholarly classifications. Their comparison of ANES survey data against the newspaper coverage of three 1998 gubernatorial races reveals asymmetrical impressions of campaign tone, as well as inconsistent perceptions among voters within each state. However, although they attribute differences among voters to partisanship, political efficacy, and above all, political sophistication, Sigelman & Kugler stop short of prescribing a solution for this problem.

The challenge lies in reconciling the rigors of systematic social science research with the reality that voters consume political information erratically and with bias. Perhaps our definition of tone should go beyond written transcripts of televised ads and include the visual images and accompanying music often aimed at triggering voters' emotions. For example, Brader's (2005, 2006) analysis of emotional appeals reveals frequent, deliberate attempts to evoke enthusiasm or fear among the electorate through the strategic use of visual images, music, and voiceover tone. These components do not change the ads' message but rather clarify that message so that it "strikes an emotional chord," as his title suggests. Unlike analyses that measure tone in terms of ad content, Brader measures the ads' intended emotional outcome. He describes enthusiastic ads—such as Ronald Reagan's famous 1984 "Morning in America" spot—as those which appeal to hope, enthusiasm, sentimentality, and national pride. Such ads aim to capture the attention of a candidate's base by motivating participation and ultimately turnout. By contrast, fear ads do not automatically elicit a transfer of allegiance from one candidate to another, but rather trigger anxieties among undecided or weakly committed voters,

causing a re-evaluation of both candidates' track records and capabilities.

But do purely factual, unemotional ads even exist? Doesn't every ad seek to arouse either confidence in its sponsor or anxiety about the competition? Even "objectively" pointing out an opponent's weak record on crime is an effort to create anxiety about his or her ability to address the issue in the future. Similarly, highlighting one's own proven track record is intended to instill confidence and security. Brader (2006) maintains that, though rare, "unimpassioned ads" do exist—if only so that their sponsors may appear to function above the political fray. But isn't that an enthusiastic appeal, painting oneself as a candidate of whom we can all feel proud? Brader points to Ross Perot's 1992 advertisements and 30-minute "infomercials" as an example of such authentic communications, yet Perot's ads tapped into a disgust and frustration with the Republican status quo and garnered enthusiasm for a more responsible fiscal conservatism—so much enthusiasm that he received more votes than any third-party candidate in U.S. history. Even his dry format triggered emotions among the electorate. To be fair, Brader concedes that "ads are almost never stripped completely of emotional content"; he categorizes as "unimpassioned" those ads which do not *overtly* appeal to voters' emotions, acknowledging that minimal residual sentiments are inevitable (Brader 2005).

Every candidate's goal is to win. Therefore, just as candidates understand that it is unrealistic to expect all voters to make decisions on the merit of facts and the content of campaign communications alone, scholars must appreciate the role of audiovisual elements in emotionally priming vote choice. This is not to suggest that political decisions are fundamentally purely emotional, but rather that emotional cues can take the voter to a place that facts alone may not always reach. As Brader explains, emotional elements may reinforce content, but they do not override it. Ultimately, Brader's research reminds us that researchers who code the tone of a political ad or campaign based solely on the direction of the candidates' appeals, without

also considering the role of emotional overlays, are accepting a broad and less precise definition of campaign negativism, one that can be potentially misleading. On the other hand—a very important hand—simple directional coding of campaign appeals can be accomplished very easily and reliably. The more subjective and complicated a coding scheme, the more difficult (and thus expensive) it is to use, and typically the less reliable are its results.

THE DECISION TO ATTACK: WHEN, HOW, AND WHO?

When does a candidate decide to inform voters of his or her opponent's alleged negative attributes, despite the risk of alienating potential and existing supporters? The decision to attack is a political calculation, based on the presumption that its execution will damage the intended target more than it will jeopardize the status of the candidate sponsoring the attack.

Skaperdas & Grofman (1995) use formal theory to identify factors that motivate a candidate to go negative. Assuming that candidates are aware of their relative status in the polls and that future vote shares remain indeterminate until the day of the election, and most importantly that campaign attacks provide a net benefit to the attacking candidate, Skaperdas & Grofman examine the use of negative advertising in both two-candidate and three-candidate contests. In a two-candidate race, they find that the frontrunner typically uses a greater proportion of positive ads and his opponent more frequently goes negative. However, as the gap between the two candidates narrows, the leading candidate's use of negative advertising will increase in order to maintain his or her frontrunner status. Multi-candidate races are less straightforward. A candidate who is trailing by a definitive margin and thus only serving as a "spoiler" is predicted to run positive ads exclusively. Further, in any three-way race, no candidate should launch negative attacks against the weaker of the two opponents. Instead, all attacks should be directed toward the frontrunner or, in the case of the frontrunner him-

herself, against the greatest threat. This was evident in the 1992 presidential election: George H.W. Bush and Bill Clinton focused their attacks on each other, rather than against third-party candidate Ross Perot. However, as Perot began to chip away at Bush's base toward the end of the campaign, he found himself on the receiving end of Bush-sponsored attacks, which is counter to the predictions of the model.

The rebuttal calculus also matters. How does a candidate's decision to go negative without provocation compare to his willingness to *retaliate* against a negative attack, in light of poll positioning, closeness of the race, or the number of opponents? Why, for example, did John Kerry fail to rebut a series of negative ads throughout August 2004? Where is the line between staying above the fray and being an opponent's punching bag? There are costs either way. On the one hand, a candidate's failure to refute attacks can leave the information environment dominated by the charges from the opponent. On the other hand, responding to the allegations gives them credibility, and it can also put the candidate on the defensive and take him or her off message, which can prove equally damaging (Kern 1989, Jamieson & Campbell 1983). After the candidate is attacked, the campaign faces not only a question of strategy (whether to go negative) but also a question of tactics (how).

Harrington & Hess (1996) disagree with the assertion that a candidate's willingness to attack reflects his or her status in the polls. Instead, their theoretical model suggests that a candidate's personality and ideology are stronger predictors. Any connection between the use of negative ads and poll position, they maintain, is purely coincidental; both are driven by a candidate's weakness in terms of personal attributes. Advertising tactics, they argue, are driven by the ideological position of undecided voters rather than the relative strength of an opponent. Candidates will use positive messages to reposition themselves as ideologically compatible with undecided voters and will likewise use negative ads to create an ideologically based wedge between their opponent and potential swing voters.

Inherent to this theory is an assumption that the candidate with the more desirable personal attributes or larger “valence index score” will present the more positive campaign, regardless of his position in the polls.

In order to reconcile these opposing arguments and identify whether actual campaign activity follows theoretical prediction, Theilmann & Wilhite (1998) ran a series of experiments using members of the American Association of Political Consultants, exposing this unique sample to a series of hypothetical campaign scenarios and assessing their recommendations regarding the willingness to attack. Theilmann & Wilhite report confirmatory evidence for Skaperdas & Grofman’s theory of relative strength, but they find no connection between personal attributes and negative ads, as suggested by Harrington & Hess. They also examine the role of partisanship, testing Ansolabehere & Iyengar’s (1995) hypothesis that positive ads are more effective with Democratic voters whereas negative ads more often resonate with Republicans. The Republican consultants in their sample were indeed more open to implementing attack strategies than were their Democratic counterparts overall, more inclined to employ them early in a campaign, and more likely to use them when a client is comparatively underfunded.

Intrigued by this decision-execution calculus, Buell & Sigelman (2008) examined all major party presidential tickets between 1960 and 2004 to identify how competing tickets decide not only when, but also how, to attack. Buell & Sigelman find support for some of Skaperdas & Grofman’s predictions, but this support is primarily limited to races they consider blowouts—the elections of 1964, 1972, and 1984. Although the Skaperdas-Grofman model can play a useful organizing function, when put to empirical test it “served primarily as a foil for one contrary finding after another.” The model “abstracts away too much and thereby ignores a multitude of factors that figure importantly in candidates’ decisions to go negative” (Buell & Sigelman 2008, pp. 441–42).

Rather than examining the campaign in its entirety, Haynes & Rhine (1998) focus on those variables—namely “competitive positioning, reward factors, and media related conditions”—that predict negative communications in the early primary season, a time when the intraparty candidate field is crowded with ideologically similar candidates, and lack of funding may prevent some candidates from purchasing any airtime. Using content analysis, they discover that underfunded candidates seek to increase name recognition through personality-driven, self-promotional (positive) ads, and utilize a variety of “intermediated” vehicles—including press releases, television interviews, stump speeches, and debates—to launch attacks on their opponents. These alternative media provide a low-cost opportunity for lower-tier candidates to garner free (some would call this *earned*) media coverage by comparing themselves to the frontrunners while presenting a positive and more dignified image through advertising.

All of the empirical data on the nature of campaigns reported thus far have come from studies of relatively few campaigns, typically at the presidential level. In contrast, Lau & Pomper (2001, 2004) examine the tone of the campaigns of virtually every contested U.S. Senate election between 1988 and 2002. They propose seven hypotheses about when candidates will go negative. In particular, they propose that candidates who are behind, candidates in close elections, challengers, candidates with relatively few campaign resources, Republicans, males, and candidates whose opponents are attacking them have higher levels of negativism in their campaigns. Lau & Pomper report significant support for all seven of these hypotheses at the bivariate level. Only three of these hypothesized effects maintain their strength in a multivariate analysis, however. Controlling on the other explanations, Republicans, candidates with less money than their opponents, and candidates whose opponents are attacking them are all significantly more likely to attack their opponents. Particularly noteworthy in Lau & Pomper’s findings is the almost one-to-one

correspondence between the tone of a candidate's campaign and that of his or her opponent. These days, attacking your opponent in a senate election is almost guaranteed to provoke an approximately equivalent counterattack.

In an unpublished manuscript, "Going Negative in a New Media Age: Congressional Campaign Websites, 2002–2006," Druckman, Kifer, and Parkin find similar patterns in political attacks seen on the internet. Having analyzed a dataset that includes 730 candidate web sites spanning three campaign cycles, Druckman et al. concur that competition drives negativity, but they provide a somewhat different explanation for this now familiar result. First, Druckman et al. argue that the audiences for campaign web sites are highly self-selective, falling into two distinct groups: journalists and politically active supporters of a candidate. Drawing on the emotions literature (Brader 2006, Marcus et al. 2000), Druckman et al. go on to argue that negative information raises anxiety levels and encourages voters to seek out additional information about all candidates and engage in compensatory decision making, wherein candidates are compared against one another across a variety of dimensions in order to "choose the best alternative, regardless of cognitive demands." Because calm voters remain satisfied with the status quo, incumbents prefer to avoid campaign attacks except in highly competitive races. The same is true of both frontrunners and members of the majority party. By contrast, challengers, nonfrontrunners, and members of the minority party not only have an incentive to go negative, but ultimately must go negative if they are to convince the electorate that a change of course is necessary. Nonetheless, as Druckman et al. and others have pointed out, the greater the competitiveness of the race, the more likely it is that *both* candidates will go negative. Ultimately competitiveness trumps incumbency, frontrunner, and majority-party status.

A review of the circumstances under which candidates go negative would be incomplete without assessing whether candidates adopt different message strategies depending on their

gender and that of the opposition. Kahn and Kenney (Kahn 1993, Kahn & Kenney 2000) find that gender does indeed factor into a candidate's decision to attack. However, their research has yielded somewhat inconsistent results. Kahn's 1993 study examining the campaign messages of 38 candidates for U.S. Senate in the 1984 and 1986 races revealed that although both male and female candidates prefer policy-based messages over candidate-oriented appeals, men make use of negative ads less frequently and are particularly reluctant to go negative against female opponents, for fear of being perceived as "beating up on a woman" (Kahn 1993, p. 491). She argued that regardless of candidate status (incumbent or challenger), the messages delivered by female candidates tend to reflect those of underdogs, exhibiting a greater proportion of negativity than their opponents demonstrate.

Less than a decade later, Kahn & Kenney (2000) again compared the communicative styles of male and female senatorial candidates, now spanning the 1988, 1990, and 1992 campaigns. They found female candidates consistently *less* likely than male candidates—regardless of candidate status—to go negative. Identifying five categories of negative ads (attack, comparative, negative-trait, negative-issue, and criticism), Kahn & Kenney found that male challengers consistently launch more negative appeals than their female counterparts, and in particular are 10% more likely to run negative-issue ads. This finding is consistent with the bivariate results reported by Lau & Pomper (2001, 2004). When closeness of race is taken into account, the differences are even more stark; the frequency of male-sponsored attacks greatly increases with competitiveness, whereas female candidates are more consistent. The two studies are not entirely contradictory, however. Kahn and Kenney confirm that male candidates remain far more aggressive toward other males than toward female opponents, attacking 56% of the time as opposed to 39%.

Kahn and Kenney's inconsistent findings, and the null effect of gender in a multivariate analysis reported by Lau & Pomper

(2001, 2004), cast doubt on any speculation that female candidates adhere to a softer style and more stereotypically feminine cultural norms while male candidates are more aggressive and cut-throat. This skepticism is supported by Procter et al. (1994), who examine the use of negative advertising in eight 1990 U.S. Senate campaigns and eight gubernatorial campaigns from the same year, all of which featured a female candidate against a male opponent. Procter et al. find that male and female candidates do not differ in their use of negative advertising. Female candidates are no less likely to strike first, do not manipulate ad formats to soften their own images, and do not shy away from attacking their opponent's personal ethics or issue positions. Procter et al.'s findings illustrate that decisions to go negative and decisions concerning how to strike are not reflections of a candidate's gender but rather of a formulaic calculus that is unaffected by gender. We should mention one important caveat here: Research on the contrasting campaign styles of male and female candidates when they oppose each other suffers from the absence of an important comparison group, female-versus-female races. There are still too few examples of statewide races where major party female candidates oppose each other.

DO NEGATIVE CAMPAIGNS WORK?

If you ask observers of the American political scene to think of a negative political advertisement, they might mention the infamous Willie Horton ads from the 1988 presidential election campaign, or remember seeing Michael Dukakis riding around in a tank during that same campaign. Or they might come up with the Swift Boat Veterans ads from the 2004 presidential campaign. Somewhat older voters might remember the daisy ad from the 1964 presidential campaign. In each case, the target of those attacks—Michael Dukakis in 1988, John Kerry in 2004, Barry Goldwater in 1964—lost the election. These impressions are perfectly consistent with the conventional wisdom about

negative campaigning: It works. This is practically a mantra among political practitioners (Kamber 1997, Swint 1998) and has reverberated throughout the scholarly literature as well.

The only reason a rational candidate should choose any campaign strategy is because that candidate believes the strategy will win votes. Thus, if attacking one's opponent is an effective campaign strategy, it must result in the attacker receiving more votes than he or she would have if some other (less negative) campaign strategy had been adopted. Assuming the vote choice is little more than selecting the more highly evaluated candidate, attacks that lower mean evaluations of the target of those attacks would constitute additional evidence for the effectiveness of negative campaigns. However, there is always a danger that evaluations of the attacking candidate will also decrease (a so-called backlash effect: see, for example, Garramone 1984, Roeser & Sande 1993). Candidates who attack their opponent might be perceived as mean-spirited or nasty, particularly by voters who were taught "if you can't say anything nice, don't say anything at all." But even if backlash effects against the sponsors of campaign attacks are frequent occurrences, attacking one's opponent would still be an effective strategy if evaluations of the target of the attacks went down more than evaluations of the sponsor, such that the net differential evaluation of the two competing candidates worked to the advantage of the attacker.

Lau et al. (2007) have just published a meta-analytic review of the social science literature on the effects of negative political campaigns (see Lau et al. 1999 for an earlier review). A meta-analysis is a quantitative synthesis of research findings from independent studies on a particular topic, where the reported findings themselves, rather than the raw data on which they are based, are the data analyzed. A meta-analysis is as comprehensive as possible, including both published and unpublished studies so that well-known publication biases (e.g., only publishing papers with statistically significant findings), or the prejudices of the reviewer (e.g., deciding a particular study is "junk" and therefore should not be included in the review), do not influence

the outcome of the meta-analysis. Any factor the analyst believes could influence the magnitude of the reported results (e.g., experimental versus correlation designs; the perceived quality of the reported study) can be coded and tested empirically—if there are enough relevant studies.

Lau et al. (2007) located 43 studies examining the effects of negative campaigns (or exposure to particular negative advertisements) on the actual or intended vote choice of those exposed to the campaigns/advertisements. For example, King et al. (1998) exposed subjects to either a positive or negative ad from Bill Clinton or Bob Dole near the end of the 1996 presidential campaign. Controlling for initial liking of the candidates, King et al. found that exposure to a negative ad from Clinton strongly decreased intention to vote for him, but exposure to a negative ad from Dole had no effect on likelihood of voting for him. In neither case is there any evidence that negative ads are particularly effective in winning votes. Of the 43 relevant studies, 12 reported positive effects (that is, the more negative candidate received more votes), four of which were statistically significant; 28 studies reported unfavorable effects (that is, attacking the opponent was counterproductive), six of which were statistically significant. Three studies reported that there were “no significant differences,” without providing any numbers or hints as to which direction the nonsignificant results leaned; in such cases, results can only be coded as having an effect size of 0. Looking at all of these studies together, no one could conclude that negative campaigning is a particularly effective campaign strategy.

Many factors go into a vote decision, of course, including habit, and it may be asking a lot of a typical political campaign, never mind exposure to a few ads in an experiment, to change that decision. But surely, one would think, evaluations of the targets of political attacks must go down. Lau et al. (2007) located 31 studies with relevant data. For example, Kahn & Kenney (2004) coded the tone of up to five ads aired by major party candidates in 97 competitive senate elections between 1988 and 1992.

Controlling for a large number of factors in a multivariate analysis, Kahn & Kenney report that both incumbents and challengers were liked significantly less when they were attacked by their opponent than when their opponent stayed primarily positive. Twenty-three of the 31 studies reported similar effects, 12 of which were statistically significant. Only seven studies (two of which were significant) reported the opposite effect, with evaluations of the target actually *increasing* after the attack—which might plausibly occur if many people perceived an attack to be exaggerated, false, or otherwise out-of-bounds, and responded in part by sympathetically increasing their evaluations of the target of the unjustified attacks. In any case, we are reasonably confident that attacking a political opponent is generally likely to result in lower evaluations of that candidate.

But attacking one’s opponent is not a risk-free enterprise, as evaluations of the attacker might also decrease if voters are convinced the attacks are unjustified or mean-spirited. This same study by Kahn & Kenney (2004) reported that evaluations of both incumbents and challengers decreased (the former significantly) after they sponsored attacks against their opponent. All told, 31 of the 40 relevant studies (19 of which were statistically significant) reported similar findings, whereas only six studies (two significant) found evaluations of the attacker increasing after the attacks. If anything, this backlash effect appears to be a bit stronger and more consistent than lowering evaluations of the target of the attacks. On balance, then, there is simply no support in the scientific literature for the hypothesis that negative campaigns are any more effective than any other type of campaign strategy.

To this point we have not commented on the different designs that researchers employ to study the effects of negative campaigns. We assume the reader is generally familiar with the relative advantages and disadvantages of experiments and correlational (or observational) studies. Despite what we often tell our undergraduates, it *is* often possible to make reasonably confident causal statements from correlational

data. But the very fact that we can reliably identify several factors that explain when candidates decide to go negative raises a difficult but very important methodological problem when we try to examine the success of different strategies with observational data from real election campaigns. In statistical terms, this problem arises because the choice of campaign strategy is endogenous to the subsequent outcome of the election itself, which makes it very difficult to determine the effect of the campaign on the outcome of the election. Now, how can a campaign strategy, which logically must take place prior to election day, be endogenous to (influenced by) the outcome of the election, which we only know after the campaign is over and the polls close on election day? This is possible because many of the factors that help determine the outcome of the election are widely known long before election day, and therefore are likely to influence candidates' choice of campaign strategies in the first place. Most incumbents running for re-election are better known, better liked, and better financed than their challengers, and thus can reasonably expect to win (Jacobson 2004). As we have seen, challengers and incumbents unexpectedly finding themselves in close elections are more likely to decide to go negative . . . but they are also more likely to lose. This makes it extremely difficult to determine if such a candidate lost *because* of her choice of campaign strategy, or in spite of it. This statistical problem does have a statistical solution: instrumental variables or two-stage least squares. As Bartels (1991) warns, however, this solution is no better than the quality of the instruments. See Lau & Pomper (2002) for further discussion of this issue and details on one reasonably successful attempt to provide such instruments.

Does this methodological problem account for the lack of empirical evidence for the efficacy of negative campaigning? Probably not; in their meta-analytic review of the literature, Lau et al. (2007) report that experiments (which avoid the methodological problem discussed above by randomly assigning campaign

strategies to different candidates) were no more likely to produce positive results than were studies using various nonexperimental designs where the endogeneity problem must be confronted. Twelve of the 25 experimental studies showed positive effects (that is, the more negative candidate received more votes), which were statistically significant in four of the studies; three of the 13 negative findings were significant, and the overall mean effect was -0.02 . Again, there is no way one could conclude that negative ads or negative campaigns are any more or less effective than more positive ads or campaigns. Still, this methodological issue must be kept in mind when reviewing the available literature or when thinking about the success of any particular negative campaign.

At this point one may reasonably ask why, if attacks are not particularly effective, political consultants continue to urge candidates to attack their opponents. There are several quick answers. First, as discussed above, few political campaigns are in fact overwhelmingly negative. We are more likely to hear about negative campaigns in the media, but relatively few ads are entirely negative, and most campaigns, on balance, are probably more positive than negative. Second, to say that all or even most political attacks are not effective is not to say that attacks are never effective. Even if all elections involved only two candidates, half of all political campaigns would have to fail. Rarely are political consultants trained as scientists, and they have few incentives—and even less time and money—to study the efficacy of their actions. Like most of us, they are likely to repeat what they have done in the past. Furthermore, the time horizon of a consultant for judging the success of a particular ad campaign is probably shorter than that of a researcher, who inevitably is not looking at the campaign until the election is over. All of these factors combine to make negative campaigns a steady part of the political scene, even if they sometimes, or even most times, do not ultimately produce the results that are desired by the candidates who employ them.

THE DEMOBILIZATION HYPOTHESIS

The scholarly debate over the impact of negative advertising has centered on the demobilization hypothesis, developed by Ansolabehere, Iyengar, and colleagues (Ansolabehere et al. 1994, Ansolabehere & Iyengar 1995). Seeking to establish a causal link between the presumed rise of negative advertising and the steadily declining turnout in national elections, Ansolabehere et al. suggested three reasons why negative campaigns might lower turnout. First, negative campaigns could be successful in lowering the probability of voting for the target of the attacks without simultaneously increasing the probability of voting for the sponsor of the attacks, resulting in a decision to stay home on election day. Second, negative campaigns could lower evaluations of both the target and the sponsor of the attacks, thus decreasing the probability of voting for either of them. Third, negative campaigns could influence turnout indirectly through an effect on civic attitudes such as trust in government and political efficacy, which in turn are often associated with turnout.

Ansolabehere et al. (1994, 1995) tested these hypotheses by conducting a series of controlled experiments on the effects of exposure to a single positive or negative political ad. Participants, all of whom were drawn from the Los Angeles area, were shown a 15-minute local newscast and systematically exposed to one of two experimental political ads (either positive or negative) embedded midway through the broadcast. While manipulating tone by changing key words within the script, both ads used the same announcer and were visually and thematically identical. In order to accurately reflect real-world ad exposure, Ansolabehere et al. used authentic political ads spanning a variety of California campaigns, including the 1990 governor's race, the 1992 senate races (primaries and general election), and the 1993 Los Angeles mayor's race. Upon comparing ad exposure to each individual's intention to vote and general level of political participation, Ansolabehere et al. found that participants exposed to the

negative ad were 5% less likely to intend to vote in the upcoming election compared to those exposed to a positive ad. A subsequent study, combining analysis of raw aggregate turnout and ballot roll-off data from all thirty-four 1992 senate races with a tone-specific content analysis of the news coverage of those campaigns, yielded similar results. (Roll-off occurs when an individual only votes in the race at the top of the ballot—president—and abstains from voting in the senate election that was on the same ballot.) Ansolabehere et al. go on to report that political independents seem particularly susceptible to demobilization from exposure to negative ads, and further report a general increase in political cynicism and lower sense of political efficacy among those exposed to negative ads. In concluding that exposure to negative advertising does in fact reduce turnout and erode confidence in the political process, Ansolabehere et al. warn that because “candidates with sufficient resources can, through the use of negative messages, keep voters away from the polls,” the phenomenon poses a unique threat to democracy (Ansolabehere et al. 1994, p. 835).

This is wonderful research, with clear theoretical predictions and amazingly consistent results across two very different research designs, and it did what all wonderful research should do: generate a lot more research. Finkel & Geer (1998) were among the first to respond. While acknowledging the logic of Ansolabehere et al.'s three reasons why negative campaigns might decrease turnout, Finkel & Geer counter with three reasons why negative campaigns might actually *stimulate* turnout. First, negative advertising often conveys large amounts of policy and retrospective performance information to voters, and knowledgeable citizens are more likely to participate. Second, negativity effects could make attack ads more useful than comparable positive ads in helping voters discriminate between the two candidates, and people with distinct alternatives are more likely to vote. Third, negative ads may be more likely to produce strong emotional responses that can get people to care more about the outcome of the election, which again should increase turnout.

Finkel & Geer (1998) then turn to an empirical investigation of whether the negativism of presidential election campaigns affects turnout. They measured the tone of presidential election campaigns by the percentage of negative themes in the campaign ads aired by the major party candidates between 1960 and 1992. Finkel & Geer use this variable to predict turnout in a thoroughly specified multivariate equation. Controlling for the standard demographic and political variables, they find no systematic relationship between turnout and tone. Thus, here is one study that adopted a different methodology and studied a different type of elections, but found no support for demobilization due to exposure to negative campaigns.

The major weakness of this study, as Finkel & Geer (1998) readily admit, is that campaign tone only varies across election years, providing very gross measurement of the crucial independent variable. When this study was published, there were no better data available. Geer & Lau (2006) try to overcome this limitation with a Bayesian analysis that treats these yearly averages as a baseline around which one can estimate plausible variations of tone across different electoral contexts (states)—and find very similar results.

But there have been many more tests of the demobilization hypothesis. The meta-analytic review of the literature by Lau et al. (2007) mentioned above identified 55 additional tests of demobilization. Of the results of those tests, 25 were in the demobilization direction (nine statistically significant results) and 29 were in the mobilization direction (eight of which were significant). This hypothesis has been such a focus of research in the field that we want to very carefully illustrate the lack of any clear findings. **Figure 3** displays the adjusted effect size for all 56 findings reported in the social science literature. An effect size is a way of expressing the results of many different studies in the same metric. For the prototypic study, it represents the difference in turnout between a group exposed to a positive ad (or campaign) and another group exposed to a negative ad (or campaign), expressed in standard deviation

units. These raw effect sizes can then be adjusted for sampling error (so that studies with larger sample sizes count more than studies with smaller sample sizes) and unreliability of measurement of the dependent variable. By meta-analytic standards, there is one very large positive effect in this literature, one moderately large positive effect, one moderately large negative effect, and 53 findings that hover a little above or a little below zero but do not amount to much. In sum, the research literature provides no general support for the hypothesis that negative political campaigning depresses voter turnout. In fact, if you had to bet on one or the other, mobilization is a bit more likely than demobilization.

Some of the individual findings in this literature are based on very large samples. As is always the case with statistics, particular findings can be statistically significant but not represent very large effects. For example, the experimental work reported by Ansolabehere & Iyengar (1995) is easily statistically significant with a *t* value of -2.56 , but because it is based on a sample of 2216 subjects, its effect size is only -0.10 . At the same time, the accompanying analysis of 34 senate election campaigns from 1992 (Ansolabehere et al. 1999) has a larger *t* value (-3.11), but this represents a *much* larger effect (-1.27) because it is based on only 34 cases. Of course we might trust a data point based on >2000 cases a lot more than one based on 34 cases, and it is conventional in the meta-analytic literature to adjust all effect sizes for sampling error, which has the effect of weighting findings based on larger samples relatively more than those based on a smaller number of cases. Lau et al. (2007) report the raw, unadjusted averages along with adjusted means; never do the various measures differ by very much.

It could still be the case that although voters are not turned off by all campaign negativism, when legitimate criticism crosses the line into mudslinging, demobilization occurs. But when exactly does a critical ad cross the line? Goodman (1995), writing in *Campaigns and Elections*, a magazine for political consultants, attempts to clarify the point: “Does it feel

negative? Is it mean-spirited? Is it crudely produced? Does it go beyond the pale of good taste and appropriate manners? Cross any of these lines, and you've crossed the viewer" (p. 23). Besides the fact that production sophistication is hardly an operational measure, the author's description of inappropriate negative advertising reflects another problem: Like the perception of campaign negativism, the "point of no return" is very much in the eye of the beholder. Although we can agree that presentations of documented facts are fair game, we are also aware that talented strategists are capable of manipulating facts in such a way that they distort the truth. Even topics that may be considered "off limits" (for example, those listed in the bottom panel of **Table 1**) can be framed in such a way that they are perceived by some voters to be relevant.

Yet Kahn & Kenney (1999) maintain that voters know the difference. They examine three distinct measures of the nature of the 1990 senate election campaigns: the tone of up to five ads aired by each candidate, the negativism of newspaper coverage of the campaign, and, of particular concern here, a judgment by the media consultants involved in these campaigns of whether the campaign degenerated into mudslinging. Kahn & Kenney use all three of these measures in a multivariate equation predicting turnout. Controlling on a variety of demographic, psychological, and political orientations, Kahn & Kenney find that both the negativism of the candidates' political ads and the amount of criticism they received in newspaper coverage were associated with higher turnout. Campaigns involving mudslinging, however, had significantly lower turnout. More fine-grained analyses suggest that political independents and those with little interest in politics were particularly affected by mudslinging. Although other researchers have had difficulty replicating this finding (see Brooks & Geer 2007, Jackson & Sides 2006), Kahn & Kenney's research provides some evidence that voters are capable of discerning malice from legitimate critiques, and that it is the former that turns voters off from politics, whereas legitimate, campaign-relevant

negative information mobilizes participation, particularly among independents and those who are generally more interested in or more knowledgeable about the political process.

METHODOLOGICAL CONCERNS

The major methodological issues plaguing the study of political campaigns can be placed into three broad categories: the lack of good data on the very nature of political campaigns, the difficulties of measuring exposure to political campaigns, and the statistical assumptions that are sometimes unwittingly made in our analysis of campaign effects. We are making progress on all of these fronts, however. The designs that researchers use to study political campaigns are one important vehicle for addressing some of these problems.

Data

As discussed above, the most fundamental problem in the study of political campaigns involves data: the lack, until very recently, of any good evidence on exactly what candidates actually do when they are running for office. With recent technological advances this information is becoming easier to gather, but there are still major parts of campaigns that remain largely hidden from the researcher's view. Before we had more objective measures of the actual content of political campaigns, some researchers asked survey respondents to recall an advertisement from some campaign, code the nature of the remembered ad, and then try to judge the impact of that ad on the voter. Alternatively, researchers asked respondents to rate how positive or negative a particular campaign has been. But memory is notoriously unreliable, and we have already seen that subjective perceptions of how negative a campaign is suffer from serious partisan biases. Noisy data about a phenomenon are always better than no data at all, but more objective, systematically gathered measures should almost always be preferred to less reliable subjective measures. Today there should be no reason to have to rely on memory

of political advertisements from any campaign as the primary measure of that campaign's tone.

We have already mentioned the Campaign Media Analysis Group (CMAG), which since 1998 has tracked every political advertisement aired on television in the 75 largest media markets (which include 80% of the population) in the United States.³ Among campaign professionals, the television and radio advertisements are often called the "air war." But there is also a "ground war," which consists of the speeches candidates deliver and campaign visits they make to different constituencies, the bulk mailings they send out, the billboards, lawn signs, bumper stickers, and get-out-the-vote phone calls. Any candidate blessed with sufficient resources to advertise on television also typically invests heavily in different aspects of the ground war. In fact, because ground war activities are not included in the definition of electioneering communication by the Bipartisan Campaign Reform Act of 2002, there are now incentives for parties to shift their campaign activities toward the ground (Dwyre & Kolodny 2003). It is usually possible to obtain a fairly complete record of the campaign stops a candidate makes from media accounts (see Shaw 1999), but we have not seen anyone systematically collect this information except at the presidential level. Magelby et al. (2006) describe a new data collection effort in which a random sample of registered voters were asked to collect every bit of campaign mail and log every political phone call and personal contact they received during the last three weeks of the 2004 presidential election (see also Hillygus & Monson 2006). But this is only one election, and the focus is on the presidential campaign. Today we must also speak of the "cyber war," the candidates' official web pages and the less official pages of their

supporters (Cornfield 2004; Druckman et al., unpublished manuscript). Limiting attention to only one aspect of a political campaign provides at best an incomplete image, and at worst a misleading picture, of the entire campaign. Yet few researchers have the resources to consider more than one or (rarely) two of these different aspects of a campaign. It is vitally important that any researcher who gathers data on the nature of any of the multiple components of a political campaign make that data available to other researchers studying campaigns.

Exposure

In gauging the effects of any aspect of a campaign, one also must be careful to distinguish the content or tone of the campaign from its volume. In using the CMAG data to measure negative advertising campaigns, for example, some researchers have mistakenly employed the number of negative ads shown in a media market as their primary measure of campaign negativism. Without a comparable measure of the number of positive ads shown in that same market, this approach confounds the nature of the advertising campaign with its magnitude, as Dan Stevens details in his unpublished manuscript, "More than Volume: Proportion, Volume, and the Effects of Exposure to Negative Advertising." The resulting regression coefficient will be estimating the effect of one more negative ad, all else (in the equation) equal, that is, irrespective of the number of positive ads that have also been shown. This makes little sense. We recommend devising separate measures of the nature or tone of the campaign (e.g., the proportion of all ads shown in a particular market that are negative) and of its volume (the total number of ads shown in that market).

Even if we have a very good idea about how much a candidate advertises in a particular media market, we still do not know how many campaign messages any individual in that media market has been exposed to. There is a large random element to individual exposure levels, but there is a larger nonrandom component to

³The Wisconsin Political Advertising Project also has CMAG data from the 1996 presidential election, but for no other races that election year. CMAG also records radio ads and the text of many local and cable news programs (to the extent their content has been transcribed for the hearing impaired), but these data are not stored at Wisconsin and are thus much more difficult for political scientists to obtain. We have never seen anyone utilize these other CMAG data.

exposure because most of us are creatures of habit. Some of us watch three or four hours of television almost every day whereas others watch an hour or less. Some of us watch television in the morning while getting ready for the day to start, others in the afternoon before the kids come home from school, and others exclusively during the prime-time hours in the evening. Freedman & Goldstein (1999) have devised survey-based measures of typical viewing habits that can be combined with objective measures of the number and mix of ads that were shown at various times of day. This is analogous to, although more complicated than, asking survey respondents which (if any) newspaper they typically read before estimating the impact of editorial endorsements on the vote choice. Best practice now dictates trying to combine some measure(s) of the content of a media message with measures of its likelihood of being heard, either through global measures of volume or, ideally, with the addition of more idiosyncratic measures of individual viewing habits.

Statistical Assumptions

Whenever we study a random sample of voters but measure media content at some aggregated level (such as the media market), we introduce another source of potential bias to our studies: data clustering. Inevitably with such research designs, individual respondents are nested within media markets, which violates basic assumptions of our most common multivariate statistical techniques (Steenbergen & Jones 2002). Unfortunately, following “best practices” in media research does not yet include explicitly modeling the multilevel structure of our data, and the conclusions from that research must consequently be somewhat suspect. One recent study found that for the most part, the results of multilevel analysis converge with those of conventional methods (Lau & Steenbergen 2006), but this result could be due to the small intraclass correlations observed in the three datasets examined in that paper. How

general this result will prove to be is still an open question.

We have already discussed the problem of the endogeneity of the decision of how to campaign with the outcome of the election, which makes causal statements about the likely effects of those campaign strategies much more difficult to make. Holding that problem aside, there is also the issue that every effort by one candidate to convince citizens to vote for her is typically offset by the efforts of another candidate trying to convince those same citizens to vote for him. Cast in this light, it may seem miraculous whenever we can detect *any* effect of a campaign on its outcome. Fortunately, there is still a great deal of variance in the conduct of campaigns that gives us leverage to estimate their effects. Even if candidates had the same amount of money to spend (as presidential candidates do in the United States, if they accept public financing), they do not have to spend it in the same way, nor in the same localities. Competing candidates typically do not try to deliver the same message or focus on the same issues in their advertising, nor are individual candidates required to deliver the same speeches or use the same advertisements in every locale within their constituency. Such variance is great for analytic purposes, but we must remember to take it into account in our data-gathering efforts. The Federal Election Commission, for example, has since its inception kept very good records on how much money candidates and parties raise, who they raise it from, and how much they spend. But there is almost no reporting on where the money is spent.

The study of political campaigns is made even more difficult by the fact that in most elections there are multiple races on the ballot. Every two years all citizens face an election selecting their member of the House of Representatives, but usually there is at least one and sometimes more simultaneously occurring statewide races for senator, governor, lieutenant governor, and so on. And every four years the presidential election is also on the

ballot. How much information are we losing by the almost universal practice of focusing on only one of these campaigns and ignoring every other race on the ballot? If our goal is to study the effectiveness of a particular campaign strategy, the practice might be justified. If instead we are interested in the indirect and presumably unintentional effects of negative campaigns on the political system as a whole (e.g., political efficacy, turnout), the practice seems much harder to justify. Most voters do not make race-by-race decisions of whether to vote—they make one universal decision to go to the polls or stay home, and to fully understand that universal decision, we must consider all of the political stimuli to which they have been subjected.

Research Designs

All of the methodological issues raised so far apply to the problems of studying real-world political campaigns. Most of these problems are solved by experiments, although now the natural focus becomes studying the effects of particular campaign events (e.g., individual campaign ads) rather than the effects of more extended campaigns. This review is not the place to discuss the relative merits of experimental and nonexperimental designs, or even the desirability of conducting “realistic” as opposed to more artificial or “ideal setting” experiments. Some degree of artificiality is inevitable in almost any controlled environment, but particularly in a discipline such as political science, many of whose practitioners have not been trained as experimental scientists, the experimental methodology developed by Iyengar & Kinder (1987) is about as good as it gets, providing high levels of internal and external validity. In the Iyengar & Kinder protocol, a carefully controlled experimental manipulation is unobtrusively delivered in a very familiar setting—for example, an experimental ad is inserted into the normal commercial break in a tape of an actual television program.

Experiments may not be as useful in trying to understand the decisions made by candidates

and their political consultants during an actual political campaign, however. If this is one’s goal, then longitudinal designs, which gather information (e.g., relative standing in the polls) and measure behavior (the volume and tone of campaigning in different media markets) at multiple points in time across an entire campaign, are crucial. We as scientists can wait to analyze the success of any advertising campaign until after it is over, but campaigners must make their decisions during the heat of battle, so to speak, and undoubtedly under conditions of considerable uncertainty. Fortunately the CMAG data are collected with very precise time stamps, so we know exactly when and where different ads are shown. It is typically much harder to get comparably precise measures of the public’s response to any advertising campaign, but many campaigns conduct tracking polls that provide exactly those data—if we can get a hold of them. Developing personal relationships with political consultants who often manage high-level political campaigns can prove invaluable to such an enterprise [see Sides (2006) for a recent study utilizing this type of data].

CONCLUSION: ARE NEGATIVE CAMPAIGNS BAD FOR DEMOCRACY?

The literature on the use and utility of negative campaigns is extensive, including descriptive accounts, systematic experimentation, survey research analysis, and formal theoretical prescriptions. Although we have not had the space to present a comprehensive review of all of this research, we do believe we have presented a thorough one. Descriptive accounts, many of which are by political consultants rather than academics, address the strategic concerns guiding the decision to attack. Empirical investigations include both systematic manipulations of hypothesized causal factors, analysis of media content—real, as well as fictitious—and survey analysis of voter reactions to actual political campaigns.

Yet there is nothing in what we have presented so far that would allow us to say that

negative campaigns are inherently good or bad. Many recognize that voters are capable of separating vicious attacks from legitimate comparisons but also appreciate that even justifiable attacks are not always effective, nor are they always desirable. The standards for making such judgments should be the public's interests, not the candidates'. Elections are, after all, about choices, not courtesy (Lau & Pomper 2004), so questions about the effectiveness of negative campaigns become irrelevant.

The literature does reflect an overall indictment of baseless attacks and unwarranted character assaults. Beyond that, however, exists a spectrum of views. The widespread concern that negative campaigns can reduce turnout appears unwarranted. There is less concern but stronger evidence that negative campaigns can deleteriously affect the political system in other ways, for example, by reducing political efficacy and decreasing trust in government (Lau et al. 2007). But negative campaigns also stimulate interest and learning about the candidates and issues in an election.

Is it unfair to objectively point out an opponent's shortcomings? Conversely, are so-called "positive" ads that boast of a candidate's strength to the point of exaggerated self-aggrandizement fair, simply because they are not negative? Mayer (1996) points out that no candidate can speak honestly and substantively about what he intends to do upon election without drawing a contrast to his opponent and explaining why he is a superior choice over the other candidate. If we assume that voters want to know why a given candidate is better suited or more qualified than his or her opponent, we cannot legitimately fault a candidate for tactfully juxtaposing him- or herself against that opponent.

Geer (2006) takes this argument even farther. All advertisements, by their very nature, exaggerate the truth. But this point applies equally to positive campaigns as well as to negative. We cannot expect candidates to discuss their own policy failures and shortcomings, any more than we can expect Chevrolet to tell us about the problems with its cars, or Hoover

to tell us its vacuum cleaners sometimes fail to suck. That is the job of the opposition, in politics as well as in commerce. Geer argues that democracy *requires* negative campaigning if citizens are to become fully informed about the choices they face.

Although Ansolabehere & Iyengar (1995) conclude their book with policy recommendations aimed at limiting the extent of negative campaigning, we see no need to impose any further restrictions on negative advertising. Voters are clearly capable of distinguishing legitimate criticism from unfair attacks and punish candidates who are out of line. We do approve of recent legislation requiring federal candidates to appear in and publicly approve of the ads they sponsor, as this increases accountability in election campaigns and exposes those individuals who would prefer to shield themselves from boomerang effects by allowing surrogates to do the attacking. Of course what this change has really done is push most of the attacking off to party-sponsored ads, but as long as voters can hold the parties accountable for their actions, that too is fine.

Although it is difficult to convincingly document the widely presumed increase in negative campaigning over recent decades, no one is suggesting the practice is on the decline, and there is little doubt that negative campaigns are an important part of the current political scene. Thus we are sure that negative campaigning will remain fertile ground for political science research for years to come. Future research should explore the impact of negative ads that are sponsored by "independent" (504) groups and individuals other than the candidates themselves. It would be good to examine the extent to which the messenger matters, and how attacks from third parties are weighted relative to those launched personally by candidates. How does the timing of attacks during a campaign, or the medium through which attacks are delivered, influence their effectiveness? Is the internet a medium where almost anything goes, and candidates can "get away with" more vicious attacks of their opponents? What is the role of music, production quality, or particular

types of images in the effectiveness or effects of negative advertisements? Despite the volume of research already available on the decision to attack and its impact on political behavior, a myriad of questions remain for future research to address.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

LITERATURE CITED

- Ansolabehere S, Iyengar S. 1995. *Going Negative*. New York: Free Press
- Ansolabehere S, Iyengar S, Simon A. 1999. Replicating experiments using aggregate and survey data: the case of negative advertising and turnout. *Am. Polit. Sci. Rev.* 93:901–9
- Ansolabehere S, Iyengar S, Simon A, Valentino N. 1994. Does attack advertising demobilize the electorate? *Am. Polit. Sci. Rev.* 88:829–38
- Bartels L. 1991. Instrumental and “quasi-instrumental” variables. *Am. J. Polit. Sci.* 35(July):777–800
- Brader T. 2005. Striking a responsive chord: how political ads motivate and persuade voters by appealing to emotions. *Am. J. Polit. Sci.* 49(Apr.):388–405
- Brader T. 2006. *Campaigning for Hearts and Minds*. Chicago: Univ. Chicago Press
- Brooks D, Geer J. 2007. Beyond negativity: the effects of incivility on the electorate. *Am. J. Polit. Sci.* 51(Jan.):1–16
- Buell E, Sigelman L. 2008. *Attack Politics: Negativity in Presidential Campaigns since 1960*. Lawrence: Univ. Press Kansas
- Cornfield M. 2004. *Politics Moves Online: Campaigning and the Internet*. New York: Century Found. Press
- Dwyre D, Kolodny R. 2003. National parties after BCRA. In *Life after Reform: When the Bipartisan Campaign Reform Act Meets Politics*, ed. M Malbin, pp. 95–112. Lanham, MD: Rowman & Littlefield
- Finkel S, Geer J. 1998. A spot check: casting doubt on the demobilizing effect of attack advertising. *Am. J. Polit. Sci.* 42:573–95
- Freedman P, Goldstein K. 1999. Measuring media exposure and the effects of negative campaign ads. *Am. J. Polit. Sci.* 43:1189–208
- Freedman P, Wood W, Lawton D. 1999. Do’s and don’ts of negative ads: what voters say. *Campaigns Elections* 20:20–25
- Garramone G. 1984. Voter response to negative political ads. *J. Q.* 61:250–59
- Geer J. 2006. *In Defense of Negativity: Attack Ads in Presidential Campaigns*. Chicago: Univ. Chicago Press
- Geer J, Lau R. 2006. A new approach for estimating campaign effects. *Br. J. Polit. Sci.* 35:269–90
- Goodman A. 1995. Producing TV: a survival guide. *Campaigns Elections* 16:22–24
- Harrington J, Hess G. 1996. A spatial theory of positive and negative campaigning. *Games Econ. Behav.* 17:209–29
- Haynes A, Rhine S. 1998. Attack politics in presidential nomination campaigns: an examination of the frequency and determinants of intermediated negative messages against opponents. *Polit. Res. Q.* 51:691–721
- Hillygus S, Monson J. 2006. *Campaign microtargeting and presidential voting in 2004*. Presented at Annu. Meet. Am. Polit. Sci. Assoc., 102nd, Philadelphia
- Iyengar S, Kinder D. 1987. *News that Matters*. Chicago: Univ. Chicago Press
- Jackson R, Sides J. 2006. Revisiting the influence of campaign tone in senate elections. *Polit. Anal.* 14(Spring):206–18
- Jacobson G. 2004. *The Politics of Congressional Elections*. New York: Longman. 6th ed.
- Jamieson KH, Campbell C. 1983. *The Interplay of Influence: Mass Media and their Publics in News, Advertising, Politics*. Belmont, CA: Wadsworth
- Jamieson KH, Waldman P, Sherr S. 2000. Eliminate the negative? Categories of analysis for political advertisements. In *Crowded Airwaves*, ed. JA Thurber, CJ Nelson, DA Dulio, pp. 44–64. Washington, DC: Brookings Inst.

- Kahn KF. 1993. Gender differences in campaign messages: the political advertisements of men and women candidates for U.S. Senate. *Polit. Res. Q.* 46:481–502
- Kahn KF, Kenney P. 1999. Do negative campaigns mobilize or suppress turnout? Clarifying the relationship between negativity and participation. *Am. Polit. Sci. Rev.* 93:877–89
- Kahn KF, Kenney P. 2000. How negative campaigning enhances knowledge of senate elections. In *Crowded Airwaves*, ed JA Thurber, CJ Nelson, DA Dulio, pp. 65–95. Washington, DC: Brookings Inst.
- Kahn KF, Kenney P. 2004. *No Holds Barred: Negativity in U.S. Senate Campaigns*. Upper Saddle River, NJ: Pearson Educ.
- Kamber V. 1997. *Poison Politics: Are Negative Campaigns Destroying Democracy?* Cambridge, MA: Perseus
- Kanouse D, Hanson L Jr. 1971. Negativity in evaluations. In *Attribution: Perceiving the Causes of Behavior*, ed. EE Jones, DE Kanouse, HH Kelley, RE Nisbett, S Valins, B Weiner, pp. 47–62. Morristown, NJ: General Learning
- Kern M. 1989. *Thirty Second Spots: Political Advertising in the Eighties*. New York: Praeger
- Kernell S. 1977. Presidential popularity and negative voting: an alternative explanation of the midterm congressional decline of the president's party. *Am. Polit. Sci. Rev.* 71(Mar.):44–66
- King E, Hendersen R, Chen H. 1998. *Viewer response to positive vs negative ads in the 1996 presidential campaign*. Presented at Annu. Meet. Midwest Polit. Sci. Assoc., Chicago
- Lau R. 1982. Negativity in political perception. *Polit. Behav.* 4(Dec.):353–78
- Lau R. 1985. Two explanations for negativity effects in political behavior. *Am. J. Polit. Sci.* 29(Feb.):119–38
- Lau R, Pomper G. 2001. Negative campaigning by U.S. Senate candidates. *Party Polit.* 7(Jan.):69–87
- Lau R, Pomper G. 2002. Effectiveness of negative campaigning in U.S. Senate elections. *Am. J. Polit. Sci.* 46:47–66
- Lau R, Pomper G. 2004. *Negative Campaigning: An Analysis of U.S. Senate Elections*. New York: Rowman & Littlefield
- Lau R, Sigelman L, Heldman C, Babbitt P. 1999. The effects of negative political advertisements: a meta-analytic assessment. *Am. Polit. Sci. Rev.* 93:851–75
- Lau R, Sigelman L, Rovner IB. 2007. The effects of negative political campaigns: a meta-analytic reassessment. *J. Polit.* 69(Nov.):1176–209
- Lau R, Steenbergen M. 2006. *Reconsidering "best practices" in political communications research*. Presented at Annu. Meet. Midwest Polit. Sci. Assoc., 64th, Chicago
- Magelby D, Monson JQ, Patterson K. 2006. *Mail communications in political campaigns: the 2004 campaign communications survey*. Presented at Annu. Meet. Midwest Polit. Sci. Assoc., Chicago
- Mansnerus L. 2005. Forrester tosses his "Hail Mary" ad. *NY Times*, Oct. 23, p. NJ14
- Marcus G, Neuman WR, MacKuen M. 2000. *Affective Intelligence and Political Judgement*. Chicago: Univ. Chicago Press
- Mayer W. 1996. In defense of negative campaigning. *Polit. Sci. Q.* 111:437–55
- Proctor D, Schenck-Hamlin W, Haase K. 1994. Exploring the role of gender in the development of negative political advertisements. *Women Polit.* 14(2):1–22
- Roose N, Sande G. 1993. Backlash effects in attack politics. *J. Appl. Soc. Psychol.* 23(Aug.):632–53
- Shaw D. 1999. The effect of TV ads and candidate appearances on statewide presidential votes, 1988–1996. *Am. Polit. Sci. Rev.* 93(June):345–61
- Sides J. 2006. *Swing your partner? The importance of interactions in campaigns*. Presented at Annu. Meet. Am. Polit. Sci. Assoc., Washington, DC
- Sides J, Grossmann M, Trost C, Lipsitz K. 2003. *Candidate attacks and voter aversion: the uncertain link between negativity and campaign satisfaction*. Presented at Annu. Meet. Am. Polit. Sci. Assoc., Philadelphia
- Sigelman L, Kugler M. 2003. Why is research on the effects of negative campaigning so inconclusive? Understanding citizens' perception of negativity. *J. Polit.* 65:142–60
- Skaperdas S, Grofman B. 1995. Modeling negative campaigning. *Am. Polit. Sci. Rev.* 89:49–61
- Steenbergen M, Jones B. 2002. Modeling multilevel data structures. *Am. J. Polit. Sci.* 46(Jan.):218–37
- Swint K. 1998. *Political Consultants and Negative Campaigning: The Secrets of the Pros*. Lanham, MD: Univ. Press Am.
- Theilmann J, Willhite A. 1998. Campaign tactics and the decision to attack. *J. Polit.* 60:1050–62
- West D. 1993. *Air Wars: Television Advertising in Election Campaigns, 1952–1992*. Washington, DC: CQ Press
- West D. 2005. *Air Wars: Television Advertising in Election Campaigns, 1952–2004*. Washington, DC: CQ Press. 4th ed.




Figure 1

Positive, negative, and contrast televised political ads in recent U.S. federal elections.

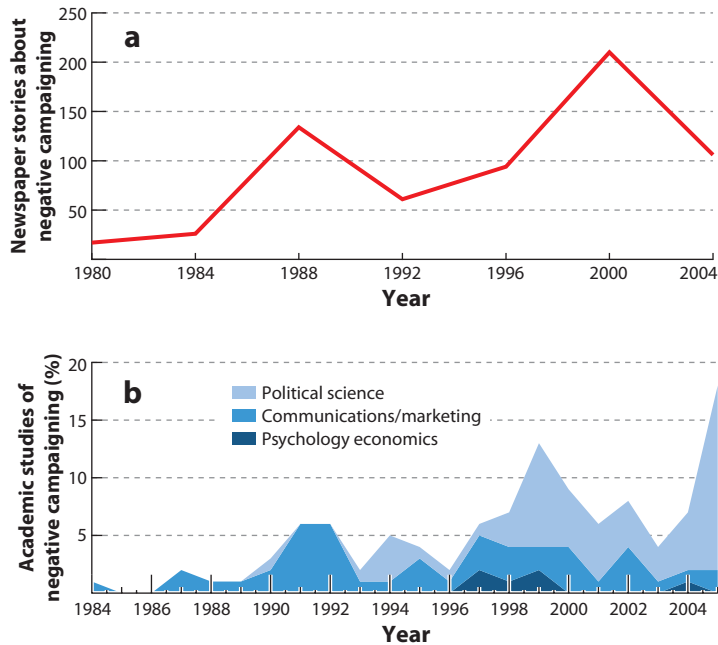


Figure 2

(a) Number of newspaper stories about negative political campaigns. (b) Academic studies of negative campaigning, by discipline.

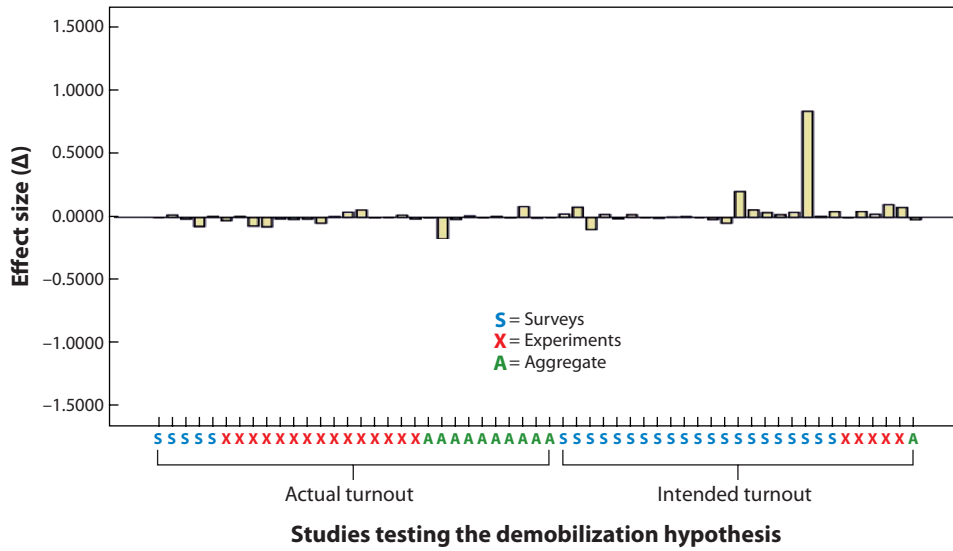


Figure 3

Do negative campaigns demobilize the electorate? This figure displays the effect size associated with every study testing the demobilization hypothesis, adjusted for sampling error and measurement unreliability (Lau et al. 2007). Data have been sorted by whether the dependent variable is *actual* (reported) vote or *intended* vote, and within that distinction, by methodology. Negative effect sizes are consistent with the demobilization hypothesis.



Contents

A Conversation with Robert A. Dahl <i>Robert A. Dahl and Margaret Levi</i>	1
Neorepublicanism: A Normative and Institutional Research Program <i>Frank Lovett and Philip Pettit</i>	11
Domestic Terrorism: The Hidden Side of Political Violence <i>Ignacio Sánchez-Cuenca and Luis de la Calle</i>	31
Women in Parliaments: Descriptive and Substantive Representation <i>Lena Wängnerud</i>	51
Self-Government in Our Times <i>Adam Przeworski</i>	71
Social Policy in Developing Countries <i>Isabela Mares and Matthew E. Carnes</i>	93
Variation in Institutional Strength <i>Steven Levitsky and María Victoria Murillo</i>	115
Quality of Government: What You Get <i>Sören Holmberg, Bo Rothstein, and Naghmeh Nasiritousi</i>	135
Democratization and Economic Globalization <i>Helen V. Milner and Bumba Mukherjee</i>	163
Has the Study of Global Politics Found Religion? <i>Daniel Philpott</i>	183
Redistricting: Reading Between the Lines <i>Raymond La Raja</i>	203
Does Efficiency Shape the Territorial Structure of Government? <i>Liesbet Hooghe and Gary Marks</i>	225
Bargaining Failures and Civil War <i>Barbara F. Walter</i>	243
Hobbesian Hierarchy: The Political Economy of Political Organization <i>David A. Lake</i>	263

Negative Campaigning <i>Richard R. Lau and Ivy Brown Rovner</i>	285
The Institutional Origins of Inequality in Sub-Saharan Africa <i>Nicolas van de Walle</i>	307
Riots <i>Steven I. Wilkinson</i>	329
Regimes and the Rule of Law: Judicial Independence in Comparative Perspective <i>Gretchen Helmke and Frances Rosenbluth</i>	345
Field Experiments and the Political Economy of Development <i>Macartan Humphreys and Jeremy M. Weinstein</i>	367
Laboratory Experiments in Political Economy <i>Thomas R. Palfrey</i>	379
Field Experiments on Political Behavior and Collective Action <i>Eline A. de Rooij, Donald P. Green, and Alan S. Gerber</i>	389
Experiments on Racial Priming in Political Campaigns <i>Vincent L. Hutchings and Ashley E. Jardina</i>	397
Elections Under Authoritarianism <i>Jennifer Gandhi and Ellen Lust-Okar</i>	403
On Assessing the Political Effects of Racial Prejudice <i>Leonie Huddy and Stanley Feldman</i>	423
A “Second Coming”? The Return of German Political Theory <i>Dana Villa</i>	449
Group Membership, Group Identity, and Group Consciousness: Measures of Racial Identity in American Politics? <i>Paula D. McClain, Jessica D. Johnson Carew, Eugene Walton, Jr., and Candis S. Watts</i>	471
Opiates for the Matches: Matching Methods for Causal Inference <i>Jasjeet Sekhon</i>	487
Indexes	
Cumulative Index of Contributing Authors, Volumes 8–12	509
Cumulative Index of Chapter Titles, Volumes 8–12	511

Errata

An online log of corrections to *Annual Review of Political Science* articles may be found at <http://polisci.annualreviews.org/>