

Breaking Empirical Deadlocks in the Study of Partisanship:

An Overview of Experimental Research Strategies

Donald P. Green

Department of Political Science

Columbia University

September 26, 2011

Abstract: The vast literature on party identification has gradually become bogged down over disputes about how to interpret observational data. This paper proposes the use of experimental research designs to shed light on the responsiveness of party identification to short term forces such as retrospective performance evaluations. Examples of recent field experiments are used to illustrate two types of experimental designs and the assumptions on which they rest.

Paper prepared for presentation at the CISE-ITANES Conference “Revisiting Party Identification: American and European Perspectives.” LUISS-Guido Carli, Roma, October 7-8, 2011. The author is grateful to the Institution for Social and Policy Studies at Yale, which furnished the replication data analyzed in this paper. Comments may be directed to Donald.P.Green@gmail.com.

Introduction

The vast behavioral literature on party identification has been propelled by a series of methodological innovations. The initial conceptualization of party identification as an enduring attachment that shapes the way in which voters view political figures and issues (Belknap and Campbell 1952; Campbell et al. 1960) was prompted by the growth and development of survey research in the early 1950s, and theoretical refinements followed as surveys became more nuanced and widespread (Campbell et al. 1966; Butler and Stokes 1969; Jennings and Neimi 1981). During the mid-1970s, nonrecursive statistical models became part of the political science toolkit, and a torrent of studies called into question the assumption that causation flows in one direction from party attachments to issue positions (Jackson 1975; Franklin and Jackson 1983), performance evaluations (Fiorina 1981; Brody and Rothenberg 1988), and candidate evaluations (Page and Jones 1979). By the mid-1980s, political scientists had become deeply skeptical of the view that identification with party is an unmoved mover, developing early in life and responding weakly to short-term changes in the political environment. These simultaneous equations models came under criticism in the wake of another methodological development, the analysis of covariance structures as a means of addressing biases due to measurement error. Response error was said to produce a variety of statistical artifacts, leading scholars to exaggerate the rate of partisan change (Achen 1975; Green and Palmquist 1994) and the responsiveness of partisanship to short-term changes in the way that voters evaluate incumbent performance and candidates' issue stances (Green and Palmquist 1990) in a variety of cross-national settings (Schickler and Green 1997). The most recent methodological innovation was the analysis of aggregate survey data, made possible by the accumulation of several decades of quarterly polling data by commercial and news organizations (MacKuen, Erikson, and Stimson 1989). This evidence was initially interpreted as demonstrating the malleability of partisanship

in the wake of economic fluctuations and scandals, although subsequent work that took sampling variability into account tempered this conclusion (Green, Palmquist, and Schickler 1998; Green and Schickler 2009).

Although each wave of methodological innovation introduced new evidence into debates about the nature and origins of party attachments, there remained considerable uncertainty about how to interpret the evidence and accompanying methodological debate. The study of partisanship currently finds itself in a state of deadlock between theoretical perspectives that emphasize the stability of partisan identities (and social identities more generally) in polities where the parties and their social constituencies are stable (Green, Palmquist, and Schickler 2002) and theoretical perspectives that regard partisanship as a running tally of past performance evaluations (Fiorina 1981; MacKuen, Erikson, and Stimson 1998), a summary of expectations about future performance (Achen 2002), or a manifestation of voters' ideological proximity to the parties (Franklin and Jackson 1983)

How might researchers break this deadlock? Many of the central debates ultimately come down to questions of identification and causal inference. The reason methodological debates about two-way causal flows, measurement error, and other specification issues have played such a prominent role in the literature on party identification is that the evidence base is almost entirely drawn from nonexperimental research. Cross-sectional surveys, panel surveys, and aggregate time-series furnish the data analyst with variation in partisanship and variation in the putative causes of partisanship. What to make of the correlation between these two sets of variables hinges on the substantive assumptions that researchers bring to bear when analyzing the data. For example, in cross-sectional analysis (e.g., Jackson 1975), the identifying assumption is that certain demographic variables predict issue stances but are unrelated to

omitted causes of party identifications. In panel analysis, the core assumption is only slightly weaker: subjects' background attributes and prior attitudes are related to current partisanship only insofar as they influence contemporary issue stances and performance evaluations (e.g., Brody and Rothenberg 1988). In time-series analysis, the identifying assumptions are somewhat more complex because they involve a range of propositions about how partisanship and short-term forces are measured over time and how the dynamics of each series are modeled (Box-Steffensmeier and Smith 1996; Erikson, MacKuen, and Stimson 1998; Green, Palmquist, and Schickler 1998). Suffice it to say that each of the competing modeling approaches involves strong and untestable modeling assumptions. New statistical techniques (e.g., matching) that introduce untestable assumptions of their own are unlikely to advance this literature.

During the past decade, largely in response to the kinds of identification problems just mentioned, a new methodological innovation has taken root in the social sciences. Increasingly, researchers in political science and economics have turned to randomized experiments in order to facilitate causal inference. Experimental designs by no means eliminate problems of inference, but they nonetheless represent an important advance that, at a minimum, calls attention to the subtle issues of identification and interpretation. This paper discusses a pair of essays that illustrate two broad classes of experimental designs. The first addresses the question of what kinds of stimuli cause people to alter their partisan attachments; the second addresses the question of what downstream consequences follow from an exogenously-induced change in partisanship. We begin by introducing the logic of inference that underlies randomized experiments, discuss the identification strategies that underlie each essay, and suggest how an experimental agenda might advance the literature on party identification.

Inference from Direct and Downstream Experiments

Randomized experiments – and research designs that attempt to approximate random assignment – are often explicated in terms of a notational system that has its origins in Neyman (1923) and is usually termed the “Rubin Causal Model,” after Rubin (1990). For each individual i let $Y_i(0)$ be the outcome if i is not exposed to the treatment, and $Y_i(1)$ be the outcome if i is exposed to the treatment. The treatment effect is defined as:

$$\tau_i \equiv Y_i(1) - Y_i(0). \quad (1)$$

In other words, the treatment effect is defined as the difference between two potential states of the world, one in which the individual receives the treatment, and another in which the individual does not. Extending this logic from a single individual to a set of individuals, we may define the average treatment effect (ATE) as follows:

$$ATE \equiv E[\tau_i] = E[Y_i(1)] - E[Y_i(0)], \quad (2)$$

where $E[.]$ indicates an expectation over all subjects. Although experimental investigation may serve many purposes, one principal aim is to estimate the ATE, the average effect of introducing some sort of information, policy, or incentive.

In an actual experiment, we observe subjects in either their treated or untreated states. Let D_i denote the treatment status of each subject, where $D_i=1$ if treated and 0 if not. The observed potential outcomes are $Y_i(D_i)$, which is just another way of writing the outcomes that result when D_i is the input. The difference in expected outcomes among those we treat and those we do not treat may be expressed as

$$E[Y_i(1)|D_i = 1] - E[Y_i(0)|D_i = 0], \quad (3)$$

where the notation $E[A_i|D_i = B]$ means the average value of A_i among those subjects for which the condition $D_i = B$ holds. For example, one could compare average outcomes (party identification scores) among those who evaluate the economy positively ($D_i = 1$) to average outcomes among those who evaluate the economy negatively ($D_i = 0$).

In a typical observational study, of course, the observed difference in partisanship between those who evaluate the economy positively or negatively does not, in expectation, reveal the average causal effect of economic perceptions. We observe average outcomes for the treated subjects in their treated state and average outcomes of the untreated subjects in their untreated state. To see how this quantity is different, in expectation, from the ATE, we rewrite equation (3) as:

$$E[Y_i(1) - Y_i(0)|D_i = 1] + \{E[Y_i(0)|D_i = 1] - E[Y_i(0)|D_i = 0]\}. \quad (4)$$

In other words, the expected difference in outcomes of the treated and untreated can be decomposed into the sum of two quantities: the average treatment effect for a subset of the subjects (the treated), and a selection bias term. The selection bias term (in braces) is the difference between what the outcome $Y_i(0)$ would have been for those who are treated had they not been treated and the value of $Y_i(0)$ observed among those who were not treated. The threat of selection bias arises whenever systematic processes determine which people receive treatment. In this example, if people choose the sorts of economic news they read and remember, expected $Y_i(0)$ potential outcomes may be quite different among those who evaluate the economy differently.

Random assignment solves the selection problem. When random assignment determines which treatment each subject receives, D_i is independent of potential outcomes. Those randomly

selected into the treatment group have the same expected outcomes in the treated state as those randomly assigned to remain untreated (control group):

$$E[Y_i(1)|D_i = 1] = E[Y_i(1)|D_i = 0] = E[Y_i(1)]. \quad (5)$$

By the same token, those randomly assigned to the control group have the same expected $Y_i(0)$ outcomes as those assigned to the treatment group:

$$E[Y_i(0)|D_i = 0] = E[Y_i(0)|D_i = 1] = E[Y_i(0)]. \quad (6)$$

Equations (5) and (6) reveal why, when subjects are randomly treated, the selection bias term vanishes and the difference of treatment and control group means measures the ATE. This identification result can be shown by substituting equations (5) and (6) into equation (1):

$$E[Y_i(1)|D_i = 1] - E[Y_i(0)|D_i = 0] = E[Y_i(1)] - E[Y_i(0)]. \quad (7)$$

This proof demonstrates an attractive property of randomized experiments. At the same time, it glosses over two implicit assumptions. One assumption, which plays a minor role in the analysis that follows, is the stable unit treatment value assumption (Rubin 1990), which stipulates that potential outcomes do not depend on which subjects are assigned to treatment. This assumption is jeopardized, for example, when the treatment administered to one subject affects the outcomes of other subjects. More pertinent to our discussion below is the exclusion restriction assumption (Angrist, Imbens, and Rubin 1996), which requires that outcomes respond solely to the treatment itself and not the random assignment or other backdoor causal pathways that are set in motion by the random assignment. For example, we must assume that when we randomly assign economic evaluations, we are not inadvertently deploying other treatments, such as information about the party platforms on environmental issues.

Readers may be wondering whether an experiment could feasibly assign how people evaluate the economy. The answer is probably not, and we must therefore introduce another layer of notation to describe the imperfect translation of intended treatments into actual treatments. Let $Z_i = 1$ if a subject is *assigned* to the treatment group, and $Z_i = 0$ if the subject is assigned to the control group. In experiments with full compliance, all those assigned to the treatment group ($Z_i = 1$) also receive the treatment ($D_i = 1$), and all those assigned to the control group ($Z_i = 0$) are untreated ($D_i = 0$). In experiments with some degree of noncompliance, $D_i(Z_i) \neq Z_i$. Encouragement designs, for example, attempt to induce some subjects to take the treatment D_i but recognize that there may be some subjects who will fail to do so or who will take the treatment even when not encouraged.

In the context of experiments that encounter noncompliance, the exclusion restriction holds that $Y_i(D_i, Z_i) = Y_i(D_i)$ for all D_i and Z_i . In other words, potential outcomes respond solely to actual treatment, not assigned treatment. Consider a recent experiment by Middleton (2011) that randomly encourages some subjects to read upbeat news stories about the economy (Z_i) in an effort to change their assessment of national economic conditions (D_i), which in turn may affect partisanship (Y_i). The causal effect of interest is the influence of D_i on Y_i , but D_i itself is not randomly assigned. The exclusion restriction holds that assignment Z_i has no influence on Y_i except insofar as it affects D_i . The encouragement is assumed to affect partisanship only insofar as the encouragement changes assessments of national economic conditions.

In order to recover the causal effect of D_i on Y_i using an encouragement design, we need one further assumption known as monotonicity (Angrist, Imbens, and Rubin 1996). Describing this assumption requires a bit more terminology. Depending on the way their received

treatments potentially respond to treatment assignment, subjects may be classified into four types, Compliers, Never-takers, Always-takers, and Defiers. Compliers are subjects who take the treatment if and only if assigned to the treatment. For this group $D_i(1) - D_i(0) = 1$. Never-takers are those who are always untreated no matter their assignment: $D_i(1) = D_i(0) = 0$. Conversely, Always-takers are those who are always treated no matter their assignment: $D_i(1) = D_i(0) = 1$. Defiers are those who take the treatment if and only if they are assigned to the control group: $D_i(1) - D_i(0) = -1$. The monotonicity assumption rules simply stipulates that there are no Defiers. In context of our running example, when assigned to receive upbeat economic news, everyone's economic assessments either remain unchanged or becomes more buoyant.

Under the stable unit treatment value, exclusion restriction, and monotonicity assumptions, one can identify the ATE among Compliers, also known as the Complier average causal effect (CACE). This quantity is estimated by dividing two quantities. The numerator in equation (8) is the average outcome in the assigned treatment group minus the average outcome in the assigned control group; the denominator is the observed rate of treatment in the assigned treatment and control groups:

$$\widehat{CACE} = \frac{\hat{E}[Y_i|Z_i=1] - \hat{E}[Y_i|Z_i=0]}{\hat{E}[D_i|Z_i=1] - \hat{E}[D_i|Z_i=0]} \quad (8)$$

This ratio is equivalent to the estimate generated by a 2SLS regression of Y_i on D_i using Z_i as an instrumental variable. Because the denominator is a difference, this ratio is consistent but not unbiased and becomes undefined when the treatment rate in the two experimental groups is the same. Precise estimation requires a substantial difference in treatment rates, a point that has special importance for the analysis of downstream experiments.

A downstream experiment (Green and Gerber 2002) is one in which an initial randomization causes a change in an outcome, and this outcome is then considered a treatment affecting a subsequent outcome. For example, in Middleton's (2011) study of news coverage on economic assessments, subjects in an internet survey were assigned to read newspaper coverage of the 2008 economic crisis. Random assignment produced a change in economic evaluations. A downstream analysis might examine the consequences of changing economic evaluations on party identification. This analysis parallels an encouragement design in terms of its underlying assumptions (stable unit treatment value, exclusion restriction, monotonicity), mode of analysis (instrumental variables regression), and causal estimand (the CACE). Of special importance is the exclusion restriction, which holds that exposure to news stories had no effect on party identification through paths other than economic evaluations. When these assumptions are met, the experimenter obtains consistent estimates of the ATE among Compliers, who are in this case those whose economic evaluations are favorable if and only if they are exposed to the news stories. In order to estimate the CACE with reasonable power, there must be ample numbers of Compliers, which is to say that the news stories must have a sizable impact on economic evaluations.¹

In sum, random assignment allows researchers to sidestep the selection problem, but important assumptions remain. Both full-compliance and encouragement designs force the researcher to impose exclusion restrictions. Encouragement designs require the additional assumption of monotonicity and confine the causal estimand to the average treatment effect among Compliers. Whether one can safely generalize from the ATE among Compliers to the

¹ The usual rule of thumb is that the t-ratio from a regression of D_i on Z_i should be greater than $\sqrt{10} \approx 3.2$ (Staiger and Stock 1997).

ATE for other subgroups is an open question that may be addressed empirically through replication using different sorts of encouragements.

From the standpoint of estimation, this framework departs markedly from the usual way in which researchers typically analyze observational data. Here, subjects are compared according to their experimental assignments, not according to the treatments they actually receive. Precise estimation requires that the assigned treatments bear a reasonably strong relationship to the treatments that subjects actually receive. In other words, the use of instrumental variables regression to estimate the CACE requires an experimental design that generates ample numbers of Compliers.

In order to see these assumptions and design considerations in action, we next consider a pair of recent experiments. The first illustrates an experiment that administers a treatment thought to have a direct effect on party attachments. The second illustrates the logic of a downstream experiment.

Engineering a Direct Effect on Party Identification: Chong et al. (2011)

Chong, De La O, Karlan, and Wantchekon (2011) report the results of a field experiment conducted in Mexico shortly before its 2009 municipal elections. Their intervention follows in the wake of a federal audit of municipal governments. These audits graded municipal governments according to whether they had accounting irregularities indicative of corruption; the auditors also noted whether local administrators had failed to spend federal grant money, suggesting a low level of administrative competence. The researchers conducted a precinct-level leafleting campaign designed to publicize some aspect of the auditors' reports. Some 1,910 precincts were randomly selected to a control group that received no leaflets. Three random

subsets of 150 precincts apiece each received one type of treatment flyer. The first treatment publicized the degree to which the municipality failed to spend federal grant funds. The second publicized the failure to spend grant funds that were supposed to aid the poor. A third graded the municipality according to the amount of evidence of corruption.

Much of the authors' report focuses on how precinct-level vote outcomes changed in the wake of the leafleting campaign; for our purposes, the relevant part of the study examines the effects of the intervention on individual-level attitudes of 750 respondents who were sampled from 75 of the precincts and surveyed two weeks after the election. Since Mexican elected officials are forbidden from seeking reelection, voter displeasure cannot be directed at incumbent candidates; the relevant target is the incumbent party. Chong et al. find that negative report cards addressing corruption (but not failure to spend grant money) significantly diminish respondents' approval of the incumbent mayor and identification with the incumbent's political party.² Unfortunately, no follow-up surveys were conducted to assess the extent to which the effects persisted beyond two weeks. Nevertheless, the study remains one of the first experiments to show that party attachments change when performance evaluations are altered exogenously.³

Let's now consider the study from the standpoint of the core assumptions discussed in the previous section. The exclusion restriction in this instance stipulates that random distribution of corruption-related leaflets influences outcomes because it provides evaluative information about

² The authors took care to allow for precinct-level clustering when calculating the standard error of this estimated treatment effect. For other methodological commentary on this study, see Gerber and Green

³ Given the sheer number of studies on the topic of party identification, it is surprising how few studies have attempted to influence party identification via an experimental manipulation. One rare exception is Cowden and McDermott (2000), which reports the results of a series of laboratory studies that sought to influence party attachments though, among other things, role-playing exercises in which undergraduate subjects were asked to take a pro- or ant-Clinton position. None of their interventions succeeded in changing party attachments.

incumbent performance. The authors present convincing evidence that the leaflets did tarnish the image of the incumbents who were accused of corruption, and precinct-level votes for incumbents accused of corruption were lowered significantly. As for the theoretical assumption that random assignment does not affect outcomes, it seems there are few backdoor paths that could explain the effect on partisanship: the leaflets were distributed toward the end of the campaign period, preventing incumbents from responding to the messages; the leaflets themselves did not mention political parties; and the post-election surveys did not prime the respondents to think about the leaflets they might have received. The Chong et al. design represents an instructive example of an experimental study that measures the extent to which party identification responds to a theoretically informative, real-world intervention. Although more research of this kind needs to be done before one can draw robust conclusions about party attachments in Mexico or elsewhere, this study seems to suggest that performance-related information regarding corruption has a short-term effect on partisanship, while somewhat more issue-related information concerning spending had negligible effects.

Encouraging Downstream Change in Party Identification: Gerber, Huber, and Washington (2010)

In the context of the hotly contested presidential primaries of 2008, Gerber, Huber, and Washington (2010) conducted an experiment in which they sought to intensify the partisan attachments of self-identified independents. In January of 2008, as the presidential primaries of both parties were intensifying, the authors conducted a survey of registered voters in Connecticut who, when registering, declared themselves unaffiliated with any political party. This declaration rendered them ineligible to vote in the upcoming presidential primaries. Among those who declared themselves to be independents (including those who “lean” toward the

Democrats or Republicans when asked a standard follow-up question about which party they feel closer to), half were randomly selected to receive a letter a week or two later informing them that they must register with a party in order to vote in that party's presidential primary election on February 5th. The letter also included a registration form enabling them to register with a party. In June, respondents were re-interviewed and asked about their party identification, as well as their issue stances and other evaluations.

This experiment parallels the encouragement design described earlier. The pool of experimental subjects comprises self-described independents who were interviewed in January. Random assignment (Z_i) determines which of the subjects is sent a letter. The letter is literally an encouragement to register with a political party. Although the letter might ordinarily be considered the treatment in a standard design, the treatment in the downstream experiment (D_i) is whether the subject actually registers as a Democrat or Republican. (The authors discuss other potential outcomes variables, such as whether subjects vote in the presidential primaries; what follows is a greatly simplified version of their analysis that conveys the basic logic of the design.) Some members of the control group register without encouragement; some members of the treatment group fail to register despite encouragement.

The mismatch between assigned and actual treatment prevents us from estimating the ATE for the sample as a whole; instead, we must set our sights on estimating the ATE for Compliers, those who register with a major party if and only if encouraged. In order to identify the CACE, we must assume monotonicity, or the absence of Defiers. In this case, Defiers are those who would register with one of the two major parties if and only if they are assigned to the control group. Intuition suggests that few voters are so hostile to form letters from afar that they

would cancel their plans to register with a major party if encouraged to do so. Monotonicity appears to be a plausible assumption here.

Under monotonicity (no Defiers), those who register with a major party in the control group are Always-Takers, and those who register in the treatment group are a combination of Always-Takers and Compliers. Since the treatment and control groups are selected randomly, in expectation they should have the same shares of Always-Takers and Compliers. Thus, the share of Compliers can be estimated by subtracting the party registration rate (7.23%) in the control group ($N = 346$) from the party registration rate (13.61%) in the treatment group ($N = 360$). This estimate ($0.1361 - 0.0723 = 0.0639$) forms the denominator of the estimator in equation (8). The t-ratio for this estimated effect is 2.78, which falls a bit short of the recommended threshold for avoiding the weak instruments problem, but using the full sample of subjects (rather than just those reinterviewed in June) leaves no doubt about the robustness of the relationship. For these 2,348 subjects, the t-ratio is 5.48.

The numerator of equation (8) is the observed difference in outcomes, in this case, identification with a major party when re-interviewed several months later. Identification could be measured in various ways; for purposes of illustration, we will use the convention of measuring partisan strength by folding the 7-point party identification scale at the center (pure independent) and counting independent leaners as 1, weak partisans as 2, and strong partisans as 3. Using this scoring method, partisan strength averaged 1.0361 in the treatment group, as compared to 0.9624 in the control group. In other words, assignment to receive a letter boosted the apparent probability of identifying with a party by $1.0361 - 0.9624 = 0.0737$ scale points. Putting the numerator and denominator together gives us the instrumental variables regression estimate of the CACE:

$$\widehat{CACE} = \frac{\hat{E}[Y_i|Z_i=1] - \hat{E}[Y_i|Z_i=0]}{\hat{E}[D_i|Z_i=1] - \hat{E}[D_i|Z_i=0]} = \frac{1.0361 - 0.9624}{0.1361 - 0.0723} = \frac{0.0737}{0.0639} = 1.153 . \quad (9)$$

This estimate suggests that among Compliers, those who register with a party if and only if encouraged to do so, the act of registering with a party increases partisan strength by 1.153 scale points.⁴ The magnitude of this effect is not trivial: in their pre-election round of interviews with registered voters who were not registered with a party (including respondents who were not part of the letter experiment because they were weak or strong partisans), the average level of partisan strength was 1.01 with a standard deviation of 0.85.

Before drawing substantive inferences based on this estimate, let's first evaluate the plausibility of the exclusion restriction in this application. One of the many strengths of the Gerber, Huber, and Washington (2010) design is the lengths to which the authors went to anticipate exclusion-restriction concerns, which the authors discuss in detail (pp.737-741). Clearly, the letters (Z_i) influenced party registration (D_i) and partisan strength (Y_i). The question is whether $Y_i(D_i) = Y_i(D_i, Z_i)$, an assumption that would be violated if potential outcomes (partisan strength) respond not only to whether people register with a party but also to whether they receive a letter. The letters themselves were designed to be empty of partisan content; they simply remind voters of the administrative fact that a change of registration will be necessary if they want to participate in an upcoming election. In terms of measurement procedures, the authors took care to assess outcomes in the June survey in ways that preserved the symmetry between treatment and control groups, avoiding any questions that would prompt members of the treatment group to recall the letter or the circumstances surrounding their change in registration. In terms of substantive confounders, it is possible that the letters piqued voters'

⁴ The authors present evidence to show that this relationship is statistically significant when one controls for pre-election survey covariates. For ease of presentation, I have constructed my dependent variable somewhat differently but find the same basic results.

interest in the campaign, so that even if they did not change their registration, their partisan attachments were altered. This backdoor pathway from Z_i to Y_i seems unlikely, and the authors find no evidence in June that subjects in the treatment group were any more interested or hungry for political information (p.739).

If we accept the exclusion restriction, two issues of interpretation remain. The first is whether one can generalize from the estimated ATE for Compliers to causal effects for other subjects, contexts, and interventions. Would the results be the same if one's treatment caused every person who was registered but unaffiliated with a party to change their party registration? This question is best settled by follow-up experiments that assess whether the results depend on number and frequency of encouragements (which will determine the proportion of Compliers) or the particular arguments that are used in the encouragement. The same goes for experimenting with different contexts: instead of offering voters a chance to vote in both parties' contested primaries, what about circumstances in which only Republican candidates are vying for the nomination?

Another question of interpretation is what to make of the effect of changing registration. A variety of hypotheses could be adduced: a public declaration of a partisan identity changes the way one regards oneself, sets in motion a search for information to justify one's partisan choice, or causes political campaigns to make increased efforts to mobilize and persuade (p.737). Each of these subsidiary hypotheses has testable implications, and the authors investigate whether subjects in the treatment and control group evaluate partisan figures differently or have different types of interactions with political campaigns. They find that partisan evaluations do change concomitantly with changes in party identification (p.735), but there is no apparent relationship between the treatment and contact with campaigns or other manifestations of greater interest in

issues or information. Over the course of a few months, change in partisanship seems to have coincided with changes in partisan attitudes but not changes in behaviors such as searching for information or discussing politics with others.

We say “coincided” because one cannot distinguish the causative effects of partisanship from the effects of all of the other changes that were set in motion by the letter. The authors note that “receipt of the letter informing the recipients about the need to be affiliated with a party in order to vote in that party’s primary increased partisan identity, partisan registration, voter turnout, and partisan evaluations of political figures” (p.737). With just a single randomly assigned treatment (the letter), one cannot separately identify the effects of each intervening variable. For example, one cannot separately identify the effects of registration and the effects of actually voting; voting is just one of the many possible by-products of registration. If one wanted to isolate the effect of registration per se, a different design would be needed – perhaps encourage unaffiliated voters to re-register with a party shortly after the primary has passed in order to estimate the effect of (solely) registering with a party? Conversely, one could determine whether voting per se increases partisanship by urging people to vote using nonpartisan messages (see Green and Gerber 2008). The single-factor encouragement used in this study paves the way for more elaborate encouragement designs that aim to identify distinct sources of partisan change.

Discussion

The two studies summarized above provide a template for advancing our understanding of party identification. The Chong et al. (2011) study is an example of how one might fruitfully study causes of partisan change by deploying an array of different kinds of interventions. In that

study, information about corruption in municipal government cause voters to subsequently change their party attachments, while information about public spending did not. The Gerber et al. (2010) study is an example of deploying a treatment that in itself has no partisan content and functions solely to facilitate behaviors that are believed to reinforce partisanship. This style of intervention-oriented research could be expanded to include information about the parties' policy stances, their financial backers, their level of support among different segments of the electorate, and so forth. A combination of treatments could be designed to test competing theories about how party identities are formed. One kind of treatment might be designed to affect retrospective performance evaluations, while another might be crafted to alter perceptions of the parties' platforms or support among voters with different social identities. What makes this approach distinctive is that it attempts to mint partisans through random interventions rather than to observe passively the partisan changes that occur on their own.

Both experiments illustrate how this approach might be deployed in a field setting (perhaps as a by-product of a broader field experiment), but the basic design applies also to laboratory research (Cowden and McDermott 2000) or lab-like survey research (Middleton 2011). One could imagine a lab or on-line study in which subjects are pre-screened for weak partisan attachments, randomly exposed to theoretically-inspired appeals that are designed to move them closer to a political party. For example, one could imagine a "social identity" video that explains what sorts of people favor the Democratic and Republican parties and a competing "spatial proximity" video that explains the ideological stances of the party with respect to several leading issues. Indeed, one can even imagine a vacuous "feel-good" video that deploys a slogans and attractive imagery while endorsing one of the parties – in this case, the same video could be adapted to support each party. The main practical constraints are the need to expose the control

group to something that is vaguely similar (but not party-focused) so that subjects in both groups have similar suspicions about what the study is about given the need to reinterview subjects at some later point in time to see whether enduring changes in partisanship took root.

More challenging is the task of designing experiments to test the effects of partisan attachments on other attitudes and behavior. For example, partisanship is said to alter issue stances, economic evaluations, and interest in political news. In an ideal design, a randomly assigned intervention would affect party attachments without directly affecting these outcomes.⁵ This exclusion restriction obviously rules out the use of economic news as an inducement to identify with the allegedly more competent party. Developing effective interventions that seem to satisfy the exclusion restriction may require a fair amount of trial-and-error. Social scientists are as yet unaccustomed to developing interventions that successfully change partisanship; the two experiments discussed above are important first steps in that direction.

⁵ Satisfying this exclusion restriction is a drawback to using what would otherwise be attractive natural experiments, such as the Vietnam Draft Lottery. See Erikson and Stoker (2011).

References

- Achen, Christopher H. 1975. Mass Political Attitudes and the Survey Response. *American Political Science Review* 69: 1218-31.
- Achen, Christopher H. 2002. Parental Socialization and Rational Party Identification. *Political Behavior* 24: 141-170.
- Angrist, Joshua D., Guido W. Imbens, and Donald D. Rubin. 1996. Identification of causal effects using instrumental variables. *Journal of the American Statistical Association* 91:444–55.
- Belknap, George, and Angus Campbell. 1952. Political Party Identification and Attitudes toward Foreign Policy. *Public Opinion Quarterly* 15: 601-623.
- Box-Steffensmeier, Janet M., and Renée M. Smith. 1996. The Dynamics of Aggregate Partisanship. *American Political Science Review* 90: 567-80.
- Brody, Richard A., and Lawrence S. Rothenberg. 1988. The Instability of Partisanship: An Analysis of the 1980 Presidential Election. *British Journal of Political Science* 18: 445-65.
- Butler, David and Donald Stokes. 1969. *Political Change in Britain*. London: Macmillan.
- Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes, 1960. *The American Voter*. New York: Wiley.
- Campbell, Angus, Philip E. Converse, Warren E. Miller, and Donald E. Stokes. 1966. *Elections and the Political Order*. New York: Wiley.
- Chong, Alberto, Ana De La O, Dean Karlan, and Leonard Wantchekon. 2011. Information Dissemination and Local Governments' Electoral Returns: Evidence from a Field Experiment in Mexico. Unpublished manuscript.

- Converse, Philip E. and Roy Pierce. 1992. Partisanship and the Party System. *Political Behavior* 14: 239-259.
- Cowden, Jonathan A., and Rose M. McDermott. 2000. Short-term Forces and Partisanship. *Political Behavior* 22: 197–222.
- Erikson, Robert S., Michael B. MacKuen, and James A. Stimson. 1998. What Moves Macropartisanship? A Response to Green, Palmquist, and Schickler. *American Political Science Review*. 92: 901-12.
- Erikson, Robert S., and Laura Stoker. 2011. Caught in the Draft: The Effects of Vietnam Lottery Status on Political Attitudes. *American Political Science Review* 105: 221-237.
- Fiorina, Morris P. 1981. *Retrospective Voting in American National Elections*. New Haven, CT: Yale University Press.
- Franklin, Charles H., and John E. Jackson. 1983. The Dynamics of Party Identification. *American Political Science Review* 77:957-73.
- Gerber, Alan S., and Donald P. Green. 2012. *Field Experiments: Design, Analysis, and Interpretation*. New York: W.W. Norton.
- Gerber, Alan S., Gregory A. Huber, and Ebonya Washington. 2010. Party Affiliation, Partisanship, and Political Beliefs: A Field Experiment. *American Political Science Review* 104: 720-744.
- Green, Donald P., and Alan S. Gerber. 2002. The Downstream Benefits of Experimentation. *Political Analysis* 10(4): 394-402.

- Green, Donald P., and Alan S. Gerber. 2008. *Get Out The Vote: How to Increase Voter Turnout*, Second Edition. Washington, D.C.: Brookings Institution Press.
- Green, Donald P., and Bradley Palmquist. 1990. Of Artifacts and Partisan Instability. *American Journal of Political Science* 34: 872-902.
- Green, Donald P., and Bradley Palmquist. 1994. How Stable is Party Identification? *Political Behavior*. 16: 437-466.
- Green, Donald P., Bradley Palmquist, and Eric Schickler. 1998. Macropartisanship: A Replication and Critique. *American Political Science Review*. 92: 883-99.
- Jackson, John E. 1975. Issues Party Choices, and Presidential Votes. *American Journal of Political Science* 19: 161-185.
- Jennings, Kent M. and Richard G. Niemi. 1981. *Generations and Politics: a panel study of young adults and their parents*. Princeton, NJ: Princeton University Press.
- MacKuen, Michael B., Robert S. Erikson, and James A. Stimson. 1989. Macropartisanship. *American Political Science Review* 83: 1125-42.
- Middleton, Joel. 2011. *On the Foundations of Economic Voting: An Experimental Study*. Unpublished doctoral dissertation, Department of Political Science, Yale University.
- Neyman, Jerzy. 1990. On the Application of Probability Theory to Agricultural Experiments. Essay on Principles. Section 9. *Statistical Science* 5:465-480.
- Page, Benjamin I., and Calvin C. Jones. 1979. Reciprocal Effects of Policy Preferences, Party Loyalties, and the Vote. *American Political Science Review* 73:1071-89.

Rubin, Donald. 1990. Formal Mode of Statistical Inference for Causal Effects. *Journal of Statistical Planning and Inference* 25:279-92.

Schickler, Eric, and Donald P. Green. 1997. The Stability of Party Identification in Western Democracies: Results from Eight Panel Surveys. *Comparative Political Studies* 30: 450-83.

Staiger, Douglas, and James H. Stock. 1997. Instrumental variables regression with weak instruments. *Econometrica* 65:557–586.