

Redistricting and the Personal Vote: When Natural Experiments Are Neither Natural Nor Experiments*

Jasjeet S. Sekhon[†] and Rocío Titiunik[‡]
Associate Professor Postdoctoral Fellow
Travers Dept. of Political Science Dept. of Political Science
UC Berkeley University of Michigan

12/13/2009 (19:26)

*A previous version of this paper was circulated under the title “Exploiting Tom DeLay: A New Method for Estimating Incumbency Advantage and the Effect of Candidate Ethnicity on Turnout” and received the 2008 Robert H. Durr Award for best paper applying quantitative methods to a substantive problem during the 66th Midwest Political Science Association Annual National Conference, April 3-6, 2008. For valuable comments we thank Steve Ansolabehere, Jake Bowers, Daniel Enemark, Bob Erikson, Shigeo Hirano, Luke Keele, Gary King, Walter Mebane, Jr., Rebecca Morton, Eric Schickler, Jonathan Wand, and participants of the Society for Political Methodology’s Annual Summer Meeting 2007, CELS 2009 and seminar participants at Berkeley, Columbia, Michigan, NYU, Princeton, Yale and UCSD. We thank Sam Davenport in the Texas Legislative Council, Nicole Boyle in the California Statewide Database, Gary Jacobson and Jonathan Wand for providing data. Sekhon also thanks the MIT Department of Political Science for hospitality during the summers of 2007 and 2008. All errors are our responsibility.

[†]<sekhon@berkeley.edu>, <http://sekhon.berkeley.edu/>, Survey Research Center, 2538 Channing Way, UC Berkeley, Berkeley, CA 94720

[‡]<titiunik@umich.edu>, <http://www-personal.umich.edu/~titiunik>, Center for Political Studies, ISR, University of Michigan, P.O. Box 1248, Ann Arbor, MI 48106-1248

Abstract

Although natural experiments are increasingly prominent in the social sciences, they often have more in common with traditional observational studies than with randomized experiments. We illustrate our argument by examining the use of redistricting to estimate the personal vote. Strikingly, even if voters were redistricted randomly, previous uses of redistricting would not identify the causal effect of interest. We also find that the redistricting process is sufficiently nonrandom as to require significant covariate adjustment to overcome confounding. To avoid these difficulties, we propose a new design for estimating the personal vote and the partisan incumbency advantage that relies on the implementation of multiple redistricting plans. Analyzing data from U.S. House elections in California and Texas, we find that there is a large partisan incumbency advantage in both states but that the effect of the personal vote is zero in Texas and small in California.

Keywords: Incumbency, Redistricting, Natural Experiments, Potential Outcomes

JEL Classifications: P16, D72

1 Introduction

The number of political scientists who use natural experiments has increased considerably in recent years (Dunning 2008). Such experiments have been used by scholars in a wide variety of fields, including political participation (Brady and McNulty 2004; Huber and Arceneaux 2007; Krasno and Green 2008; Lassen 2005), elections and electoral competitiveness (Carman, Mitchell, and Johns 2008; Gordon and Huber 2007), political psychology (van der Brug 2001; Buetler and Marechal 2007), ethnic politics (Abrajano, Nagler, and Alvarez 2005), comparative politics (Posner 2004; McCauley and Posner 2007), and bureaucracy (Whitford 2002).

Given the inherent difficulties of making causal inferences in the social sciences, it is unsurprising that natural experiments have become popular. Social scientists often wish to estimate the causal effects of policies, institutions, and individual actions that they do not control. Experimental manipulations are rare. And in observational studies, the behavior of political actors plausibly depends on unobserved factors.

Natural experiments offer a path to overcome these difficulties by possibly providing researchers with an *exogenous* manipulation of a cause. The condition of exogeneity is crucial: a valid natural experiment generates variation in a cause that is independent of all other competing factors that may affect the outcome of interest. If this condition is not satisfied, the hoped for natural experiment has less in common with an actual experiment than with an observational study. For example, we might hypothesize that incumbent retirement can be used as a natural experiment to estimate the incumbency advantage, but this approach will fail if incumbents who have reason to expect that they will perform poorly in the following election are more likely to retire.¹

Researchers generally recognize that exogeneity is a crucial condition for their natural experiments to be valid and often spend substantial time justifying the exogeneity of the “naturally occurring” variation they exploit. But less attention is usually given to another crucial feature of natural experiments, namely, the extent to which the manipulation, even if exogenous, is the manipulation

¹A more promising strategy is to use unforeseeable retirements (such as those caused by death), since these are arguably independent of factors affecting expected electoral performance—see Cox and Katz (2002) for a discussion of this strategy.

needed to identify the causal effect *of interest* as opposed to some *other* causal effect. Although this issue is critical in any research design, it is more pressing in a natural experiment where, by construction, the manipulation is outside the control of the researcher and is therefore typically only indirectly related to the effect of interest. Thus, although arbitrary events or interventions may be appealing because they are a source of exogenous variation, they can force the researcher to redefine the causal effect of interest, often without the researcher noticing or acknowledging that this redefinition has taken place. To state it another way: even if one grants that the exogeneity assumption holds, does the natural experiment identify the causal effect of interest?

In this paper, we reanalyze the use of redistricting as a natural experiment to identify the portion of the incumbency advantage that stems from an incumbent's personal appeal with her constituents, a research design first proposed by Ansolabehere, Snyder, and Stewart (2000).² Even frequent critics of observational studies such as Green and Gerber (2008) have noted that this is a well designed natural experiment. The design exploits the fact that after redistricting most incumbents face districts that contain a combination of old and new territory, and hence face a combination of old and new voters. The authors analyzed U.S. House elections from 1872 to 1990 and compared an incumbent's vote share in the new part of the district with her vote share in the old part of the district, finding an average personal vote of about 4 percent.³

Desposato and Petrocik (2003) employ the Ansolabehere, Snyder, and Stewart design to estimate the personal vote in California using block-level data, and they obtain similar average estimates of the personal vote for members of Congress.⁴ Carson, Engstrom, and Roberts (2007) also use the design to estimate the personal vote in late-nineteenth-century House elections (1872–1900). They estimate the personal vote to be about 2.5% and find that nearly all of the incumbency advantage during this time-period can be attributed to the personal vote.

²The authors sometimes refer to this electoral advantage as the “personal vote” and other times as the benefits of “homestyle” or as “direct office holder benefits.”

³This figure corresponds to the 1972 – 1988 period. Consistent with previous literature, their estimate of the personal vote is smaller for earlier periods.

⁴Desposato and Petrocik (2003) allow for heterogeneous treatment effects. They find that in the average neighborhood (with 50% Democratic registration) Republican incumbents lost about 5% with new voters and Democratic incumbents lost about 4%.

The proposal by Ansolabehere, Snyder, and Stewart to use a natural experiment to identify the personal vote adds to the broader literature on whether the incumbency status of legislators in the United States affects their electoral outcomes, one of the most studied topics in electoral politics.⁵ Despite the vast amount of work, the literature faces formidable methodological challenges. Erikson (1971) was the first to recognize that traditional measures of incumbency advantage such as sophomore surge and retirement slump could be severely biased. Gelman and King (1990) provided a formal analysis of these problems and developed an alternative estimator. Cox and Katz (2002) showed that most extant estimators of the incumbency advantage assume that there is no strategic exit of incumbents and no strategic entry of challengers. Zaller (1998) demonstrated that high reelection rates for incumbents may be explained simply by survival bias.

The innovative design proposed by Ansolabehere, Snyder, and Stewart is an effort to surmount some of these methodological difficulties. Unfortunately, although using redistricting as a natural experiment to identify the personal incumbency advantage has some appealing properties, we show that the design fails to recognize crucial difficulties in both the exogeneity of the intervention and the appropriateness of its manipulation to identify the causal effect of interest, resulting in invalid inferences. First, we show that redistricting is not an exogenous manipulation but (at most) a conditionally exogenous one, that is, redistricting is exogenous only after conditioning on appropriate covariates. The conditioning set implicit in prior work is theoretically implausible because it fails to include crucial covariates, and it empirically fails a placebo test which our design passes. Second, we demonstrate that even if the assumption of exogeneity is relaxed to conditional exogeneity and a valid conditioning set is used, previous uses of redistricting do not identify a causal effect because the wrong potential outcomes are used. This leads to the striking result that even if voters were assumed to be redistricted at random, the design would result in biased estimates.

To overcome these difficulties, we propose a new method for estimating a portion of the incumbency advantage using a natural experiment which relies upon the successive implementation

⁵Studies of the effects of incumbency for legislative offices include, among many others, Alford and Brady (1989), Ansolabehere, Brady, and Fiorina (1988), Ansolabehere and Snyder (2002), Breaux (1990), Born (1979), Cox and Katz (1996), Cox and Morgenstern (1993), Ferejohn (1977), Fiorina (1977), Jacobson (1987), Krashinsky and Milne (1993), Krehbiel and Wright (1983), Mayhew (1974), Nelson (1979), Payne (1980), Rush (1993).

of multiple redistricting plans. Our research design uses the correct potential outcomes and makes a selection on observables assumption which passes a crucial placebo test which previous designs fail. We estimate it with data from congressional elections in Texas, where two different redistricting plans were successively implemented in 2002 and 2004. When multiple redistricting plans are not available, we propose a “second-best” design, which is similar in spirit to the original design proposed by Ansolabehere, Snyder, and Stewart (2000). We estimate this second design using data from congressional elections in California and Texas. In all cases, we estimate the effects of incumbency using Genetic Matching (Diamond and Sekhon 2005; Sekhon In Press) to achieve covariate balance.

Even after correcting the methodological problems in the original design, using redistricting as an identification strategy for estimating the personal vote still leads to some conceptual ambiguities. For example, even when new voters do not have the same history with the incumbent as old voters, they may nonetheless respond to the incumbency cue. Voters know who the incumbent is, separate from anything that the incumbent does such as provide constituency service and pork, take public positions, etc. As Mayhew (1974, 313) notes: “It is possible that incumbents have been profiting not from any exertions of their own but from changes in voter attitudes. . . . Voters dissatisfied with party cues could be reaching for any other cues that are available in deciding how to vote. The incumbency cue is readily at hand.”

The rest of the paper is organized as follows. In the next section we discuss previous uses of redistricting as a natural experiment, and propose our research design. We focus on the exogeneity of the intervention in Subsection 2.1 and on the identification of the causal effect of interest in Subsection 2.2. In Section 3 we discuss the substantive and theoretical implications of rethinking the design. Section 4 describes the data used in our empirical application. We present the empirical results in Section 5, and conclude in Section 6.

2 Research Design

In this section, we examine in detail the conditions that must hold for the variation introduced by congressional redistricting to identify (a portion of) the causal effect of incumbency. Redistricting induces variation in at least two dimensions: a time dimension, as voters vote both before and after redistricting, and a cross-sectional dimension, as some voters are moved to a different district while others stay in the district they originally belonged to. We are interested in learning about the incumbency advantage by comparing the behavior of voters who are moved to a new district (*new voters*) to the behavior of voters whose district remains unchanged across elections (*old voters*).

As first recognized by Ansolabehere, Snyder, and Stewart (2000), the attractiveness of redistricting as a research design relies upon the fact that old voters and new voters face the same electoral environment in the elections following redistricting. Indeed, the design was not originally intended to capture the overall incumbency advantage, since important sources of this advantage such as challenger quality and the incumbency cue are held constant in both parts of the district. The comparison between old and new voters was therefore intended to isolate the fraction of the overall incumbency advantage that stems solely from the incumbent’s homestyle, although, as we will see in Section 3, such a clear “personal vote” interpretation may not be possible.

Interpretation issues aside, in this section we show that while using redistricting as an empirical strategy to identify a portion of the incumbency advantage is promising, its correct implementation requires a more careful consideration than previously thought. First, the redrawing of district boundaries cannot be considered exogenous unless one conditions on crucial covariates that have been ignored in previous designs. Second, even if redistricting were randomly assigned, previous designs would yield *biased* estimates.

2.1 Redistricting as an exogenous intervention

We now show that redistricting does not satisfy the exogeneity condition that is necessary for unbiased estimation of the incumbency advantage, because voters who are redistricted vote for

their incumbent at lower rates *even before redistricting occurs*. Moreover, this bias remains even after conditioning on presidential vote, the single covariate included in previous designs.

Figure 1 illustrates that the naive comparison of voters who are moved to a new district with voters who remain in their old district is likely to find a *spurious* personal vote. This figure shows the empirical Quantile-Quantile (QQ) plots of the baseline vote share received by the incumbent U.S. House member in the election before redistricting, comparing units⁶ that were to be redistricted to a different incumbent in the following election to units that were to remain with the same incumbent after redistricting. For Texas, we use 2002 as the baseline year, and compare 2002 incumbent vote shares between units that will be moved in the 2004 redistricting and units that will remain with their old incumbent after the 2004 redistricting.⁷ For California, the baseline year is 2000, and we compare 2000 incumbent vote shares between units that will be moved in the 2002 redistricting and units that will remain with their old incumbent after the 2002 redistricting.⁸ Figure 1(a) shows the QQ plot for California, while Figure 1(b) shows the QQ plot for Texas. In both states, the empirical quantiles of the baseline incumbent vote share for the units which are to remain with the same incumbent after redistricting are everywhere larger than the empirical quantiles of the units whose incumbent is to change after redistricting, indicating that those units to be redistricted vote for their *old* incumbent at a systematically *lower* rate than units whose incumbent will not change. In other words, there is a bias in the selection of units to be moved, with units with a *lower* incumbent vote share in the election before redistricting being *more* likely to be moved to a different incumbent when redistricting is implemented.

Moreover, this bias remains even after both types of voters are matched on their partisan attachments as measured by presidential vote shares. Figures 1(c) and 1(d) show the same QQ plots than figures 1(a) and 1(b), respectively, but this time the QQ plots are produced after units whose

⁶The unit of analysis is the Voting Tabulation District for Texas, and the 2000 census block for California. Details are provided in Section 4.

⁷For Texas, using 2000 as the baseline instead of 2002 does not change the direction of the results, although the differences become smaller. This is to be expected given the less partisan nature of the 2000 redistricting when compared to the 2002 redistricting.

⁸Note that using 2000 as the baseline year for California is the only option, as this state, unlike Texas, implements only one redistricting plan (in 2002) after the 2000 census.

incumbent does not change are matched to units whose incumbent does change on their Democratic share of the two-party presidential vote, which is the measure of “normal vote” used by Ansolabehere, Snyder, and Stewart (2000). As can be seen, restricting the comparison to units whose normal vote is on average identical does not eliminate the bias. Even when new voters vote for the Democratic presidential candidate at the same rate as old voters, they still vote for their old incumbent at a lower rate. But note that the outcome examined is the incumbent vote share *before* redistricting occurs, and so the effect is zero by construction.

If at least part of this tendency of new voters to vote for their incumbent at a lower rate persists in the future, comparing old voters and new voters will be biased towards finding a positive personal vote even when there is none. The plots in Figure 1 suggest that if the bias is severe enough, previous uses of redistricting could be estimating a positive personal vote even if the true effect were zero. The evidence presented in Section 5.1 shows that the bias is indeed severe.

2.2 Redistricting as an identification assumption for the effect of interest

We now show that the challenges involved in using redistricting as a natural experiment are not limited to issues of exogeneity. By analyzing the redistricting manipulation, we demonstrate that even if voters were assumed to be redistricted randomly, the manner in which previous work has used redistricting to identify the personal vote results in biased estimates because the wrong counterfactuals (i.e., potential outcomes) are used.

We divide our argument in two parts. First, we show that although new voters are naturally defined as the voters whose district changes between one election and another, there is an ambiguity in the way in which old voters are defined, as these could be either the electorate of the district to which new voters are moved (henceforth *new neighbors*), or the electorate of the district to which new voters belonged before redistricting occurred (henceforth *old neighbors*). Second, independently of how old voters are defined, we show that only under strong assumptions does the difference in the behavior of old voters and new voters identify the effect of incumbency. We propose different research designs to address these issues.

We discuss both arguments by means of the following thought experiment, which we illustrate in Figures 2(a) and 2(b). We imagine that just before election t a redistricting plan randomly redraws the boundaries of an arbitrary district (referred to as district A), in such a way that some voters that used to be in this district are randomly chosen and moved to a new district (referred to as district B). From the point of view of district B 's incumbent, in the first election after redistricting (referred to as election t) voters that come from district A are *new voters* and voters that were originally in B are *old voters*. In principle, it seems natural to compare how differently these two groups vote for the incumbent and attribute the difference to a personal incumbency advantage, since both types of voters face not only the same incumbent but also the same challenger, the same campaign, etc.⁹

Moreover, the assumption of random redistricting seems to make this comparison even more attractive. Since randomization (if successful) ensures exchangeability between treatment and control units, we may be tempted to claim that in this hypothetical case B 's old voters are guaranteed to be valid counterfactuals for B 's new voters. But a crucial feature of this experiment prevents this claim from being true: while this randomization guarantees that voters who stay in A (A 's old voters) and voters who leave A (B 's new voters) are exchangeable, randomization says nothing about the exchangeability of B 's new voters and B 's old voters. In the absence of redistricting, B 's new voters would have been in a different district than B 's old voters and therefore nothing ensures that B 's old voters are a good counterfactual for what would have happened to the new voters in the absence of redistricting, precisely because in the absence of redistricting both groups of voters would not have been in the same district at all.

In other words, the fact that B 's new voters are originally in a different district than B 's old voters implies that both types of voters have different histories—that is, at election $t - 1$, B 's old and new voters may have faced incumbents who belonged to different parties, or candidates who were of different qualities, or campaigns that were managed in different ways, etc. Since these factors

⁹Although this thought experiment places constraints on which precincts may move, the results are general. That is, the conclusions are the same if we assume that every precinct in every district in the state has a positive probability of moving to any other district. However, the notation and discussion become unwieldy.

are likely to affect how new voters react to their new incumbent, one needs a design that balances these covariates between treated and control groups. The crucial point is that the randomization we are considering does not guarantee balance in the covariates related to the history of new voters and new neighbors and hence, without further assumptions, it is not appropriate to estimate the effects of incumbency.¹⁰

Formally, let T_i be equal to 1 if precinct i is moved from one district to another just before election t and equal to 0 if precinct i is not moved to a different district before election t , and let D_i be equal to 1 if precinct i has new voters in its district at election t and equal to 0 if precinct i has no new voters in its district at election t . Let $Y_0(i, t)$ be the outcome attained by precinct i if $T_i = 0$ and $D_i = 0$ (the precinct is not moved and does not have new neighbors, i.e., these are voters who stay in A after redistricting), let $Y_1(i, t)$ be the outcome attained by precinct i if $T_i = 0$ and $D_i = 1$ (the precinct is not moved and has new neighbors, i.e., these are voters who are in B before and after redistricting), and let $Y_2(i, t)$ be the outcome attained by precinct i if $T_i = 1$ and $D_i = 1$ (the precinct is moved and has new neighbors, i.e., these are voters who are moved from A to B).¹¹ Of course, the fundamental problem of causal inference is that for every precinct we observe only one of its three potential outcomes. This is, we only observe the *realized* outcome, defined as

$$Y(i, t) = Y_0(i, t) \cdot (1 - T_i) \cdot (1 - D_i) + Y_1(i, t) \cdot (1 - T_i) \cdot D_i + Y_2(i, t) \cdot T_i \cdot D_i \quad (1)$$

As is common with observational studies, we focus on the average treatment effect on the treated (ATT). Given the set-up of our hypothetical experiment, the ATT can be defined in two

¹⁰Moreover, unless we assume that the distribution of voter level characteristics such as race and partisanship between district A and B are the same at $t - 1$, randomization does not ensure that these individual level characteristics will be balanced between *old* and *new* voters.

¹¹The potential outcome when $T_i = 1$ and $D_i = 0$ is not defined because it is not possible to be moved from one district to another and not to have new neighbors.

different ways:

$$ATT_0 \equiv E [Y_2(i, t) - Y_0(i, t) | T_i = 1, D_i = 1] \quad (2)$$

$$ATT_1 \equiv E [Y_2(i, t) - Y_1(i, t) | T_i = 1, D_i = 1] \quad (3)$$

It can be shown that the following condition is sufficient for ATT_0 to be identified:¹²

$$E [Y_0(i, t) | T_i = 1, D_i = 1] = E [Y_0(i, t) | T_i = 0, D_i = 0] \quad (4)$$

Similarly, it can be shown that the following condition identifies ATT_1 :

$$E [Y_1(i, t) | T_i = 1, D_i = 1] = E [Y_1(i, t) | T_i = 0, D_i = 1] \quad (5)$$

In words, Assumption (4) says that voters who stay in A and voters who are moved from A to B would have attained the same average outcomes if they hadn't been moved and if they had not received new neighbors in their districts. Assumption (5), on the other hand, states that voters who are originally in B and voters who are moved from A to B would have attained the same average outcomes if A 's voters would not have been moved and B 's voters would not have received new neighbors.

This makes clear that randomization does not imply that B 's old voters are a valid counterfactual for B 's new voters: while randomization, if successful, ensures that Assumption (4) is satisfied (and hence that the average treatment effect defined by Equation (2) is identified), randomization does *not* imply Assumption (5). In other words, randomization ensures exchangeability between the set of voters for which $(1 - T_i) \cdot (1 - D_i) = 1$ (i.e., voters who stay in A after redistricting) and the set of voters for which $T_i \cdot D_i = 1$ (i.e., voters who are redistricted from A to B), but not between the latter set of voters and the set of voters for which $(1 - T_i) \cdot D_i = 1$ (i.e., voters who are originally in B).

¹²For a formal treatment of these and related assumptions, see, for example, Heckman, Ichimura, and Todd (1998).

Indeed, a close examination of Assumption (5) reveals that it is a rather peculiar requirement since, as shown in Figure 2(a), in the absence of redistricting voters in A would have been in a different district than voters in B . The assumption that they would have attained the same average outcomes is a very strong one precisely because in the absence of redistricting these voters would have been in completely different districts facing different incumbents and challengers.

Henceforth, we will refer to the design that uses new neighbors as counterfactuals as the “second-best one-time redistricting design”.

2.2.1 Making the most of old and new neighbors

We have shown that under randomization the group guaranteed to be a valid counterfactual for the new voters is not the new neighbors (i.e., the electorate in the new district to which new voters are moved), but rather the old neighbors (i.e., the voters that are left behind in the new voters’ original district). However, the question arises of whether the one-time redistricting design that uses old-neighbors as counterfactuals is appropriate to estimate the personal vote. On the one hand, using old neighbors ensures that both new and old voters are from the same district and hence from the same population at baseline. But this design also introduces important sources of heterogeneity, since it compares voters who at election $t - 1$ are in the same district (and hence face the same electoral environment) but who at election t are in different districts (and hence face a different incumbent, a different challenger, a different campaign, etc.). In principle, one could restrict the universe of the comparison to reduce this heterogeneity (for example, one could restrict the old and new district to have the same incumbent’s party and the same challenger’s quality). However, there is a crucial difficulty in adopting this approach, as in order to induce homogeneity one would have to condition on characteristics of the environment *after* redistricting, and since these characteristics are likely to have been affected by redistricting itself one runs the risk of introducing post-treatment bias. We therefore conclude that the one-time redistricting design that uses old-neighbors as counterfactuals is not appropriate to estimate the effect of incumbency status

on electoral outcomes.¹³

We have yet to establish a design that is both valid and appropriate for estimating the personal vote. In the next subsection we propose what we consider to be the first-best design to estimate the incumbency advantage using redistricting. But before turning to this design, we consider additional methodological issues that arise if one decides to implement the second-best one-time redistricting design despite its difficulties.

Since Assumption (5) is not valid even with random assignment, we define a weaker version of this assumption:

$$E [Y_1(i, t) | T_i = 1, D_i = 1, X] = E [Y_1(i, t) | T_i = 0, D_i = 1, X], \quad (6)$$

where X is a vector of observable characteristics. Assumption (6) can be shown to identify ATT_1 conditional on X and is considerably weaker than Assumption (5). Thus, if one were still interested in using B 's original voters as counterfactuals despite the methodological difficulties, one could attempt to find the subpopulation of B 's old voters who are most similar to the new voters on some set X of observable characteristics and use these as counterfactuals, under the assumption that once the joint distribution of X is equated among new voters and new neighbors, their average potential outcomes would have been identical in the absence of redistricting. But note that Assumption 6 defines a selection on observables assumption which is not guaranteed to hold even under random assignment!

To complicate things further, if Assumption 6 were true this approach may still not result in unbiased estimates, because the distribution of X between B 's old and new voters is not guaranteed to be equal *even if conditional on X both groups of voters would have attained the same average outcomes in the absence of redistricting.*¹⁴ The reason is that the support of the distribution of X among B 's new voters may be different from the support of the distribution of X among B 's

¹³Note, however, that this design could be used to estimate how voters react to a change in the race or ethnicity of their incumbent, since in this case one wishes to consider the different electoral environments which incumbents of different races or ethnicities bring about.

¹⁴See Heckman, Ichimura, Smith, and Todd (1998) for a formal proof that the lack of common support introduces bias.

old voters, a concern that becomes all the more relevant given that B 's old and new voters were originally in different districts. In sum, the fact that new voters and new neighbors are never in the same population at baseline may imply that both groups are different by construction, and hence that unbiased estimates may not be achieved even if a strong identifying condition is assumed to hold.

The second-best one-time redistricting design introduces a lack of common support by construction on covariates that are related to the previous history in the district. For example, new voters may have been moved from a Hispanic to a white incumbent, from a Democratic to a Republican incumbent, from a female to a male incumbent, or from a moderate to an extreme incumbent, while new neighbors by definition would face no such variation in the characteristics of their incumbent (assuming the incumbent runs in both elections). Since different previous histories will likely affect voters' behavior differently, having balance on these history-related covariates is crucial to identify the causal effect of interest. Hence, the second-best design must be modified so that balance on these covariates is achieved.

2.2.2 Consecutive redistricting: the first-best design

We propose a different design. We consider a modification of the thought experiment introduced above, and imagine that after some voters are randomly moved from district A to B (and after election t takes place), another random redistricting plan is implemented right before election $t + 1$ so that some voters who were in district A until after election t are randomly chosen and moved to district B . At $t + 1$, there are three types of voters in district B : voters who always belonged to B , voters who became part of district B just before election t (henceforth *early new voters*), and voters who became part of district B just before election $t + 1$ (henceforth *late new voters*). The design is illustrated in Figures 2(c) and 2(d). In this case, the most natural way to estimate the causal effect of incumbency is to compare early new voters to late new voters, as not only do they both face the same electoral environment at election $t + 1$, but they also have the same electoral environment up to election $t - 1$, which implies that their histories are the same except for the fact

that early new voters are moved to the new district one election earlier than late new voters. We call this the “first-best two-time redistricting design”, as it is free from the complications that arise in the two alternatives considered above.

To formally establish the parameter identified by the first-best two-time redistricting design, let $W_{i,t+1}$ be equal to one if precinct i is moved from district A to district B at election $t + 1$, and $W_{i,t+1}$ be equal to zero if precinct i is moved from A to B at election t and remains in B at election $t + 1$. In other words, $W_{i,t+1}$ is a late new voter treatment indicator, where late new voter is defined as voting in B for the first time at election $t + 1$. Letting $Y_0(i, t + 1)$ denote the outcome of i at election $t + 1$ if $W_{i,t+1} = 0$ and $Y_1(i, t + 1)$ denote the outcome of i at election $t + 1$ if $W_{i,t+1} = 1$, we define the parameter of interest, ATT_2 , as

$$ATT_2 \equiv E [Y_1(i, t + 1) - Y_0(i, t + 1) | W_{i,t+1} = 1] \quad (7)$$

which is identified under

$$E [Y_0(i, t + 1) | W_{i,t+1} = 1] = E [Y_0(i, t + 1) | W_{i,t+1} = 0] \quad (8)$$

In words, ATT_2 is identified if late new voters and early new voters would have attained the same average outcomes if they both had been in the new district for exactly two elections. Below, we will show that randomization under this design together with an assumption of stationarity guarantees that Assumption (8) holds.

Since we assumed that both groups of voters are in the same district at election $t - 1$, and that just before election t the set of voters for which $W_{i,t+1} = 0$ is randomly chosen and moved to district B , we have

$$E [Y_0(i, t - 1) | W_{i,t+1} = 1] = E [Y_0(i, t - 1) | W_{i,t+1} = 0] \quad (9)$$

Assumption (9), implied by randomization, guarantees that both groups of voters have the same pre-treatment average outcomes at $t - 1$. But Assumption (9) does not imply Assumption

(8), hence we need to add an assumption to the first-best two-time redistricting design in order to obtain exchangeability at election $t + 1$. We make the following additional assumption:

$$E [Y_0(i, t + 1) - Y_0(i, t - 1) | W_{i,t+1} = 1] = E [Y_0(i, t + 1) - Y_0(i, t - 1) | W_{i,t+1} = 0] \quad (10)$$

Assumptions (9) and (10) together imply Assumption (8). If late new voters are randomly chosen *and* early new voters and late new voters would have followed the same path between election $t - 1$ and election $t + 1$ if they both had spent election t and election $t + 1$ in the new district, ATT_2 is identified. As before, since in practice district boundaries are not randomly modified, in order to achieve identification of the parameters of interest in this design we must make the assumption that, *conditional on certain observable characteristics*, late new voters are exchangeable with early new voters. This is undoubtedly a strong assumption, but is plausible considering that we use the same data which participants in the redistricting battles fed into their computer programs to design their various redistricting plans.

The first-best two-time redistricting design proposed here allows us to implement a crucial placebo experiment to test the validity of the identification strategy, because we observe precincts which will be redistricted before they are redistricted. The placebo test examines precincts which will be redistricted (or not) at election $t + 1$ but which are in the same district in elections t , $t - 1$, $t - 2$, etc. Calling those to be redistricted in election $t + 1$ “treated” and those who will not be redistricted in election $t + 1$ “controls”, we can arbitrarily denote $t - 1$ to be the baseline year, and our placebo test is that in t there should be no significant difference between the outcomes of our treated and control groups.

Note that in the placebo test treatments and controls are always in the same district. As such, this test can be used to validate the first-best design but not the second-best design, because in the former design treatment and control precincts are in the same district before redistricting, while in the latter precincts are in the same district after redistricting but in a *different* district *before* redistricting.

Before presenting the estimated effects for California and Texas using the different designs described above, we discuss the interpretation of the redistricting estimand.

3 Personal Vote, Cues, or Learning?

Ansolabehere, Snyder, and Stewart (2000) proposed to use redistricting as an identification strategy for the *personal vote*, this is, for the electoral advantage that an incumbent acquires by providing casework, bringing federal resources to the district, taking positions that match voters' tastes, etc. The authors distinguished the personal incumbency advantage from the incumbency advantage that stems from candidate quality and from incumbency as a cue, and claimed that redistricting could only identify the personal vote but not the last two sources of incumbency advantage because old and new voters face both the same candidates and the same cues. However, a careful analysis reveals that this design may fail to capture some aspects of the personal vote and, contrary to previous interpretations, may reveal a great deal about the importance of using incumbency as a cue.

The comparison between new voters and old voters in a given district will only capture those elements of the personal vote that stem from the personal relationship that the incumbent has established with her constituents over the years, but will miss the electoral advantage that stems from the resources associated with being an incumbent, since these resources can, at least in principle, be targeted to both new and old voters. For example, the incumbent can exploit her "reputation" not only among old voters but also among new voters. And the incumbent can deploy resources such as franking privileges, campaign funds, etc., to both old voters and new voters alike. Hence, by construction, this design cannot provide information about how the incumbent's vote share is affected by these kinds of activities.

It follows that if old voters and new voters were found to vote for the incumbent at the same rate, this should not be interpreted as evidence that there is no personal incumbency advantage, but rather as evidence that the incumbent's history with her constituents does not translate into

a significant electoral advantage. This relationship between history with the incumbent (or lack thereof) and electoral performance could be explained in more than one way. We discuss some of these alternative explanations here, without subscribing to any one in particular.

Some authors would say that the difference between early new voters (ENV) and late new voters (LNV) represents how accurately both types of voters identify the incumbent's type or ability. For example, in the context of Ashworth and Bueno de Mesquita's (2008) model, we can think that early new voters receive more informative (lower variance) signals about the incumbent's ability (and possibly the challenger's) than late new voters, since they spend an extra period with the incumbent. Having more informative signals, early new voters are more likely to correctly identify the higher quality candidate who, under typical selection mechanisms (candidates who become incumbents are in equilibrium of higher quality than candidates who do not), is the incumbent. In this framework, observing no difference in the voting rates for the incumbent between ENV and LNV can be attributed to the degree of information contained in the signals about candidates' abilities being equivalent for both types of voters. Conversely, observing that ENV vote for the incumbent at higher rates could be interpreted as their having more informative signals than LNV. In section 5, we show that when the party of the incumbent remains the same before and after redistricting, ENV and LNV vote for the incumbent at similar rates. But when the party of the incumbent changes, ENV vote for the incumbent at higher rates than LNV. According to this model, this difference could be explained by the degree of information contained in the signal: when the movement involves no change in party, the signal observed by ENV and LNV has approximately the same amount of information, because a large number of factors remain unchanged before and after redistricting. When the movement involves a change in party, the extra period that ENV spend with the incumbent results in a more informative signal at $t+1$ than what LNV are able to learn in only one period, causing a significant difference in voting behavior since ENV are more likely to identify the higher quality candidate, which is the incumbent.¹⁵

A different, political psychology approach would explain a difference between ENV and LNV

¹⁵Assuming that entry costs for challengers are not excessively large. See Ashworth and Bueno de Mesquita (2008).

by appealing to voter learning rather than to revelation of information about candidates' types. Marcus and Mackuen (1993) argue that when voters experience a mismatch between what they have come to expect from the environment and what the environment is, ongoing activity is inhibited and attention is shifted towards the intrusive source, making voters less reliant on habit. In other words, anxiety (caused by novelty or threat) stimulates voters' attention and releases them from habit and standing decisions.¹⁶ This view might interpret the change in incumbent that occurs after redistricting as such a mismatch between expectations and the environment, prompting voters to rely less on habitual voting cues, such as incumbency. This would predict that late new voters would vote for the incumbent at a lower rate than early new voters: they may be relying less on incumbency for their vote than voters who were moved a period before because they are trying to turn to more candidate-specific information for judgment as opposed to merely incumbency status. And if this information takes time to collect and/or reveals unfavorable features about the incumbent, one would expect LNV to vote at a lower rate for their incumbent.

According to this interpretation, the difference that we see between same-party and different-party movements after redistricting could be attributed to the intensity of the threat or anxiety induced by a change in incumbent. When the incumbent changes but the party remains the same, voters may not perceive the change as too threatening and may continue to base their voting decisions on the same cues and rules that they used before redistricting. But when the party of the incumbent also changes, voters may face a more radical mismatch between what they had come to expect and the environment, and may decide to rely less on incumbency as a cue and engage in a more elaborate learning process which may take some time and/or reveal information less favorable to the incumbent. The crucial issue is that although old and new voters are exposed to the same incumbency cue, they may differ in how much they rely on it for making their voting decisions. The fact that the incumbency cue is held constant does not imply that old and new voters will interpret or react to this cue in identical ways.

In this section we have offered two alternative substantive interpretations of what is being esti-

¹⁶See also Valentino, Hutchings, Banks, and Davis (2008) for a recent review of this view and experimental evidence on the effects of anxiety on information seeking and learning.

mated when redistricting is used as an identification strategy for the personal vote. One implication of both alternatives is that finding a zero “personal vote” effect need not be any more surprising than finding a positive effect as is common in the literature. For example, a zero effect does not imply that incumbents lack an electoral advantage over challengers due to deploying more resources or scaring off high-quality challengers. The design does allow one to measure the importance of the incumbency cue and possibly how quickly voters learn the type or quality of the candidates in their new districts both when the partisanship of incumbents is held constant and when it is not.

4 Empirical Application: California and Texas

We implement the first-best design in Texas and the second-best design in both Texas and California. We analyze congressional elections between 1998 and 2006. Our choice of Texas is motivated by the availability of data at the Voting Tabulation District (VTD) level, which allows us to track the same geographical unit over time, and by the consecutive congressional redistricting plans implemented in 2002 and 2004, which give us the unique opportunity of implementing the first-best two-time redistricting design described in Section 2.2. Our choice of California is motivated by the availability of data at the census block level, which, as in Texas, allows us to track the same geographical unit over time, and by the fact that Desposato and Petrocik (2003) estimated Ansolabehere, Snyder, and Stewart’s (2000) design in California during the 1990s.

Texas’ congressional districts were redrawn once after the reapportionment that followed the 2000 census, and again before the 2004 elections in a highly controversial mid-decade plan engineered by former Republican House Majority Leader Tom DeLay (Bickerstaff 2007).¹⁷ We define *late new voters* as voters who were in a given district in the 2000 and the 2002 elections and in a different district in the 2004 election, and *early new voters* as voters who were in the same district as late new voters in the 2000 election but in the 2002 and 2004 elections were in the district where late new voters were moved to in 2004.

¹⁷See Appendix A for a detailed description of Texas and California redistricting plans.

For Texas, data on electoral returns were collected from the Texas Legislative Council (TXLC) at the VTD level. VTDs are census blocks grouped to approximate voting precincts as closely as possible, providing a link between census data and electoral data.¹⁸ Since there is a one-to-one mapping between VTDs and 2000 census blocks, we are able to track the electoral returns of the same geographical unit over time. Election returns reported by the TXLC include congressional, state house, state senate, U.S. Senate, and presidential elections. Data files also include total and Hispanic voter registration, voter turnout, and candidate information including name, party affiliation, race, ethnicity and incumbency status.

For California, data on electoral returns were collected from the Statewide Database (SWDB) at the 2000 census block level.¹⁹ As in Texas, using 2000 census blocks as the unit of analysis allows us to track the electoral returns of the same geographical unit over time. Electoral returns include congressional, state house, state senate, U.S. senate, presidential, and other elections. The data also include registration and turnout for different age groups and party affiliations. The roster of congressional candidate and incumbents was obtained directly from the California Secretary of State, and data on race and ethnicity were obtained from the *Hispanic Americans in Congress* website, maintained by the Library of Congress, and the Congressional Research Service Report for Congress 2008.

We added data on challengers' quality to both the California and Texas datasets.²⁰ We also merged data from the 2000 census at the VTD-level for Texas and at the census block level for California. For Texas, census data from Summary File 1 was easily obtained at the VTD level by aggregating census blocks; for California, we merged block-level data directly since our unit of analysis is the census block. Census data from Summary File 3 was converted to the VTD-level for Texas and to the block-level for California.²¹ Variables include population by age, white,

¹⁸For details about how VTDs are constructed, data sources, and other issues regarding data construction, see Texas Legislative Council (2000, 2001).

¹⁹Data for 1998 and 2000 were directly obtained at the block level, while data for 2002 through 2006 were obtained at the precinct level and converted to 2000 census block level using conversion files provided by the Statewide Database.

²⁰Challenger quality data were kindly provided by Gary C. Jacobson.

²¹The assignment of Summary File 3 variables to blocks and VTDs is only approximate because the smallest geographical unit for which Summary File 3 variables are reported is the block-group level.

black and Hispanic population, and population by language spoken at home, employment status, place of birth, and education level. Every VTD and block in each dataset was assigned to the congressional district it belonged to in each general election between 1998 and 2006, according to the redistricting plan that was effective at the time of each election in each state.

As mentioned in Section 2.2.1, the second-best one-time redistricting design must condition on crucial covariates related to voters' previous history in their old districts. For this reason, when estimating the personal vote using this design, we restrict our analysis to movements between incumbents of the same race, ethnicity, and gender. Our substantive results are unchanged if we don't restrict the analysis in this way. Our analysis also excludes incumbents whose original districts are modified so radically that the share of old voters in the newly redrawn district is almost zero.

5 Results

5.1 Placebo Test

As discussed in our research design section, because redistricting involves the nonrandom assignment of blocks and VTDs to Congressional districts, a selection on observables assumption must be made in order to make progress. Fortunately, a placebo test is available to check this assumption for one of the states considered.

Our placebo test uses data from Texas, where the multiple redistricting plans allow us to implement our first-best design. The test examines VTDs which will be redistricted (or not) in 2004 but which are in the same district in 1998, 2000 and 2002. We assume that those to be redistricted for the 2004 election are treatment and those who will remain are control, and arbitrarily denote 2000 to be the baseline year. As explained in Section 2.2, our placebo test is that in 2002 there should be no significant difference between our treated and control groups in vote intention.

Our dataset allows us to draw on a rich set of covariates based on electoral returns, registration files, and census data. Table 1 provides the covariates we use to perform the matching for our placebo test. Like past work, we use past presidential vote returns, but we also use data from

statewide offices, registration figures, past turnout numbers, and the past vote for the Democratic Party's House candidate. Note that because both the treated and control units in this placebo test are drawn from the same Congressional district (as they are in our first-best design), we by definition match on the party of the incumbent, the historical quality of challengers, and other aspects of past races at the local, statewide and national level as experienced by the VTDs we are matching.

The variables listed in Table 1 were chosen on a priori theoretical grounds because we believed them to be theoretically important. This is the set of variables that we used in the first specification of the placebo test, and no further modification of the set was necessary to pass the test. For completeness, we then also matched on other variables drawn from the census (such as the percentage of the voting eligible population which is Hispanic, black, white, native born, and who speaks English), but these extra covariates were not necessary to reliably pass the placebo test.

We match on the covariates using Genetic Matching (GenMatch), a nonparametric matching method which algorithmically maximizes the balance of observed covariates between treated and control groups (Diamond and Sekhon 2005; Sekhon In Press). GenMatch is a generalization of propensity score and Mahalanobis distance matching that uses a genetic algorithm to optimize the balance of observed covariates given the data, and it does not depend on estimating a propensity score. In this analysis, the data is restricted to ensure common support for each covariate.

The balance statistics in Table 1 show the excellent balance that Genetic Matching was able to find post matching. The mean differences between treatment and control groups, the maximum differences in the empirical QQ-plots, and the significance of the differences greatly shrank post matching in every case. The smallest bootstrapped Kolmogorov-Smirnov (KS) p-value post matching is 0.235 and the second smallest is 0.601 while all of the pre-matching p-values are significant at 0.00.

The first row of Table 2 presents results for the 2002 incumbent vote proportions, and it shows that the estimate for the matched data is statistically indistinguishable from zero and substantively small (0.00245). This is not a case of the confidence interval simply being large: the point estimate is tiny and the confidence interval tight.

Genetic Matching estimates of vote proportions and confidence intervals from Hodges-Lehmann Interval Estimation are presented in all tables.²² Our substantive results are unchanged if either bivariate overdispersed GLM models are estimated on the matched data²³ or if Abadie-Imbens standard errors (Abadie and Imbens 2006) are used instead.

Figure 3 plots the QQ-plot for the incumbent vote in 2002 between treatment and control groups. The figure visually presents the results in the first row of Table 2. It is clear that the result for incumbent vote is zero, as it should be in this placebo test.

Presidential vote, which is the conditioning variable used by Ansolabehere, Snyder, and Stewart (2000), is not sufficient to satisfy this placebo test. We present detailed results for using the 2000 presidential vote as the estimate of the normal vote, but we have also conducted placebo tests using means and medians of a number of presidential elections. Figure 4(a) presents the balance on Presidential vote in 2000 matching only on this variable. As can be seen, balance is excellent. Figure 4(b) presents the QQ-plot for the estimand in question, House vote in 2002. Unlike the case for the rich conditioning set, there is a significant treatment effect. Row (2) in Table 2 (Data 1) shows the formal results for this additional placebo test.²⁴

Notwithstanding the excellent balance displayed in Figure 4, it may be argued that balance is not good enough on past Presidential vote. The QQ-plot for Presidential vote in Figure 4 corresponds to a mean difference of only 0.00165, a maximum difference of 0.021 and a bootstrap KS p-value of 0.958. However, if a tight caliper is used, even better balance can be obtained. Using this caliper, balance on presidential vote improves to a mean difference of 0.000156, a maximum difference of 0.000362 and bootstrapped KS p-value of 1. But as the Data 2 results in Table 2 show, even when balance of Presidential vote is improved this much, balance on the House vote proportion in 2002 is still poor, as there is still a significant effect for incumbent vote in 2002

²²Rosenbaum (2002) provides details, and Hill and Reiter (2006) provide a simulation study comparing the performance of Hodges-Lehmann intervals relative to other methods of interval estimation for treatment effects using matching.

²³See McCullagh and Nelder (1989) for a general treatment of overdispersed models for binary data, and Mebane and Sekhon (2004) for a discussion of overdispersion in aggregate election data.

²⁴Conditioning on party registration, as done by Desposato and Petrocik (2003), instead of or in addition to the presidential vote is not sufficient to pass the placebo test.

($p=0.000$). And the point estimate does not significantly change, even with the caliper.

5.2 Difference Between Old Voters and New Voters When Party Remains The Same

The results for Texas are displayed in Table 3. Rows (1) and (2) present results from the second-best one-time redistricting design as originally proposed by Ansolabehere, Snyder, and Stewart (2000) for the treatment effect of the incumbent changing, but the party remaining the same. This is the design which fails the placebo test because it only conditions on past presidential vote. Using this design, there is a 5.8% difference in the vote share of Texas old and new voters in 2004 with a p -value of 0.000. This point estimate and the confidence interval of 4.7% to 7.0% is similar to the average national estimate provided by Ansolabehere, Snyder, and Stewart (2000) in the modern era. And it is somewhat larger than Desposato and Petrocik's (2003) estimate from California. These results show that, in Texas, the second-best one-time redistricting design that conditions only on past presidential vote estimates a positive and highly significant effect of the same magnitude as found in the existing literature. The results also show that this difference remains significant in 2006, even though its magnitude is lower than in 2004.

Rows (5) and (6) in Table 3 present the Texas results from our first-best two-time redistricting design when the party of the incumbent remains the same. The difference in the vote shares of late new voters and early new voters is estimated to be statistically indistinguishable from zero for both 2004 and 2006. Note that our point estimates are also extremely small. For example, for 2004, the estimated vote proportion is 0.00637 and for 2006 it is 0.00843.

The third and fourth rows in Table 3 present our estimates from the second-best one-time redistricting design, but this time conditioning on the same covariates we used in our first-best design, with the addition of variables which attempt to measure details of the House election at baseline in 2000 and in 1998. In particular, we add Jacobson's challenger quality measures in 2000 and in 1998. As we extensively discussed in the research design section, this design is not as compelling as the previous one. Nonetheless, it allows us to use more data (434 observations to estimate a single parameter).

As is made clear in rows (3) and (4) of Table 3, when conditioning on a large set of covariates, all of our estimates of the difference in the incumbent vote shares between new neighbors and new voters in Texas are extremely small and all are *substantively* and *statistically* indistinguishable from zero. The largest absolute value of the point estimate is 0.00472 (incumbent vote in '06). This is of course insignificant, but even if it were significant, it would not be a substantively meaningful effect.

The results for same-party movements for California are displayed in Table 4. Since in this state there is only one redistricting plan implemented during the 2000 decade, we cannot apply our first-best two-time redistricting design nor is a placebo test available to validate the design. We can only use the second-best one-time redistricting design in California.

Rows (4) through (6) in Table 4 show the same-party results for the second-best design that conditions on a large set of covariates, as in in rows (3) and (4) of Table 3. Contrary to the results for Texas, in California there is a statistically significant difference between old and new voters as estimated by this design. But although this difference is significant, the results are about half the size of those usually found in the literature. In 2002, the difference in the incumbent vote shares between old and new voters is 2.2%, which is significantly smaller than the 4 to 5 percentage average results found by Desposato and Petrocik (2003) in California in the 1990s. Consistent with the Desposato and Petrocik results, row (1) in Table 4 shows that when we only condition on presidential vote, we estimate a 4.1% ($p=0.000$) effect in 2002 between old and new voters when the party does not change. This presidential-vote only result is also similar to the 4% nation-wide estimates of Ansolabehere, Snyder, and Stewart (2000). Therefore, although we find a significant 2.2% effect in California, it is notably smaller than that found using the standard approach.

When conditioning on all key covariates, the results for 2004, the second election after redistricting, are similar to the 2002 results in magnitude and significance. But in the 2006 election, the initial electoral advantage enjoyed by the incumbent among old voters decreases. The effect appears to switch sign, but this is not a robust result. As shown in row (6) of Table 4, the effect is substantively small, just -0.72% , and the effect is not significant when alternative statistical tests

are used. For example, the p-value of a simple difference-in-means t-test is 0.1698.

5.3 Difference Between Old Voters and New Voters When Party Changes

Rows (7) and (8) in Table 3 present the results for Texas from our preferred two-time redistricting design, but when the party of the incumbent changes. Here we find that there is a significant and large difference in the incumbent vote shares of early and late new voters. Early new-voter VTDs vote for the incumbent party at a much higher rate than late new-voter VTDs: 11.9%. But by the time that late new-voter VTDs have been in the Congressional district for a term, this effect drops to about 3.9%.

The California results for movements to an incumbent of the opposite party using the second-best design that conditions on a large set of covariates are shown in rows (7) through (9) of Table 4. The incumbent vote share among old voters is larger than among new voters, and the difference is large and statistically significant. In 2002, the difference in the incumbent vote shares between old and new voters is about 10%. The figure for 2004 is very similar, although the effect decreases to about 3% by 2006, two elections after redistricting. But these results must be interpreted with caution. As mentioned above, the lack of multiple redistricting plans in California forces us to consider only the second-best design, which becomes particularly problematic when considering movements between incumbents of the opposite party. The reason is that when movement occurs across parties, it is problematic to condition on incumbent vote shares at baseline. Since the previous incumbent vote was for an incumbent of a different party, it is a poor idea to condition on this variable because it probably would introduce bias. Voting for a Democratic incumbent is not comparable to voting for a Republican incumbent. The two-time redistricting design does not suffer from this issue.

6 Conclusion

The use of so called natural experiments to estimate causal effects has recently become popular. Although natural experiments offer significant advantages, they do not possess key benefits of actual experiments and hence require careful theoretical and statistical work to make valid inferences. First, there exists the obvious problem that natural experiments do not have random assignment with a known probability distribution, and selection on observables is a strong assumption that is difficult to justify. As in the current example, rarely can natural experiments be used without significant covariate adjustment.

A less often noted but crucial issue with natural experiments is that they may not identify the causal effect of interest even if the manipulation they are based on is assumed to be exogenous. When designing a controlled randomized experiment, researchers *a priori* design the study so that randomization will ensure the identification of the causal effect of interest. But with natural experiments, one must find a way to connect the manipulation done by someone else to the question of interest. As we have seen, although it initially seemed that one-time redistricting allowed for a straightforward estimation of the personal vote, the manipulation is far from providing the appropriate counterfactual.

Our analysis has also shown that a benefit of natural experiments that should not be underestimated is that there *is* a manipulation, and as such one can think through the process by which the manipulation was assigned. For example, it is easier to determine what is post- and what is pre-treatment than with the observational designs more commonly used, certainly easier than with cross-sectional data that lacks any manipulation. And placebo tests may be available to help determine if the selection on observables assumption is plausible.

Substantively, we find that when the party of the incumbent remains the same, in Texas old voters vote for the incumbent at the same rate as new voters, even though using the design as originally proposed by Ansolabehere, Snyder, and Stewart (2000) we estimate an effect of about 5.8%. Since this design uses the incorrect potential outcomes and fails a placebo test, our results show that there is a methodological bias in existing published estimates of the personal vote (e.g.,

Ansolabehere, Snyder, and Stewart 2000; Carson, Engstrom, and Roberts 2007; Desposato and Petrocik 2003). In California, we estimate a small effect of 2.2%. It is unclear if the different results in Texas and California are the consequence of actual differences between the states or because the assumptions fail in California. Unlike in Texas, it is not possible to conduct a placebo test in California so it is difficult to distinguish the explanations. If the differences are real, one conjecture is that the Texas results are the consequence of one of the most partisan redistricting plans in recent history, and that the California results are the consequence of one of the most pro-incumbent protection redistricting plans of the 2000 cycle (Bickerstaff 2007). In either case, our estimated effects when the incumbent party stays constant are much smaller than those previously reported in the literature.

The incumbency cue appears to provide the necessary information for old and new voters to vote in a similar fashion when the partisanship of the incumbent does not change. As discussed in Section 3, our results do not imply that incumbents do not have an electoral advantage over challengers. But they do imply that the incumbency cue is important. When the partisanship of the incumbent does not change, the incumbency cue provides the information needed for new voters and old voters to vote alike in Texas and only slightly differently in California. When the partisanship of the incumbent changes, the incumbency cue is no longer sufficient. In this case, new voters support their new incumbent substantially less than voters with more experience with the incumbent. The effect is large in both Texas and California. In California, the effect when the party of the incumbent changes is almost five times larger (10%) than when it does not (2.2%). While in Texas, the effect is 12% when the party of the incumbent changes, and zero when it does not.

Our zero (in Texas) or small (in California) estimates when the incumbent party is unchanged are consistent with theoretical arguments that existing positive estimates of incumbency advantage are plagued by selection problems (Cox and Katz 2002; Zaller 1998). And our finding of a large effect when the party of the incumbent changes is consistent with the results of Lee (2008), who finds a significant party incumbency advantage using a regression discontinuity design. The fact

that two different quasi-experimental methods that make very different identification assumptions find the same results increases our confidence. With observational data, where strong assumptions must necessarily be made, such corroboration is essential.

A Texas and California redistricting plans in the 2000s

Texas implemented six different congressional district plans between 1990 and 2006.²⁵ After the reapportionment that followed the 1990 census, the districts enabled by the old C001 plan were redrawn. The 1992 elections were held under the new districts enacted by plan C657, which remained in effect until the 1996 primaries. In August 1996, 13 of Texas' 30 congressional districts were redrawn. The new plan, C746, was used in the 1996 general election and it remained in effect during the 1998 and 2000 elections.

In 2001, after the reapportionment following the 2000 census that created two new congressional seats, the Texas Legislature was in charge of redrawing the senate, house, congressional, and State Board of Education districts during the regular session of the 77th Legislature. But the plans failed to be considered by the full Senate and the full House, and the legislature adjourned without enacting new districts. A number of congressional proposals were submitted to state and federal courts. Finally, on November 14, 2001, the U.S. District Court issued an order adopting new congressional districts (Plan C1151) for the 2002 elections.

But plan C1151 was only in effect for the 2002 elections. In 2003, Republican majority leader Tom Delay led an effort to enact a new congressional district plan, with the objective of maximizing the number of Texas' Republicans elected to Congress in the 2004 and subsequent elections. After a legislative battle that included Democratic lawmakers massively fleeing to New Mexico and Oklahoma to avoid quorum, the new plan (Plan C1374) was passed in October, 2003. The 2004 primaries and general election were held under this new plan. Congressional districts were redrawn one more time in 2006.

In California, there was only one redistricting plan implemented in the 2000s. The districts in effect during the 1990s were redrawn by the 2001 redistricting plan, which was enacted in two separate bills in September 2001. Bill AB632 established Senate and Congressional districts, and bill SB802 established Assembly and Board of Equalization districts.

²⁵See the Texas Legislative Council's Redistricting website <http://www.tlc.state.tx.us/redist>, Texas Legislative Council (2000), and Texas Legislative Council (2001) for details about Texas' redistricting plans during the 1990s and 2000s.

References

- Abadie, Alberto, and Guido Imbens. 2006. "Large Sample Properties of Matching Estimators for Average Treatment Effects." *Econometrica* 74: 235–267.
- Abrajano, Marisa A., Jonathan Nagler, and R. Michael Alvarez. 2005. "A Natural Experiment of Race-Based and Issue Voting: The 2001 City of Los Angeles Elections." *Political Research Quarterly* 58(2): 203–218.
- Alford, John R., and David W. Brady. 1989. "Personal and Partisan Advantage in US Congressional Elections." In *Congress Reconsidered 4th edition*, ed. Lawrence C. Dodd, and Bruce I. Oppenheimer. Washington, D.C.: CQ Press.
- Ansolabehere, Stephen, and James Snyder. 2002. "The Incumbency Advantage in U.S. Elections: An Analysis of State and Federal Offices, 1942-2000." *Election Law Journal* 1(3): 315–338.
- Ansolabehere, Stephen, David W. Brady, and Morris P. Fiorina. 1988. "The Vanishing Marginals and Electoral Responsiveness." *British Journal of Political Science* 22(1): 21–38.
- Ansolabehere, Stephen, James M. Snyder, and Charles Stewart. 2000. "Old Voters, New Voters, and the Personal Vote: Using Redistricting to Measure the Incumbency Advantage." *American Journal of Political Science* 44(1): 17–34.
- Ashworth, Scott, and Ethan Bueno de Mesquita. 2008. "Electoral Selection, Strategic Challenger Entry, and the Incumbency Advantage." *Journal of Politics* 70(4): 1006–1025.
- Bickerstaff, Steve. 2007. *Lines in the Sand*. Austin, Texas: University of Texas Press.
- Born, Richard. 1979. "Generational Replacement and the Growth of Incumbent Reelection Margins in the US House." *American Political Science Review* 73(3): 811–817.
- Brady, Henry, and John McNulty. 2004. "The Costs of Voting: Evidence from a Natural Experiment."

- Breaux, David. 1990. "Specifying the Impact of Incumbency on State Legislative Elections." *American Politics Quarterly* 18(3): 270–286.
- Buetler, Monika, and Michel A. Marechal. 2007. "Framing Effects in Political Decision Making: Evidence from a Natural Voting Experiment."
- Carman, Christopher, James Mitchell, and Robert Johns. 2008. "The unfortunate natural experiment in ballot design: The Scottish Parliamentary elections of 2007." *Electoral Studies* 27(3): 442–459.
- Carson, Jamie L., Erik J. Engstrom, and Jason M. Roberts. 2007. "Candidate Quality, the Personal Vote, and the Incumbency Advantage in Congress." *American Political Science Review* 101(2): 289–301.
- Cox, Gary W., and Jonathan N. Katz. 1996. "Why Did the Incumbency Advantage in U.S. House Elections Grow?" *American Journal of Political Science* 40(2): 478–497.
- Cox, Gary W., and Jonathan N. Katz. 2002. *Elbridge Gerry's Salamander: The Electoral Consequences of the Reapportionment Revolution*. New York: Cambridge University Press.
- Cox, Gary W., and Scott Morgenstern. 1993. "The Increasing Advantage of Incumbency in the U.S. States." *Legislative Studies Quarterly* 18(4): 495–514.
- Desposato, Scott W., and John R. Petrocik. 2003. "The Variable Incumbency Advantage: New Voters, Redistricting, and the Personal Vote." *American Journal of Political Science* 47(1): 18–32.
- Diamond, Alexis, and Jasjeet S. Sekhon. 2005. "Genetic Matching for Estimating Causal Effects: A General Multivariate Matching Method for Achieving Balance in Observational Studies."
- Dunning, Thad. 2008. "Improving Causal Inference: Strengths and Limitations of Natural Experiments." *Political Science Quarterly* 61(2): 282–293.

- Erikson, Robert S. 1971. "The advantage of Incumbency in Congressional Elections." *Polity* 3(3): 395–405.
- Ferejohn, John A. 1977. "On the Decline of Competition in Congressional Elections." *American Political Science Review* 71(1): 166–176.
- Fiorina, Morris P. 1977. "The Case of the Vanishing Marginals: The Bureaucracy Did It." *American Political Science Review* 71(1): 177–181.
- Gelman, Andrew, and Gary King. 1990. "Estimating Incumbency Advantage without Bias." *American Journal of Political Science* 34(4): 1142–1164.
- Gordon, Sandy, and Greg Huber. 2007. "The Effect of Electoral Competitiveness on Incumbent Behavior." *Quarterly Journal of Political Science* 2(2): 107–138.
- Green, Donald, and Alan Gerber. 2008. "Field Experiments and Natural Experiments." In *The Oxford Handbook of Political Methodology*, ed. Janet M. Box-Steffensmeier, Henry E. Brady, and David Collier. New York: Oxford University Press pp. 357–381.
- Heckman, James J., Hidehiko Ichimura, and Petra Todd. 1998. "Matching as an Econometric Evaluation Estimator." *Review of Economic Studies* 65(2): 261–294.
- Heckman, James J., Hidehiko Ichimura, Jeffrey Smith, and Petra Todd. 1998. "Characterizing Selection Bias Using Experimental Data." *Econometrica* 66(5): 1017–1098.
- Hill, Jennifer, and Jerome P. Reiter. 2006. "Interval Estimation for Treatment Effects Using Propensity Score Matching." *Statistics in Medicine* 25: 2230–2256.
- Huber, Gregory A., and Kevin Arceneaux. 2007. "Identifying the Persuasive Effects of Presidential Advertising." *American Journal of Political Science* 51(4): 957–977.
- Jacobson, Gary C. 1987. "The Marginals Never Vanished: Incumbency and Competition in Elections to the U.S. House of Representatives." *American Journal of Political Science* 31(1): 126–141.

- Krashinsky, Michael, and William J. Milne. 1993. "The Effects of Incumbency in U. S. Congressional Elections, 1950-1988." *Legislative Studies Quarterly* 18(3): 321–344.
- Krasno, Jonathan S., and Donald P. Green. 2008. "Do Televised Presidential Ads Increase Voter Turnout? Evidence from a Natural Experiment." *Journal of Politics* 70(1): 245–261.
- Krehbiel, Keith, and John R. Wright. 1983. "The Incumbency Effect in Congressional Elections: A Test of Two Explanations." *American Journal of Political Science* 27(1): 140–157.
- Lassen, David D. 2005. "The Effect of Information on Voter Turnout: Evidence from a Natural Experiment." *American Journal of Political Science* 49(1): 103–118.
- Lee, David S. 2008. "Randomized Experiments from Non-random Selection in U.S. House Elections." *Journal of Econometrics* 142(2): 675–697.
- Marcus, George E., and Michael B. Mackuen. 1993. "Anxiety, enthusiasm, and the vote: The emotional underpinnings of learning and involvement during presidential campaigns." *American Political Science Review* 87(3): 672–685.
- Mayhew, David R. 1974. "Congressional Elections: The Case of the Vanishing Marginals." *Polity* 6(3): 295–317.
- McCauley, John F., and Daniel N. Posner. 2007. "African Borders as Sources of Natural Experiments."
- McCullagh, Peter, and John A. Nelder. 1989. *Generalized Linear Models*. New York: Chapman & Hall.
- Mebane, Walter R. Jr., and Jasjeet S. Sekhon. 2004. "Robust Estimation and Outlier Detection for Overdispersed Multinomial Models of Count Data." *American Journal of Political Science* 48(2): 391–410.
- Nelson, Candice J. 1979. "The Effect of Incumbency on Voting in Congressional Elections." *Political Science Quarterly* 93: 665–678.

- Payne, James L. 1980. "The Personal Electoral Advantage of House Incumbents, 1936-1976." *American Politics Research* 8(4): 465–482.
- Posner, Daniel N. 2004. "The Political Salience of Cultural Difference: Why Chewas and Tumbukas Are Allies in Zambia and Adversaries in Malawi." *American Political Science Review* 98(4): 529–545.
- Rosenbaum, Paul R. 2002. *Observational Studies*. 2nd ed. New York: Springer-Verlag.
- Rush, Mark E. 1993. *Does Redistricting Make a Difference? Partisan Representation and Electoral Behavior*. Baltimore: Johns Hopkins University Press.
- Sekhon, Jasjeet S. In Press. "Matching: Multivariate and Propensity Score Matching with Automated Balance Search." *Journal of Statistical Software* .
- Texas Legislative Council, Research Division. 2000. *Guide to 2001 Redistricting*. Austin, Texas: Texas Legislative Council.
- Texas Legislative Council, Research Division. 2001. *Data for 2001 Redistricting in Texas*. Austin, Texas: Texas Legislative Council.
- Valentino, Nicholas A., Vincent L. Hutchings, Antoine J. Banks, and Anne K. Davis. 2008. "Is a Worried Citizen a Good Citizen? Emotions, Political Information Seeking, and Learning via the Internet." *Political Psychology* 29(2): 247–273.
- van der Brug, Wouter. 2001. "Perceptions, Opinions and Party Preferences in the Face of a Real World Event: Chernobyl as a Natural Experiment in Political Psychology." *Journal of Theoretical Politics* 13(1): 53–80.
- Whitford, Andrew B. 2002. "Decentralization and Political Control of the Bureaucracy." *Journal of Theoretical Politics* 14(2): 167–193.
- Zaller, John. 1998. "Politicians as Prize Fighters: Electoral Selection and Incumbency Advantage." In *Party Politics and Politicians*, ed. John Geer. Baltimore: Johns Hopkins University Press.

Table 1: Balance for Placebo Test Covariates for Texas

Variable	Before Matching			After Matching		
	mean diff	D-statistic	KS-pvalue	mean diff	D-statistic	KS-pvalue
Dem Pres. vote share '00	.0447	.100	0.00	.00459	.0337	0.953
Dem House vote share '00	.159	.305	0.00	.00693	.0344	0.678
Dem House vote share '98	.127	.340	0.00	.00585	.0368	0.996
Dem Senate vote share '00	.0426	.120	0.00	.00576	.0317	0.846
Dem Governor vote share '98	.0305	.0974	0.00	.00510	.0241	0.942
Dem Att. Gen. vote share '98	.0353	.141	0.00	.00683	.0358	0.868
Dem Comptroller vote share '98	.0304	.208	0.00	.00499	.0373	0.994
Voter turnout '00	.0331	.102	0.00	.00607	.0327	0.943
Voter turnout '98	.028	.199	0.00	.0111	.0378	0.235
Registration '00	.0308	.157	0.00	.00736	.0608	0.601

The mean differences are the simple differences between treatment and control, the D-statistic is the largest difference in the empirical QQ-plot on the scale of the variable, and the KS-pvalue is from the bootstrapped Kolmogorov-Smirnov test.

Table 2: Placebo Tests for 2002 Incumbent Vote Share in Texas

		Estimate	95% CI		p.value
Matching on all key covariates					
(1)	Incumbent vote '02	0.00245	-0.00488	0.00954	0.513
Matching on past presidential vote only					
(2)	Incumbent vote '02 (Data 1)	0.0237	0.0178	0.0294	0.000
(3)	Incumbent vote '02 (Data 2)	0.0285	0.0160	0.0413	0.000

Genetic Matching estimates of vote proportions. There are 474 observations for matching on all key covariates, and 2666 and 412 observations for matching on past presidential vote only for Data 1 and Data 2, respectively.

Table 3: Results for Texas

	Estimate	95% CI		p.value
Same-Party, Second-Best One-Time Redistricting Design				
Matching on Past Presidential Vote Only				
(1) Incumbent vote '04	0.0579	0.0470	0.0700	0.000
(2) Incumbent vote '06	0.0104	0.00317	0.0176	0.005
Matching on All Key Covariates				
(3) Incumbent vote '04	0.00214	-0.00807	0.0124	0.690
(4) Incumbent vote '06	0.00472	-0.00539	0.0149	0.378
Same-Party, First-Best Two-Time Redistricting Design				
(5) Incumbent vote '04	0.00637	-0.00428	0.0177	0.254
(6) Incumbent vote '06	0.00843	-0.00938	0.0258	0.457
Different-Party, First-Best Two-Time Redistricting Design				
(7) Incumbent vote '04	0.119	0.0595	0.191	0.0000
(8) Incumbent vote '06	0.0389	0.00973	0.0692	0.0106


Genetic Matching estimates of vote proportions. For the same-party second-best design, there are 412 observations when matching only on presidential vote and 434 when matching on all key covariates. For the first-best design, there are 166 observations in the same-party case and 70 observations in the different-party case.

Table 4: Results for California


	Estimate	95% CI		p.value
Same-Party, Second-Best One-Time Redistricting Design				
Matching on Past Presidential Vote Only				
(1) Incumbent vote '02	0.0409	0.0405	0.0413	0.000
(2) Incumbent vote '04	0.0090	0.0087	0.0093	0.000
(3) Incumbent vote '06	0.0031	0.0028	0.0035	0.000
Matching on All Key Covariates				
(4) Incumbent vote '02	0.0219	0.0171	0.0266	0.000
(5) Incumbent vote '04	0.0240	0.0195	0.0284	0.000
(6) Incumbent vote '06	-0.0072	-0.0129	-0.0015	0.012
Different-Party, Second-Best One-Time Redistricting Design				
(7) Incumbent vote '02	0.1025	0.0925	0.1122	0.000
(8) Incumbent vote '04	0.1020	0.0932	0.1109	0.000
(9) Incumbent vote '06	0.0307	0.0186	0.0428	0.000

Genetic Matching estimates of vote proportions. There are 3526 observations for the same-party design and 1394 observations for the different-party design.


Figure 1: QQ Plots of Baseline Vote Share for Incumbent House Member
California and Texas




(a) CA 2000, unconditional



(b) TX 2002, unconditional

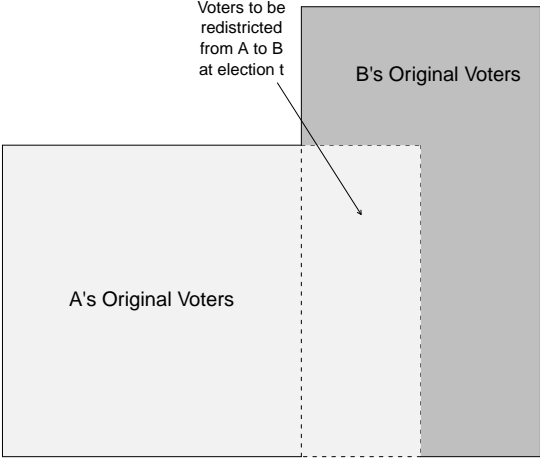



(c) CA 2000, matched on 2000 presidential vote

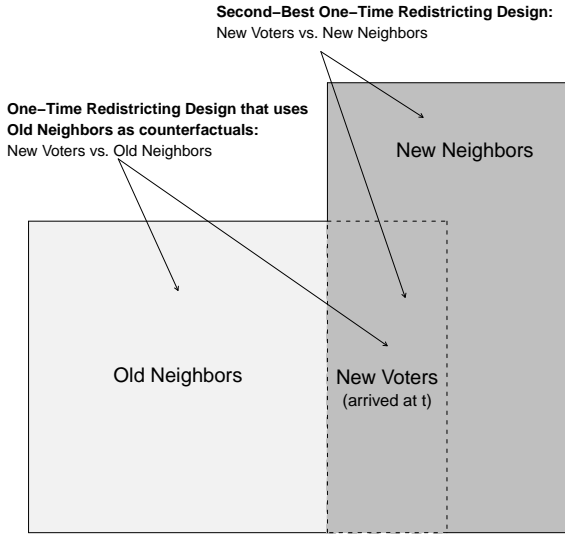


(d) TX 2002, matched on 2000 presidential vote

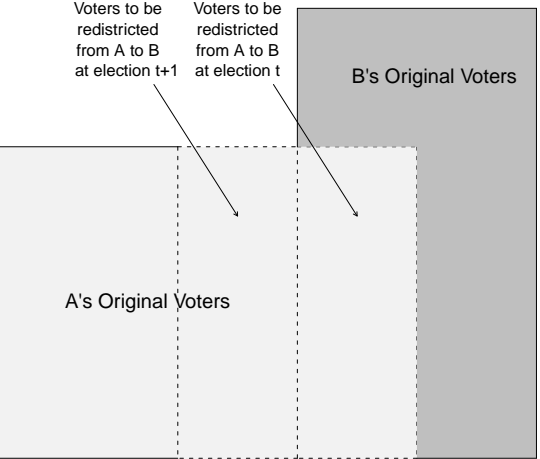
Figure 2: Illustration of Research Designs



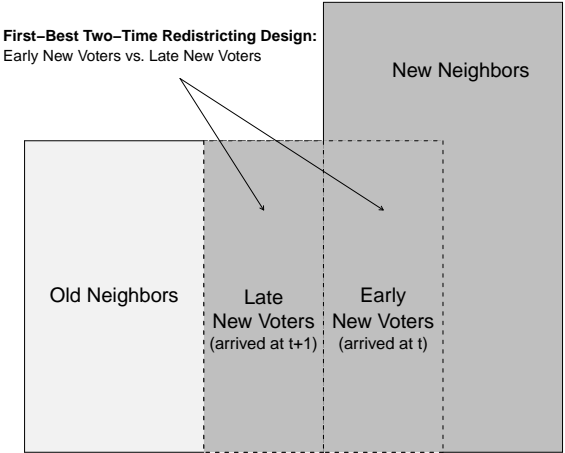
(a) Before One-Time Redistricting



(b) After One-Time Redistricting



(c) Before Two-Time Redistricting



(d) After Two-Time Redistricting

Figure 3: QQ Plot for Placebo Test Conditioning on All Key Covariates

2002 Vote for the Incumbent House Member






Figure 4: QQ Plots for Placebo Test Conditioning Only On Presidential Vote



(a) 2000 Presidential Vote (Baseline)



(b) 2002 Vote for the Incumbent House Member (Placebo Test)