

# Panel Effects in the American National Election Studies

Larry M. Bartels  
Princeton University

Parallel panel and fresh cross-section samples in recent National Election Study surveys provide valuable leverage for assessing the magnitude of biases in statistical analyses of survey data due to panel attrition and panel conditioning. My analyses employing a variety of typical regression models suggest that substantial panel biases are likely to be fairly rare in these data, even when panel and cross-section respondents have markedly different characteristics. However, two of the dependent variables considered here—campaign interest and turnout—do appear to be sufficiently sensitive to panel effects to warrant significant discounting or adjustment of panel data. I propose adjustments for panel effects in both cross-sectional and dynamic analyses, based upon variants of the “fractional pooling” (Bartels 1996) and “two-stage auxiliary instrumental variables” (Franklin 1990) methods.

## 1 Introduction

The most powerful way to incorporate a dynamic element in survey research is to interview the same individuals at two or more points in time and attribute observed changes in their attitudes or behavior to the effects of intervening processes or events. This so-called *panel design* provides more direct evidence of change than the primary alternative, the *longitudinal* (or *rolling cross-section*) *design*, in which change is inferred from comparisons of different survey respondents at different times rather than from comparisons of the same respondents at different times.

Panel designs have been especially prominent in the field of electoral studies. Significant campaign studies based primarily upon election-year panel surveys include the classic Columbia studies of the 1940 and 1948 presidential campaigns (Lazarsfeld et al. 1948; Berelson et al. 1954), the parallel television-era panel studies by Patterson and McClure (1976) and Patterson (1980), and several studies based upon the 1980 American National Election Study (NES) election-year panel (Markus 1982; Bartels 1993; Finkel 1993).<sup>1</sup> The

---

*Author's note:* An earlier version of this report was presented at the 1998 summer meeting of the Political Methodology Group in La Jolla, California. I am grateful to Kathryn Cirksena and Charles Franklin for helpful discussions of issues addressed here, to anonymous referees for constructive criticism, and to Princeton University's Woodrow Wilson School of Public and International Affairs for financial support. Tables A1 through A12 and replication data sets are available from the *Political Analysis* website.

<sup>1</sup>The Columbia team's 1940 Erie County study represents an especially sophisticated and remarkably self-conscious panel design, including (for example) parallel fresh cross sections to gauge the impact of the interviews themselves on the panel respondents. Unfortunately, the resulting data do not seem to have been subjected to concerted methodological analysis of the sort attempted here.

recent efforts of Just et al. (1996) suggest that panel designs continue to appeal to scholars interested in capturing campaign effects.

The immediate stimulus for the research reported here is the heavy reliance on panel data in recent American National Election Studies.<sup>2</sup> The design employed in NES surveys in every presidential election year since 1952 has included preelection interviews with a national sample of respondents in the 2 months before the election and postelection reinterviews with as many as possible of those same respondents. However, my focus here is on longer panels, in which the same respondents are reinterviewed over the course of an entire election year or even over a multiyear election cycle. The first four decades of NES surveys included three of these more ambitious panel studies—4-year three-wave panels in 1956–1958–1960 and 1972–1974–1976, and an election-year panel in 1980.

Recent NES surveys have more routinely incorporated significant panel components. A 1990–1991–1992 “Gulf War” panel produced 1359 panel respondents in the 1992 NES survey (55% of the total 1992 sample), the 1994 NES survey included 759 panel respondents (42% of the total sample) originally interviewed in 1992, and the 1996 NES survey included 1316 panel respondents (77% of the total sample) originally interviewed in 1992 or 1994. The proposed design for the 2000 NES survey calls for a similarly heavy reliance on panel reinterviews with respondents originally interviewed in 1998.

These design choices reflect a growing appreciation of the potential power of panel designs. However, the added power of panel designs comes with some associated inferential costs. The primary questions addressed here are how to assess the magnitude of those costs and how to minimize the disadvantages of panel designs while maximizing their inferential payoff.

One of the unexploited benefits of the NES investment in panel surveys over the last decade is that it provides an unusually strong data base for analysis of methodological problems that arise in all panel surveys. Most importantly, the emphasis of NES on facilitating panel analysis while preserving the integrity of a time series extending over 50 years has produced a series of data sets with parallel panel and fresh cross-section components. As shown below, such data sets provide important opportunities for assessing and, possibly, mitigating the inferential pitfalls of panel data.

While the specific problems and opportunities posed by recent NES panel surveys provide a starting point for my research, the questions addressed here are potentially of much broader significance, since panel surveys figure prominently throughout the social sciences—for example, in the Panel Study of Income Dynamics in economics (Beckett et al. 1988). Care is always required in applying methodological lessons from one setting more generally. Nevertheless, my hope is that both the assessments offered here of the magnitude of biases in panel data and the methods outlined here for addressing those biases will stimulate comparable analyses by social scientists using panel data in a variety of other substantive settings.

## 2 Advantages of Panel Data

The variety of significant panel surveys conducted in political science and related fields should be sufficient to suggest that researchers perceive important advantages in panel

---

<sup>2</sup>The American National Election Studies (NES) is a national social science data resource based in the University of Michigan’s Institute for Social Research and funded by the National Science Foundation. NES and its organizational predecessors at the University of Michigan have conducted surveys in connection with American national elections since 1948. Project information and data are available from the NES website, <http://www.umich.edu/~nes/>. The NES project is supervised by a national Board of Overseers representing the scholarly community. Although I have been a member of the NES Board of Overseers since 1992 and chair of the Board since 1997, the research reported here was conducted in my capacity as an independent scholar, and the NES organization bears no responsibility for my analysis or conclusions.

data by comparison with cross-sectional data. Most important, in my view, are the following.

(a) Panel data facilitate adjustments for measurement error in survey responses using Wiley–Wiley (1970) or other measurement models. Given the embarrassingly low reliability of typical survey measures of attitudes and behavior, adjustments for measurement error are often essential for achieving plausible statistical estimates of causal effects (Achen 1975, 1983). Panel data are neither necessary nor sufficient for producing such adjustments, but they represent the most common basis for measurement error corrections in analyses of survey data. For example, I have used data from the 1980 NES election-year panel survey to document biases of 50% or more—both positive and negative—in a variety of estimates of partisan and media exposure effects in a presidential campaign setting (Bartels 1993).

(b) Panel data permit direct analyses of opinion change in which prior opinions appear as explanatory variables. Given the stability of typical political opinions (at least, after adjustment for measurement error)—and their modest correlations with relevant explanatory variables—direct measurement of prior opinions substantially increases the efficiency of statistical estimation and provides crucial perspective on the relative importance of new and preexisting attitudes and perceptions.

(c) Panel data facilitate analyses in which relevant explanatory variables are measured outside the immediate political or social setting affecting the attitudes or behavior to be explained. Given the susceptibility of attitudes and behavior to rationalization (for example, in terms of causally irrelevant themes prominent in campaign discourse), prior measurement of potential explanatory variables provides important insights regarding the causal priority of specific political attitudes and perceptions. Thus, for example, Zaller and Hunt (1995) used data from the 1990–1992 NES panel survey to show that support for Ross Perot in the 1992 presidential election was strongly correlated with attitudes about the budget deficit in 1992 but virtually uncorrelated with attitudes about the budget deficit among the same respondents in 1990, suggesting (at least to me) that concern about the deficit was primarily a *rationalization* for supporting Perot rather than a *reason* for doing so. I have used an essentially similar approach to assess the extent to which the impact of partisanship on voting behavior is exaggerated by the short-term responsiveness of expressed partisanship to current vote intentions (Bartels 1998).

### 3 Pitfalls of Panel Data

While panel surveys have several significant methodological virtues, they also have some notable potential drawbacks. The most serious of these is that the sample surveyed in successive panel waves may become increasingly unrepresentative of the original population, either because some of the original respondents are not reinterviewed (panel attrition) or because the experience of being interviewed itself affects the subsequent behavior and responses of those who are reinterviewed (panel conditioning). In either case, the observed opinion change in the surviving panel sample may provide a biased estimate of the corresponding opinion change in the relevant population.

#### 3.1 Panel Attrition

Failure to reinterview some panel respondents (as a matter of survey design or, more typically, due to refusal, moving, death, or other causes) may result in selection bias if this *panel attrition* is correlated with substantively relevant characteristics of the respondents (Heckman 1976). Panel reinterview rates in carefully conducted surveys can be quite high—for example, 75% through three waves of the 1980 NES election-year panel, 78% over the 2 years of the 1990–1992 NES panel, and 90% over the 6 weeks or so between

NES pre- and postelection interviews. When as many as 80 or 90% of the original panel respondents are successfully reinterviewed, it seems likely that the respondents who drop out of the panel would have to be quite different from the survivors in order for their absence to produce serious selection bias.

In any case, well-developed econometric techniques exist for analyzing data in the presence of selection bias (Heckman 1979; Achen 1986). However, I am not aware of any effort to gauge the magnitude of selection bias specifically attributable to panel attrition using these techniques. The closest analogue to such an effort is Brehm's (1993) analysis of selection bias due to nonresponse in cross-sectional surveys. Brehm used data from the 1986 and 1988 NES surveys to estimate a variety of familiar regression models (of turnout, candidate evaluation, economic voting, issue preferences, and so on) with and without Heckman-style corrections for selection bias. He concluded (Brehm 1993, 158) that most of the models "escaped with only small changes to the coefficients," but also provided some examples of more substantial biases.<sup>3</sup>

### 3.2 *Panel Conditioning*

Selection bias due to panel attrition is a fairly well-understood phenomenon, and Heckman-style corrections have been successfully applied to a variety of more or less similar problems. However, the situation with respect to the more general problem of *panel conditioning* is a good deal less happy. Here, the pitfalls facing analysts of panel data are less well understood and have received much less attention from methodologists.

The basic problem is clear enough: the unusual experience of answering arcane questions about politics, economics, or social issues for 20 min, 60 min, or more may spur ordinary people to behave in unusual ways, especially if they have any inkling that the experience is likely to be repeated at some point in the future. If participating in one survey—or anticipating subsequent reinterviews—causes respondents to change their attitudes or behavior (for example, by paying more attention to news, discussing a political campaign with their friends and family, or registering to vote), their subsequent survey responses will be biased in comparison with those of the population from which they were originally sampled.

The most straightforward way to investigate the potential effects of panel conditioning is to compare panel responses with responses from a fresh cross section unexposed to previous interviews. Unfortunately, such a comparison requires costly additional data collection, and the resulting data are not themselves amenable to panel analysis.<sup>4</sup> Moreover, a comparison of panel respondents with a fresh cross section will shed light only on cross-sectional differences between the two samples; estimating effects of panel conditioning on the *dynamics*

<sup>3</sup>The two most notable biases in Brehm's (1993) analysis were a significant overestimate of the intercept in a model of turnout and a significant underestimate of the impact of political information (as measured by familiarity with the candidates) in a model of congressional voting. The latter bias is especially worrisome given the prominence of political information as a conditioning variable in recent political studies. Moreover, Traugott's (1996) analysis of panel attrition in the 1995 NES Pilot Study, which reinterviewed respondents from the 1994 NES survey, documented a significant relationship between political information and panel attrition: of the 1994 respondents included in the 1995 Pilot Study sample, 81% of the most informed third, but only 72% of the middle third and 51 percent of the least informed third, were successfully reinterviewed. It is unclear to what extent the same differential panel attrition is reflected in regular NES production surveys, which have somewhat higher overall reinterview rates.

<sup>4</sup>Of course, the fresh cross section might also serve as the first wave of a new panel, as in the 1992 American National Election Study or, more systematically, in the Panel Study of Income Dynamics. Such interleaved panel designs are potentially quite powerful but require large-scale, long-term planning and execution.

of opinion change (for example, on the stability of opinions over time) will require more complicated analysis.

In most panel surveys, all of the surviving panel respondents have been exposed to the same sequence of previous interviews, and comparable data from a fresh cross section are unavailable. In that case, even a rough assessment of the consequences of panel conditioning is impossible.<sup>5</sup> Fortunately, NES data from parallel panel and fresh cross-section samples interviewed at the same time by the same interviewers using the same survey instrument allow for an unusually straightforward and powerful assessment of the effects of both panel attrition and panel conditioning. Here, I use those data to assess the magnitude of panel effects in the NES surveys and to explore new methods for analyzing data subject to panel effects.

#### 4 Treating Panel Effects as a Nuisance in Cross-Sectional Analyses

Analysts interested in examining cross-sectional relationships may often be uninterested in the fact that respondents were interviewed at an earlier time, except insofar as the earlier interviews affect the observed relationship in the current cross section, either by biasing the composition of the sample (panel attrition) or by causing respondents to alter their attitudes or behavior (panel conditioning). For example, many analysts have used the data from the 1992 and 1996 NES surveys considered here as self-contained cross sections (or in longitudinal studies combining data from several NES surveys), without reference to the panel structure of the 1992 and 1996 surveys. Analyses of this sort will be subject to more or less serious biases, depending upon the nature and magnitude of the panel effects produced by panel attrition and panel conditioning.

The only way to guard completely against such biases is to embed a full-scale model of panel effects within the substantive cross-sectional analysis. In this section, I outline such a model and describe how it can be estimated using parallel panel and fresh cross-section data. I also explain why actually estimating such a model may often be a bad idea—because doing so allocates too many of the available data to estimating panel effects and too few to estimating the substantive effects of interest. Finally, I introduce two simpler approaches to the estimation of panel effects, either or both of which may provide an attractive compromise between estimating an overly elaborate model of panel effects and ignoring them altogether.

I assume for concreteness that the substantive model of interest is a standard regression model of the form

$$\mathbf{y} = \mathbf{X}\beta + \varepsilon \quad (1)$$

but with the added complication that the dependent variable  $\mathbf{y}$  is observed only for cross-section respondents (i.e., those who have not been previously interviewed). For panel respondents, we observe instead

$$\mathbf{y}^p = \mathbf{y} + \omega \quad (2)$$

where

$$\omega = \mathbf{X}\alpha + \nu \quad (3)$$

<sup>5</sup>In more complicated panel designs, some respondents may have been exposed to fewer previous interviews than others, and these differences may provide some leverage for estimating conditioning effects. However, the problem remains of how to extrapolate from observed differences between respondents with different numbers of previous interviews to the desired hypothetical situation of *no* previous interviews.

represents the combined effects of panel attrition and panel conditioning on the observed panel responses  $\mathbf{y}^p$ .

The model of panel effects in Eq. (3) is quite flexible, in that it allows these effects to be correlated with any or all of the substantive explanatory variables in the analysis. Moreover, two apparent restrictions on the generality of this model can be readily overcome. First, although the model as written implies that the explanatory variables are not themselves subject to panel effects, situations in which they are can be accommodated (at the cost of some additional complexity) by adding a “reduced form” in which these effects are traced back to exogenous variables (such as demographic characteristics) which can themselves be plausibly assumed to be immune to panel effects. Second, although the model as written implies that panel effects must be linear in the explanatory variables, nonlinear functions of these (and other) variables (such as the inverse Mills ratio of the fitted value from an auxiliary probit analysis, as in a Heckman-style model of panel attrition) can simply be added as additional explanatory variables subject to the usual identification conditions.

Substituting Eqs. (1) and (3) into Eq. (2) and stacking the distinct models for cross-section and panel respondents produces the combined model

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon} \quad (4a)$$

$$\mathbf{y}^p = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\varepsilon} + (\mathbf{X}\boldsymbol{\alpha} + \boldsymbol{\nu}) \quad (4b)$$

for the cross-section and panel respondents, respectively. This formulation emphasizes the inferential leverage provided by the availability of parallel data from a fresh cross section, since it would clearly be impossible to distinguish the substantive effects  $\boldsymbol{\beta}$  from the panel effects  $\boldsymbol{\alpha}$  using the panel data and Eq. (4b) alone. However, when both panel and fresh cross-section data are available, both sets of parameters may be estimated, either by separate regressions using the cross-section and panel data (with constraints imposed to reflect the fact that the same  $\boldsymbol{\beta}$  parameters appear in both regressions) or by a single joint regression using the combined data (with interactions between the explanatory variables and an indicator variable for panel respondents, and with an allowance for heteroskedasticity to reflect the different error structures assumed for the cross-section and panel models).<sup>6</sup>

The problem with this approach is that it will usually be too risk-averse. Indeed, it is not hard to see that estimating the model in Eqs. (4a) and (4b) as it stands with the combined data from the panel and the fresh cross section produces exactly the same estimate of  $\boldsymbol{\beta}$  as if we used the fresh cross section alone to estimate Eq. (4a). Any differences between the observed relationships in the cross section and the panel are captured in the vector  $\boldsymbol{\alpha}$  and interpreted as panel effects. In effect, this approach allocates *all* of the panel data to estimating panel effects and *none* to estimating the substantive effects of interest.

Clearly, more fruitful use of panel data in cross-sectional analyses will require some simplifying assumptions about the nature of potential panel effects. Here, I consider two possible simplifying strategies, each of which represents a potentially attractive middle-ground between the very general treatment of panel effects implied by Eq. (3) and the laissez-faire approach of simply ignoring the possibility of panel effects altogether.

---

<sup>6</sup>If the random components of substantive responses ( $\boldsymbol{\varepsilon}_i$ ) and panel effects ( $\boldsymbol{\nu}_i$ ) are assumed to be uncorrelated, as seems natural, the model in Eqs. (4a) and (4b) implies that the overall disturbance variance must be greater for panel respondents than for cross-section respondents. As it happens, this is not always the case in the analyses presented below, which suggests that a model more complicated than the one in Eqs. (4a) and (4b) may be necessary to capture more subtle aspects of panel conditioning.

#### 4.1 A “Semipooled” Model

One approach, which I refer to as a *semipooled* strategy, involves pooling the panel and cross-sectional data, but with some allowance for possible panel effects. This approach imposes a priori restrictions on the very general treatment of panel effects implied by Eq. (3) but does not ignore the possibility of panel effects altogether. The Heckman model of selection bias is one familiar example of a semipooled model, in the sense that all of the potential panel effects are captured (or not) in a single summary variable (the inverse Mills ratio of the fitted value from a first-stage probit analysis) representing the expected value of the disturbance term for surviving panel respondents.

Here, I consider a simpler semipooled strategy—including a dummy variable for panel respondents as an additional explanatory variable in each regression analysis. The coefficients on this dummy variable will capture any differences between the mean values of each dependent variable in the panel and cross section that cannot be accounted for by differences in the values of the other explanatory variables, but will not allow for the possibility of differential panel effects among panel respondents with different characteristics.

Whether this particular semipooled strategy will be better or worse than any other depends, of course, on the nature of the panel effects actually present in the data. For example, a Heckman-style correction may be preferable to a simple dummy variable for panel observations *if* most of the panel effects in any particular instance are due to panel attrition rather than to panel conditioning and *if* the specific functional form imposed in the first stage accurately reflects the actual process of panel attrition. If these conditions are not satisfied, the simple dummy variable strategy proposed here may be preferable to the more complicated Heckman-style correction procedure.

#### 4.2 A “Fractionally Pooled” Model

A second approach, which I refer to as a *fractionally pooled* strategy, involves pooling the panel and cross-section data, but with the panel data discounted by some fixed proportion to reflect a prior belief that they are generated by a somewhat different causal process than the one of substantive interest—in the language of Eq. (4b), by the parameter vector  $(\beta + \alpha)$ , representing substantive effects plus panel effects, rather than by the parameter vector  $\beta$ , representing substantive effects alone. If the panel effects are small in magnitude in comparison with the other sources of error in parameter estimation from the panel data, the appropriate discount rate will be modest; however, if the panel effects are likely to be substantial, then the panel data should be discounted heavily by comparison with the cross-section data.<sup>7</sup>

The main uncertainty in the application of fractional pooling in the present context is how to choose an appropriate discount factor to reflect the potential biases resulting from panel effects. Intuition or previous experience with more or less similar data and models may suggest an appropriate discount factor, or a range of possible discount factors may be employed in systematic sensitivity testing. Alternatively, data from the panel and a parallel cross section may be combined to estimate an appropriate discount factor in the spirit of “empirical Bayes” estimation.

<sup>7</sup>Bartels (1996) described the rationale for fractional pooling in a broad range of cases of potential parameter heterogeneity and, also, presented a Bayesian interpretation of the resulting parameter estimates and a series of empirical examples.

The primary virtue of fractional pooling in this context is that it provides a way to deal with situations where panel effects cannot be directly estimated, either because no data are available from a parallel fresh cross section or because the available data will not support the estimation of a suitably general model of panel effects. In the latter situation, the analyst's choice will be between alternative simplifications, and fractional pooling may be preferable to a semipooled model or some other more or less ad hoc approach. In the former situation, the analyst's choice will be between doing something and doing nothing, and doing *something* to allow for the potential impact of panel biases may be preferable to doing *nothing*, even if the precise value of an appropriate discount factor must be more or less uncertain.

## 5 Assessing the Magnitude of Panel Effects in NES Surveys

Having introduced a general model of panel effects as well as two simpler alternative models, my aim in this section is to use these models to characterize the magnitude of panel biases in recent NES surveys. Along the way, I also hope to shed some light on the relative utility of these alternative models in dealing with panel effects of the sort considered here.

The first step in my analysis is simply to examine the nature and magnitude of differences in some potentially relevant demographic characteristics, attitudes, and behavior between cross-section and panel respondents in the 1992 and 1996 NES surveys. To that end, Table 1

**Table 1** Comparison of cross-section and panel respondents, 1992 and 1996 NES surveys

	1992			1996		
	<i>X-section</i>	<i>Panel</i>	<i>Diff.</i>	<i>X-section</i>	<i>Panel</i>	<i>Diff.</i>
Age (years)	44.6	46.7	+2.1 (0.7)	43.9	48.6	+4.7 (1.0)
Education (years)	13.2	12.5	-0.7 (0.1)	13.2	13.4	+0.2 (0.2)
Income (percentile)	53.0	47.5	-5.5 (1.1)	46.4	50.4	+4.0 (1.5)
Black (%)	12.7	12.9	+0.2 (1.3)	14.6	11.3	-3.3 (1.9)
Female (%)	53.2	53.6	+0.4 (2.0)	60.6	53.6	-7.0 (2.8)
Married (%)	53.4	55.3	+1.9 (2.0)	50.5	55.4	+4.9 (2.8)
Conservative ideology (-1 to 1)	0.047	0.053	+0.006 (0.016)	0.043	0.097	+0.054 (0.023)
Republican partisanship (0 to 1)	-0.070	-0.119	-0.049 (0.027)	-0.137	-0.097	+0.040 (0.040)
Trust in government (0 to 1)	0.440	0.444	+0.004 (0.007)	0.436	0.433	-0.003 (0.010)
Follow public affairs (0 to 1)	0.623	0.600	-0.023 (0.013)	0.539	0.595	+0.056 (0.019)
Campaign interest (0 to 1)	0.613	0.603	-0.010 (0.014)	0.479	0.532	+0.053 (0.020)
Presidential turnout (%)	75.2	72.7	-2.5 (1.9)	68.8	75.3	+6.5 (2.7)
Presidential vote (% Rep.)	43.4	40.0	-3.4 (2.7)	36.4	43.3	+6.9 (3.9)



summarizes a variety of differences in the composition of the cross-section and panel samples separately for each of those two surveys.

The most striking aspect of the comparisons presented in Table 1 is the difference between the 2 election years. The panel respondents in 1996 were older, richer, whiter, more interested in politics, more conservative, and more Republican than respondents in the 1996 fresh cross section—all differences consistent with standard suppositions about the effects of panel attrition and panel conditioning. In contrast, the panel respondents in 1992 were actually poorer, less well educated, less interested in politics, and less Republican than respondents in the 1992 fresh cross section.<sup>8</sup> The only statistically “significant” difference between cross-section and panel respondents that appears consistently in both years is with respect to age, and that difference is largely a mechanical function of the aging of panel respondents between the time they first entered the sample and the time they were reinterviewed.

The second step in my analysis is to assess the impact of panel attrition and panel conditioning on parameter estimates from a variety of typical regression analyses. To that end, I used data from the 1992 and 1996 NES surveys to estimate a total of 88 distinct parameters in 12 different regression and probit analyses. The dependent variables in these analyses include ratings of Bill Clinton’s “morality”; correct relative placements of the presidential candidates on the NES’s seven-point ideology scale; perceptions of whether the nation’s economy had gotten better, gotten worse, or stayed the same over the past year; interest in the campaign; reported turnout (measured after the election); and reported presidential vote (also measured after the election). Each of these six analyses is repeated in identical form using the 1992 and 1996 NES data. Taken together, these 12 analyses and 88 regression coefficients are intended to be roughly representative of a broad range of typical uses of NES data.

The results of all 12 analyses are reported in Tables A1 through A12, available at the *Political Analysis* website. For purposes of illustration, Table 2 reproduces the results for my analysis of campaign interest in 1992. The table presents results for three distinct model specifications:

- (1) an “unconstrained model” with separate estimates for the effect of each explanatory variable among cross-section and panel respondents;
- (2) a “semipooled model” with a single set of estimates for the effects of the explanatory variables, but with an additional indicator variable distinguishing panel respondents; and
- (3) a “pooled model” with a single set of estimates for the effects of the explanatory variables and no distinction between cross-section and panel respondents.

Columns 2–4 in Table 2 present the results for the unconstrained model—separate sets of parameter estimates and goodness-of-fit statistics for the cross-section and panel respondents, plus a column indicating the differences between these estimates. The fifth and sixth columns in the table present the results for the semipooled and pooled models, respectively.

One way to gauge the “significance” of the differences between cross-section and panel results in Table 2 is through a formal likelihood-ratio test of the pooled specification in

<sup>8</sup>These counterintuitive differences may reflect the fact that the NES sampling frame changed between 1990, when the 1992 panel respondents were originally interviewed, and 1992. Thus, comparisons between the panel and cross-section respondents in 1992 require an additional assumption that this change in the sampling frame did not itself produce systematic differences in responses beyond those captured by observed differences in respondents’ demographic characteristics.

**Table 2** Campaign interest, 1992: Regression parameter estimates, with standard errors in parentheses (preelection respondents)

	<i>X-section</i>	<i>Panel</i>	<i>Difference</i>	<i>Semipooled</i>	<i>Pooled</i>
Age	0.00307 (0.00060)	0.00245 (0.00054)	-0.00062 (0.00081)	0.00272 (0.00040)	0.00272 (0.00040)
Education	0.0440 (0.0043)	0.0231 (0.0037)	-0.0209 (0.0057)	0.0324 (0.0028)	0.0324 (0.0028)
Income	0.058 (0.041)	0.071 (0.037)	0.013 (0.055)	0.065 (0.028)	0.065 (0.027)
Black	-0.012 (0.031)	-0.028 (0.028)	-0.016 (0.041)	-0.019 (0.020)	-0.019 (0.020)
Female	-0.014 (0.020)	-0.055 (0.018)	-0.041 (0.027)	-0.037 (0.014)	-0.037 (0.014)
Partisan strength	0.204 (0.030)	0.214 (0.027)	0.010 (0.041)	0.209 (0.020)	0.209 (0.020)
Days before election	-0.00246 (0.00053)	-0.00144 (0.00050)	0.00102 (0.00073)	-0.00200 (0.00036)	-0.00200 (0.00036)
Intercept	0.362 (0.042)	0.391 (0.038)	0.029 (0.057)	0.383 (0.029)	0.384 (0.028)
Panel indicator	—	—	—	0.002 (0.014)	—
Adjusted $R^2$	0.17	0.11	0.14	0.13	0.13
SE of regression	0.335	0.331	0.333	0.334	0.333
$N$	1126	1359	2485	2485	2485

which those effects are ignored (presented in the sixth column in the table) against the unconstrained specification in which they are not (presented in columns 2–4). The results of these tests for each of the regression models in Tables A1 through A12 (six distinct models in each of the 2 election years) are presented in the form of  $p$  values in the second column in Table 3.

In general, these tests provide rather little evidence of “significant” panel biases. Only 2 of the 12  $p$  values are less than 0.10 (whereas 1 would be expected simply by chance), and only 1 is less than 0.05. To the extent that these results are typical, they suggest that the null hypothesis of no panel effects will only rarely be rejected by an analyst employing conventional hypothesis-testing technology. It may be worth noting that neither of the 2 cases among the 12 in Table 3 for which the null hypothesis of no panel effects can be “rejected” at the 0.10 level—the analyses of Campaign Interest and Presidential Turnout based upon data from the 1992 survey—correspond to variables with “significant” marginal differences between panel and cross-section respondents in Table 1; the associated  $t$  values for the differences between cross-section and panel mean values are  $-0.7$  and  $-1.3$ , respectively. On the other hand, several of the variables that *did* show “significant” marginal differences in Tables 1 (including Campaign Interest and Presidential Turnout in 1996) produced little or no evidence of parameter bias in Table 3. This discrepancy highlights the fact that statistically “significant” biases in panel composition may not produce “significant” biases in parameter estimates, while apparently representative panel samples *may* produce such biases.

The remaining columns in Table 3 present  $p$  values for similar likelihood-ratio tests involving the semipooled models with dummy variables for panel respondents. The fourth


**Table 3** Tests for panel effects in 1992 and 1996 NES surveys:  $p$  values for likelihood-ratio tests based upon pooled, semipooled, and unconstrained parameter estimates in Tables A1–A12

	<i>Pooled vs. unconstrained model</i>	<i>Pooled vs. semipooled model</i>	<i>Semipooled vs. unconstrained model</i>
1992			
Clinton “morality” ratings	0.382	0.448	0.339
Correct relative ideological placements	0.598	0.086	0.898
Economic perceptions	0.507	0.359	0.489
Campaign interest	0.016	0.887	0.009
Presidential turnout	0.065	0.786	0.038
Presidential vote	0.886	0.332	0.916
1996			
Clinton “morality” ratings	0.475	0.466	0.427
Correct relative ideological placements	0.303	0.708	0.216
Economic perceptions	0.552	0.247	0.597
Campaign interest	0.239	0.164	0.295
Presidential turnout	0.539	0.439	0.488
Presidential vote	0.261	0.051	0.534

column reports  $p$  values for likelihood-ratio tests of the semipooled model versus the unconstrained model. In effect, these test statistics index the inferential cost of reducing the general model of panel effects in Eq. (3) to the much simpler model with a single dummy variable for panel respondents. Two of these 12  $p$  values are less than 0.05, suggesting that for these 2 analyses the more elaborate unconstrained model captures statistically “significant” panel effects that are not captured by the semipooled model. However, the range of  $p$  values for all 12 analyses provides rather little evidence that the more elaborate model of panel effects is likely to be necessary in typical applications.

On the other hand, the third column in Table 3 reports  $p$  values for likelihood-ratio tests of the pooled model with no allowance for panel effects versus the semipooled model with a single coefficient to capture panel effects. These test statistics index the inferential benefit of including a dummy variable for panel respondents rather than ignoring completely the panel structure of the data. None of the 12  $p$  values is less than 0.05, and only 2 are less than 0.10. Thus, while there is rather little indication in Table 3 that the more general model of panel effects in Eq. (3) will often be worth estimating in preference to the semipooled model, there is equally little indication that the semipooled model will often be worth estimating in preference to a fully pooled model with no distinction between panel and cross-section respondents. In most cases, with data like these, analysts employing conventional hypothesis tests will be led simply to ignore the panel structure of the data.

Another, somewhat more concrete way to assess the differences between cross-sectional and panel parameter estimates in Tables A1 through A12 is by examining the distribution of  $t$  statistics for the differences between these estimates. If panel attrition and panel conditioning had no effect, the parameter estimates for cross-section and panel respondents would be identical within sampling error, and the distribution of  $t$  statistics for the differences between them would have zero mean and unit variance. If panel attrition and/or panel conditioning produced significant differences between the cross-section and the panel parameter estimates, the result would be to produce larger (in absolute value)  $t$  statistics



**Fig 1** Distribution of  $t$  statistics for estimated panel effects in NES surveys.

than would be expected simply on the basis of sampling error. Thus, the extent of departure from the null distribution provides a useful gauge of the magnitude of panel effects in these analyses.

The distribution of the 88  $t$  statistics for the differences between cross-section and panel parameter estimates in Tables A1 through A12 is displayed in Fig. 1, in the upper panel as a histogram and in the lower panel as a quantile-normal plot.<sup>9</sup> It is clear in both panels that there are more extreme  $t$  statistics in the observed distribution than would be expected simply on the basis of sampling variability—but not many more. For example, 32 of the 88  $t$  statistics in Fig. 1 are larger than 1 in absolute value, whereas 28 would be expected by chance, and 5 are larger than 2 in absolute value, whereas 4 would be expected by chance. The single dramatic outlier in the observed distribution of  $t$  statistics is the value of  $-3.6$  for

<sup>9</sup>To detect any systematic tendency for the panel data to understate or overstate the corresponding effects in the fresh cross section, the sign of the relevant  $t$  statistic is reversed in Fig. 1 when the cross-sectional parameter estimate is negative; this adjustment has no effect on the subsequent calculations, which are based upon the *squared*  $t$  statistics. As it turns out, there may be a slight tendency for the panel coefficients to underestimate the corresponding cross-sectional effects: the average value of the 88  $t$  statistics (with adjusted signs) is  $-0.153$ , with a standard error of 0.117.

**Table 4** Relative efficiency of panel data for cross-sectional analysis, 1992 and 1996 NES surveys

	<i>Average estimated panel effect (root mean squared <math>t</math>)</i>	<i>Average estimated MSE/variance for panel data (<math>1/\lambda^2</math>)</i>	<i>Implied relative value of panel data (<math>\lambda</math>)</i>
1992			
Clinton "morality" ratings	1.07	1.323	0.869
Correct relative ideological placements	0.45	-0.761	— <sup>a</sup>
Economic perceptions	0.98	0.903	1.052
Campaign interest	1.53	4.115	0.493
Presidential turnout	1.59	4.404	0.477
Presidential vote	0.66	-0.252	—
Total	1.13	1.640	0.781
1996			
Clinton "morality" ratings	1.11	1.473	0.824
Correct relative ideological placements	1.01	0.167	2.448
Economic perceptions	0.92	0.621	1.269
Campaign interest	1.14	3.172	0.561
Presidential turnout	1.01	2.597	0.620
Presidential vote	1.16	1.651	0.778
Total	1.07	1.621	0.785
Overall total	1.10	1.631	0.783

<sup>a</sup> Out-of-bounds estimate.

the difference reported in Table 2 between cross-section and panel estimates of the impact of education on campaign interest in 1992. In general, however, there is no evidence in Fig. 1 of widespread biases in parameter estimates due to panel effects.

The same distribution of  $t$  statistics for panel effects presented in Fig. 1 is summarized in a different way in the second column in Table 4, which reports the root mean squared  $t$  statistic for panel effects in each of the 12 regression models presented in Tables A1 through A12. Here, the extent to which the observed root mean squared  $t$  statistics exceed 1.0 reflects the extent to which panel effects, rather than mere sampling variability, produce deviations from the corresponding cross-section estimates.

For 2 of the 12 regression models I consider—the model of campaign interest in 1992 presented in Table 2 and the model of presidential turnout in 1992—the root mean squared  $t$  statistics are more than 50% larger than would be expected on the basis of sampling variability alone. In two other cases the observed  $t$  statistics for panel effects are significantly *smaller* in magnitude than expected, with average values of 0.45 and 0.66; for the other eight regression models, the observed average values are fairly close to the null value of 1.0. The overall averages for the 2 election years (each based upon 44 separate parameter estimates) are quite similar—1.13 in 1992 and 1.07 in 1996—despite the significant differences in the characteristics of panel respondents in the two surveys reported in Table 1.

While these results provide some sense of the prevalence of panel biases in the various regression models considered here, they provide no clear indication of the *magnitude* of the inferential costs associated with those biases. One important conceptual advantage of the "fractional pooling" approach introduced above is that it focuses attention clearly on

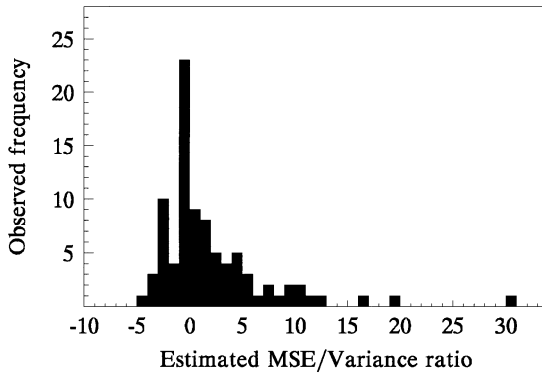
the issue of inferential cost by summarizing the impact of panel effects in a single discount factor to be applied to the panel data. While in typical applications the relevant discount factor must be assigned on a priori grounds, the availability of parallel panel and fresh cross-section data in the cases considered here makes it possible to estimate an appropriate discount factor *ex post*. Such an exercise may be useful in two respects: for summarizing clearly the inferential cost of panel effects in a situation where those costs can be ascertained in some detail and for providing general guidance to analysts using panel data in many other situations where the nature and magnitude of panel biases cannot be directly ascertained.

The key issue in the choice of an appropriate discount factor for panel data is an assessment of the extent to which panel biases inflate the errors in parameter estimates arising simply from sampling variability. The comparisons of parameter estimates from panel and fresh cross-section data presented in Tables A1 through A12 provide the raw material for such an assessment. In particular, the magnitude of the estimated panel biases (represented by the squared *difference* between panel and cross-section estimates for each parameter in these tables) can be compared with the magnitude of sampling variability (represented by the *variance* of each parameter estimate estimated from the panel data alone). The ratio of mean squared error (squared bias plus variance) to variance is inversely related to the weight  $\lambda$  attached to the panel data in the fractional pooling approach: for an appropriate choice of  $\lambda$ , the expected ratio of mean squared error to variance is  $1/\lambda^2$  (Bartels 1996). When panel biases are small in magnitude relative to sampling variance, this ratio will be close to 1, and the appropriate weight  $\lambda$  will also be close to 1. When panel biases are large in magnitude relative to sampling variance, this ratio will be large, and the appropriate weight will be close to 0.

Our aim, then, is to calculate, for each of our 88 parameters, the ratio of total error variance in the panel estimate (that is, squared panel bias plus sampling variability) to sampling variance alone (which reflects sources of uncertainty unrelated to panel effects). However, the estimates of panel bias for each parameter presented in Tables A1 through A12 are just that—*estimates*—and ignoring that fact will significantly overstate the contribution of panel biases to total error variance in the parameter estimates based upon panel data.<sup>10</sup> Fortunately, we also have standard errors for the estimates of panel bias, and we can use those standard errors to correct (in expectation) our calculation of the contribution of panel bias to total error variance for each parameter estimate.

Unfortunately, the individual estimates of panel bias in Tables A1 through A12 are quite imprecise, while the underlying distribution of *true* squared panel biases is left-truncated at 0 and strongly right-skewed. The result is that the distribution of parameter-by-parameter *estimates* of the ratio of mean squared error to variance for the panel data, displayed in Fig. 2, includes a large number of out-of-bounds estimates (that is, *negative* estimates of mean squared error!) along with a few very large estimates (implying squared panel biases of the order of 15 to 30 times as large as the nominal variances of the parameter estimates based upon the panel data). Obviously, ignoring the *negative* estimates while taking the *positive* estimates at face value would bias upward our overall estimate of the impact of panel effects. Thus, my approach here is simply to report the *average* estimated ratios for each of my 12 regression analyses in the third column in Table 4, along with overall averages for each election year and for both years combined.

<sup>10</sup>The expected value of each squared panel bias *estimate* is equal to the *true* squared panel bias plus the variance of the panel bias estimate. The larger the variance of the panel bias estimate (relative to the true bias), the more the squared panel bias *estimate* will tend to overstate the corresponding *true* squared panel bias.



**Fig 2** Distribution of estimated MSE/variance ratios for panel data in NES surveys.

For 2 of the 12 analyses, the result of this calculation is an illogical *negative* estimate of the ratio of mean squared error to variance for the panel data, and in 3 other instances the estimated ratio is positive but less than 1.0, implausibly implying that the panel data in these cases are relatively more valuable than the cross-sectional data for estimating purely cross-sectional relationships. Obviously, these calculations must be taken with a large grain of salt. Thus, so too must the remaining seven cases, in which the average estimated ratios of mean squared error to variance for the panel data range from 1.3 to 4.4, implying weights for the panel data (presented in the fourth column in Table 4) ranging from 0.477 to 0.869.

While it would clearly be a mistake to place too much confidence in these results, they do seem plausible in two key respects. First, the overall estimates of the panel weight  $\lambda$  implied by these calculations—0.781 in 1992 and 0.785 in 1996—suggest that, on average, biases produced by panel attrition and panel conditioning reduced the inferential value of panel data by about 20% across the range of analyses considered here. Second, the estimates suggest that panel biases were concentrated in the analyses which seem most likely on a priori grounds to be sensitive to the potential effects of panel attrition and panel conditioning—the analyses of campaign interest and turnout. The estimated relative value of panel data for the 28 distinct parameters in my analyses of campaign interest and turnout is 0.585, implying an appropriate discount rate of about 40% by comparison with fresh cross-section data. The corresponding estimate of the relative value of panel data for the 60 distinct parameters in my analyses of Clinton “morality,” correct ideological placements, economic perceptions, and vote choice is 0.987, implying an appropriate discount rate of only 1% by comparison with fresh cross-section data. These contrasting results suggest that the inferential costs of panel biases are likely to be quite modest in most analyses of the sort considered here but that the relative value of panel data may be dramatically reduced in analyses where the dependent variable of interest is strongly related to processes of panel attrition or panel conditioning.

## 6 Adjusting for Panel Effects in Dynamic Analyses

So far, my analysis has addressed the situation in which panel effects merely complicate cross-sectional analyses; the implicit assumption has been that analysts have no use for the panel data in their own right. Obviously, if that were the only situation to arise in practice, there would be no reason to gather panel data in the first place. It is the power of panel data to facilitate dynamic analyses that may (or may not) warrant paying the inferential

price associated with panel effects examined in the preceding section. Thus, my aim in this section is to propose an adjustment for panel effects in dynamic analyses.

I assume, as before, that the second wave of the panel is paralleled by a separate fresh cross section.<sup>11</sup> By analogy with Eq. (1), our substantive dynamic model will be a regression model of the form

$$\mathbf{y} = \mathbf{X}\boldsymbol{\gamma} + \mathbf{Z}\boldsymbol{\psi} + \boldsymbol{\delta} \quad (5)$$

where matrix  $\mathbf{Z}$  includes lagged variables (possibly including  $\mathbf{y}_{t-1}$  and  $\mathbf{X}_{t-1}$ ) potentially observable in a previous panel wave. However, we continue to assume that our observed data  $\mathbf{y}^p$  for panel respondents are subject to the panel effects specified in Eq. (3), so that the combined model for cross-section and panel respondents analogous to Eqs. (4a) and (4b) is

$$\mathbf{y} = \mathbf{X}\boldsymbol{\gamma} + \mathbf{Z}\boldsymbol{\psi} + \boldsymbol{\delta} \quad (6a)$$

$$\mathbf{y}^p = \mathbf{X}\boldsymbol{\gamma} + \mathbf{Z}\boldsymbol{\psi} + \boldsymbol{\delta} + (\mathbf{X}\boldsymbol{\alpha} + \boldsymbol{\nu}) \quad (6b)$$

The problem with this formulation is that the cross-section and panel data cannot be readily combined [as they could be in the purely cross-sectional analysis based on Eqs. (4a) and (4b)], because *neither* Eq. (6a) nor Eq. (6b) can be estimated independently. Eq. (6b) suffers from the same identification problem as Eq. (4b)—since (at least some of) the same contemporaneous explanatory variables may have both substantive effects and panel effects, it is impossible to distinguish the parameters  $\boldsymbol{\gamma}$  and  $\boldsymbol{\alpha}$  using the panel data. And Eq. (6a), unlike the corresponding Eq. (4a), suffers from the problem of missing observations on the lagged variables  $\mathbf{Z}$ , since there are no prior observations for precisely those respondents whose contemporaneous data are uncontaminated by panel effects.

My solution to this problem is based upon the observation that the panel effects hypothesized in the dynamic Eq. (6b) are just the same panel effects hypothesized in the cross-sectional Eq. (4b), where it *was* possible to estimate them using a comparison between the panel and the fresh cross-section data. Given consistent estimates of the panel effect parameters  $\boldsymbol{\alpha}$  from analyses like those presented in Tables A1 through A12, we can “panel adjust” the observed values of  $\mathbf{y}^p$  in the panel data by subtracting the expected value of panel effects for each observation, producing

$$\mathbf{y}^p - \mathbf{w} = \mathbf{X}\boldsymbol{\gamma} + \mathbf{Z}\boldsymbol{\psi} + [\boldsymbol{\delta} + (\mathbf{w} - \boldsymbol{\omega})] \quad (7)$$

where

$$\mathbf{w} = \mathbf{X}\boldsymbol{\alpha} \quad (8)$$

is a consistent estimate of the panel bias  $\boldsymbol{\omega}$  in Eq. (3). The “panel-adjusted” values of  $\mathbf{y}$  can then simply be regressed on the contemporaneous and lagged explanatory variables  $\mathbf{X}$  and  $\mathbf{Z}$  to produce consistent estimates of the substantive parameters  $\boldsymbol{\gamma}$  and  $\boldsymbol{\psi}$ .

<sup>11</sup>If the only available data are panel data, there seems to be little hope of constructing any explicit adjustment for potential panel effects. Even in that case, however, it seems desirable in principle to discount the available data to allow for biases in the estimated parameters due to panel effects, just as in the case of “fractional pooling” of panel and cross-section data described above. The only difference is that in this case, all the data will be discounted by the same amount, so the only impact on the results of the analysis will be to generate larger standard errors, reflecting the additional uncertainty attributable to unknown panel biases.



This two-step estimation strategy is just a slightly convoluted variation of Franklin's (1990) two-stage auxiliary instrumental variables estimator, in which the auxiliary instrumental variable  $\mathbf{w}$  is constructed jointly from the main (panel) data set and an auxiliary (cross-section) data set rather than from the auxiliary data set alone, and in which the coefficient on this auxiliary instrumental variable is constrained to equal unity by the identity of hypothesized panel effects in Eqs. (4b) and (6b).

Simply regressing the panel-adjusted dependent variable on the explanatory variables in the second-stage regression will underestimate the standard errors of the second-stage parameter estimates, since the ordinary regression calculations take no account of sampling error in the first-stage estimated panel effects.<sup>12</sup> The easiest way to deal with this complication is to resample the entire estimation procedure, including both the first-stage estimation of panel effects and the second-stage regression. The panel-adjusted parameter estimates and standard errors reported below were calculated using a jackknife estimator in which 10 random subsets of the original data were omitted 1 at a time.<sup>13</sup> Due to sampling variability in the first-stage estimates of panel effects, the resulting second-stage standard errors are, on average, 10 to 20% larger than those from the corresponding ordinary regression models.

Obviously, any adjustment for panel effects can have relatively little impact when those effects are themselves negligible in magnitude, as they are in most of the analyses reported in Tables A1 through A12. Thus, for purposes of illustration, I focus here on the single most egregious case of apparent panel effects in those analyses, the model of campaign interest in 1992 reported in Table 2. That analysis suggested that the impact of education on interest in the campaign was only about half as great for panel respondents as for those in the fresh cross section ( $t = -3.6$ ), that females in the panel were relatively less interested in the campaign ( $t = -1.5$ ), and that interest increased less perceptibly over the 2 months between Labor Day and Election Day among panel respondents ( $t = 1.4$ ). The root mean squared  $t$  statistic for all eight parameter estimates in Table 2 (reported in Table 4) is 1.53, and the  $p$  value for the general likelihood-ratio test of no panel effects (reported in Table 3) is 0.016, much smaller than for any of the 11 other models considered here.

To gauge the impact of these apparent panel effects, I estimated three dynamic models of interest in the 1992 presidential campaign using 1990 and 1992 data for the 1992 panel respondents. The simplest of these includes interest in the 1990 midterm campaign as the sole explanatory variable; the second model adds 1990 measures of partisan strength and turnout as additional explanatory variables; the third model also adds the three variables associated with significant panel effects in Table 2—education, sex, and the date of the 1992 preelection interview. Each of these models is estimated once using the unadjusted campaign interest variable for the 1992 panel respondents and a second time using “panel-adjusted” values of 1992 campaign interest derived from the comparison of panel and fresh cross-section respondents in Table 2. The results of these analyses are presented in Table 5.

The only important difference between the ordinary regression and the “panel-adjusted” results for the first two models considered in Table 5 is that, in each case, correcting for panel effects reduces the intercept level of campaign interest by about 0.05 on the zero-to-one scale. This difference suggests, quite plausibly, that campaign interest was artificially

<sup>12</sup>This is a standard complication in two-stage estimators, including Franklin's (1990) two-stage auxiliary instrumental variables estimator, Heckman's (1976) two-stage estimator for selection bias, and classical two-stage least squares.

<sup>13</sup>Miller (1974) provided a standard survey of jackknife estimators. Efron and Gong (1983) and Mooney and Duval (1993, pp. 22–27) compared the logic and inferential strengths and weaknesses of jackknife estimators with those of bootstrap and other related resampling estimators.

**Table 5** Dynamic analyses of campaign interest with and without panel adjustments, 1992 NES survey: Regression parameter estimates, with standard errors in parentheses;  $N = 1359$  (panel respondents)

	<i>Ordinary regression</i>	<i>Panel adjusted<sup>a</sup></i>	<i>Ordinary regression</i>	<i>Panel adjusted</i>	<i>Ordinary regression</i>	<i>Panel adjusted</i>
1990 campaign interest	0.436 (0.023)	0.448 (0.016)	0.381 (0.027)	0.388 (0.021)	0.370 (0.027)	0.358 (0.016)
1990 partisan strength			0.083 (0.027)	0.070 (0.024)	0.088 (0.027)	0.078 (0.021)
1990 turnout			0.054 (0.019)	0.066 (0.026)	0.044 (0.019)	0.048 (0.024)
Education (years)					0.0137 (0.0031)	0.0285 (0.0051)
Female					-0.045 (0.017)	-0.005 (0.018)
Days before election					-0.00167 (0.00046)	-0.00346 (0.00093)
Intercept	0.404 (0.014)	0.353 (0.024)	0.352 (0.019)	0.305 (0.026)	0.425 (0.025)	0.409 (0.024)
Adjusted $R^2$	0.20	0.21	0.21	0.22	0.23	0.29
SE of regression	0.313	0.324	0.311	0.322	0.307	0.307

<sup>a</sup> “Panel adjusted” entries are jackknife parameter estimates and standard errors based on the estimated panel effects in Table 2.

inflated in the panel—despite the fact (reported in Table 1) that the 1992 panel respondents were, on average, slightly *less* interested in the campaign than the corresponding fresh cross-section respondents. In this respect, at least, the panel adjustment seems to make a real, and reasonable, difference.

In the third model reported in Table 5, most of this difference is absorbed in the estimated effects of education, sex, and interview date, all of which are substantially different in the ordinary regression and “panel-adjusted” analyses. These differences clearly reflect the attribution of panel effects based upon the comparison of panel and cross-section respondents in Table 2. The effect of the panel adjustment is to roughly double the apparent impact of education and interview date and to erase the apparent difference in campaign interest between female and male respondents.

## 7 Discussion

The results presented in Table 5 provide some indication that the “panel-adjusted” estimation strategy proposed here may be of value in situations where dynamic analyses would otherwise be plagued by significant panel biases. However, it is worth bearing in mind that the analysis of campaign interest presented in Tables 2 and 5 represents the most dramatic example of apparent panel biases among the twelve analyses considered here. There is no reason to believe that my adjustment for panel effects would have as much impact in more typical analyses.

Indeed, the more general implication of my research is that such significant panel biases are likely to be relatively rare in practice, at least for data and models of the sort considered here. The results presented in Table 4 for a range of typical analyses using NES survey

data suggest that panel effects probably reduced the inferential value of panel data (by comparison with fresh cross-section data) by an average of about 20%. But this overall average is an unduly pessimistic summary of the inferential cost of using panel data in cross-sectional analyses of the sort considered here, since it obscures a sharp distinction between the four analyses of campaign interest and turnout, where the appropriate discount rate appears to be of the order of 40%, and the eight analyses of Clinton traits, ideological placements, economic perceptions, and vote choices, where the estimated discount rate is essentially zero.

Of course, any estimate of the inferential cost of using panel data must be weighed against the benefits of panel data summarized in Section 1. While I have made no corresponding effort to quantify the advantages of panel data, it seems clear that, in most applications along the lines of those presented here, the inferential advantages of panel data will easily outweigh the rather modest inferential cost of panel effects implied by my analysis. In these cases, there is little to be lost and much to be gained by making panel interviews an increasingly common component of survey research.

The problem, obviously, is to judge which data sets and analyses are likely to be subject to more substantial panel biases. My analyses of campaign interest and turnout suggest that panel effects may create significant difficulties even for analyses of NES data when our dependent variables of interest are especially likely to be related to survey participation. In such cases, it seems prudent to take some account of potential panel biases by discounting panel data, by explicitly estimating panel effects, or both.

Surveys with different substantive foci, lower reinterview rates, shorter time spans between interviews, or more intensive conditioning experiences than the NES surveys considered here may be subject to more frequent and more serious panel biases. The only way to tell is by careful comparison of panel and fresh cross-section data along the lines suggested here. Thus, an important practical implication of my analysis is that panel surveys should routinely include parallel fresh cross-section components, to provide a solid basis for assessing and, if necessary, adjusting for biases arising from panel attrition and panel conditioning.

## References

- Achen, Christopher H. 1975. "Mass Political Attitudes and the Survey Response." *American Political Science Review* 69:1218–1223.
- Achen, Christopher H. 1983. "Toward Theories of Political Data." In *Political Science: The State of the Discipline*, ed. Ada W. Finifter. Washington, DC: American Political Science Association.
- Achen, Christopher H. 1986. *The Statistical Analysis of Quasi-Experiments*. Berkeley: University of California Press.
- Bartels, Larry M. 1993. "Messages Received: The Political Impact of Media Exposure." *American Political Science Review* 87:267–285.
- Bartels, Larry M. 1996. "Pooling Disparate Observations." *American Journal of Political Science* 40:905–942.
- Bartels, Larry M. 1998. "Partisanship and Voting Behavior, 1952–1996," Unpublished manuscript. Princeton, NJ: Department of Politics and Woodrow Wilson School of Public and International Affairs, Princeton University.
- Beckett, Sean, William Gould, Lee Lillard, and Finis Welch. 1988. "The Panel Study of Income Dynamics After Fourteen Years: An Evaluation." *Journal of Labor Economics* 6:423–445.
- Berelson, Bernard R., Paul F. Lazarsfeld, and William N. McPhee. 1954. *Voting: A Study of Opinion Formation in a Presidential Campaign*. Chicago: University of Chicago Press.
- Brehm, John. 1993. *The Phantom Respondents: Opinion Surveys and Political Representation*. Ann Arbor: University of Michigan Press.
- Efron, Bradley, and Gail Gong. 1983. "A Leisurely Look at the Bootstrap, the Jackknife and Cross-Validation." *American Statistician* 37:36–48.
- Finkel, Steven E. 1993. "Reexamining the 'Minimal Effects' Model in Recent Presidential Campaigns." *Journal of Politics* 55:1–21.

- Franklin, Charles H. 1990. "Estimation Across Data Sets: Two-Stage Auxiliary Instrumental Variables Estimation (2SAIV)." *Political Analysis* 1:1–24.
- Heckman, James J. 1976. "The Common Structure of Statistical Models of Truncation, Sample Selection, and Limited Dependent Variables and a Simple Estimator for Such Models." *Annals of Economic and Social Measurement* 5:475–492.
- Heckman, James J. 1979. "Sample Selection Bias as a Specification Error." *Econometrica* 47:153–161.
- Johnston, Richard, André Blais, Henry E. Brady, and Jean Crête. 1992. *Letting the People Decide: Dynamics of a Canadian Election*. Montreal: McGill–Queen's University Press.
- Just, Marion R., Ann N. Crigler, Dean E. Alger, Timothy E. Cook, Mantague Kern, and Darrell M. West. 1996. *Crosstalk: Citizens, Candidates, and the Media in a Presidential Campaign*. Chicago: University of Chicago Press.
- Lazarsfeld, Paul F., Bernard Berelson, and Hazel Gaudet. 1948. *The People's Choice*. New York: Columbia University Press.
- Markus, Gregory B. 1982. "Political Attitudes During an Election Year: A Report on the 1980 NES Panel Study." *American Political Science Review* 76:538–560.
- Miller, R. G. 1974. "The Jackknife: A Review." *Biometrika* 61:1–15.
- Mooney, Christopher Z., and Robert D. Duval. 1993. *Bootstrapping: A Nonparametric Approach to Statistical Inference*. Newbury Park, CA: Sage.
- Patterson, Thomas E. 1980. *The Mass Media Election: How Americans Choose Their President*. New York: Praeger.
- Patterson, Thomas E., and Robert D. McClure. 1976. *The Unseeing Eye: The Myth of Television Power in National Elections*. New York: G. P. Putnam's Sons.
- Traugott, Santa. 1996. "Effects of Panel Attrition and Forms Assignment in the 1995 Pilot Study." Technical Report 52, The National Election Studies. Ann Arbor: Institute for Social Research, University of Michigan.
- Wiley, David E., and James A. Wiley. 1970. "The Estimation of Measurement Error in Panel Data." *American Sociological Review* 35:112–117.
- Zaller, John, and Mark Hunt. 1995. "The Rise and Fall of Candidate Perot: The Outsider Versus the Political System—Part II." *Political Communication* 12:97–123.