Political Analysis Advance Access published October 21, 2013

Political Analysis (2013) pp. 1–18
doi:10.1093/pan/mpt020

Looking Beyond Demographics: Panel Attrition in the ANES and GSS

Laura Lazarus Frankel
Department of Political Science, Duke University, 140 Science Drive, Box 90204, Durham, NC 27708

D. Sunshine Hillygus
Department of Political Science, Duke University, 140 Science Drive, Box 90204, Durham, NC 27708
e-mail: hillygus@duke.edu (corresponding author)

Edited by Jasjeet Sekhon

Longitudinal or panel surveys offer unique benefits for social science research, but they typically suffer from attrition, which reduces sample size and can result in biased inferences. Previous research tends to focus on the demographic predictors of attrition, conceptualizing attrition propensity as a stable, individual-level characteristic—some individuals (e.g., young, poor, residentially mobile) are more likely to drop out of a study than others. We argue that panel attrition reflects both the characteristics of the individual respondent as well as her survey experience, a factor shaped by the design and implementation features of the study. In this article, we examine and compare the predictors of panel attrition in the 2008–2009 American National Election Study, an online panel, and the 2006–2010 General Social Survey, a face-to-face panel. In both cases, survey experience variables are predictive of panel attrition above and beyond the standard demographic predictors, but the particular measures of relevance differ across the two surveys. The findings inform statistical corrections for panel attrition bias and provide study design insights for future panel data collections.

There has been increasing use of longitudinal surveys in recent decades throughout the social sciences, medicine, government, and business. The National Science Foundation's three major recurring surveys—the American National Election Study (ANES), the General Social Survey (GSS), and the Panel Survey on Income Dynamics (PSID)—now all contain a multi-wave panel component. Panel surveys offer unique benefits for observational studies; because they track the same respondents over time, they offer more causal leverage than a cross-sectional survey (Bartels 2006). And yet, this advantage comes with a downside in the form of panel attrition.[1] That is, not all respondents in Wave 1 will continue to participate in subsequent waves. This attrition disturbs the sampling design, reduces effective sample sizes, and, if correlated with the outcomes of interest, can bias substantive results (Groves and Couper 1998).

Although there exist numerous ways to account and correct for attrition bias, the first step for any of these methods is to better understand the attrition process. In other words, who drops out and why? Previous research predicting panel attrition focused primarily on the demographic characteristics of the respondents; those most mobile and difficult to reach—the young, minorities, and

Authors' Note: Laura Lazarus Frankel is a Ph.D. candidate in political science at Duke University. D. Sunshine Hillygus is an Associate Professor of Political Science at Duke University. An earlier version of this paper was presented at the Annual Meeting of the American Association of Public Opinion Research. Replication data are available on the *Political Analysis* Dataverse.

[1]Panel conditioning is another concern of panel studies, but is not one considered here. See Bartels (1999) and Clinton (2001) for discussions of panel conditioning in the American National Election Study and Knowledge Networks panels, respectively.

those living alone—are the most likely to drop out of a panel (Behr, Bellgardt, and Rendtel 2005). And to the extent that attrition is addressed using post-stratification weights or controls within a statistical model, these demographics are assumed to fully account for the attrition process. Yet, the broader literature on nonresponse has shown that respondents' propensity to answer a question or questionnaire depends on their survey experience, a factor shaped by the design and implementation features of the study (Frankel 1980; Sharp and Frankel 1983; Loosveldt and Carton 2001).

In this article, we move beyond basic demographics to examine a broader range of predictors of panel attrition in the 2008–2009 ANES and the 2006–2010 GSS Panel Studies. As the cornerstone surveys of political science and sociology, it is especially important to identify the predictors of attrition in their recent panel additions. Because these surveys differed in survey mode—the ANES panel was a self-administered online survey, whereas the GSS was face-to-face—the comparison provides additional insights about the various ways survey experience can be shaped by survey design.

Our analysis finds that, in the GSS, respondents interviewed by more experienced interviewers are less likely to attrite. Additionally, respondents rated as more cooperative by their interviewer at Wave 1 are less likely to attrite. In the ANES panel, where the online mode means there is no interviewer, we find that interaction with the equipment matters. Individuals who required the use of the clunky MSN TV2 interface because they lacked home Internet access were more likely to attrite, as are those who reported a negative survey experience in response to an open-ended question inviting comments about the Wave 1 survey. Finally, the results suggest that an individual's motivation for participating in the study in the first place is related to his or her probability of attriting, with those motivated strictly by the monetary incentives more likely to drop out.

1 Background

Although there is considerable recognition of the potential problems of panel attrition, the way it is reported and handled varies widely. It is often ignored entirely.[2] To be sure, some dedicated research has found that panel attrition did not substantially bias many of the results in a number of "gold standard" surveys (Fitzgerald, Gottschalk, and Moffitt 1998; Bartels 1999). For example, in an analysis of the 1992–1996 ANES panel, Bartels (1999) concludes that there is "rather little evidence of 'significant' panel biases" (10).[3] Yet, there are a number of reasons to believe it is the right time for reassessment. First, overall survey response rates have significantly decreased in the years since those studies (Keeter et al. 2006), so we might expect changes in attrition rates as well.[4] Second, panel studies have become more common in the social sciences, in large part reflecting the emergence of online panels like GfK Knowledge Networks (formerly Knowledge Networks) and YouGov (formerly Polimetrix).[5] It is well known that survey mode has implications for unit and item nonresponse, so we might also expect differences in the rate and predictors of attrition in an online mode (de Leeuw 1992; Dillman et al. 2009). The ANES still uses the face-to-face mode in its traditional pre/post-election survey, but a multi-wave panel component added in 2008 had an

[2]In fact, it can be difficult even to determine the extent of panel attrition in many studies. Although initial response rates are routinely provided, there is no standardized reporting of respondent retention (Atrostic et al. 2001). For instance, changes in sample sizes across waves could reflect changes in panel eligibility, sampling design, or panel attrition. The distinction between loss of sample due to ineligibility and attrition of eligible cases is important since they can reflect very different processes. Because the ANES defines eligibility differently than many other panel studies, we cannot distinguish if a respondent dropped out, died, or moved somewhere without Internet access.

[3]He does note that panel surveys with different design features could well suffer from more problematic panel biases and finds that remaining panelists have higher levels of political interest and turnout.

[4]The expected impact of initial nonresponse on attrition is unclear. On the one hand, we might expect panel attrition to have worsened given the additional respondent burden of multiple interviews. On the other hand, lower initial response rates might weed out those most likely to attrite.

[5]These panels involve the recruitment and maintenance of a cohort of respondents for surveys that are fielded online. Online recruitment methods vary from simply asking respondents to opt in (nonprobability) to probability sampling that uses RDD or address-based methods in order to contact respondents and recruit them for participation in the panel. Once in the panel, respondents participate in a wide array of surveys available to the panel participants. Although the jury is still out on best sampling, recruitment, and maintenance procedures for online panels (see AAPOR 2010 Report on online panels), as well as new sources of error and bias they may create, it is clear that their use continues to increase.

online panel design fielded by GfK Knowledge Networks.[6] In contrast, the 2006–2010 GSS was an in-person interview. It is worth highlighting that the ANES and GSS both incorporate a number of design features to explicitly mitigate against panel attrition—incentives, refusal conversions, and the like—so we might expect less attrition bias than surveys with greater budgetary and time constraints. Nonetheless, understanding the panel attrition process in these high-quality panel surveys should inform the conduct and analysis of the broader set of panel studies being used in political science.

Although there are a wide range of statistical techniques available to account for panel attrition (for an overview, see Vandecasteele and Debels 2007), political scientists most often ignore it, conducting a complete case analysis on the subset of respondents who completed all panel waves (e.g., Wawro 2002). In doing so, scholars are making an assumption that panel attrition occurs randomly. In the language of the missing-data literature (Little and Rubin 2002), any complete-case descriptive analysis assumes the missing data—unit nonresponse in subsequent survey waves, in this case—are missing completely at random (MCAR). That is, no observed or unobserved data exist that could systematically predict or account for this missingness. Unfortunately, this assumption is almost always unfounded. Countless analyses have found that panel attrition is related to a variety of respondent characteristics (e.g., Behr et al. 2005).

Even within the context of a multivariable analysis, complete-case analysis still assumes that attrition occurs randomly within each class of the covariates. Yet, the variables scholars include in multivariable models are rarely motivated by concerns about panel attrition and are often limited to a small set of demographic and political characteristics. Resulting estimates can still be biased if omitted variables are predictive of attrition. It is clear, then, that there is a need for a better understanding of the attrition process. That is also true, perhaps even more so, for those using statistical techniques that explicitly correct for panel attrition bias. The most common correction for panel attrition is post-stratification weighting to known population distributions, such as from the Current Population Survey or American Community Survey (e.g., Henderson et al. 2010). This approach assumes that the variables used in creating the weights—typically a limited number of demographic benchmarks—fully account for any bias introduced by panel attrition. Likewise, statistical correction using multiple imputation (e.g., Honaker and King 2010) or selection models (e.g., Brehm 1993) must also make explicit assumptions about the attrition process, so it remains critical to understand and gather data measuring the predictors of attrition.[7] Finally, and perhaps most importantly, understanding who is most likely to attrite can inform design decisions for future surveys to help prevent attrition in the first place. Thus, the basic motivation for this article is to identify and explore the correlates of panel attrition in the ANES and GSS. We look to the more general research on survey nonresponse for possible predictors of panel attrition that extend beyond the standard demographic and political controls included in many multivariable models and typically used in constructing post-stratification weights.

Broadly speaking, much of the previous research has emphasized that panel attrition is related to a respondent's ability and motivation to participate. A large body of scholarship has demonstrated that characteristics like country of origin, income, education, gender, and race are predictive of attrition (Gray et al. 1996; Fitzgerald, Gottschalk, and Moffitt 1998; Loosveldt, Pickery, and Billiet 2002; Behr, Bellgardt, and Rendtel 2005; Lynn et al. 2005; Watson and Wooden 2009). Individuals who are more socially engaged and residentially stable—homeowners and those with children—are more likely to remain in a panel study, whereas younger respondents and those who live alone are more likely to drop out (Groves and Couper 1998; Lipps 2007; Uhrig 2008; Watson and Wooden 2009). Civic engagement and interest in the survey topic are also predictive of attrition (Traugott and Morchio 1990; Traugott and Rosenstone 1994; Loosveldt and Carton 1997;

[6]It is worth noting that although GfK Knowledge Networks utilizes probability-based sampling procedures—formerly RDD and now address-based sampling using the United States Postal Service's Delivery Sequence File—they also maintain respondents in panels for many months and frequently survey them. The ANES panelists, however, were all new recruits and did not participate in other surveys while members of the ANES panel.

[7]See Deng et al. (2013) for a discussion of model-based and multiple imputation approaches to correcting for nonignorable panel attrition that could be used with the ANES and GSS.

Lepkowski and Couper 2002; Loosveldt, Pickery, and Billiet 2002; Voogt 2005; Smith and Son 2010). Measures of political engagement and political interest, in particular, can be predictive of attrition in surveys on all topics, but are especially predictive of attrition in surveys with political content (Brehm 1993; Traugott and Rosenstone 1994; Bartels 1999; Burden 2000; Voogt and Saris 2003; Olson and Witt 2011). For example, Olson and Witt (2011) find that political interest has been consistently predictive of retention in the ANES time series from 1964 to 2004.

Yet, the broader literature on survey nonresponse also emphasizes the importance of the survey experience in collecting quality responses. This is true both in the research on item nonresponse, in which a selected respondent fails to respond to some, but not all, questions, and unit nonresponse, in which a selected respondent fails to participate in the entire study or survey wave (see Groves et al. [2002] for a comprehensive accounting of causes of survey nonresponse). We expect the same to hold true in the case of panel attrition.

Scholarship on survey methodology has emphasized the importance of viewing the survey process as a form of conversation, rather than a rote standardized process of question-asking yielding data. Interviews vary by respondent, interviewer, and context, and this can affect the data collected. Respondents are more likely to respond when they feel heard, understood, and comfortable in the survey process, much like we would expect in a conversational setting.

In line with viewing the interview as a form of conversation, much of the survey experience research focuses on interviewer effects, as interviewers represent at least one half of the conversation. Research has demonstrated that the race and gender of an interviewer can affect respondents' likelihood to participate in a survey (Groves, Cialdini, and Couper 1992) and to provide answers to individual questions, especially items directly connected to race and gender (Davis 1997; Groves et al. 2009). In terms of unit response, survey field managers often contend that middle-aged female interviewers have higher response rates, although the empirical evidence is mixed (Campanelli and O'Muircheartaigh 1999). In addition, more experienced interviewers have been found to be more skillful at eliciting initial participation in a survey as well as responses to sensitive items once the survey begins (Couper and Groves 1992; Groves and Couper 1998; Durrant et al. 2010).

Beyond interviewer effects, research also demonstrates the importance of respondent satisfaction and comfort throughout the interview process for maintaining high-quality data and minimizing nonresponse. Survey design and complexity, interview length, and respondents' reports of satisfaction with the survey all affect survey nonresponse and data quality (e.g., Burchell and Marsh 1992). In addition, respondents who have a generally less cooperative disposition provide lower-quality data (e.g., higher item nonresponse) even if they complete the survey. Even accounting for political and demographic characteristics, some respondents simply have a lower propensity to cooperate (Brehm 1993). For example, in adjusting for initial unit nonresponse in the 1980, 1984, and 1988 ANES, Peress (2010) estimates each individual's "response propensity" based on number of contact attempts, interviewer-rated cooperativeness, and interest in the survey. Finally, it is well known that incentives can increase response rates (Singer 2002), but there is some question as to whether monetary incentives create an expectation about future payment that could potentially reduce future cooperation (Martin, Abreu, and Winters 2001). Certainly, it is clear that those motivated by civic duty or altruistic motivations are among the most likely to respond (Groves and Couper 1998).

Although the bulk of the evidence on survey experience and respondent disposition has focused on initial nonresponse or item nonresponse, there is every reason to expect these factors to matter in a longitudinal context as well. As Lynn et al. (2005) note, "The sample member has had the experience of the previous survey and so is in a position to know what to expect. They can make an evaluation of whether to participate in the survey based on how they felt about the previous survey" (25). Indeed, there is limited research showing some of these factors to be predictive of panel attrition. Research specifically focused on panel attrition has found that satisfaction with the Wave 1 survey is associated with attrition, although such measures are only rarely available (Kalton et al. 1990; Traugott and Rosenstone 1994).[8] Groves and Couper (1998) use a

[8]In terms of respondent burden, higher levels of reported understanding at Wave 1 reduce attrition (Kalton et al. 1990; Loosveldt and Carton 2001; Lepkowski and Couper 2002; Olsen 2005; Lipps 2010; Smith and Son 2010). Some research

labor-intensive process of coding interviewers' field notes to locate respondents' negative utterances about the Wave 1 survey, for example. Measures of respondent cooperation in Wave 1 are also related to attrition (Kalton et al. 1990; Traugott and Rosenstone 1994; Smith and Son 2010; Olson and Witt 2011). A handful of studies have examined interviewer effects on panel attrition, focusing on the importance of providing consistency in interviewer assignment across waves (Smith and Davis 1992; Zabel 1998; Campanelli and O'Muircheartaigh 1999; Hill and Willis 2001).[9] Clearly, many expectations about possible determinants of panel attrition have been laid out in previous research, but such analyses are often limited to technical reports for a particular survey, are scattered across various disciplines, or are focused on isolating the effect of just one factor or another in a single discipline-specific survey. Thus, our key contribution is to bring together insights from these various sources in examining panel attrition in the ANES and GSS.

2 Panel Attrition in the ANES and GSS

In the expanding world of panel data, two surveys stand out for political scientists and sociologists: the General Social Survey (GSS) and the American National Election Survey (ANES). The GSS has recently moved to a full rotating panel design, and the ANES added a multi-wave panel component in 2008 as a supplement to the face-to-face, pre/post design of the time series.

The GSS is a nationally representative recurring survey started in 1972 to study societal change in the United States.[10] Beginning in 2006, the GSS transitioned from a repeated cross-sectional design to a three-wave rotating panel, so that each survey round now includes a new cohort completing their first wave simultaneously with the previous cohorts completing their second or third wave. We use the 2006–2008–2010 completed panel in this analysis.[11] The ANES time series have been conducted since 1948, typically through in-person interviewing, in a pre/post-election design. In 2008, the ANES added a multi-wave panel component using the online survey firm GfK Knowledge Networks. The panel was conducted online using newly recruited panel members to the probability-based Knowledge Networks panel.[12] This panel was maintained separately from the larger Knowledge Networks panel for the entirety of the study. The new ANES panel component was composed of twenty-one waves that were fielded during the course of 2008 and 2009, with ten of those waves including political content as developed by the ANES and eleven others composed of content funded by outside investigators.

In examining the predictors of attrition, there are several distinct design features of the ANES and GSS that could be relevant. The first and most important difference between the ANES and the GSS panels has to do with survey mode. The GSS is primarily fielded through face-to-face interviews, whereas the ANES was fielded online.[13] In terms of thinking about how survey experience shapes and affects respondents' decisions about returning to participate in subsequent waves, differences in mode must play a key role in determining what criterion respondents base their overall

indicates that longer interviews may increase a respondent's likelihood of attriting (Frankel 1980; Olson and Witt 2011), whereas other studies have failed to demonstrate this relationship (Hill and Willis 2001). We found no consistent relationship between interview length and panel attrition in the ANES or GSS studies.

[9]The 2006–2010 GSS panel did not make an explicit effort to assign the same interviewer to the same respondent across waves, and lacks data on whether matches were made by chance.

[10]The survey is conducted by the National Opinion Research Center (NORC) and is composed of four parts: one set of permanent items that all respondents answer and three sets of rotating items, of which respondents only answer a subset. Since 1994, the survey has used a biennial split-ballot design, in which the survey is fielded every two years and the sample is broken up into three subsamples. Each subsample receives the same set of permanent items as well as two-thirds of the set of rotating items. This is relevant to our analysis because we restricted variable selection to items asked of the entire sample.

[11]The GSS uses multistage area probability sampling from the NORC national sample frame, which is constructed using postal and census data, as well as geocoding. The 2006 GSS had a response rate of 71%.

[12]GfK Knowledge Networks now primarily uses address-based sampling using the USPS's Delivery Sequence File rather than RDD. However, all respondents who participated in the ANES panel were recruited to the online panel using a stratified list-assisted RDD sample of landline telephones and a subsequent telephone recruitment survey. The minimum response rate (AAPOR RR1) at the panel recruitment stage was 26%. The estimated response rate (AAPOR RR3, based on observed eligibility among known cases) at the recruitment stage was 42%.

[13]A small number of GSS interviews were fielded by phone (7% in Wave 1); controlling for them did not change the results.

experience on. Second, the substantive content of the surveys are different: the GSS includes a wide variety of measures for social attitudes and outcomes, whereas the ANES is political in its focus. Third, the complete GSS panel consists of three waves over five years, whereas the 2008–2009 ANES included twenty-one waves completed over the course of two years. Although a higher number of waves may lead to respondent fatigue and higher rates of attrition, the longer time span between waves in the GSS may render respondents less likely to return or more difficult to contact in subsequent waves. The final difference between the two studies is in the panel follow-up design: once a respondent attrited in Wave 2 in the GSS, they could not participate in Wave 3; ANES respondents who completed Wave 1 were invited to complete each subsequent (ANES-content) wave, even if they had failed to complete a previous wave.

This brings us to the definition of "attrition" itself, and how we should think about measuring and predicting it. Unfortunately, there are no standardized attrition rate formulas, so in looking at any given panel survey it is not always clear if reductions in sample size between waves reflect respondents dropping out, changes in eligibility, varying definitions of "eligibility" between panel studies, or changes to the sampling design. The ability for respondents to return after skipping a wave complicates things further. In the GSS, attrition is defined as respondents to Wave 1 who, though eligible, failed to complete Wave 3. This excludes those who became ineligible because they moved out of the United States, died, or became institutionalized (108 cases). Using this definition, the GSS attrition rate was 33% (Smith and Son 2010). Twenty-one percent of Wave 1 respondents dropped out by Wave 2, and the remaining attriters dropped out between Wave 2 and Wave 3.

In the ANES, our analysis focuses on respondents who completed Wave 1 but did not complete Wave 5.[14] Using this definition, the attrition rate is 32%, and corresponds with the 68% Wave 5 "retention rate" reported in the ANES methodology report (Debell, Krosnick, and Lupia 2010). There are, however, several ways in which the survey design of the ANES makes it more challenging to operationalize and study panel attrition. For instance, the ANES experimented with varying incentives during the survey and across respondents; held certain response periods open longer for some waves and some respondents; and chose not to invite some respondents to participate in off-panel waves in order to increase the likelihood of their participation in ANES-content waves. Moreover, whereas the GSS considers respondents who die, move out of the country, or become institutionalized during the study as ineligible for subsequent waves, the ANES defines eligibility only at recruitment. In other words, the ANES does not distinguish those who refuse to take part in a subsequent wave from those who are no longer part of the original sample frame. And, in contrast to the GSS, respondents were allowed to answer subsequent ANES-content waves even if they failed to answer earlier ones. To be sure, all of these efforts were implemented to improve the validity and reliability of the final data, but they do complicate analysis of panel attrition.

Critically, Wave 5 was selected because it is the latest wave in the study before a new incentive experiment and structure was added to the design.[15] Following Wave 5, incentives of $30 or $50 were offered to one hundred randomly selected attriters to complete Wave 6 (with promises to make similar payments for completing each subsequent wave). Because of the success of the experiment, the incentive program was later expanded to include all those who had attrited by Waves 7–9, with payments continuing through the rest of the study. To avoid conflation of the observed attrition with this design change, our analysis looks at attrition in Wave 5.[16] Whereas 32% of those who completed Wave 1 failed to complete Wave 5, attrition from Wave 1 to Wave 21 was only 30%, reflecting the fact that some respondents returned after receiving the increased incentives.[17] Overall, 48% of respondents who completed the final wave failed to participate in at least one of the previous waves.

[14]Replication data are available on the *Political Analysis* Dataverse.

[15]For more information about the incentive experiment, see Debell, Krosnick, and Lupia (2010).

[16]Of the first five waves, only Waves 1 and 2 contained ANES content. Wave 5 focused on the topic of science and environmental policy attitudes and was sponsored by Jon Miller of Michigan State University.

[17]The ANES design included the addition of a second cohort of respondents in Wave 9; these respondents have not been included in any of our calculations.

Variations in fielding periods to try to maximize response rates to the ANES-content waves also made it more difficult to clearly identify attriters. Debell, Krosnick, and Lupia (2010) explain that "we held some waves open longer than others and we did not invite every panelist to all of the non-political waves . . . panelists who were slow to complete Wave 4 were given all of the month of May to complete the April (Wave 4) survey, and then were invited to complete Wave 6 (the June ANES wave) at the beginning of June, skipping Wave 5" (68). So, although our definition of attrition maintains consistency with the reported retention rate, and fits with the ANES's eligibility definition, it has the odd feature of treating as attriters some individuals who were not invited to complete Wave 5 because they were slow to complete Wave 4.[18] To be sure, many of those slow to respond never ended up responding—thus, confirming them as attriters—but others did. Reassuringly, however, we find nearly identical results if we replicate our analysis defining attriters as panelists who completed Wave 1 but failed to complete Wave 4 and Wave 5.[19] We simply want to acknowledge that the particularities of the survey design and the sheer number of survey waves mean that there are dozens of different ways to define attrition in the ANES. Indeed, the "Methodology Report and Users Guide for the 2008–2009 ANES Panel Study" (Debell, Krosnick, and Lupia 2010) understandably only looks at panel attrition in later ANES-content waves (Waves 11 and Wave 17) since they are likely of greater interest to most ANES users. We believe the complexities outlined here highlight the need for a clear, consistent, and precise definition of "attrition" across longitudinal surveys.

Which respondents from Wave 1 were most likely to attrite? We first look at the demographic differences between the attriters and nonattriters in each of the surveys. Table 1 reports the average values of each of these variables measured at Wave 1 for two different groups: respondents who stayed in the panel and those who dropped out.[20] Although the differences are not always large, those who attrite are more likely to be young, non-white, less educated, and less connected to their communities. Perhaps the most striking difference is on foreign-born and foreign language preference, perhaps reflecting difficulty in completing an English-language interview.[21]

Although a demographic comparison is often the extent of panel attrition analysis conducted by many scholars, we have argued that nondemographic factors might also be related to panel attrition. In going beyond these standard variables, we look to available measures in each survey to capture aspects of survey experience and respondent disposition. Because the GSS is a face-to-face survey, the primary concept of survey experience involves interaction with the interviewer in Wave 1. Thus, we use a measure for years of interviewer experience at NORC. Building on the existing research that links interviewer experience with other forms of nonresponse, more experience may enable the interviewer to provide a better experience or cultivate a better rapport with a respondent at Wave 1, thus facilitating a more positive experience for the respondent. We also include an

[18]Of the 1623 respondents who completed Wave 1, 1342 were invited to Wave 4, and 1219 were invited to participate in Wave 5. There were 118 cases who eventually completed Wave 4 after the Wave 5 field period had begun. Unfortunately, the publicly released data set does not indicate if a respondent was invited to Wave 5, although that information is available for each of the ANES-content waves.

[19]We also replicate our analysis coding as attriters those respondents who *ever* failed to complete a wave during the entire course of the study. Although there are a few minor differences, the key survey experience measures remain significant predictors of attrition.

[20]The measure for education is a numeric scale ranging from 1 to 5, where 1 = No high school diploma, 2 = A HS diploma, 3 = Some college, 4 = A college degree, and 5 = A graduate degree. Income is also a scale ranging from 1 to 18, corresponding to standard breaks in income categories. Other variables are indicators. The measure for language difficulty/shallow community roots was less feasible to standardize. In the ANES model, this concept is measured as a dummy variable that takes the values of 0 and 1 based upon whether a respondent indicated that they would prefer to complete the survey in a non-English language. In the GSS, this variable is also a dummy that takes on the values of 0 and 1 depending on whether or not a respondent was born outside the United States. Both measures have been used in the literature and, we believe, measure similar concepts of community connectedness and ability to interface with the survey with which a respondent is presented. All variable values were measured at Wave 1 or, in the case of the ANES, possibly derived from information provided in the profile survey. We restrict our analysis to respondents who completed the profile survey, but did not consider this to be a "baseline" survey because it was sometimes completed after Wave 1.

[21]Indeed, recognizing its strong relationship with panel attrition, the GSS uses foreign-born status in the construction of the post-stratification weights provided in the data set.

Table 1 Mean demographic comparisons of attriters and nonattriters in GSS and ANES

	GSS		*ANES*	
	Nonattriters	*Attriters*	*Nonattriters*	*Attriters*
Age	**46**	**42**	**53**	**48**
Female	0.59	0.56	0.56	0.57
Non-white	**0.25**	**0.31**	**0.15**	**0.24**
Education (1–5)	**2.90**	**2.67**	3.39	3.31
Income (1–18)	10.52	10.08	12.19	12.18
Owns home	**0.70**	**0.60**	**0.82**	**0.75**
Employed	0.60	0.64	**0.64**	**0.67**
Number of children under 6	0.19	0.22	0.17	0.30
Lives alone	**0.25**	**0.31**	0.19	0.16
Foreign born	**0.10**	**0.20**		
Non-English preference			**0.05**	**0.12**

Note. Bold type indicates a significant difference between groups at $p < 0.05$.

indicator for whether an interviewer is female, as female interviewers are believed to produce higher response rates, perhaps due to comfort and rapport, as well.[22]

Turning to respondent disposition, we include an interviewer rating of the respondent's cooperativeness at Wave 1—coded 1 for those reported to be friendly or cooperative and 0 for those reported restless, impatient, or hostile. We also consider item nonresponse in Wave 1, with the expectation that respondents who refuse to answer a Wave 1 question might be more likely to refuse to participate in a future survey wave.[23] Finally, we include an indicator for whether the respondent voted in the previous election. Although the GSS is not primarily a political survey, it could be that voter participation is a proxy for more general civic interest, including survey participation.

Of course, as an online survey administered without an interviewer, there are not similar experience measures in the ANES. In a web context, interface and interaction with technology supplants the interaction with the interviewer in the face-to-face context.[24] Thus, we create an indicator variable for whether a respondent required GfK Knowledge Networks to provide them with MSN TV2 in order to participate in the survey.[25] Respondents who "interact" with the technology administering the survey on a regular basis, as would be the case if they already had Internet in the home, may have a better "rapport" with this interviewing medium than those needing to familiarize themselves with the MSN TV2 technology and interface for the first time. We expect that respondents who require this technology may have greater difficulty and less comfort with the interview

[22]We also included a variable for interviewer race and whether the interviewer and respondent were of the same race. Neither was statistically significant, and neither changed our results in any substantive way. In the literature, race of interviewer seems to be more related to sensitive items and item bias, rather than overall response, so we do not include it in the model. Likewise, length of the original interview was not a statistically significant predictor of attrition and, as such, was dropped from the final model. We believe that interview length may actually be somewhat conflicted or cross-cutting in its relationship with attrition. A long interview could signal a poor experience for the respondent if she was bored, disengaged, or wanted the interview to be shorter, thus making her less likely to complete Waves 2 and/or 3. However, a long interview could also be a sign of positive interviewer/respondent rapport, leading respondents to feel heard, understood, and engaged, thus making it more likely that a respondent completes Waves 2 and 3. More work is needed in this area.

[23]The variable is an indicator if the respondent refused to answer any of the following questions asked of the entire sample: happiness, ideology, party identification, self-identified class, and religious attendance.

[24]In a methodological report on attrition in Knowledge Networks online panels, Dennis and Li (2003) note that a respondent's existing access to computer and the Internet affects levels of attrition in this context. They further stipulate that more work is needed to understand the connection between respondent access to and use of technology and likelihood to attrite. We hope that this analysis begins that exploration.

[25]We also considered measures of familiarity with computers and frequency of computer use, but neither variable was significant and did not change the MSN TV2 effect.

instrument, which could, in turn, affect their probability of attriting.[26] In addition, we used responses to an open-ended question at the end of Wave 1 that asked respondents for any comments they had about the survey.[27] These qualitative responses were recoded into a set of indicators for whether a respondent reported a positive experience, a negative experience, some other comment, or no response at all. We include the measure for negative experience in the model, as well as a control for providing no response to this question, and expect that respondents who said something negative or refused to answer the question will be more likely to attrite.[28] It is worth noting that this measure differs from the cooperation variable we have for the GSS because the cooperation measure is reported by the interviewer, whereas satisfaction is reported by the respondent. Further, cooperation is typically viewed as a dispositional measure—respondents begin the survey at some existing level of cooperation. Satisfaction, on the other hand, is a function of the interview process. Although these two concepts are related—satisfaction during the survey may affect cooperation as the interview goes on and initial disposition may likewise affect cooperation—they are distinct concepts in theory, measurement, and effect.

The ANES data set also offers unique measures about the respondent's motivation for joining the survey. The profile survey asked respondents to select from a list all of the reasons they decided to participate in the study.[29] A follow-up open-ended question asked for other reasons they joined the panel. Based on the combination of responses to the closed and open-ended questions, we created an indicator for those reporting *only* a monetary incentive as their motivation for participating. We expect that respondents motivated to participate only by money, in contrast to those who reported other motivations, such as wanting to have a voice, will be more likely to attrite.[30] Given the political content of the survey, we also examine the relationship between political interest and attrition. Finally, we look at three measures of a respondent's propensity to cooperate: the number of call attempts required to reach the respondent in the initial telephone recruitment survey, the number of days it took the respondent to complete Wave 1, and whether the respondent skipped questions in Wave 1 of the survey.[31] Table 2 compares attriters and nonattriters across these various measures.

Since many of these factors might themselves be related to the demographic characteristics of the respondents, we estimate a logit model predicting attrition that includes the full set of possible predictors. Again, the dependent variable in the General Social Survey model is an indicator if a Wave 1 respondent was eligible but failed to complete Wave 3. The dependent variable in the ANES model is an indicator if a Wave 1 respondent failed to complete Wave 5. The results of our logit models predicting attrition are reported in Tables 3 and 4.[32] For comparison, we report the results for both a restricted model that includes only demographic variables (column 1) and a fully specified model (column 2) that also includes measures of survey experience and respondent

[26]Although this variable might be related to SES, we control for education and income in the model. As an additional robustness check, we included an indicator for having a home Internet connection in the GSS model, but find that it is not a statistically significant predictor of attrition in the face-to-face GSS—as our theory regarding the import of Internet access for predicting attrition in the ANES would predict. The Internet question was asked only of a subset of respondents so was not included in the final model specification.

[27]The prompt read, "If you have any comments about any part of the survey, please write them below."

[28]By including no response as a control, we set our reference category to be those who reported a positive response at Wave 1 or provided other types of feedback that was neither positive nor negative.

[29]Multiple selections were allowed. Response options included: "I wanted an MSN TV2 device, I wanted access to the Internet, I wanted to have a voice in how decisions were made, I thought it would be interesting, I thought it would be fun, I thought it would be educational, None of these reasons."

[30]In the multivariate model, we also include a measure capturing respondents who reported no motivation as a control so that our reference category is those who indicated yes to any motivation besides or in addition to monetary incentives.

[31]Days to complete was calculated by subtracting the date the survey wave was completed from the date GfK Knowledge Networks emailed the invitation to complete the Wave 1 interview. Seventy-five percent of the sample completed the survey within the first week, but 1% of the sample took more than fifty days to respond. Because of extreme right skew in both call attempts and days to complete, these measures were censored at the 95th percentile.

[32]Our analysis uses Wave 1 weights that incorporate the sampling design and post-stratification on the basis of nonresponse to Wave 1. Although 1623 respondents completed Wave 1, our models are restricted to those who also completed the profile (89% of Wave 1 respondents) because we include profile variables as covariates.

Table 2 Survey experience and respondent disposition: mean comparisons of attriters and nonattriters in GSS and ANES

	GSS		*ANES*	
	Nonattriters	*Attriters*	*Nonattriters*	*Attriters*
Voted in 2004	**0.70**	**0.59**		
Interviewer experience (years)	**3.96**	**3.50**		
Female interviewer	0.81	0.78		
Cooperative respondent	**0.97**	**0.93**		
Skipped Wave 1 question	**0.06**	**0.09**	**0.03**	**0.07**
Political interest (1-5)			**3.68**	**3.48**
Required MSN TV2			**0.10**	**0.14**
Monetary motivation only			0.03	0.05
No motivation reported			0.01	0.01
Negative experience reported			0.07	0.07
Days to complete survey			**4.72**	**7.14**
No experience reported			0.07	0.08
Number of call attempts			**4.47**	**5.76**

Note. Bold type indicates a significant difference between groups at $p < 0.05$.

disposition. To help with substantive interpretation of the fully specified model, we report first differences in column 3.[33]

Looking first at the demographic measures in the ANES model in Table 3 largely confirms the findings of previous research. Respondents who are non-white and have limited English skills are more likely to attrite, whereas respondents with higher levels of education are less likely to attrite.

A comparison with the demographic effects in the GSS model in Table 4 reveals a couple of interesting differences that could be attributable to survey mode. First, living alone is significant in the GSS model, but not the ANES model (although it is in the expected direction). Living alone is believed to affect propensity to attrite because it serves as a proxy for social connectedness, regularity of lifestyle, and general social stability. Because the logistics of a face-to-face survey are more reliant on residential stability and a lifestyle that facilitates interview scheduling, such a measure might be more consequential for a face-to-face interview than a mobile online survey, which can be completed in the wee hours of the morning at any location. The effect of having young children at home also differs across the two surveys in the full models. Having young children at home increases the likelihood of attrition in the ANES but not in the GSS. Interestingly, this ANES finding is the opposite of what the literature generally tells us about the effects of young children in the home—typically another indicator of social connectedness and stability and believed to decrease a respondent's likelihood of attriting (e.g., Uhrig 2008). We suspect, however, that this measure is also an indicator of time availability and distraction which, again, would have differing effects on attrition in different survey modes. Respondents to the online ANES panel had the flexibility of completing the survey at any time, but they had to actually make the time to do it. Parents of young children might find that more difficult to do or could be distracted while trying to complete the survey. In contrast, in a face-to-face survey, the interview is scheduled and distractions may be less likely in a more formal interview setting.

It is clear that in both the GSS and the ANES models, the standard variables used to predict attrition are important, but they do not alone account for an individual's likelihood of dropping out. Including a number of survey experience and respondent disposition measures improves the model fit and identifies a number of significant predictors of attrition. Looking first at the GSS, we find that interviewer characteristics are related to attrition. Interviewer experience is inversely

[33]First differences are calculated for the 95th minus 5th percentiles, holding other variables at their means or modes using Zelig (Imai, King, and Lau 2006).

Table 3 Logit model of attrition in ANES

	Demographic model	*Full model*	*First differences (95th-5th percentile)*
Intercept	0.20	−0.65	
	(0.32)	(0.41)	
Age	−0.01*	−0.01*	−0.10
	(0.00)	(0.00)	[−.20, −.01]
Female	0.07	0.16	0.03
	(0.12)	(0.13)	[−0.02, 0.08]
Non-white	0.55*	0.52*	0.11
	(0.15)	(0.16)	[0.04, 0.18]
Education	−0.24*	−0.16*	−0.09
	(0.06)	(0.07)	[−0.15, −0.02]
Income	0.03	0.05*	0.12
	(0.02)	(0.02)	[0.03, 0.23]
Own home	−0.21	−0.02	−0.00
	(0.15)	(0.16)	[−0.06, 0.05]
Employed	−0.29*	−0.36*	−0.07
	(0.14)	(0.14)	[−0.12, −0.01]
Number of children under 6	0.25*	0.17	0.03
	(0.09)	(0.10)	−[0.00, 0.08]
Non-English preference	0.52*	0.64*	0.14
	(0.22)	(0.22)	[0.03, 0.24]
Lives alone	0.07	0.09	0.02
	(0.21)	(0.22)	[−0.06, 0.10]
Political interest		−0.21*	−0.12
		(0.06)	[−0.19, −0.05]
Required MSN TV2		0.74*	0.16
		(0.19)	[0.07, 0.24]
Monetary motivation only		0.63*	0.14
		(0.27)	[0.02, 0.26]
No motivation reported		1.71*	0.40
		(0.69)	[0.10, 0.62]
Negative experience reported		0.63*	0.14
		(0.27)	[0.01, 0.28]
No experience reported		0.19	0.03
		(0.18)	[−0.03, 0.10]
Days to complete survey		0.06*	0.27
		(0.01)	[0.18, 0.36]
Number of call attempts		0.01	0.04
		(0.02)	[−0.05, 0.13]
Refused Q in Wave 1		0.66*	0.14
		(0.28)	[0.02, 0.27]
N	1407	1407	
AIC	1639.38	1577.1	
log *L*	−775.69	−708.57	

Note. *Indicates significance at $p < 0.05$. First difference calculated for the 95th minus 5th percentiles, holding other variables at their means or modes.

related to likelihood of attrition. Holding all else constant, respondents who received the most experienced interviewers (95th percentile is 11.8 years of experience) were predicted to be ten percentage points less likely to attrite on average than those interviewed by the least experienced interviewers (5th percentile is three months of experience). The effect of having a female interviewer, while in a direction consistent with our theory, is not significant—possibly because it lacks much variation. Nearly 80% of Wave 1 interviewers were female, which may, in part, reflect an effort by the GSS to preempt adverse interviewer effects on nonresponse.

Table 4 Logit model of attrition in GSS

	Demographic model	*Full model*	*First differences (95th-5th percentile)*
Intercept	−0.16	1.50*	
	(0.35)	(0.61)	
Age	−0.01	−0.01*	−0.13
	(0.01)	(0.01)	[−.25, −.01]
Female	0.08	0.17	0.04
	(0.14)	(0.15)	[−0.03, 0.10]
Non-white	0.13	0.09	0.02
	(0.17)	(0.18)	[−0.06, 0.09]
Education	−0.25*	−0.21*	−0.20
	(0.07)	(0.07)	[−0.30, −0.06]
Income	0.04	0.03	0.11
	(0.02)	(0.02)	[−0.02, 0.25]
Own home	−0.23	−0.15	−0.04
	(0.16)	(0.17)	[−0.11, 0.04]
Employed	−0.06	−0.08	−0.02
	(0.16)	(0.17)	[−0.09, 0.05]
Number of children under 6	0.00	−0.00	−0.01
	(0.12)	(0.13)	[−0.06, 0.05]
Foreign born	0.84*	0.82*	0.19
	(0.21)	(0.23)	[0.08, 0.30]
Lives alone	0.44*	0.44*	0.09
	(0.17)	(0.18)	[0.02, 0.18]
Interviewer experience		−0.05*	−0.10
		(0.02)	[−0.20, −0.01]
Female interviewer		−0.22	−0.05
		(0.18)	[−0.13, 0.03]
Cooperative respondent		−1.35*	−0.31
		(0.45)	[−0.50, −0.10]
Voted in 2004		−0.09	−.02
		(0.18)	[−0.10, 0.05]
Refused Wave 1 Q		0.48	0.11
		(0.34)	[−0.03, 0.28]
N	1056	989	
AIC	1293.1	1209.4	
log *L*	−602.56	−540.7	

Note. *Indicates significance at $p < 0.05$. First difference calculated for the 95th minus 5th percentiles, holding other variables at their means or modes.

In addition, one of our measures for respondent disposition, the interviewer-rated measure of respondent cooperation, is a strong predictor of attrition. The model predicts that cooperative respondents are thirty-one percentage points less likely to attrite than uncooperative ones. The effect of refusing to answer a question in Wave 1 is in the expected direction, but does not reach traditional levels of statistical significance. Likewise, voters are less likely to attrite, but that effect is not statistically significant.

It is worth highlighting that the inclusion of these additional variables has little impact on the relationship between attrition and the demographic measures. In other words, respondent disposition and survey experience are not adequately captured by the usual demographic measures.

Turning to the ANES, we find that respondent interaction with the survey interface matters. Those requiring MSN TV2 were significantly more likely to attrite. Holding all else constant, individuals who required an MSN TV2 unit are predicted to be sixteen percentage points more likely to attrite than those who did not. As another measure of survey experience, respondents who offer a negative comment about the survey were predicted to be fourteen percentage points more likely to attrite than those who offered a positive or neutral comment.

In addition to the survey experience, dispositional factors are related to the probability of attrition. Politically interested respondents were less likely to drop out, whereas those motivated only by monetary incentive were more likely to attrite. Those who said money was their only motivation were predicted to be fourteen percentage points more likely to attrite on average. Coupled with the evidence about the success of the incentive payment starting after Wave 5, this finding suggests that money might well induce an individual to initially participate in the study, but continued participation requires further inducements.[34] Those who failed to offer any motivation for participating in the study—a measure of item nonresponse in the profile survey—were also significantly more likely to drop out. Likewise, those who refused to answer a survey question in Wave 1 were fourteen percentage points more likely to attrite. The effect of item nonresponse is not significant in the GSS, a difference that could potentially be attributed to survey mode. Scholars have found that item nonresponse tends to be higher in self-administered surveys than in interviewer-administered surveys because interviewers can clarify questions and encourage responses (Couper and Miller 2009). Thus, if item nonresponse is a signal of disinterested respondents, it is harder to give that signal in a face-to-face survey than in an online survey. We also find that the more days it took an individual to complete the first wave after being invited, the more likely she was to drop out of the study by Wave 5. Those who took the longest (95th percentile was twenty-one days) were twenty-seven percentage points more likely to attrite on average than those who answered the survey on the first day it was available. These findings point to possible ways that survey designs might be adapted during data collection to reduce the potential for panel attrition bias. For example, more experienced interviewers might be strategically assigned to less cooperative respondents for the follow-up interviews.

The only respondent disposition measure unrelated to attrition was the number of call attempts for the initial recruitment survey. One possible explanation for this null finding is that those respondents who were difficult to reach were more likely to drop out between the recruitment survey and Wave 1, and thus are not included in the survey. For the subset of respondents who actually completed Wave 1, we would expect disposition at and experience with Wave 1 to be more relevant. This again highlights the differences between the survey designs because the GSS did not utilize a separate recruitment, profile, and Wave 1 survey.

Overall, then, these results demonstrate that missing data due to panel attrition are not missing completely at random in the GSS and ANES. Moreover, an individual's propensity to drop out is not simply a function of demographic characteristics, but is also predicted by his or her disposition at and experience with Wave 1. What is more, that experience is itself shaped by the design of the survey, with different aspects of the survey design having differential importance across the two surveys. Given the large number of differences in the design of the two surveys (timing, incentives, mode, etc.), we cannot definitively identify the cause of the differences, but modal differences seem like a plausible explanation. This analysis demonstrates the importance of both accounting for potential biases that may arise from panel attrition, and also informing and motivating these accountings with a broader theory of survey experience, nonresponse, and the interviewing process.

3 Attrition Bias Correction Example

One obvious question that arises from the above analysis is how we might use the findings from our analysis to assess panel attrition bias and correct for it. Although not the primary focus of the article, we walk through an example of such a correction to help illustrate the value of our analysis. We look more closely at a few measures in the ANES to see if we find different estimates when we use complete case analysis, thereby assuming that attrition occurs randomly, than when we correct for the systematic patterns of panel attrition identified in the analysis above.

[34]Interestingly, a model predicting attrition in the final wave of the study finds that those motivated by money were no more likely to attrite from the final wave, perhaps reflecting the success of the large incentive for completing later waves.

There are a multitude of different correction approaches available, including weighting, selection models, and multiple imputation, among others—and there is a rich literature comparing each of their merits and limitations (Vandecasteele and Debels 2007; Deng et al. 2013). In this example, we rely on a multiple imputation approach in which we use the variables from our fully specified model of attrition to inform the imputation of a handful of substantive measures from Wave 5 of the ANES panel survey. As will be recalled, Wave 5 of the ANES panel survey was an off-ANES wave and, thus, does not contain the standard ANES political content. The questionnaire was largely focused on topics related to scientific knowledge and attitudes, but we identified a handful of items that could be of interest to political scientists: news readership, campaign interest, concern about global warming, and Internet use.[35]

With multiple imputation, values for missing data—from panel attrition, in this case—are repeatedly simulated by sampling from predictive distributions of the missing values to create multiple data sets. Point and variance estimates of interest with each data set can then be combined using simple formulas developed by Rubin (1987). In essence, patterns of systematic missingness in the data are modeled as a function of the other data present and the relationships between them. Because our goal here is simply to illustrate the predictive power of survey experience variables for modeling attrition, we use only the variables in our models above to inform the imputation of the Wave 5 values for those who attrited.[36] We then compare the complete case and imputed means for each of the variables.[37]

The results are reported in Table 5. The first column reports the estimate from a complete case analysis—the estimated mean for those who completed Wave 1 and Wave 5. The second column is the imputed mean value of the attriters, and the final column is the corrected mean that accounts for the attrited respondents. Thus, the difference between column 3 and column 1 represents the bias correction. The differences are quite small across these items, although accounting for nonrandom panel attrition does slightly change some of our estimates. For example, the complete case analysis of campaign interest finds that 48% of respondents were very interested in the primary elections, but correcting for attrition reduces that estimate to 46% because attriters are predicted to be less interested. Likewise, failing to account for attrition overestimates daily newspaper readership by two percentage points and overestimates time spent on the Internet. Correcting for panel attrition bias has no effect on our estimate of concern about global warming, however.

Although simple mean comparisons find only small biases on these particular variables, panel attrition can also affect the relationships between variables (Bartels 1999). To explore that possibility, we focus on the measure of campaign interest. We estimate a logit model of campaign interest that includes age, sex, race, education, and income as independent variables. We compare the results of this model using the imputed data that corrects for panel attrition with those of the complete case subset, in which panel attrition is ignored. Despite the small differences found in looking at the mean estimates of campaign interest, the multivariate analysis leads to different conclusions regarding the relationship between sex and campaign interest depending on the panel attrition approach used. The complete case analysis finds that women are significantly more likely than men to report being very interested in the campaign ($p < 0.05$). In contrast, the attrition-corrected model finds no difference in the predicted campaign interest of men and women ($p < 0.16$)—the sex variable is no longer significant, despite the increased sample size that comes from including the imputed attriters.[38]

[35]The imputation for this analysis was done using the Amelia package in R (Honaker, King, and Blackwell 2011).

[36]Typically, analysts want to use as much information as is available—often all other variables in the data set. Certainly, a substantive motivation would warrant closer attention to the specific correlates of the dependent variable (Pasek et al. 2009).

[37]Imputation was performed five times, and the means reported are the mean value of the resulting five means. We exclude cases with missingness on any other variable from our model and reported means. Scholars would typically want to simultaneously correct for both item and unit nonresponse. We do not do that here simply to cleanly illustrate the correction for panel attrition bias alone—we want the imputation to correct for bias from panel attrition, not missingness from item nonresponse.

[38]The confidence intervals of the specific predictions overlap, but the important point here is that different conclusions would be drawn in a complete case than in an attrition-corrected model.

Table 5 ANES: multiple imputation based on full model versus listwise deletion

		Imputed	
	Complete-case mean	*Attriters mean*	*Corrected mean*
Very interested in primary election	0.481	0.438	0.462
Reads news daily	0.373	0.313	0.353
Hours of Internet use (0-99)	7.7	6.8	7.4
Very concerned with global warming	0.200	0.199	0.200

To be sure, this cursory comparison of complete case and attrition-corrected analyses in the ANES finds rather small biases due to panel attrition. On the one hand, it is somewhat reassuring for political scientists that the differences are not large for these items. However, there are a number of reasons that these particular examples might give unfounded reassurance about the threats of panel attrition. The extent to which panel attrition biases any given estimate depends on both the rate of attrition and the size of the difference between attriters and nonattriters. The American National Election Study invested considerable resources in designing and adapting the survey to mitigate against bias from panel attrition, so we might expect higher attrition in studies with greater budgetary constraint, less focus on preventing attrition, or simply with higher rates of attrition overall. We might also find greater bias in particular subsets of the data with higher attrition rates (e.g., racial minorities). More importantly, the difference between attriters and nonattriters can vary widely depending on the particular outcome of interest. For example, there may be large differences between the levels of political interest of attriters and nonattriters, but this difference might matter more for a researcher interested in vote turnout than for one focused on health outcomes. Finally, for the sake of clarity, whereas most imputation is based on every available variable—often hundreds of pieces of information—our imputation was based *only* on the few variables from our panel attrition model (and imputed only for panel attrition, not item nonresponse). This decision may have resulted in imputed values reflecting smaller differences between attriters and nonattriters than actually exist.

Again, the primary contribution of this article is to demonstrate that panel attrition reflects both the characteristics of the individual respondent as well as her survey experience, which depends on the survey design. As such, there is not a single one-size-fits-all solution for correcting bias from panel attrition. Bias corrections should vary depending on the particular substantive variable of interest as well as the design features of the data being analyzed. One important implication is that future data collection can use this knowledge to reduce potential panel attrition bias before the data are collected. For example, understanding the influence of interviewer characteristics on attrition likelihood might shape how interviewers are assigned to particular households. Understanding the influence of interest in and satisfaction with Wave 1 on attrition likelihood might shape the content of the questionnaire. And, certainly, our results make clear that more and better paradata are necessary to understand the way survey experience influences panel attrition.

4 Discussion

The results we have presented demonstrate the inadequacy of standard demographic and political variables (as are so often used in constructing post-stratification weights or as model controls) for accounting for panel attrition in the ANES and GSS. We find that survey experience and respondent disposition are as important as these more traditional measures in predicting who is most likely to drop out. The comparison between the ANES and GSS also suggests that survey mode influences panel attrition by shaping the respondent's survey experience. In the GSS, the interaction between the respondent and interviewer predicted the respondent's likelihood of attriting, whereas the respondent's interaction with the technology predicted attrition in the web-based ANES. Failing to account for this systematic variation in panel attrition risks biased estimates of many outcomes of interest.

Our point here is not to consider the benefits and limitations of various bias correction approaches or to focus solely on the two surveys used as examples, but rather to inform any corrections that are made (and to highlight the need for correction!). There is every reason to believe that the import of nontraditional predictors of attrition extends beyond the ANES and GSS. The results of this analysis also offer some guidance to future data-collection efforts, making clear the importance of collecting and using paradata. Analysts too often overlook survey administration variables when answering substantive questions. And for scholars commissioning original surveys, these results illustrate the need to collect not just the respondents' answers to survey questions but also information about the interviewer, the difficulty of connecting with the respondent, and respondent satisfaction with the survey experience. These findings also suggest that as survey designs become more complicated (with the increase in mixed-mode methods and the introduction of new technologies), it will become more important than ever to inform designs with a theory of survey participation—and a recognition that the survey design shapes survey experience—rather than simply assuming that some individuals are participators and others are not. To think about attrition without also considering the contours of an individual survey experience is to ignore the human respondent. Our analysis was necessarily guided by the data that were available to us, but the results do suggest the need to build a more robust theory of panel attrition, within larger survey response theories, and the need to collect information about the respondent's survey experience and her assessments of that experience. As panel data continue to grow in popularity and availability, we must ensure that we measure, control for, and learn about the varied effects of panel attrition.

Funding

This research is supported in part by NSF Graduate Research Fellowship Grant 1106401 and NSF Grant SES-10-61241. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

Atrostic, B. K., Nancy Bates, Geraldine Burt, and Adriana Silberstein. 2001. Nonresponse in US government household surveys: Consistent measures, recent trends, and new insights. *Journal of Official Statistics, Stockholm* 17:209–26.

Bartels, Larry M. 1999. Panel effects in the American National Election Studies. *Political Analysis* 8:1–20.

———. 2006. Three virtues of panel data for the analysis of campaign effects. In *Capturing campaign effects*, eds. Henry E. Brady and Richard Johnston, 134–63. Ann Arbor: University of Michigan Press.

Behr, Andreas, Egon Bellgardt, and Ulrich Rendtel. 2005. Extent and determinants of panel attrition in the European Community Household Panel. *European Sociological Review* 23:81–97.

Brehm, John. 1993. *The phantom respondents: Opinion surveys and political representation*. Ann Arbor: University of Michigan Press.

Burchell, Brendan, and Catherine Marsh. 1992. The effect of questionnaire length on survey response. *Quality & Quantity* 26(3):233–44.

Burden, Barry C. 2000. Voter turnout and the National Election Studies. *Political Analysis* 8(4):389.

Campanelli, Pamela, and Colm O'Muircheartaigh. 1999. Interviewers, interviewer continuity, and panel survey nonresponse. *Quality & Quantity* 33(1):59–76.

Clinton, Joshua D. 2001. Panel bias from attrition and conditioning: A case study of the knowledge networks panel. In *AAPOR 55th Annual Conference*.

Couper, Mick P., and Peter V. Miller. 2009. Web survey methods: Introduction. *Public Opinion Quarterly* 72:831–5.

Couper, Mick P., and Robert M. Groves. 1992. The role of the interviewer in survey participation. *Survey Methodology* 18(2):263–78.

Davis, Darren W. 1997. Nonrandom measurement error and race of interviewer effects among African Americans. *Public Opinion Quarterly* 61(1):183–207.

de Leeuw, Edith Desiree. 1992. *Data quality in mail, telephone and face to face surveys*. Amsterdam, The Netherlands: TT Publikaties.

Debell, Matthew, JonA. Krosnick, and Arthur Lupia. 2010. Methodology report and users guide for the 2008–2009 ANES Panel Study. Technical Report, Stanford University and University of Michigan.

Deng, Yiting, D. Sunshine Hillygus, Jerome P. Reiter, Yajuan Si, and Siyu Zheng. 2013. Handling attrition in longitudinal studies: The case for refreshment samples. *Statistical Science* 28(2):238–56.

Dennis, J. Micahel, and Rick Li. 2003. Effects of panel attrition on survey estimates. Paper presented at Annual Meeting of the American Association for Public Opinion Research, Nashville, TN.

Dillman, Don A., Glenn Phelps, Robert Tortora, Karen Swift, Julie Kohrell, Jodi Berck, and Benajmin L. Messer. 2009. Response rate and measurement differences in mixed-mode surveys using mail, telephone, interactive voice response (IVR), and the Internet. *Social Science Research* 38(1):1–18.

Durrant, Gabriele B., Robert M. Groves, Laura Staetsky, and Fiona Steele. 2010. Effects of interviewer attitudes and behaviors on refusal in household surveys. *Public Opinion Quarterly* 74(1):1–36.

Fitzgerald, John, Peter Gottschalk, and Robert. Moffitt. 1998. An analysis of sample attrition in panel data: The Michigan Panel Study of Income Dynamics. Technical Report, National Bureau of Economic Research.

Frankel, Joanne. 1980. Measurement of respondent burden: Study design and early findings. Technical Report. Washington, DC: Bureau of Social Science Research.

Gray, Rebecca, Pamela Campanelli, Kavita Deepchand, and Patricia Prescott-Clarke. 1996. Exploring survey non-response: The effect of attrition on a follow-up of the 1984–85 Health and Life Style Survey. *Statistician* 45(2):163–83.

Groves, Robert M., Don Dillman, John L. Eltinge, and Roderick J.A. Little. 2002. *Survey nonresponse*. New York: Wiley.

Groves, Robert M., Floyd J. Fowler, Mick P. Couper, James M. Lepkowski, and Eleanor Singer. 2009. *Survey methodology*, vol. 561. Hoboken, NJ: John Wiley & Sons.

Groves, Robert M., and Mick P. Couper. 1998. *Nonresponse in household interview surveys*. New York: Wiley.

Groves, Robert M., Robert B. Cialdini, and Mick P. Couper. 1992. Understanding the decision to participate in a survey. *Public Opinion Quarterly* 56(4):475.

Henderson, Michael, D. Sunshine Hillygus, and Trevor Tompson. 2010. "Sour grapes" or rational voting? Voter decision making among thwarted primary voters in 2008. *Public Opinion Quarterly* 74(3):499–529.

Hill, Daniel H., and Robert J. Willis. 2001. Reducing panel attrition: A search for effective policy instruments. *Journal of Human Resources* 36(2):416–38.

Honaker, James, and Gary King. 2010. What to do about missing values in time-series cross-section data. *American Journal of Political Science* 54(2):561–81.

Honaker, J., G. King, and M. Blackwell. 2011. Amelia II: A program for missing data. *Journal of Statistical Software* 45(7):1–47.

Imai, Kosuke, Gary King, and Olivia Lau. 2006. *Zelig: Everyone's statistical software*. R package version, 2–7. http://gking.harvard.edu/zelig (accessed February 2013).

Kalton, Graham, James Lepkowski, G. E. Montanari, and Dalisay Maligalig. 1990. Characteristics of second-wave nonrespondents in a panel survey. In *Proceedings of the Section on Survey Research Methods, American Statistical Association*, 462–7.

Keeter, Scott, Courtney Kennedy, Michael Dimock, Jonathon Best, and Peyton Craighill. 2006. Gauging the impact of growing nonresponse on estimates from a National RDD Telephone Survey. *Public Opinion Quarterly* 70:759–79.

Lepkowski, James, and Mick P. Couper. 2002. Nonresponse in the second wave of longitudinal household surveys. In *Survey nonresponse*, eds. Robert M. Groves, Don Dillman, John L. Eltinge, and Roderick J. A. Little, 259–72. New York: Wiley.

Lipps, Oliver. 2007. Attrition in the Swiss Household Panel. *Methoden–Daten–Analysen* 1(1):45–68.

———. 2010. Effects of different incentives on attrition and fieldwork effort in telephone household panel surveys. *Survey Research Methods* 4(2):81–90.

Little, Roderick J. A., and Donald B. Rubin. 2002. *Statistical analysis with missing data*. Hoboken, NJ: Wiley-Interscience.

Loosveldt, Geert, and Ann Carton. 1997. Evaluation of nonresponse in the Belgian Election Panel Study 91–95. In *Proceedings of the Survey Research Methods Section, American Statistical Association*.

———. 2001. An empirical test of a limited model for panel refusals. *International Journal of Public Opinion Research* 13(2):173.

Loosveldt, Geert, Jan Pickery, and Jaak Billiet. 2002. Item nonresponse as a predictor of unit nonresponse in a panel survey. *Journal of Official Statistics, Stockholm* 18(4):545–58.

Lynn, Peter, Nicholas Buck, Jonathan Burton, Annette Jackle, and Heather Laurie. 2005. *A review of methodological research pertinent to longitudinal survey design and data collection*. Working Paper, Institute for Social and Economic Research, University of Essex.

Martin, Elizabeth, Denise Abreu, and Franklin Winters. 2001. Money and motive: Effects of incentives on panel attrition in the survey of income and program participation. *Journal of Official Statistics, Stockholm* 17(2):267–84.

Olsen, Randall J. 2005. The problem of respondent attrition: Survey methodology is key. *Monthly Labor Review* 128:63–71.

Olson, Kristen, and Lindsey Witt. 2011. Are we keeping the people who used to stay? Changes in correlates of panel survey attrition over time. *Social Science Research* 40(4):1037–50.

Pasek, Josh, Alexander Tahk, Yphtach Lelkes, Jon A. Krosnick, B. Keith Payne, Omair Akhtar, and Trevor Tompson. 2009. Determinants of turnout and candidate choice in the 2008 US presidential election: Illuminating the impact of racial prejudice and other considerations. *Public Opinion Quarterly* 73:943–94.

Peress, Michael. 2010. Correcting for survey nonresponse using variable response propensity. *Journal of the American Statistical Association* 105(492):1418–30.

Sharp, Laurie M., and Joanne Frankel. 1983. Respondent burden: A test of some common assumptions. *Public Opinion Quarterly* 47(1):36.

Singer, Eleanor. 2002. The use of incentives to reduce nonresponse in household surveys. In *Survey nonresponse*, eds. M. Robert, Don Dillman, John L. Eltinge, and Roderick J. A. Little, 163–77. New York: Wiley.

Smith, Tom W., and Jaesok Son. 2010. An analysis of panel attrition and panel change on the 2006–2008 General Social Survey Panel. Technical Report, NORC, University of Chicago.

Smith, Tom W., and James A. Davis. 1992. The 1992 General Social Survey reinterview. GSS Methodological Report 75. Technical Report, NORC, University of Chicago.

Traugott, Santa, and Giovanna Morchio. 1990. Assessment of bias due to attrition and sample selection in the ANES 1989 Pilot Study. ANES Technical Report Series, No. NES010154. Technical Report, Center for Political Studies.

Traugott, Santa, and Steven J. Rosenstone. 1994. Panel attrition among the 1990–1992 panel respondents. ANES Technical Report Series, No. NES010163. Technical Report, Center for Political Studies.

Uhrig, S. C. Noah. 2008. The nature and causes of attrition in the British Household Panel Survey. Technical Report, Institute for Social and Economic Research, University of Essex.

Vandecasteele, Leen, and Annelies Debels. 2007. Attrition in panel data: The effectiveness of weighting. *European Sociological Review* 23(1):81–97.

Voogt, Robert J. J. 2005. An alternative approach to correcting response and nonresponse bias in election research. *Acta Politica* 40(1):94–116.

Voogt, Robert J. J., and Willem E. Saris. 2003. To participate or not to participate: The link between survey participation, electoral participation, and political interest. *Political Analysis* 11(2):164–79.

Watson, Nicole, and Mark Wooden. 2009. Identifying factors affecting longitudinal survey response. In *Methodology of longitudinal surveys*, 157–81. West Sussex, UK: John Wiley & Sons.

Wawro, Gregory. 2002. Estimating dynamic panel data models in political science. *Political Analysis* 10:25–48.

Zabel, Jeffrey E. 1998. An analysis of attrition in the panel study of income dynamics and the survey of income and program participation with an application to a model of labor market behavior. *Journal of Human Resources* 33(2):479–506.