Panel Bias from Attrition and Conditioning: A Case Study of the Knowledge Networks Panel*

Joshua D. Clinton
Department of Political Science
Stanford University

May 25, 2001

*The author would like to thank Mick Couper and Karol Krotki for their helpful comments. A version of this paper was presented at the 2001 AAPOR Conference in Montreal, Canada.

Abstract

With the recent rise of Internet based public opinion studies, panel studies have been utilized with much more frequency. Although the benefits of panel methods are well known (e.g., Sharot 1991), using a panel risks bias in two ways. First, since panels rely on re-interviewing panelists, systematic panel attrition can produce a panel that is unrepresentative of the target population. Second, interviewing and re-interviewing panelists may change the opinions/behaviors of the panelists – creating unrepresentative panelists. To investigate the prevalence and impact of these possible biases, I investigate a panel that is particularly suspect to these sources of bias – the panel of Knowledge Networks. Knowledge Networks' panelists are not only given an interactive TV appliance and Internet access, but they are also surveyed weekly. In this paper I both examine the extent (and effect) of panel attrition in Knowledge Networks's panel over a 7 month period, as well as report the results of an experiment designed to isolate the possible opinion/behavior changes introduced by panel participation. I find little evidence of either type of bias in the Knowledge Networks panel.

The study of panel methodology is important for two reasons. First, panel studies provide public opinion researchers with a valuable source of data that qualitatively differs from that offered by cross-sectional data. As Kish notes, "Only panels permit studies of individual changes; these may be needed not only for counting the frequency of changes, but also for research on the dynamics of causation and relationships"(Kish, 265). Second, given the recent rise of surveying via the internet, many companies' research methodologies rely on the creation of internet-enabled panels.

There are two important methodological issues associated with the utilization of panels: panel attrition and panel conditioning. Panel attrition refers to the possibility that panel members of certain demographics or at-

titudes may disproportionately opt out of the panel, leaving an unrepresentative panel. The concern of conditioning is roughly summarized by the proposition that panelists may change their attitudes and/or behaviors as a result of being on the panel.

Either panel attrition or conditioning produces severe problems for inference. If membership on the panel changes either the composition or the responses of the panel in a systematic manner, then the panel's responses are not unbiased estimates of the national population. As a result, absent correction, it will be impossible to correctly infer the opinions of the national population from the panel's responses.

Given the nature of the research, it is scientifically impossible to definitively and generically dismiss the possibility of panel bias. First, statistically speaking, the only inference possible is that there is no evidence of a systematic change. However, because both panel attrition and panel conditioning can occur simultaneously, it is not clear whether the correct inference from this observation is: neither panel attrition nor conditioning affect the panel's validity, or the effects of panel attrition and panel conditioning are offsetting. Although the substantive conclusion about the panel's validity is identical in either case, the causal story is very different.[1]

A second limitation arises because the results are based on a single instrument in a single issue area. I attempt to measure opinions that would

[1]These concerns are similar to those voiced by Biemer (1988) in regards to mode effect studies.

plausibly (if not probably) change as a result of being on the panel, but as is true with all such research, it is not clear how the results generalize – either across survey topics, or across different panels.

A related concern is that the panel examined differs from traditional panel studies in slight, but important, ways. Panel surveys are traditionally utilized to examine change over time.[2] As a result, panelists are repeatedly asked identical questions to measure change. Although this approach certainly enables the documentation of individual changes in opinion, the source of a change in a panelist's answer is not entirely certain. For example, it could be due to: a genuine opinion change, opinion/behavioral changes due to the fact that the panelist knows that they are part of an ongoing opinion study, and/or the fact that the panelist has seen the question previously (and possibly thought more about the question as a result). Given that Knowledge Networks' panelists are interviewed on different topics, it is doubtful that any resulting differences are due to repeat measures effects (i.e., the effect of asking identical questions at different times).

With these caveats, the results of this paper suggest that there is very little evidence to sustain the claim that a panel design produces biased population estimates - at least in terms of the panel and topic examined. This finding is surprising and unexpected given that the Knowledge Networks

[2]In contrast, Knowledge Networks uses a panel primarily to economize resources. By utilizing a panel, the recruitment cost is minimized due to the cost savings brought about by "re-using" panelists. Furthermore, since demographic measures about each panelist are collected only once, future surveys condition on this information – thereby shortening the length of interviews.

panel has several characteristics that would presumably make it more susceptible than normal to panel bias.

I establish this result as follows. In section 1, I summarize the possible sources of panel bias. In section 2, I outline the methodology employed by Knowledge Networks – whose panel I investigate. Section 3 examines the extent and effect of panel attrition in terms of available demographic variables. I outline the experiment performed in section 4. Section 5 presents the results of the experiment, and section 6 concludes.

Although numerous studies on panel bias exist (e.g., Groves (1989)), this paper makes a valuable contribution for three reasons. First, it is the first study that investigates the bias associated with an Internet enabled panel – albiet a panel that relies on Random Digit Dialing (RDD) technologies to recruit its members. Second, the panel has two characteristics that make it especially susceptible to panel bias: every panelist is provided with an interactive TV device and free Internet access, and every panelist is surveyed on a weekly basis. Thus, we should certainly expect clear evidence of panel bias in the Knowledge Networks panel. Third, the panel examined is that of a private firm. As a result, it does not have the prestige associated with a governmental or academic source – a prestige that could provide an impetus to panelists to stay involved, thereby decreasing panel attrition.

1 The Possible Biases of a Panel

I examine two possible biases introduced by a panel design. The first potential problem is that of panel attrition. Panel attrition refers to withdrawal of panelists from the panel. The second possible source of bias, panel conditioning, results from the fact that a panelist's behavior and/or opinions may change as a result of being subjected to repeated surveys (Lazarsfeld (1940)).[3]

These two possible biases can be depicted mathematically using the framework of Groves and Couper (1998). Let $\bar{y}_r$ denote the response from all panelists, $\tilde{y}_r$ denote the "true" responses of the panelists, ϵ denote the change in responses brought about as a result of being on the panel, n be the sample size and m be the number of non-respondents to the survey. Let μ be the number of panelists who have withdrawn from the panel *and* who would have been assigned the survey if still in the panel. Finally, let $\bar{y}_{m+\mu}$ be the response that non-respondents (m) and withdrawn panelists (i.e., μ) would have provided.

$$\bar{y}_r = [\tilde{y}_r + \epsilon] + \frac{m+\mu}{n+\mu}[\bar{y}_{m+\mu} - \bar{y}_r]$$

Panel attrition creates bias by increasing the non-response rate by $\frac{m+\mu}{n+\mu} -$

[3]Although it is beyond the scope of the current paper to examine hypotheses about the causal process underlying such a change, many hypothesis are possible. For example, knowing that they are to be surveyed, a panelist might consume more news – and therefore be more knowledgeable than s/he "should." Alternatively, a panelist might take response cues off of a previous survey and answer subsequent surveys accordingly.

$\frac{m}{n}$, and as well as changing the response of the non-respondents from $\bar{y}_m$ to $\bar{y}_{m+\mu}$. Panel effects affect bias by changing the observed response $\tilde{y}_r$ by ϵ to $\tilde{y}_r + \epsilon$.

Although panel attrition is inevitable, the substantive question of interest is: who withdraws, and what impact does it have on the bias. By definition, we never observe the non-response bias (i.e., $[\bar{y}_{m+\mu} - \bar{y}_r]$) because we never observe the response of the non-responders. At best, we can compare the demographic composition of m to the demographic composition of μ using variables that are known for both. Assuming that $\bar{y}_\mu$ (unobserved) is a function of the observed demographic variables (or unobserved variables highly correlated with the observed demographic variables), we could then estimate $\bar{y}_\mu$, but this clearly does not provide much (if any) leverage in assessing $\bar{y}_{m+\mu}$.[4] As a result, although I address the issue of panel attrition, the conclusions drawn from the exercise must be speculative.

In terms of panel conditioning, *a priori* it is reasonable to expect to observe effects, as it is hard to believe that interviewing panelists weekly would not induce either behavioral or attitudinal changes – even though panelists are interviewed on different topics. Additionally, the effect should be increasing in panel tenure.

I term these two components of the conditioning hypothesis the "direc-

[4] A further limitation is that one is required to assume that the mapping from demographic variables for μ to $\bar{y}_\mu$ is identical to that of m to $\bar{y}_m$. To the extent that non-respondents are different from respondents (which they are by definition), this assumption is problematic.

tional," and "magnitude" change.[5] Directional change refers to a consistent and statistically significant bias in responses across time. For example, a reasonable hypothesis is that the longer a panelist has been on the panel, the more likely it is that a panelist consumes more news and thinks themselves more influential relative to a non-panelist. Thus, response for media consumption should be uniformly higher for panelists of a sufficient tenure. In terms of the framework presented above, under the directional hypothesis, ϵ is statistically significant and either strictly positive or strictly negative.

The magnitude change hypothesis postulates that changes resulting from the directional hypothesis should be a (weakly) increasing function of panel tenure. Continuing the above example, the longer a panelist has been on the panel, the more media the panelist should consume. Thus, the magnitude hypothesis predicts that ϵ is positively correlated with panel tenure.

There have been many studies that examine panel attrition and conditioning (e.g., Kalton and Citro (1993), Sen (1976)). One source of much research has been the National Longitudinal Survey. Zagorsky and Rhoton (1999,4) summarize the results of 7 studies by indicating that no "serious problems" are created by attrition. Furthermore, they themselves find only slight differences between the demographic composition of the Young and Mature Women cohorts of the National Longitudinal Study Young and the Current Population Survey demographics even after attrition rates of 47%

[5]This is very similar to the approach advocated by Biemer for mode effect research (276,1988).

(Mature) and 42% (Young) (Zagorsky and Rhoton 1999, 30). Studies examined with a shorter field period (e.g., Price and Zaller (1993), Rahn, Krosnick, and Breuning (1994)) reach similar conclusions.

In terms of estimating and quantifying the bias resulting from attrition, a recent examination used the American National Election Studies(ANES) (Bartels 2000).As a result of fitting multiple statistical models, Bartels concludes that there is "rather little evidence of 'significant' panel biases" (Bartels 10, 2000). Bartels tempers his conclusion by noting that panels that do not share similar characteristics to the ANES may not be so negligibly affected.

2 The Knowledge Networks Methodology

I examine panel attrition and the panel conditioning hypothesis using the panel of the research firm Knowledge Networks. To determine the extent and effect of panel attrition, I track the disposition of panelists recruited in a given month to determine how the demographic composition of that group changes over time. To test the panel conditioning hypothesis I perform an experiment that measures the effect of being on the panel in terms of responses to behavioral and attitudinal measures.

As briefly noted previously, there are three reasons why using the Knowledge Networks panel provides a very strong test of the biases associated with the panel effects hypothesis. First, because Knowledge Networks provides

each household with an interactive TV device on which panelists take surveys via the internet, either the interactive TV device or the internet could affect panelists' behavior/responses. Second, since Knowledge Networks panelists are surveyed weekly (on different topics), the number of surveys that a Knowledge Networks panelist sees and completes is very high. Given these two factors, if panel bias exists, we should observe it in the Knowledge Networks panel. Third, as a private firm, Knowledge Networks lacks the prestige associated with governmental or academic panels.

Knowledge Networks utilizes list-assisted RDD (Random Digit Dialing) sampling techniques on the sample frame consisting of the entire United States telephone population. The sample frame is updated quarterly. Knowledge Networks excludes only those banks of telephone numbers that have zero directory-listed phone numbers (i.e., Knowledge Networks's telephone numbers are selected from "1+ banks"). All numbers have an equal probability of selection, and the sampling is done without replacement.

After generating the initial list of telephone numbers, the sample preparation system then excludes confirmed disconnected and non-residential telephone numbers. Next, the sample is screened to exclude numbers that are not in the interactive TV Internet Service Provider network. As of the time that the experiment was performed, this resulted in the exclusion of approximately 17% of the United States population. Telephone numbers for which Knowledge Networks is able to recover a valid postal address (about 50%) are sent an advance mailing informing them that they have been selected

to participate in the Knowledge Networks panel. In addition to containing some information about the Knowledge Networks panel, the advance mailing also contains a small monetary incentive to encourage cooperation.

Following the mailing, the telephone recruitment process begins. The numbers called by interviewers consist of all numbers sent an advance mailing, as well as 50% of the numbers not sent an advance mailing. During the recruitment interview, which typically requires about 10 minutes, the interviewer informs the household member that they have been selected to join the Knowledge Networks panel, and that in return for completing a short survey weekly, the household will be given an interactive TV set-top box and free monthly internet access. All members in the household are then enumerated, and some initial demographic variables (i.e., age, gender, household mailing address) are collected. The household cooperation rate for the period examined averaged 56%.

The first survey that the panelist receives is a profile survey that collects standard demographic information. Panelists are not eligible for subsequent surveys until completing a profile survey. 65% of all households have at least member who completes the profile survey. The first panelist joined the Knowledge Networks panel in June of 1999 and as of June 31, 2000, there were approximately 100,000 panelists associated with the Knowledge Networks panel.

3 Panel Attrition

To gauge the extent of panel attrition requires examining a group of panelists over time (this is the approach of Sobol (1959), Evan (1959) and Sharot (1991)). There are two questions of interest. The first question is simply a descriptive question - what happens to a recruitment group over time (e.g., how many drop out, how many remain active)? The second question is: who drops out? The second question clearly takes precedence over the first, as unbiased panel attrition does not compromise the statistical validity of the panel. A problem arises if some population groups withdraw from the panel more frequently than others, thereby leaving the panel unrepresentative of the national population. Recalling a previous discussion, although I find no evidence of statistically significant differences in the demographic composition of withdrawn panelists and active panelists, it is not possible to conclusively infer the impact of attrition.

To determine the extent of panel attrition, I examine the disposition of recruited panelists as of July 31, 2000 by recruitment month. Table 1 below summarizes the disposition of adult panelists who have been shipped an interactive TV unit as of July 31. "Active" indicates the number of panelists who regularly take surveys. "Inactive" indicates the number of panelists who have requested a temporary hold on their survey assignment for personal reasons (e.g., vacation). "Not complete initial survey" indicates those who have not completed the initial profile survey, and are therefore not

eligible for future survey assignments. "Withdrawn" are those panelists who have dropped from the panel, either before or after the initial survey. The table is organized so that each row represents a recruitment cohort (i.e., all the members were recruited during that month). Since the earliest reported cohort was recruited in January 2000, and the latest in June 2000, it it possible to examine the changes across the cohorts to make inferences of behavior across time.[6]

Table 1: Disposition of Adult Panelists as of July 31, 2000

Recruited in	% Active	Inactive	Not Complete Initial Survey	Withdrawn	Total
January	64%	1%	25%	10%	6734
February	62%	1%	25%	12%	8216
March	60%	1%	25%	14%	8367
April	60%	1%	25%	9%	9160
May	59%	1%	32%	9%	7990
June	54%	1%	39%	5%	8866

There are four interesting aspects of Table 1. First, although "Not Complete Initial Survey" increases from the January cohort to the June cohort, this is almost entirely due to the lag associated with getting households to install an interactive TV unit and complete the initial survey. Since the data is current as of July 31, those recruited in June have had the least amount of time to get connected – hence they have the highest non-response. Second, and associated to the first point, the percent of the panel that are active and taking

[6]Although Knowledge Networks began recruitment in June 1999, the recruitment procedures used by Knowledge Networks varied between June 1999 and January 2000. To examine panel attrition holding constant the panel recruitment process, only those panelists recruited between January 2000 and June 2000 are examined.

surveys increases 10% (from 54% to 64%) over the 7 month time period examined.[7] Third, the percentage of withdrawn panelists increases around 6% over the same period. Finally, whereas the percent that are active increases gradually, the observed data suggests that those who withdraw do so almost immediately. The percent of panelists that withdraw in a recruitment cohort is relatively stable – the withdrawal rate of the recruitment cohort with two months tenure (i.e., May) is almost identical to the cohort with 7 months tenure (i.e., January).

Although interesting, these results by themselves do not provide the means to assess the statistical quality of the panel. To perform this analysis requires comparing the demographic composition of the original recruitment cohort and the demographic composition of active panelists in the cohort as of July 31. This can be done using only those demographics that are known for both active panelists, and those that are inactive, have withdrawn or never completed the initial survey.

The only variables that are known for all panelists are those that are gathered during the initial RDD household recruitment contact. Given the nature of the call, the number of demographic items collected is rather sparse. However, using these variables, it is possible to examine the demographic composition of the panel before and after attrition.

Table 2 presents the demographic compositions of the active panelists

[7]Given the relatively constant withdrawal rate, this is due to the lag associated with the households connecting the interactive TV.

in each recruitment cohort. Although not strictly the correct comparison because of the non-response bias present prior to panel attrition (as a result of either non-response on the RDD recruitment or being out of the interactive TV network), the Current Population Survey (CPS) estimates from March 2000 for the same quantities are presented for heuristic interest.

Table 2: Demographic Composition of Active Panelists

Recruited in	**Male**	**18-34**	**35-54**	**55+**
December	50%	35%	44%	21%
January	49%	36%	45%	19%
February	49%	35%	45%	20%
March	50%	35%	47%	18%
April	50%	35%	46%	18%
May	49%	38%	46%	17%
June	50%	38%	46%	18%
CPS	50%	35%	47%	19%

Inspection of Table 2 reveals that the active panel is consistently very close to the current population estimates – at least in terms of age and gender. However, to gauge the effect of panel attrition requires investigating how the demographic composition of active panelists compares to the *entire* recruitment cohort. Table 3 enables this comparison by presenting the percent change in demographic composition between the entire recruitment cohort (including those panelists that are active, inactive, withdrawn and failed to complete the initial survey), and the active members in the recruitment cohort.

Table 3: Change in Demographic Composition of Recruitment Cohort

Recruited in	Male	18-34	35-54	55+
January	1%	2%	1%	-3%
February	1%	2%	2%	-4%
March	1%	2%	3%	-4%
April	2%	0%	2%	-3%
May	1%	0%	4%	-5%
June	1%	2%	1%	-4%

Inspection of Table 3 shows that the primary source of attrition are panelists over the age of 55.

Since Knowledge Networks conducts its surveys via interactive TV and the internet, examination of internet penetration is another useful benchmark. The question of interest is – are withdrawn panelists systematically more or less "wired" than those who remain? Unlike the demographic groups, stable baseline estimates of internet penetration are not available. As a result, it is not entirely clear how the internet composition of the panel compares the the national population.

Table 4 presents the individual level internet penetration percentage by recruitment cohort. Since this information is collected at the time of recruitment, and since national internet penetration is increasing, the June cohort necessarily has a higher penetration rate than the January cohort.

Table 4: Computer Possession By Recruitment Cohort

Recruited in	No Computer	Computer (no Internet)	Computer (Internet)
January	32%	15%	53%
February	30%	14%	55%
March	30%	15%	56%
April	28%	13%	59%
May	29%	13%	57%
June	27%	14%	60%

There is no systematic change between active and inactive panelists in terms of computer possession.

Tables 2 thru 4 indicate that although only 54% to 64% of the recruited panelists are "active" as of July 31, the panel remains largely representative of the national population in terms of the variables examined. However, the result is actually stronger, as the results in the above tables support the conclusion that the sample is representative in terms of *any* characteristic that is correlated with age and gender.

A more in-depth comparison of demographic differences using those that have withdrawn but completed the first survey (1691) is also possible. Although it does not delineate the total impact of all sources of panel attrition (as the previous analysis did), it is informative nonetheless because it allows the examination of other important demographics.

Table 5 presents the racial and ethnic composition of active and withdrawn panelists. There is surprisingly little variation between the two groups.

Table 5: Race and Ethnic Composition of Withdrawn Versus Active Panelists

Demographic	Active	Withdrawn
Hispanic	8 %	9%
White	80%	84%
Black/African-American	10%	10%
American Indian/Alaska Native	2%	2%
Asian/Pacific Islander	4%	4%

Table 6 summarizes the income and education composition of active and withdrawn panelists. There is again no evidence that withdrawn panelists differ from active panelists in terms of education, household ownership, or household income.

Table 6: Education and Income Characteristics of Withdrawn Versus Active Panelists

Demographic	Active	Withdrawn
Less than High School degree	7%	10%
High School Degree	24%	27%
Master's,Professional or Ph.D. degree	11%	12%
Own residence	74%	78%
1999 pre-tax Household $<$ \$15K	7%	7%
1999 pre-tax Household $\in$ [\$50K,\$60K)	13%	15%
1999 pre-tax Household $>$ \$125K	7%	7%

In terms of panel attrition, this section has shown that although attrition is certainly evident in the Knowledge Networks panel, its effect in terms of the demographic composition of the panel (according to the variables examined) is not pathological. Although the lack of difference does not enable us to directly quantify (and therefore address) the possible bias caused by attrition,

the demographic similarity in recruitment cohorts prior and subsequent to attrition provides some reassurance.

There are two reasons why these findings are preliminary. First, we are limited in our ability to fully assess the effect of attrition in terms of the demographic composition of the panel by the lack of available data. Second, given the youth of the Knowledge Networks panel (i.e., the earliest recruitment cohort has a tenure of only 7 months), it is not clear how these results generalize to longer tenures. A fact that gives some confidence that the results may not change significantly is the frequency with which the panelists are surveyed. Since panelists are surveyed approximately (but never more than) once a week, remaining on the panel for 7 months implies having seen (and completed) a fairly large number of surveys.

4 Experimental Design

To test the direction and magnitude hypotheses of panel conditioning, I employ an experiment to capture the beliefs and actions of panelists of varying tenure.[8] Five groups of 800 panelists were simultaneously assigned an identical instrument dealing with politics. Group assignment was determined by the panelist's recruitment date (i.e., when the respondent joined the Knowledge Networks panel).

The five groups were: 1] recruited during June 1999 and October 1999,

[8]This is similar to the approach taken by Coombs (1973).

2] recruited in November 1999, 3] recruited in January 2000, 4] recruited in March 2000, and 5] recruited during April 17-25, 2000. In other words, groups 1-4 contain panelists that vary in panel tenure, and group 5 consists of panelists who have just joined the panel *and* never taken an Knowledge Networks survey.

The 5 groups were selected to roughly correspond to the CPS demographic estimates according to age, gender, education and race. Each household was allowed to contribute at most 1 panelist. All groups except group 5 had completed at least one previous Knowledge Networks survey, with the number of surveys completed increasing in tenure length. Thus, group 5 represents a sample with no previous Knowledge Networks survey experiences, whereas groups 1 - 4 represent (increasingly) biased samples under the panel conditioning hypotheses.

By looking at the weighted response differentials among the five groups, we can cleanly test for panel conditioning. The direction hypothesis requires a statistically significant response difference between groups 1 - 4 and group 5. The magnitude difference predicts that the biggest difference in responses should be evident between groups 1 and 5, followed by groups 2 and 5 and so on (with the smallest difference between groups 4 and 5).

Recalling that the groups' demographic characteristics are selected to correspond with current CPS estimates, I explicitly control for the possible (and independent) effect of panel attrition. As a result, differences in response distributions could result from sampling variation, non-response, and/or panel

tenure bias.

The survey instrument used in the study contains 15 questions and tracks responses to three types of questions: media consumption questions (e.g., how often the panelist watches the news or uses the internet), behavioral questions (e.g., perception of panelist influence over politics, likelihood of voting), and substantive opinions (e.g., opinion of economy, Gore and Bush). The survey was assigned to respondents twice: once on 4/23/00 and once on 4/25/00 (only for new respondents added between 4/23/00 and 4/25/00). The study ended on May 15, 2000.

The survey response rate for each group is given in Table 7 below.[9]

Two items are of note. First, the response rate for group 5 is much lower than the other groups. This is due to the fact that: 1) only group 5 had to set up an interactive TV device to access the survey, and 2) the field period for some respondents in group 5 was 2 days shorter than the other groups. Second, the response rate of group 4 is much higher than than the other groups. Although it is not possible to support the claim, a reasonable claim is that panelists who have just joined the panel may be especially anxious to participate – an anxiety that subsides over time (as is indicative of the response rates of groups 1-3).[10]

[9]Note that the cumulative response rate must take account of all sources of non-response, which occurs in three stages – recruitment, completion of the profile survey, and completion of the survey itself. For group 4, the cumulative response rate is: $.56 \times .65 \times .87 = 32\%$.

[10]The willingness could either be due to an interest in participating in surveys, or a sense of guilt/obligation caused by the free interactive TV device and internet they received.

Table 7: Survey Response Rates for Assignment Groups

Group	1	2	3	4	5
Recruit Date	June 99 to Oct 99	Nov 99	Jan 00	Mar 00	Apr 17-25 00
Assigned	800	800	800	800	800
Completed	544	587	580	692	436
Response	68%	73%	73%	87%	55%

In the analysis that follows, the data is weighted according to age, education, rural/urban, and race/ethnicity. Weighting is desirable to attempt to control for survey non-response and panel attrition.

5 Panel Conditioning

Summarizing the results presented in this section, I find very little evidence of panel conditioning in the Knowledge Networks panel. What evidence I do find is neither systematic, nor consistent, with either the magnitude or direction hypothesis. For expositional clarity, I focus on panelist behavior in the three types of questions separately.[11]. To control for any non-response effects, I use post-stratification weighting to make the responses of each group representative of the national population.

The null hypothesis of the direction hypothesis is that the distribution of responses does not vary according to panel tenure. I implement a test of this null by comparing responses for groups 1-4 and responses to group 5. To the extent that the responses are statistically identical, I conclude that there is not a systematic bias in the responses.

[11]This document presents the results for a subset of the questions asked. The unreported questions exhibit a similar response pattern and are available upon request

The magnitude hypothesis' null hypothesis is that the non-zero differences between the groups' responses is not increasing in panel tenure. Thus, the difference between groups 1 and 5 is the same as the difference between groups 4 and 5. To test this null hypothesis, I inspect the response distributions to determine if a trend in differences is evident (i.e., is the difference between groups 1 and 5 larger than than of groups 2 and 5, ..., which is larger than 4 and 5 – or *visa versa*). Note that a necessary condition for the magnitude hypothesis to be true is that the direction hypothesis is true (i.e., a statistically significant difference exists).

5.1 Media Consumption Questions

I begin by examining media consumption questions. According to the direction hypothesis, given that panelists know that they will be surveyed weekly, panelists who have been on the panel longer should be more active consumers of the news and the internet.[12] According to the magnitude hypothesis, the consumption difference should be increasing in tenure length – with those who have been on the panel the longest being the most active consumers of the media. I find that neither is a correct characterization.

[12]This is consistent with the findings of Sobol (1959,65).

Table 8: How often do you watch or read the national news?

Group	1	2	3	4	5
Every day	52%	48%	52%	57%	57%
Few times	29%	34%	29%	29%	29%
Once a week	6%	7%	6%	7%	6%
Less than once	9%	9%	10%	7%	5%
Never	4%	2%	2%	2%	3%

Inspection of Table 8 reveals that the response distributions do not yield inferences favorable to the panel conditioning hypotheses. Specifically, although there appears to be a direction difference, it is not in the anticipated direction – longer tenure length is associated with *lower* news consumption. On the basis of inspection, there is very little evidence that the responses for groups 1-4 systematically differ from those of group 5. In terms of the magnitude hypothesis, there is no obvious trend associated with tenure length, as only the "Every day" response shows any sign of systematic movement.

To formalize tests of the direction and magnitude hypotheses, I transform the response categories into a dichotomous measure. For example, the response of individual i (Y_i) is:

$$Y_i = \begin{cases} 1 & \text{if } i \text{ answers "Every day"} \\ 0 & \text{else} \end{cases}$$

I model the unweighted probability that a respondent answers "Every day" as a function of recruitment cohort. Mathematically, I estimate the following model:

$$\text{Prob}(Y_i = 1) = \alpha 1 + \alpha_1\beta_1 + \alpha_2\beta_2 + \alpha_3\beta_3 + \alpha_4\beta_4 + \epsilon_i$$

where 1 denotes the constant, and β_1 through β_4 are indicator variables denoting membership in groups 1 through 4 respectively. Assuming that ϵ is distributed logistically, this model is estimated using the standard logit model. Note that the absence of any demographic characteristics as independent variables implies that we assume that the only source of variation in the probability of watching the news every day is due to panel tenure.[13]

In this framework, testing the direction and magnitude hypothesis is straightforward. To test the direction hypothesis I examine if $\alpha_1, \alpha_2, \alpha_3$ and α_4 are all not statistically differentiable from zero. Intuitively, the interpretation is that the mean percentage of respondents who answer "Every day" is identical across recruitment groups. Testing the magnitude hypothesis is less rigorous, as I rely on inspection of the α's to see if: $\alpha_1, \alpha_2, \alpha_3$ and α_4 are all similarly signed, and also that either $\alpha_1 > \alpha_2 > \alpha_3 > \alpha_4$ or else $\alpha_1 < \alpha_2 < \alpha_3 < \alpha_4$. Note that, by definition, if the null hypothesis of the magnitude hypothesis is *not* rejected, so too must the null hypothesis of the direction hypothesis.

[13] I also estimate a model that allows responsiveness of each group to vary according to whether or not the household has a computer. Although this data is only available for groups 2-5, the (unreported) results of this analysis is identical to those reported in the text – there is no differential effect conditional on having a computer in the home.

This model specification correctly predicts 53.4% of the observations (as opposed to 53.1 % by the naive model), with the following results.[14]

Table 9: Results for "Every Day" News Consumption

Coefficient	Value	Standard Error
α	.23*	.10
α_1	-.24	.13
α_2	-.33*	.13
α_3	-.15	.13
α_4	.02	.12
PCP	53.4 %	

Inspecting the coefficients reported in Table 9 confirms the conclusions reported above. Only α_2 is statistically significant at traditional levels of significance (i.e., probability of Type I error = 5%), although α_1 is statistically significant at a Type I error level of 10%. The inference to be drawn is that there is some evidence that the oldest groups (i.e., groups 1 and 2) are more likely not to consume news every day than average. However, all other groups (i.e., groups 3,4 and 5) are all identically likely to consume news every day. As a result, although some evidence for the difference hypothesis can be found, evidence for the magnitude hypothesis is strained given that all but one of the coefficients are not statistically different from zero (or two under a weaker statistical test).

A harder test of the panel effects hypothesis involves investigation of Internet usage. In some sense this test is actually *too* strong of a test, as a

[14]In all of the models estimated, the estimated models do only fractionally better (if at all) relative to the naive model. Although not a statistically meaningful assessment of model fit strictly speaking, the relative lack of improvement is further evidence that panel tenure does not explain the observed variation.

large amount of the variation could be driven by the fact that panelists in group 5 have only just connected the interactive TV device. Additionally, panelists must be on the Internet to complete the survey.[15] Table 10 presents the weighted response distribution by recruitment cohort.

Table 10: How many days in the last week have you been on the Internet?

Group	1	2	3	4	5
None	21%	19%	18%	10%	33%
1 day	20%	14%	16%	16%	17%
2-3 days	24%	29%	30%	29%	18%
4-5 days	16%	18%	15%	19%	13%
6-7 days	19%	20%	22%	26%	19%

Proceeding directly to the statistical test, I predict the unweighted probability that an individual answers "None." Table 11 presents the resulting coefficient estimates.[16]

Table 11: Results for "None" Internet Usage

Coefficient	Value	Standard Error
α	-.87*	.11
α_1	-.54*	.15
α_2	-.75*	.15
α_3	-.72*	.15
α_4	-1.38*	.17
PCP	82.5%	

The results in Table 11 show that the direction hypothesis is indeed relevant, as α_1 through α_4 are all statistically significant and of the same sign

[15] Which gives rise to the question – why the large percentage of "None" responses. The interpretation that must be given is that these panelists exclude the current session from consideration.

[16] Despite the statistically significant coefficients, the predictions of the model are identical to that of the naive model.

(i.e., negative). Given the caveats noted above, this is unsurprising – existing panelists are more likely to have used the internet in the past week than new recruits because existing panelists would have taken surveys over the internet previously.[17]

Support for the direction hypothesis is more ambiguous. Clearly there is no monotonic trend in the coefficients. Ordering the groups by non-usage, group 5 has the highest level of non-usage, followed by 4,3,2, and 1. In other words, the most recently recruited group of existing panelists are the most active users of the internet. The likely explanation, similar to that supplied previously in explanation of group 4's high response rate, is that new panelists are excited about the Internet enabled interactive TV appliance supplied to them – using it often. However, as time passes, the novelty of the interactive TV device wears off and the panelists' non-usage increases. Regardless of the veracity of this explanation, the finding runs contrary to the expectation that providing panelists with internet access turns them into a "wired" panelists. The evidence shows that, if anything, this might be true for the first month following recruitment. However, after the initial excitement, panelists apparently resort to their traditional behavior.[18]

[17]When internet usage is allowed to vary by tenure and computer ownership, the only significant difference that emerges is that group panelists in group 3 with a computer were more likely to use the internet. Given that this effect is not systematic (i.e., there is no effect on panelists whose tenure is either younger or older than group 3), it is not clear what to make of this finding.

[18]Of course, this interpretation is a bit strained as their "traditional behavior" is not known.

5.2 Behavioral Questions

Turning to the responses to behavioral questions, were the panel conditioning hypotheses true, panelists' interest in politics should change. The direction of the change is less certain, as it is not clear whether panelists would become more or less interested in politics as a result of being surveyed on political topics. As Tables 12 through 15 below show, I find no evidence of change in either direction.

Table 12: How interested are you in politics and public affairs?

Group	1	2	3	4	5
Very interested	20%	20%	21%	24%	23%
Somewhat interested	43%	43%	40%	40%	37%
Slightly interested	28%	31%	27%	24%	30%
Not at all interested	9%	7%	10%	8%	6%

Determining the change in terms of the percent of respondents that are either "Very" or "Somewhat" interested, I find that there is no evidence for the direction (and therefore the magnitude) hypothesis. Table 13 presents the evidence.[19]

Table 13: Results for "Very" or "Somewhat" Politically Interested

Coefficient	Value	Standard Error
α	.60*	.09
α_1	-.09	.13
α_2	-.07	.13
α_3	-.06	.13
α_4	.16	.13
PCP	64.4 %	

[19]The model generates the same predictions as the naive model.

The lack of significance across all groups is immediately obvious – the implication being that panel tenure does not appear to change the percentage of panelists that express more than a slight interest in politics.[20]

The percentage of respondents reporting an intention to vote is also invariant across panel tenure.

Table 14: Are you planning to vote in the upcoming presidential election?

Group	1	2	3	4	5
Yes	74%	77%	74%	75%	72%
No	12%	10%	14%	12%	11%
Don't know	14%	12%	12%	14%	17%

Similar to the conclusion drawn from political interest, there is little evidence that intention to vote differs across the recruitment groups.[21] As only the vote intentions of group 2 varies from that of group 5 in a statistically significant manner, there is no evidence of any systematic effect of panel tenure on the respondents' intention to vote. Table 15 documents the lack of effect.[22]

[20]There is no difference in political interest in groups 2-5 conditional on panel tenure and computer ownership.

[21]As is the case for political interest, the lack of difference holds even when the intentions of panelists are allowed to vary by both panel tenure and computer ownership.

[22]Similar to the models above, the fit of the model is indentical to that of a naive model.

Table 15: Results for "Yes" Vote Intention

Coefficient	Value	Standard Error
α	1.2*	.12
α_1	-.06	.15
α_2	.08	.15
α_3	-.10	.15
α_4	-.00	.15
PCP	77.2 %	

5.3 Substantive Questions

In some sense, of the three question types I examine, differences in responses to the substantive questions are the ones that we care most about, as it is these types of questions that form the basis of much of public opinion research. Although I can still examine the direction hypothesis, examination of the magnitude hypothesis becomes more ambiguous, as it is not clear how panel membership should systematically bias responses to substantive political questions.

Regardless of this uncertainty, the results clearly show that opinion on the substantive questions does not systematically change as a function of panel tenure. Given the strength of the similarity, only the response distributions on 2 of the 7 substantive questions are provided – although the conclusions generalize to the unreported questions.

Table 16: If elected president, how good of a job would George W. Bush do as President?

Group	1	2	3	4	5
Very Good	9%	8%	8%	8%	8%
Good	24%	29%	24%	28%	23%
OK	40%	43%	42%	37%	45%
Poor	18%	16%	20%	17%	16%
Very Poor	8%	5%	7%	10%	9%

Table 17: If elected president, how good of a job would Al Gore do as President?

Group	1	2	3	4	5
Very Good	11%	7%	9%	11%	9%
Good	22%	24%	22%	21%	24%
OK	42%	41%	40%	39%	40%
Poor	14%	17%	18%	17%	20%
Very Poor	11%	12%	11%	12%	7 %

Considering the response distribution for either question shows that there is no statistically significant difference in response distributions based on tenure length. Given that the direction hypothesis is unsustainable for these questions, so too must the magnitude hypothesis.

6 Conclusion

Panel bias has many sources. Two prominent sources are attrition and conditioning. As a result of examining the panel of Knowledge Networks over a 7 month time period – a panel that would seem to be especially prone to panel bias – I find little evidence of either systematic panel attrition or panel conditioning.

The panel conditioning hypothesis consists of investigating two related concerns: first do panel members differ from non-panel members (i.e., direction hypothesis), and second, is the difference increasing in tenure length (i.e., magnitude hypothesis). To examine the existence and magnitude of these two concerns, I survey 4,000 panelists (divided into 5 demographically balanced groups based on tenure length) simultaneously with an identical instrument dealing with politics. Looking at the response distributions across the groups of varying tenure, I find little support for either proposition.

Although it is not clear how these results generalize to other panels, other issues and longer tenure lengths, there is some reason to suspect that the results are robust. Specifically, the panel of Knowledge Networks is in some sense the hardest test of the null hypothesis of no panel bias. Three reasons contribute to this: first, Knowledge Networks is a private firm, lacking the prestige associated with academic and governmental studies, second, Knowledge Networks provides every panelist with both an interactive TV device and Internet access, and third, Knowledge Networks' panelists are surveyed weekly. Individually, any one of these factors would make one believe that the panel of Knowledge Networks is susceptible to bias. That all three exist makes the repeated finding of "no difference" even more believable – and more conclusive.

7 Works Cited

Biemer, Paul P. 1988. "Measuring Data Quality," in *Telephone Survey Methodology* ed: Robert M. Groves *et. al.* John Wiley & Sons:NY.

Coombs, Lolagene. 1973. "Problems of Contamination in Panel Surveys: A Brief Report on an Independent Sample, Taiwan, 1970," in *Studies in Family Planning*, 4:10 (Oct.). pp.257-61.

Evan, William M. 1959. "Cohort Analysis of Survey Data: A Procedure for Studying Long-Terms Opinion Change," in *Public Opinion Quarterly*, 23:1 (Spring) pp.63-72.

Groves, Robert M. and Mick P. Couper. 1998 *Nonresponse in Household Interview Surveys.* John Wiley & Sons:NY.

Groves, Robert M. 1989 *Survey Errors and Survey Costs.* John Wiley & Sons:NY.

Kalton, G. and C.F. Citro. 1993 "Panel surveys: Adding the fourth dimension," in *Survey Methodology*, 19 (2): 205-215.

Lazarsfeld, Paul F. 1940 "Panel" Studies," in *Public Opinion Quarterly*, 4:1 (Mar.) pp. 122-8.

Price, Vincent and John Zaller. 1993. "Who Gets the News? Alternative Measures of News Reception and Their Implications for Research," in *Public Opinion Quarterly*, 57:2. (Summer), pp. 133-164.

Rahn, Wendy M. Rahn, Jon A. Krosnick, and Marijke Breuning. 1994. "Rationalization and Derivation Processes in Survey Studies of Political Candidate Evaluation," in *American Journal of Political Science*, 38:3. (Aug.), pp. 582-600.

Sen, A.R. 1976. "Developments in Migratory Game Bird Surveys," in *Journal of the American Statistical Assocaition*, 71:353. (March), pp.43-48.

Sharot, Trevor. 1991. "Attrition and Rotation in Panel Surveys," in *Statistician*, 40:3 pp. 325-31.

Sobol, Marion Gross. 1959. "Panel Mortality and Panel Bias," in *Journal of the American Statistical Association*," 54:285 (Mar.) 52-68.

Zagorsky, Jay and Pat Rhoton. 1999. "Attrition and the National Longitudinal Survey's Women Cohorts," Manuscript, Center for Human Resource Research, Ohio State University.