

Increasing Response Rates & Data Quality of Web Surveys: Pre-Notification and Questionnaire Paging Format

Dominic Lusinchi

Far West Research, San Francisco, California
dominic@farwestresearch.com

Abstract

Although web-based survey research of the general population has serious limitations (e.g., non-coverage, lack of a sampling plan), it has been embraced by many who do survey research. It has been shown to be particularly useful when the population surveyed is almost universally connected to the Internet, and is proficient with the medium and the tools to use the medium (computing devices: e.g., desktop computers, PDAs, etc.).

In this experiment conducted on a panel of electronics engineers and engineering managers, we explore issues that may affect response rates, and response quality. Respondents are asked to report on the process of designing an integrated circuit.

First, we examine the effect of a time-honored method on both unit response rates and speed of delivery: a random half of the panel received a notification in advance of the survey invitation, informing them that they would be asked to participate in a few days; the other half did not. Second, we analyze what effect two different formats of the same instrument have on response quality. One half of the panel was randomly assigned to a scroll-type web questionnaire, in which the instrument is loaded onto one web page and the respondent scrolls down from one question to the next; the other half of the panel received a multiple page-type questionnaire, in which the respondents click on a button to move from one question to the next. Is one type more advantageous than the other?

Measures of quality used will be the break-off rate (the proportion of respondents that abandon the survey) and item response rate (looking both at errors of omission, failing to answer a relevant question; and errors of commission, answering a question for which the respondent is not eligible).

Keywords: pre-notification, unit nonresponse, response speed, screen- vs. scroll-type web survey, item nonresponse

1. Introduction

Although the World Wide Web, as an open-access communications infrastructure, is less than a generation old, many in the field of survey research recognized very early its potential as a data collection tool. Thus, it was quickly adopted and now is part and parcel of a survey research organization's data collection tool kit.

2. Research Background

The Internet as a data collection tool shares common attributes with longer established modes: like paper-and-pencil, it is a self-administered mode; but like CATI, its flow can be entirely controlled by the survey administrator. Of course, it has features all to itself: for instance, its multi-media capabilities, *i.e.*, sound and graphics (both static and interactive) can be incorporated into the instrument.

As any survey mode, the web is prone to various sources of error. Response error or bias is a particularly vexing one: the researcher must find ways of reduce the incidence of nonresponse. Maximizing the response rate insures that the response bias (the unknown difference between respondents and nonrespondents) is less likely to affect the accuracy of the sample

estimates: the values that are reported as the survey results. Furthermore, a smaller sample size, which results from nonresponse, increases the standard error of these estimates. To be effective, techniques to increase the response rate must raise it above a threshold; small increments in the response rate may continue to yield the same bias already present, because the effort on the part of the researcher entices only a certain segment of the population.

Unit response rate is one measure of survey quality, but, by no means the only one. Item response rate is another measure: here again the researcher must find ways to minimize its impact which compounds that of unit nonresponse. Another threat to the validity of survey results is measurement error, where the respondent, willingly or unwillingly, provides a response that is incorrect.

The purpose of a pre-notification is to encourage sampled individuals to become survey participants and thus minimize the rate of nonresponse. The goal of the questionnaire format is to eliminate, or at least reduce, item nonresponse and measurement error.

2.1. Research on pre-notification

It is generally considered that providing prospective survey participants with an advance letter informing them that they have been selected for an upcoming study and letting them know when they can expect to receive the actual survey invitation has a positive impact on response rates irrespective of the data collection mode employed (in-person, telephone, or self-administered paper-and-pencil). (Linsky, 1975; Fowler 2002; see Slocum, Empey, and Swanson, 1956, for an example of a study in which an advance letter was used prior to a face-to-face interview) It is also recommended particularly when one expect a low overall response rate (<50%) because any additional contact with potential participants is likely to increase participation (Wright, 1995).

For example, in an experiment involving a random sample of university students in the Houston area, the investigator sent an advance letter to one group of students who were told of an upcoming paper survey (Stafford, 1966). The response rate of the experimental group was more than twice (43.7%) that of the control group (20.5%). A third group, which was contacted ahead of time by telephone (those were the days before telemarketers and answering machines), had the highest response rate (68.2%). In another experiment (Ford, 1967), also conducted during the course of two mail surveys, advance letters were mailed, in one survey, six days prior to survey invitation itself, and, in the second, thirteen days before. Again, the use of an advance letter showed a positive effect on response rates. However, the response rates were low (less than 40%), and the differences between the two experimental group was under seven percentage points. Similar results were obtained in a mail survey on consumer credit (Walker and Burdick, 1977): individuals that received an advanced notice (postcard or letter) had a significantly higher response rate (52.4%) than the control group (38.8%) which received no pre-notification.

Advance letters have also been used in establishment surveys. In one experiment conducted by the Bureau of Labor Statistics (Chun and Robertson, 1995), establishments that received a pre-notification had, overall, a higher response rate (52.0%) to the self-administered mail survey than firms that were not notified (40.8%). Moreover, investigators found interesting variations in the level of effectiveness of advanced letters. Among small establishments, i.e. those employing less than 50 employees, there was no significant difference between the experimental (letter) and control (no-letter) groups: their respective response rates were 58.8 percent and 60.6 percent – in fact this result went against the researchers' expectations. A similar finding was obtained from large establishments (with 250 or more employees): although the difference was larger (8.5%), it too was not significant. Advance letters were at their most effective with medium size establishments (between 50 and under 250 employees); the response rate of the experimental group was 59.4 percent, some 27 points higher than that of the control group (29.0%).

In addition, meta-analytic reviews of a large number of experiments designed to enhance the response to mail surveys reach the conclusion that pre-notification generally encourages participation on the part of sampled individuals (Heberlein and Baumgartner, 1978; Fox, Crask and Kim, 1988; Yammarino, Skinner and Childers, 1991).

There are exceptions to this general rule. For instance, Parsons and Medford (1972) found that an advance letter had no effect on response among two "homogeneous" populations: one study surveyed alumni of an MBA program at a large private university; the second study examined the leaders of two religious sects. They concluded that "an advance notice is not necessary when the sample is from a fairly homogeneous population." A somewhat similar result is reported by Wright (1995) in a study of local government officials who were recently elected for the first time: a difference of 3.9 percent separate the overall response rate of the group that received an advanced letter from that of the control group – a result that was not statistically significant. In cases as these, pre-notification adds to the cost burden of the research organization, while

yielding little in return.

Advance notification has also been used successfully when conducting telephone surveys. As with mail surveys, overall results appear to be positive but there are exceptions in this mode also. In a study using both an RDD sample and a list-assisted sample, a statewide omnibus survey conducted in Illinois (Parsons, Owens, and Skogan, 2002) found that the experimental condition (advance letter) had no effect on participation. In another experiment, reported by the same authors, with an RDD survey of the city of Chicago, they found, contrary to their first study, that the response rate was significantly higher by 5.3 percentage point among households that received the advance letter. Although, the citywide response rates were lower than they were in the statewide survey, the investigators were not able to explain why the advance letter had an effect in only one of the experiments.

A telephone survey of 13 neighborhoods in south Indianapolis (Parks, Kennedy, and Hecht, 1994) found that, generally, a pre-notification letter helps to boost the cooperation rate of sampled households. The higher cooperation rate obtained from pre-notified households varied greatly from one neighborhood to the next: from a minimum of 66.3 percent to a maximum of 80.8 percent. The investigators wondered whether the apparent effectiveness of the advance letter was not simply a reflection of the fact that people who are willing to have their phone numbers listed in the White Pages may also have a higher propensity to cooperate than people not listed. In a second study of 37 neighborhoods, including all of Indianapolis and several small municipalities, although they found that the use of an advance letter increased the cooperation rate by 9 percent overall, they also observed that some neighborhoods that had not received a pre-notification had higher cooperation rate than neighborhoods that had received an advance notice. This led the investigators to conclude that the effectiveness of pre-survey letters might be determined by the characteristic of the neighborhood, and therefore might not be necessary in all cases.

In a RDD survey of preventive health practices and risk behaviors, Link and Mokdad (2005) found that among households that had been sent an advance letter, the response rates of respondents was higher (55.2%) than that of respondents whose household had not been sent a letter (49.4%). Although, this difference is not large (5.8 points), it is significant ($p < .001$) – the sample size for this study being quite large ($n > 5200$). Moreover, the researchers found that the cost per 1,000 completed interviews was lower for the group that received the letter (\$18-574) than for the control group (\$20-011). Positive results were also reported in an epidemiologic survey of gastrointestinal illness in British Columbia conducted by telephone: those who received an advance letter had response rate that was 1.58 times higher than those who were not contacted in this way. (Majowicz et al., 2004). Similar results were reported of a telephone survey of registered voters in Arizona. (Goldstein and Jennings, 2002) Those who received a pre-notification letter had a response rate that was 12 percentage points higher than the control group. Other studies that report a positive effect from an advance letters are Dillman, Gorton Gallegos and Frey (1976) where the refusal rate of the experimental group was about half that of the control group, and Hembroff et al. (2005) which report a small but significant difference (5.4%) between those that were notified ahead of the study and those that were not.

Exceptions to the rule are reported for telephone surveys also. Woodruff, Mayer, and Clapp (2006) in a pilot survey exploring the health behaviors of teens and their parents, found only small and nonsignificant differences between the advance letter and the no-letter groups.

The effect of advance notification on *response speed* has also been explored but to a lesser extent than its impact on unit nonresponse. In a study mentioned earlier, Ford (1967) found no difference in response speed (measured as the cumulative percent of returns) between the group who had received the pre-survey letter and the group who was sent simply the survey invitation. However, Wright (1995) shows that the response speed was higher among those who received an advance notice as measured before the first reminder was sent out: the response rate of the experimental group was 56 percent compared to 41% for the control group.

2.2. Research on web-questionnaire paging format

Not surprisingly there is a smaller body of research on web questionnaire formatting than there is on the effects of pre-notification. After all the medium is a relative newcomer and its full potential as data collection tool is still being explored by survey researchers.

The research reviewed here explores the impact of two types of web survey designs:

- 1) the first type, in which the entire questionnaire is implemented in one HTML file or one screen (scroll-based);
- 2) the second type, in which the questionnaire is broken into a number of separate “pages” or “screens” (HTML files), each

page being loaded after the previous page has been submitted (screen-based).

Researchers at the University of Connecticut (Clark and Nyiri, 2001) conducted two surveys, one of undergraduates, and another of graduate students covering different topics. Their goal was to observe the effect the two survey designs might have on item nonresponse. They hypothesized that the scroll-type design would have a higher rate of item nonresponse than the screen-based design. They believed that a multi-screen design “forces respondents to focus in on the question at hand”, while a single-screen design, which lets respondents navigate through the entire survey, will increase the likelihood of respondents answering the questions out of order thereby enhancing the risk of skipping (intentionally or not) over some questions. They also expected that the screen-based design would elicit longer responses to open-ended questions.

The results of the survey of undergraduate students showed no significant difference between the two designs on item nonresponse. They found a small but significant difference on open-ended questions. The difference went against their initial assumption: on average, the length of response to open-ended questions in the scroll-based design was 1.7 words longer than in the multi-screen design.

The results of the graduate students’ survey showed that the screen-based design had a higher rate of item nonresponse and break-off than the scroll-based design.

A comparison between the two modes was also done at the University of Illinois at Chicago (Burriss et al., 2001). The topic of this study, conducted by the Survey Research Laboratory, was building safety. The researchers sampled the faculty, staff and students of the College of Urban Planning and Public Affairs.

One issue the researchers investigated was the time it took to complete the questionnaire. They found no significant difference between the two designs. They also observed that in the scroll version more respondents (37.3%) changed their answers before submitting their survey than respondents to the multi-screen version (23.9%). More multi-screen respondents (56.5%) reported reviewing their answers before submission than did the scroll-based respondents (20%). Furthermore, multi-screen respondents (69.6%) reported having more experience with web surveys than did respondents to the scroll version (52.0%). Unfortunately, the study had a relatively small sample size (about $n = 50$ for each treatment), which made it difficult to reach statistically significant conclusions.

A survey conducted among German high school students (Fuchs, 2001) also experimented with two forms of a web survey. Unlike the University of Illinois design, where each question item was set on a separate screen (for a total of 12 screens), this study examined two forms of multi-screen surveys. One version only had one question item per screen; the other version had multiple items on the screen. It was found that including multiple items on the screen required less time to complete the survey. No significant difference could be detected between the two modes on the distribution of responses to questions items using scales. Although the researchers stressed that what they presented were preliminary results, they stated that “showing several items on one screen might change the respondent’s understanding and interpretation of the question.”

In a similar study Couper, Traugott, and Lamias (2001) explored several types of web design: in one the respondents are presented with only one question per screen, in another each screen shows a set of related questions. They hypothesized that grouping related items on the same screen should increase the internal consistency (correlation) among items, and take less time to complete than if presented separately, one item per screen. The results did show that items presented on a single screen were more correlated than items provided on multiple screens, but the effect was small and did not achieve statistical significance. In contrast, their second hypothesis was supported: providing one item per screen took respondents significantly longer to complete than when they were presented with multiple items per screen: on average, over the 16 items tested, it took 26.2 seconds longer with the multiple screen format than with the single screen format.

In a survey at the University of Michigan that asked students about their use of alcohol, tobacco, and drugs, the investigators (Peytchev et al., 2006) experimented with a scrollable version of the questionnaire and a screen-based version. They explored the effect the questionnaire format had on break-off rates, completion times, data quality, as measured by the proportion of non-substantive answers (“Don’t Know”, “Refuse”), and respondent errors (answering items for which the respondent is not eligible or failing to answer eligible items). They found no significant differences in break-off rates between the two versions. They found that the scrolling version had a significantly higher overall completion time (23.53 min.) than the paging version (22.93 min.), but the difference is small enough to be of little import. The scrolling version also had slightly more non-substantive answers (3.8%) than the multi-screen version (1.0%).

3. Survey Background

The survey that was used for this experiment is part of a series of studies that take place yearly on different aspects of integrated circuit (IC or chip) design. The survey instrument was a self-administered web-based questionnaire. It contained 26 question items in all. Many of the individual questions had several entries, *i.e.*, more than one piece of information is requested per item: for example, question 2 asks respondents to provide information on the chip's attributes (number of logic gates, type of embedded memory, clock speed, etc.); or question 25 asks what software tools (classified in eleven different categories) were used for the various tasks in the process of designing the chip.

The respondents are electronics engineers or engineering managers in North America. They belong to a research panel maintained by the survey organization and from which a total of $n=601$ were selected for this particular study.

Most of them (58%) are employed by semiconductor companies (*e.g.*, AMD, Intel, National, *etc.*), the rest work for systems companies that design ICs, not for commercial sale, but for use in the products they develop and sell on the market (*e.g.*, Hewlett Packard, IBM, Xerox). Table 1 below illustrates the distribution of the sampled panel members.

Table 1: Distribution of panel members by industry

Industry	Survey Panel	
	Frequency	Percent
Semiconductor	346	57.6
Computer	143	23.8
Communications	112	18.6
Total	601	100.0

Most of the professionals on the panel have participated in at least one prior survey. Therefore, they are familiar with the *modus operandi* of the study. All studies have some basic characteristics in common: potential respondents were incentivized in the following ways: post-survey cash and non-cash incentives were promised to respondents who completed the questionnaire, *i.e.*, all potential respondents are promised \$25 cash and a chance of winning a grand prize (an electronic item like an iPod), if they fill out the survey. Thus, there is no pre-survey incentive just the promise of one; panel members who have participated in prior surveys know that the organization sponsoring the research will honor its pledge. Most panel participants are also aware that the questionnaires are very demanding in terms of the time required to complete them and the level of detailed technical information requested. Most of them have previously received follow-up email messages from the investigators requesting that they clarify and/or validate some of their answers.

This high-tech occupational group is a very challenging one to survey. The task of designing an integrated circuit is an extremely complex one. Furthermore, chips themselves have seen their complexity increase dramatically: millions of transistors are being packed on an ever shrinking silicon wafer, and drawn feature sizes are reaching into the .13 micron range ($1/1000^{\text{th}}$ the thickness of a human hair), thereby allowing ever more functionality to be embedded onto the chip. Business constraints in the form of time-to-market pressures compound these technical issues making an engineer's professional environment a very demanding one.

The questionnaire elicited quantitative information regarding an integrated circuit design that was completed in 2001. Specifically, respondents were asked to provide numerous physical quantities such as those pertaining to the characteristics and attributes of the chip (number of transistors, bytes of embedded memory, clock speed, *etc.*), the length of its design cycle, the size and composition of the team of engineers that was assembled to develop it, *etc.* The ultimate goal of the survey is to obtain parameter estimates of the "population" of chip designs in North America.

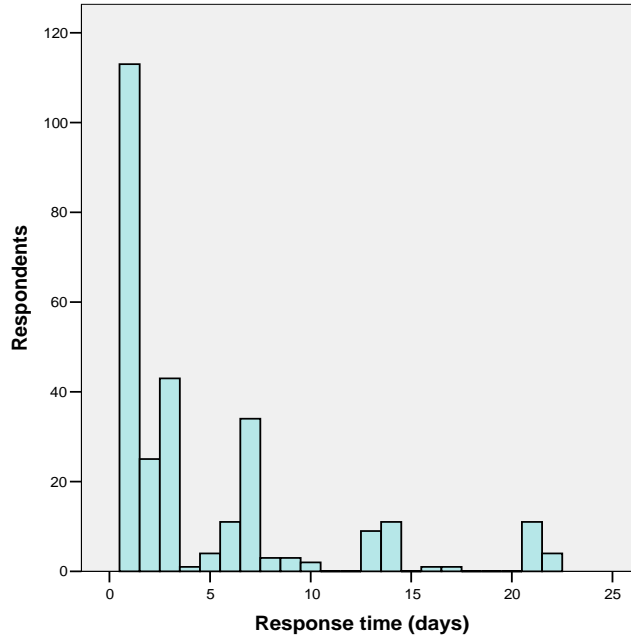
The survey was fielded over a period of about three weeks in the fall of 2001: from October 31 to November 23, to be precise. About 14 percent ($n=84$) of the original 601 emails sent to the panel to either inform members of the upcoming survey or invite them to the survey were undeliverable: in other words, the research organization received notices to that effect from the recipients' servers. It is assumed that the remaining messages were safely delivered to their intended addressee.

A total of three reminders were sent to panel members: the first was sent two days after the beginning of the field period; the second was sent four days after the first; and the third was sent a week after the second reminder or 10 days before the end of

the field period.

The median response time was 2 days, and by the end of the first week of the field period more than 80 percent of entries had been submitted. As we see in Figure 1 (n=276), the largest group of entries (41%) was submitted on the first day, *i.e.* the day the email invitation was sent out. We also notice the effects of the reminders: these are the peaks, in the histogram, on the third day, the seventh day, and the fourteenth day. The first reminder brought in forty-three respondents; the second, thirty-four; and the final one, eleven. In all, between the invitation and the three reminders, nearly three-quarters of the respondents submitted their entries; the remainder responded on days they were not contacted.

Figure 1: Distribution of response time (days)



Because of the specialized nature of the topics covered by the survey, it was necessary to make sure that potential respondents were familiar with specific areas of chip design. Therefore, the home page of the survey included a couple of screening questions. Those who did not meet the specified criteria for the survey were directed to an exit page (ineligible); the others were sent to the survey questionnaire proper (eligible).

By the end of the field period, the original sample was distributed in the following manner:

Table 2: Survey Outcome (%)

Outcome	Panel members
Ineligible	58 (9.7)
Completed	202 (33.6)
Break-off	16 (2.7)
Nonresponse	241 (40.0)
Undeliverable	84 (14.0)
Total	601 (100.0)

For this study, we measured the level of survey participation by using a “cooperation rate”, which was defined as the proportion of panel members that received the email invitation and answered the screening questions on the survey’s home page. This measure should not be construed in any way as representing a true response rate, since it does not take into account the probability of accepting the original invitation to participate in the panel. It is merely a measure of the level of

participation of those panel members who were selected for this study. The overall cooperation rate for this study was 53.4 percent, and was calculated in the following way (using the outcome numbers in Table 2):

$$COOP = \frac{C + N + I}{P - U} = \frac{202 + 58 + 16}{601 - 84} = \frac{276}{517} = 0.534$$

where:

COOP = Cooperation rate

C = Completed questionnaires

N = Ineligible (answered the screening questions and disqualified themselves)

I = Incomplete/break-off (answered too few items; or began answering the questionnaire but abandoned it)

P = Sampled panel members

U = Undeliverable (email invitation was not received by intended recipient)

Including “ineligibles” in this formula seems entirely justified. It is clear from the distribution in Table 2 that if someone comes to the survey home page, answers the eligibility questions, and meets eligibility criteria that individual is highly likely to fill out the questionnaire. Indeed, of the 218 that reported being eligible (completed and break-off) only 7 percent did not complete the questionnaire. A stricter definition of “cooperation” could eliminate the “incompletes”, the actual (n=16) and the estimated (n=58*.07=4), which would give us a rate of 49.5 percent.

Although the absolute value of the response rate to a survey is of critical importance, what this study focuses on is the differential effect, if any, of an advance notice has on participation. Of course, one would wish as high a response rate as possible; however, the latter does not depend solely on one factor such as an advance notice. What we seek to determine is whether that technique can be added to others as a reliable tool to enhance the response rates of sample surveys.

Table 3 below shows the distribution of panel members by industry at different stages of the survey process: column 2 describes the distribution of the original sample (see Table 1); column 3 illustrates the industry distribution of the panel members who were classified as respondents (as defined in the cooperation rate above); the last column (4) shows the distribution of those who completed the survey instrument. We note that the differences from one distribution to the next are rather small: the maximum difference is 2.8 percentage points (the difference, in the Semiconductor industry between the original panel and the completed questionnaires). Furthermore, none of these disparities is statistically significant.

Table 3: Distribution of panel members by industry

Industry (1)	Survey Panel (%)		
	Original (2)	Respondents (3)	Completed (4)
Semiconductor	57.6	60.1	60.4
Computer	23.8	21.7	22.3
Communications	18.6	18.1	17.3
Total	100.0	100.0	100.0
n	601	276	202

4. Method: Experimental Design

Two experiments were carried out for this research. The purpose of the first was to assess the effect that a pre-notification email would have on both survey participation and response speed (how long it takes respondents to submit their entries). This experiment was conducted with the sampled panel members: half of them (n=301) received an advance notice of the survey, the others (n=300) did not. The pre-notification email was sent on a Thursday, six days prior to the invitation email was sent off (the following Wednesday).

The second experiment implemented two versions of the web questionnaire. One version was a scrollable type instrument: in other words, the questionnaire consists of one large HTML file or screen that contains all the questions; in this mode, respondents scroll in order to follow the flow of the questionnaire. Thus, respondents are free to view the entire questionnaire at any time, if they choose to do so. The second version of the questionnaire consisted of multiple HTML files: each file or page contains one or more question item. In order to advance to the next question, respondents must click on a

“NEXT” button. This version of the survey consisted of thirty pages, including the home page and the contact information page. In contrast, the scroll-type version had two pages: the home page and the questionnaire page. The respondents assigned to this version had click once to enter the questionnaire proper and a second time to submit their responses.

The participants in this experiment were randomly assigned to one or the other version of the questionnaire: thus n=300 were selected for the scroll-type version and n=301 were sent to the multiple-page version.

The goal of this second experiment was to assess the effect of web questionnaire’s page format on data quality as measured by break-off rates (individuals who fill out part of the survey but then give up) and item response rates (items that should have been answered but are left blank).

Table 4 below describes the distribution of the panel for the two experiments at the outset of the field period.

Table 4: Distribution of Experimental Conditions (original)

	Experiment	
	Advance notice	Multipage
No	300	301
Yes	301	300
Total	601	601

After eliminating the individuals who could not be reached (n=84), panel members were distributed as described in Table 5.

Table 5: Distribution of Experimental Conditions (final)

	Experiment	
	Advance notice	Multipage
No	253	259
Yes	264	258
Total	517	517

By the end of the field period, 517 individuals were believed to have been contacted and their distributions on the experimental conditions were as described by Table 5 above. We see that in the pre-notification experiment, the treatment group was slightly larger than the control group. As for the web questionnaire format experiment, both groups were almost evenly divided.

5. Experimental Results

In this section we will review the results of the experiments that were conducted during the course of the survey. Using both descriptive and inferential statistical methods, first, we examine the outcome of the pre-notification experiment, then that of the questionnaire format experiment.

5.1 Introduction

As mentioned before, the purpose of the pre-notification experiment was two-fold: first, to determine whether individuals who are told in advance, by means of an email message, that they will be invited a few days hence to participate in a survey, have a higher propensity to respond than those who are simply invited without warning. In other words, we wanted to see, with this experiment, whether the general consensus in the literature about the benefits of pre-notification on response rates also held for email pre-notification to a web survey. The second goal of this experiment was to explore whether an advance notice also encourages invitees to submit their entries faster than individuals who did not receive such notice. This was a purely exploratory investigation since very little research, as far as I know, has tested this hypothesis (see Ford, 1967).

In the second experiment, we wanted to determine, if one web questionnaire format was superior to the other in terms of response quality. Is one type of questionnaire associated with a higher break-off rate? Is item nonresponse more prevalent in one type of format than in the other?

5.2 Advance Notice Experiment: Results

Table 6 describes the overall outcome of the pre-notification experiment. We observe that those who were sent the advance email participated at a higher rate (55.7%) than those who simply received the survey invitation (51.0%): this result at first appears to go in the direction indicated by many studies on pre-notification. However, this small difference, slightly over four percentage points, is of little consequence, aside from being not statistically significant ($X^2 = 1.144$, $df = 1$, $p = .285$). In other words, although the cost and effort of sending a pre-notification email are minimal, it seems hardly worth the bother. Although it appears to have little effect on its own, it is possible that pre-notification combined with other response rate enhancing techniques may be of some use. This is something that requires further research.

Table 6: Cooperation Rate by Advance Notice (rows %)

Advance Notice	Respondent		Total
	Yes	No	
Yes	147 (55.7)	117 (44.3)	264
No	129 (51.0)	124 (49.0)	253
Total	276	241	517

Nevertheless, it is worth examining whether the overall pattern described in Table 7, also holds when we control for industry.

Table 7: Cooperation Rate by Advance Notice & by Industry (rows %)

Industry	Advance Notice	Respondent		Total
		Yes	No	
Semiconductor	Yes	95 (59.4)	65 (40.6)	160
	No	71 (50.7)	69 (49.3)	140
	Total	166	134	300
Computer	Yes	25 (44.6)	31 (55.4)	56
	No	35 (53.0)	31 (47.0)	66
	Total	60	62	122
Communications	Yes	27 (56.3)	21 (43.8)	48
	No	23 (48.9)	24 (51.1)	47
	Total	50	45	95

There are some differences between industries. Thus, the advance notice seems to have had the largest effect in the Semiconductor industry. There, the probability of cooperating in the survey is 8.7 percentage points higher for individuals who were sent the advance email notice (59.4%) than for individuals who were simply sent the survey email invitation (50.7%). The difference in the Communications industry goes in the same way but it is slightly smaller: 7.4 percentage points separate the two groups (56.3% for the advance notice group against 48.9% for the control group). In contrast, the advance notice appears to have had a detrimental effect in the Computer industry: individuals who were *not* notified in advance of the upcoming survey cooperated at a higher rate (53.0%) than individuals who received the advance email notice (44.6%) – a difference similar to the one we encountered in the Semiconductor industry but in the opposite direction. However, none of the disparities we have described is statistically significant ($p > .05$).

Before we examine whether the pre-notification had an effect on the speed of returns, we will explore whether it had an effect at any time when the survey organization contacted potential participants. All sampled individuals who failed to respond before the end of the field period were contacted four times – if we exclude the advance notice: the first contact was the invitation to the survey (first day of field period), the second was the first reminder (third day of field period), the third was

the second reminder (seventh day of field period) and the fourth was the third reminder (fourteenth day of field period). Table 8 below summarizes the results at each contact point. It shows the cooperation rate on each day the survey organization contacted sampled individuals: on the first day that includes the entire sample; in subsequent days, it includes remaining nonrespondents. Did the individuals who were pre-notified participate at a significantly higher rate at any time?

Table 8: Cooperation Rate by Advance Notice & by Contact (rows %)

Contact	Advance Notice	Respondent		Total
		Yes	No	
First day	Yes	60 (22.7)	204 (77.3)	264
	No	53 (20.9)	200 (79.1)	253
	Total	113	404	517
1 st reminder	Yes	21 (11.1)	168 (88.9)	189
	No	22 (11.6)	168 (88.4)	190
	Total	43	336	379
2 nd reminder	Yes	20 (12.8)	136 (87.2)	156
	No	14 (8.5)	150 (91.5)	164
	Total	34	286	320
3 rd reminder	Yes	4 (3.1)	127 (96.9)	131
	No	7 (5.1)	131 (94.9)	138
	Total	11	342	269

As can be surmised from a cursory look at Table 8, the answer is “NO”. At no time did the panel members who were given an advance warning of the upcoming survey cooperate at a higher rate than the remaining panel members who were not pre-notified. If anything we would have expected them to react in larger proportion on the day of the survey invitation: in other words, because of two promptings (advance notice and survey invitation) we would have expected a larger percentage of them to participate on the first day of the field period; instead only 2.2 percentage points separate them (19.9%) from the others (17.7%). The largest difference (3.8%) between the two groups is found on the day of the second reminder but this and the other differences are only random variations we need not bother about. In other words, it takes as much effort (contacts) to convert into respondents sampled individuals who were pre-notified as it is those who were not.

We can now turn to the issue of response speed: how quickly did respondents who were notified in advance participate in the survey compared to their counterparts? After reviewing the prior results on cooperation rates, what we are going to discover next will come as no surprise.

Table 9: Response time (days) by Advance notice

Statistic	Advance Notice	
	Yes	No
Mean	4.65	4.88
Std. Dev.	5.522	5.385
Median	2.00	3.00
IQR	6.00	6.00
Min.	1.00	1.00
Max.	22.00	21.00

Table 9 presents summary statistics for both groups on the response time variable. Neither the measures of center nor those

characterizing dispersion differ much between the two groups. The advance-notice cohort has a smaller median (2 days) than the other (3 days), but their interquartile range is identical, and their range (not listed) differs by one day (21 days for the advance notice, 20 days for the other group).

Figure 2: Distribution of response time (days) by Advance notice

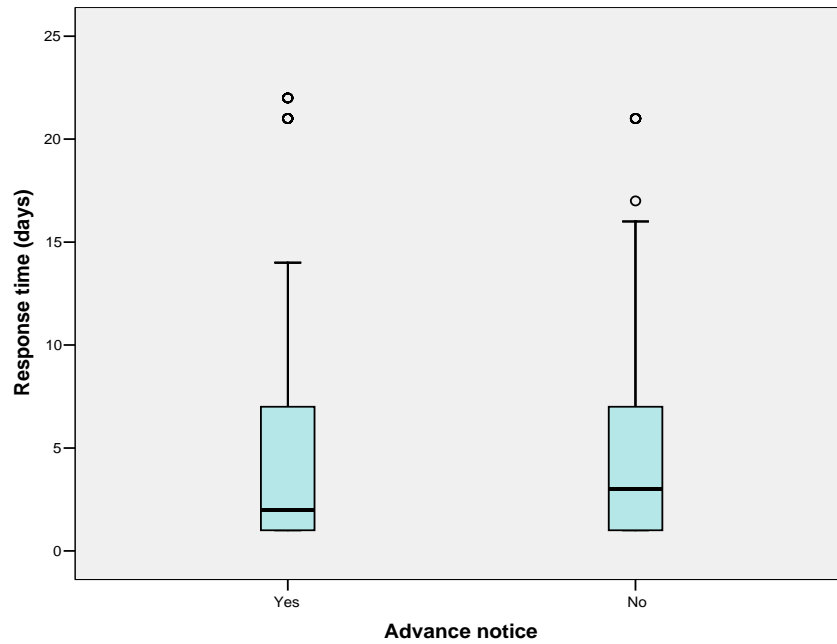


Figure 3: Cumulative returns (%) by Advance notice

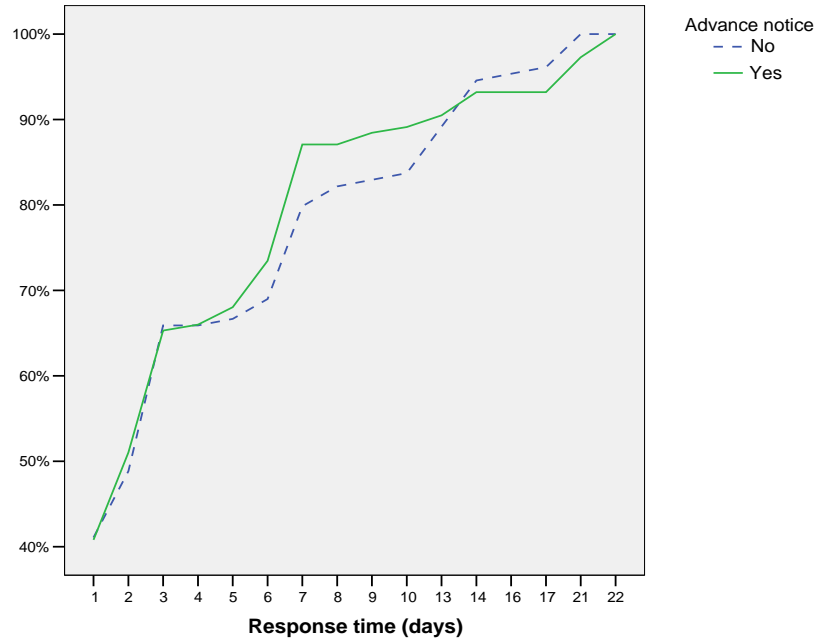


Figure 2 above provides a graphical depiction of the results presented in Table 9. We note that the distribution of response time in the two groups is quasi identical. We observe the slight difference in medians. We also see that the distribution for the no-advance notice group extends further beyond the 75th percentile, without observations becoming outliers, than does

that of the treatment group. On the other hand, the advance notice group has a few more outliers at the upper end of the distribution than the control group. Aside from this, every else looks the same.

We can follow the pattern of returns more clearly by looking at the line graph above (Figure 3). We see that up to third day, the pattern of returns for both groups is identical. Things start to diverge, in favor of the advance notice cohort, on the fourth day, and remain this way until the 13th day, the day of the third reminder. At that point the no-advance notice group takes the lead in terms of cumulative percent returns. The largest discrepancy occurs between the 7th and 10th day of field period. However none of these differences is very large: the pattern for both groups is very similar and they are very close to each other.

The results of a formal statistical test, the Wilcoxon-Mann-Whitney (W-M-W) test, confirm, if indeed it needed to, that nothing statistically significant is occurring here ($Z = .336, p = .737$). Thus, we conclude that, for this occupational group, the pre-notification email had no effect on the response (cooperation) rate or on the speed of returns – at least on its own.

5.3 Web Questionnaire Format Experiment: Results

In the second experiment, sampled individuals who chose to fill out the survey questionnaire were randomly assigned to one of two versions of its web form: in the first version, the entire instrument was loaded onto one single HTML file or web page (screen); in the second version, the questionnaire was distributed among multiple web pages: specifically, the survey’s substantive questions were placed onto twenty eight pages. In other words, in this version, respondents had to click on a “Next” button in one page to move on to the next screen to answer the next set of questions. It should be noted that respondents in the multiple-page format were not forced to answer an item, if they left that an item unanswered, nor were they notified about an incomplete answer (e.g., pop-up window) in order to move on to the next page.

The purpose of the experiment was to assess the differential effect, if any, the two versions had on the data quality that the survey organization ended up with. Data quality was measured in two ways: first by observing the proportion of individuals who abandoned the survey, the break-off rate, which is a form of delayed refusal or partial unit nonresponse; second by measuring the rate of item non-response, where an item is missing if the respondent is eligible to answer it but fails to do so, and other forms of misreporting (e.g., providing the wrong percentage when two or more items should sum to 100%).

The total sample size for the first analysis was $n=218$, of which $n=16$ were break-offs (individuals who started filling out the survey but did not complete it): thus the overall break-off rate, as mentioned before (see section 3), is 7.3 percent. Table 10 below shows the distribution of those that entered the survey by the type of questionnaire they were assigned to. As we see the sample is almost evenly distributed between the two questionnaire formats.

Table 10: Distribution of Participants by Type of Web questionnaire

Format	Participants	Percent
Multi-page	111	50.9
Single page	107	49.1
Total	218	100.0

The second analysis (item nonresponse or error) was conducted with the $n=202$ completed questionnaires. Here also the sample is fairly evenly distributed between the two types of web pages.

Table 11: Distribution of Respondents by Type of Web questionnaire

Format	Respondents	Percent
Multi-page	99	49.0
Single page	103	51.0
Total	202	100.0

Table 12 below presents the results of the experiment in terms of break-off rates. We see that if individuals have been assigned to the multi-page format of the web questionnaire, they are more likely (10.8%) to fail to complete the survey than are individuals who have been assigned to the single page format (3.7%). A chi-square test indicates that there is a statistically significant ($X^2 = 4.007, df = 1, p = .045$) association between break-off rate and questionnaire format. But the strength of the association between the two variables is small, as indicated by Cramer’s V ($V = .136, p < .05$). Furthermore,

an approximate 95% confidence interval (.17, .05) is rather wide and shows that the multi-screen version of the survey could possibly lose up to 17% of its original participants.

Table 12: Break-off Rate by Type of Web questionnaire (rows %)

Format	Break-off		Total
	Yes	No	
Multi-page	12 (10.8)	99 (89.2)	111
Single page	4 (3.7)	103 (97.3)	107
Total	16	202	218

Thus, it appears that the multi-page version of the questionnaire could have a detrimental effect on the rate of participation. However, we now need to determine what differences, if any, can be detected between the two versions in terms of error rate (primarily item nonresponse). This analysis is conducted with the n=202 participants who completed the survey. These individuals were measured on 59 question items. If, for a given item, the respondent provided an answer (item not missing) and the answer was accurate, that item for that respondent was assigned a 1 (= response), otherwise the score was 0 (= missing/error). Then, a mean score was computed across all 59 question items. This value actually measures the proportion of items correctly answered. (A related measure would have been to compute the total number of completed items. Both measures yield similar results.)

The box plot below (Figure 4) presents the distribution of the item response rate for the two experimental groups. The distributions are identical; they are both left skewed, with a few more extreme outliers (3) in the single page group than in the multi-screen group. These are participants that have a higher error rate, or, in terms of the measure presented, a lower item response rate.

Figure 4: Item response rate by questionnaire screen format (respondents)

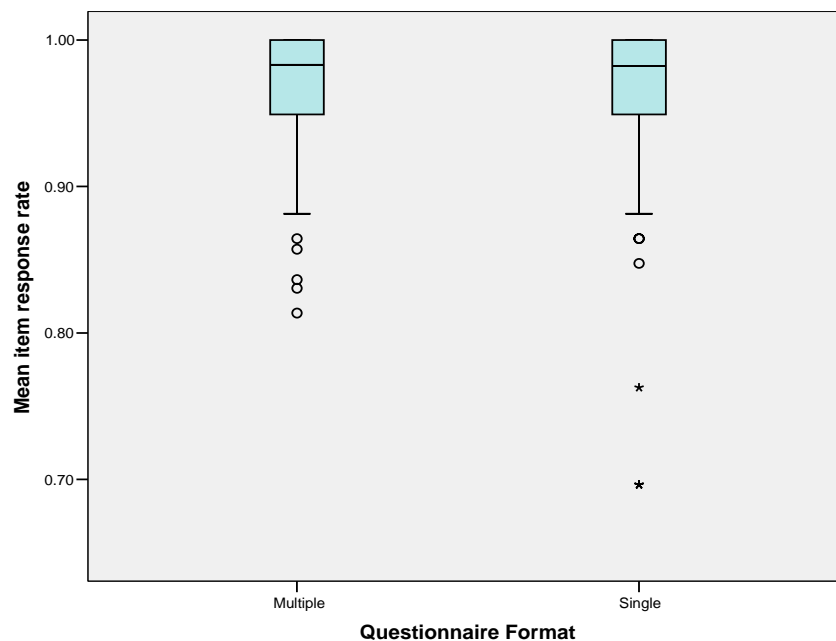


Table 13 below presents summary descriptive statistics for both experimental groups. We note, first, that, generally, response items are well answered – this is not unusual for a survey whose topic saliency is high (Donald, 1960; Murata and Gwartney, 1999): the median response rate is 98% for both groups. Due to the few extreme outliers in the lower reaches of the scale among the single-page group, the largest difference (11.8 percentage points) between the two groups on the measures presented in Table 13 is their respective minimum response rate: 69.6 and 81.4 percent, for the single-page group and the

multi-page group, respectively. There is no significant difference between the two experimental groups on their item response probabilities (W-M-W test: $Z = -.055, p > .05$).

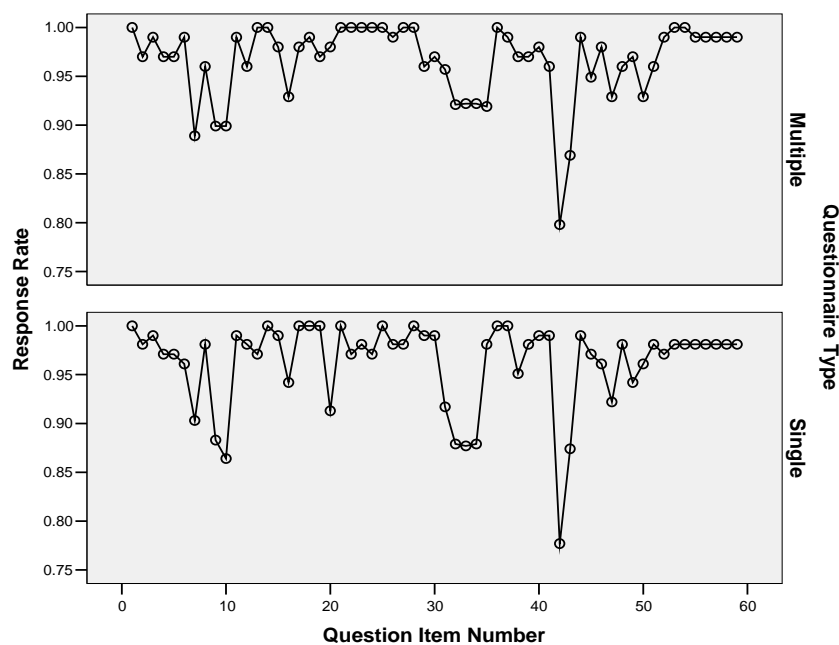
Table 13: Item response rate by questionnaire screen format (respondents)

Statistic	Screen Format	
	Multiple	Single
Mean	.967	.965
Std. Dev.	.0433	.0555
Median	.983	.982
IQR	.05	.05
Min.	.814	.696
Max.	1.000	1.000

Another measure that was examined is the proportion of respondents that completed the survey (i.e., respondents who provided an answer to all applicable questions). Overall, nearly forty-three percent of respondents completed all applicable items. Respondents in the multi-page format had a somewhat lower (41.4%) completion rate than respondents answering the single-screen questionnaire (43.7%); but this difference is neither substantively important nor statistically significant ($X^2 = .107, df = 1, p > .05$).

Clearly, the web format of the questionnaire does not appear to have had any practical or significant effect on the item response rate: both groups of respondents answer the survey as thoroughly; the format does not appear to confer any improvement to either group.

Figure 5: Item response rate by questionnaire screen format (question items)



Next, we will pursue this analysis by looking at the question items themselves: are questions less prone to error or to remain unanswered if they are presented within a multi-page questionnaire format or is this so if they are in a single screen questionnaire? Although we have seen that respondents in the two experimental groups have similar item response rates, it is possible that the response rate to some items may be significantly different. Here, the question item, rather than the respondent, is the unit of analysis. Recall that there are $n=59$ question items; what we are measuring, in this analysis, is the proportion of (eligible) respondents that answered each item adequately (did not leave it blank or made some kind of error when answering).

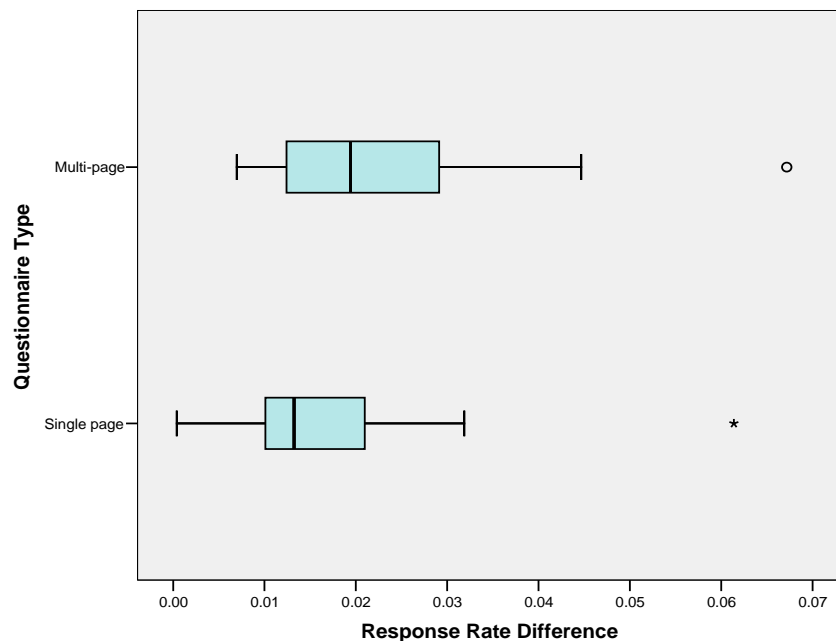
Figure 5 above, a serial plot, illustrates the pattern of responses for each of the 59 items in the two types of questionnaires. What we want to discover is whether the proportion of responses differ for the two treatments. What we see in Figure 5 are very similar patterns of response. The fact that the vertical scale has been truncated exaggerates the differences which are, generally, small – the average absolute difference between items in the two types of formats is 1.8 percentage points. Overall, response items tend to have the same or very similar response rates. Of the 59 question items, 27 have a higher response rate in the multi-screen format than in the single screen format, 26 are higher in the latter than in the former, and six items have the same response rate. We can safely say that (item) nonresponse is not related to the format in which the data collection instrument is presented. (In passing we note that there does not appear to be any “response fatigue”: i.e., an increasing trend in item nonresponse as the respondent progresses through the questionnaire.)

Table 14: Difference in response rate by questionnaire screen format (question items)

Statistic	Screen Format	
	Multiple	Single
Mean	.024	.017
Std. Dev.	.0142	.0136
Median	.019	.013
IQR	.020	.012
Min.	.0070	.0004
Max.	.0672	.0614

The average difference for those items that have a higher response rate in the multi-page format compared to the same items in the one-page format is 2.4 percentage points, whereas it is 1.7 percentage points when we compare the items with a higher response rate in the single-screen format than they have in the multi-screen format (see Table 14 above).

Figure 6: Difference in response rate by questionnaire screen format (question items)



We see in Figure 6 that the distribution of differences for the multiple-page format is shifted slightly higher than that of the single page format, which means that, generally, the differences in favor of the multi-page questionnaire are higher than the differences that express a higher response rate for the single page format. This explains also why, even though the two observations at the upper end of the scale of the box plot have nearly the same value, one of them (“o”) is a regular outlier ($Q3 + [1.5 \times IQR]$) in one distribution (single page) and the second (“*”) is an extreme outlier ($Q3 + [3.0 \times IQR]$) in the other distribution (multi-page). But, as Table 14 illustrates, these differences are minor.

If we look at the $n=59$ question items as independent observations for which two measurements have been taken, one within a single page format, and the other in a multi-page format, we can apply the Wilcoxon signed ranks test. The results show, predictably, that no statistically significant difference exists between the two types of questionnaire formats ($Z = -.845, p > .05$).

6. Discussion

In this dual experiment, our goal was to find ways of reducing two potential sources of error that arise in sample surveys: nonresponse error and measurement error.

6.1 The pre-notification experiment

The self-administered mode of survey research has been shown to yield, in general, lower response rates (Yu and Cooper, 1983) than person-mediated modes (face-to-face, telephone); although it is possible to obtain high rates (>60%) within that modality (Dillman et al., 1974 for mail surveys; Crawford et al., 2002 for web-based surveys). Ever since survey researchers discovered the negative impact nonresponse can have on the validity of survey results (circa 1930s), much effort has been put into mitigating its influence by taking measures prior to data collection (the preferred approach) and/or after the field period is over (e.g., weighting).

Thus when a study is in the planning stage, nonresponse is taken into account because it is known to be a usual, if unwelcome, part of any sample survey. Numerous techniques have been invented by survey methodologists to minimize the occurrence of nonresponse: e.g., advance contact, incentives, personalization (addressing the potential participant by name in, for example, the survey invitation letter), reminders or follow-up letters. All these devices have one purpose: motivate the sampled individual to participate in this social process called the survey.

As we have seen, a large number of studies have found that alerting potential participants of an upcoming survey tends to have a positive effect on response rates as compared to simply contacting them at the beginning of the field period. However, few, if any, studies (Linsky, 1975; Dillman et al., 1976) deal with the underlying theory related to sending advance notice. What mechanisms does an advance notice is expected to trigger which will entice an individual to become a survey respondent?

But first let's look at what the pre-notification is expected to do. A number of beneficial features have been mentioned: it introduces and personalizes the researcher, it requests the individual's cooperation, and, of course, it alerts the potential participant of the upcoming study (Linsky, 1975); it also helps in establishing the legitimacy and importance of the study (Tourangeau, Rips and Rasinski, 2000), and thus, perhaps, in minimizing the danger that the recipient will relegate the actual survey invitation to the status of "junk" mail. Intuitively, the benefits of a pre-notice would seem to be most appealing for a person-mediated survey (telephone or face-to-face), in that it allows sampled individuals to prepare themselves by being given information (e.g. purpose, importance) about the study, the survey organization, and the fact that they will be contacted by an interviewer, etc. as opposed to being confronted by an unknown entity when they receive a cold call from an interviewer or are faced with a stranger at their front door when the interviewer shows up unannounced. Potential respondents are now informed and are no longer put in a position in which they face the unknown. In contrast, because of the less threatening nature of a mail questionnaire (Dillman et al., 1976), a pre-notice does not appear to have the same importance or impact in this mode. In other words, the beneficial features mentioned earlier could be just as well those of the survey invitation in which all the relevant information about the study is given – the only exception being what is specific to an advance notice: i.e., that it alerts the prospective participant of the upcoming study.

Although the self-administered mode of data collection is an impersonal one in that the respondent never interacts with a live person (via the telephone or face-to-face), the survey organization does initiate a social process with the would-be participant. The role of the survey organization at this stage of the process (before the start of the field period) is to convince the sampled individual to participate: it is in a position not dissimilar to a political candidate who must win over the electorate, or the defense lawyer who must persuade the jury to let the accused go free, or the lover who must woo the object of his love. The advance notice is just one technique in the survey researcher's repertoire of stimuli. But why do we believe that the advance notice can become a stimulus to survey participation?

The answer to this question must be found in society's norms of behavior. Although survey research is a recognized and legitimate form of inquiry, it is intrusive: we invade a person's private space – this is very much the case with face-to-face interviews; it is so, but to a lesser extent, for the telephone interview, and even less so for the self-administered mode. In all

cases, we ask respondents to set aside time to perform a task for us. The advance notice is a way to soften the intrusion: we're intruding but at least we don't come unannounced. In other words, we are conforming to certain rules of behavior that express courtesy and respect. Because we are products of a social system, we know by experience that there are certain ways to behave that are sure to bring about a refusal (rejection) on the part of prospective participants: rudeness (deviance), for example, is simply out of the question. As survey researchers, we are not altruistic: although socially desirable actions, such as opening the door for someone else because we happen to be there first, can be rewarding in themselves, our actions (e.g. an advance notice) are motivated to yield a concrete reward (participation in our survey). We hope that in receiving our advance notice, potential respondents will recognize (social approval) that we have made a sincere effort to follow common rules of civility, and that they will reward us by becoming actual respondents. In other words, we seek social approval for our actions (e.g. an advance notice), and we hope that this social approval will translate into survey participation.

As we have seen in the experiment described above, an advance notice in this context has had no effect on the response rate. Our result is similar to that obtained by Parsons and Medford (1972) and Wright (1995). These investigators were also relying on a self-administered mode (mail survey), and they were studying "fairly homogeneous" populations. However, our situation differed from theirs in one important way. Theirs was a one-time sample, which means that the individuals they selected were unaware that they would be called upon to participate in a study. In contrast, we relied on a panel in which the vast majority of members had previously participated in prior surveys and all were aware that they would be contacted at some time to take part in a survey. Of course, in the busy life of an engineer, one can not expect that membership in a panel will be on one's mind at all times. Nevertheless, one would expect an advance notice to carry more weight, i.e. to be more necessary and useful, with a one-time sample than with a panel. But, as we discovered, in none of these cases did an advance notice make a difference. Perhaps, the reason for this is not so much that the populations that have been sampled are "homogeneous" but rather the kind of "homogeneous" populations they are. I would suggest that these populations expect to be "interviewed" as part of their job profile. This is certainly the case for clergy (Parsons and Medford, 1972) and politicians (Wright, 1995), so the use of an advance notice is probably unnecessary. In our specific case, the advance notice appears also to be irrelevant: the sampled individuals are already on the panel and they know what that entails; in a way, the original invitation to the panel takes the place of an advance notice since asking individuals to become panel members is also telling them that they should expect being contacted periodically to participate in a survey.

Membership in the panel seems the most plausible explanation for the failure of the advance notice from having an effect in this experiment. Of course, we can speculate whether the timing of the pre-notification was optimal: may be we should have sent it later or perhaps earlier; we can also wonder about the content of the advance notice: perhaps the wording could have been more convincing, more polished; finally, we can conjecture that a letter sent via postal mail might be more effective than a mere email. Of course, the latter approach would have increased the overall cost of the research. When one conducts a mail survey, one has limited options: postal mail is pretty much the unique conduit. So, in this mode, it is very important to know whether an advance notice is effective for the population sampled: including the cost of an advance letter to that of sending the survey invitation proper and the various reminders can add to the financial burden the survey organization must bear (stationery, printing, labor, postage) – especially if the effect of the pre-notification is nil. These monies could be better spent elsewhere: e.g., following-up on nonrespondents. In contrast, sending an email as an advance notice has hardly any consequence on research project's budget.

6.2 The questionnaire presentation experiment

Several issues must be taken into consideration by the survey organization when deciding which format it should present the questionnaire it will use to collect the data. Some of these issues impact the respondents, other affect the survey organization directly. For example, multi-screen web questionnaires require more programming to create them, but they allow for unambiguous branching, thus relieving respondents from having to read instructions, and thereby eliminating a place where errors can occur. Of course, it is possible to program a single screen survey for branching: what the programming does is scroll for the respondent. The choice of format (screen-vs.-scroll) is dictated by the requirements of the questionnaire: if the researcher does not want respondents to see all the questions they will be asked to answer, i.e., if the researcher does not want the respondents to preview some questions before others are answered, then the screen format is required. In other words, a screen-based instrument is best suited when the researcher does not want respondents to jump ahead or back or both: i.e., when the researcher wants complete control over the flow of the survey.

There is no definitive evidence that shows that one format is more burdensome to respondents than the other. One might think that having to click repeatedly to move along with the questionnaire, from screen to screen, could be seen as adding to the respondent's load, but then scrolling requires respondents to go from the center of the screen to its border and vice versa – respondents using a mouse with a roller would avoid this by simply activating that device to move down the page.

One issue that could turn out to be an irritant for respondents is the loading time: the time it takes a web page to load onto the respondent's computer. Although it will take more time for a one-screen survey to load initially, once completed the respondents can navigate the survey at their own pace. With a multi-screen survey, each page must be loaded in turn onto the respondent's computer, this can give rise to several problems: each upload is a point at which some technical glitch can occur; the time to upload each screen is dependent on the respondent's hardware equipment: this could unnecessarily lengthen the time the respondent takes to complete the survey. Special populations, such as the one we sampled for this survey, are unlikely to encounter these problems since they tend to be equipped with the most up-to-date hardware and software, especially at their place of work.

Our experiment in web-survey format has shown that respondents assigned to the single screen version of the questionnaire were three times more likely to complete the questionnaire than their multi-screen counterpart (see Table 12). Although these results are only marginally significant ($.05 < p < .10$), they should not be summarily dismissed. Unfortunately, one of the limitations of this study is that we did not find out the reason why certain respondents dropped out of the survey. While this occurrence (break-off) is relatively rare, it does deserve further study.

Another issue we looked at with regards to the survey's web format was response quality: measured, essentially, as item response rates. We saw that no clear advantage between the two formats could be found. One possible explanation for this has to do with the fact that the topic of the survey is highly salient to this occupational group. Hence, the subject matter of the survey, with which they are very much conversant, is more likely to capture their interest and attention, and thus they are less likely to make mistakes (e.g. miss an item).

7. Conclusion

Although an advance notice has been shown in many cases to be effective in boosting the overall response rate, it was not so in the experiment that was described here. I have speculated that the reason for this result is the fact that the would-be survey participants belong to a panel, which means that they have already given their consent to participate in surveys, and as such they expect to be called upon from time to time. To confirm this conjecture, we would have to conduct an experiment with panel members and non-panel members invited to the same survey.

Aside from the pre-notice, we used two other well-established methods to improve response rates: incentives and follow-up reminders. Yet, the overall participation was low (<50%). We do not know why (panel) nonrespondents do not participate: were some ineligible but did not bother to tell us, i.e. they came to the survey home page read the screening question but did not answer; perhaps some had bad addresses but we were not informed of it, i.e. not all servers generate an automatic message informing the sender that the recipient's address is no longer active; maybe the email message we sent did not stand out in the pile of emails in the panel member's inbox (although one would think that in addition to the advance notice, the invitation and three reminders, this would attract one's attention at some point); the panel member may have been on vacation during the survey period; some might have encountered technical difficulties in accessing the survey but did not contact the survey organization; for others, no time was available: the pressures of work were simply too high to devote any time to a survey; others still may suffer from survey fatigue (oversurveying): just because they belong to this panel does not mean they are not solicited and participate in other surveys.

Some of these questions cannot be answered but others can: the latter require a systematic study of survey nonrespondents. This is what Crawford and his colleagues (2002) did as part of their Student Life Survey at the University of Michigan. They found that the majority of the web survey nonrespondents that answered their follow-up study claim they did not recall receiving the survey invitation. Bad addresses could not explain this result since the researchers had the correct email contact for 94% of them; nor was lack of checking their email inbox since 81% of these nonrespondents checked their mail on a daily basis and 97% of them checked their mail at least once a week (the field period lasted nearly 2 months). In fact the authors do not provide an explanation for this "reason" given by nonrespondents for not participating. They do report that when asked if they would take the survey, now that they have been made aware of it, this group of nonrespondents showed a high degree of willingness to do so. This response seems to satisfy the researchers that the no-recall claim is a legitimate one. It is difficult to believe that after receiving four contacts (the survey invitation plus three reminders) from the survey organization, these nonrespondents would not be aware of the study. It seems more plausible that these individuals have some other reason for not participating but did not reveal it to the investigators. Be that as it may, the point being that it makes eminent sense for a survey organization that is repeatedly sampling the same population to find out as much as it can about it, and specifically to unearth the reasons for nonresponse. In the Crawford et al. (2002) study just mentioned the second reason

most frequently cited for nonparticipation was “no time”: these individuals expressed concern that the survey would stretch to the breaking point their already busy schedule. Perhaps the population that we sampled for the study reported in this paper contains a large proportion of individuals who have “no time” – even if they belong to a panel. But to find out we need to conduct a follow-up survey of nonrespondents.

The experiment on the web format of the questionnaire showed no difference between the two versions that were implemented. For the purpose of this particular survey a single screen layout would have done the job since no complex branching was required by the questionnaire – we would then have avoided the cost of the programming to set up multi-screen version of the survey. It is clear from this experiment that the decision regarding how the questionnaire is to be formatted should be dictated by the flow of the question items.

References

- Burris, J., et al. (2001), “An Experiment in Web Survey Design”, paper presented at the Annual Meeting of the American Association for Public Opinion Research, Montreal, Quebec.
- Chun, Y. I. and Robertson K. W. (1995), “The Effects of Advance Letters and Reminder/Thank You Letters on Reducing Nonresponse in an Establishment Survey: An Experimental Study”, *Proceedings of the Survey Research Methods Section*, American Statistical Association, 1151-1256.
- Clark, R. L., Nyiri, Z. (2001), “Web Survey Design: Comparing a Multi-Screen to a Single Screen Survey”, paper presented at the Annual Meeting of the American Association for Public Opinion Research, Montreal, Quebec.
- Couper, M. P., Traugott, M. W., and Lamias, M. J. (2001), “Web survey design and administration”, *Public Opinion Quarterly*, 65, 230-253.
- Crawford, S., McCabe, S., Couper M., and Boyd C. (2002), “From Mail to Web: Improving Response Rates and Data Collection Efficiencies”, paper presented at the International Conference on Improving Surveys, August 25-28, 2002 University of Copenhagen, Copenhagen, Denmark.
- Dillman, D. A., Carpenter, E. H., Christenson, J. A., and Brooks, R. M. (1974), “Increasing Mail Questionnaire Response: A Four State Comparison”, *American Sociological Review*, 39, 744-756.
- Dillman, D. A., Gorton Gallegos, J., and Frey, J. H. (1976), “Reducing Refusal Rates for Telephone Interviews”, *Public Opinion Quarterly*, 40, 66-78.
- Donald, M. N. (1960), “Implication of Nonresponse for the Interpretation of Mail Questionnaire Data”, *Public Opinion Quarterly*, 24, 99-114.
- Ford, N.M. (1967), “The advance letter in mail surveys”, *Journal of Marketing Research*, 4, 202-204.
- Fowler, F. J. (2002). *Survey Research Methods*. Sage Publications, 3rd edition.
- Fox, R. J.; Crask, M. R.; Kim, J. (1988) “Mail Survey Response Rate: A Meta-Analysis of Selected Techniques for Inducing Response”, *Public Opinion Quarterly*, 52, 467-491.
- Fuchs, M. (2001), “Screen Design in a Web Survey”, paper presented at the Annual Meeting of the American Association for Public Opinion Research, Montreal, Quebec.
- Goldstein, K. M. and Jennings, M. K. (2002), “The Effect of Advance Letters on Cooperation in a List Sample Telephone Survey”, *Public Opinion Quarterly*, 66, 608-617.
- Heberlein, T.A.; Baumgartner, R.M. (1978), “Factors affecting response rates to mailed questionnaires: A quantitative analysis of the published literature”, *American Sociological Review*, 43, 447-462.
- Hembroff, L. A., et al. (2005), “The Cost-Effectiveness of Alternative Advance Mailings in a Telephone Survey”, *Public Opinion Quarterly*, 69, 232-245.
- Link, M. W., Mokdad, A. (2005), “Advance Letters as a Means of Improving Respondent Cooperation in Random Digit Dial Studies”, *Public Opinion Quarterly*, 69, 572-587.
- Linsky, A. S. (1975), “Stimulating Responses to Mailed Questionnaires: A Review”, *Public Opinion Quarterly*, 39, 82-101.
- Majowicz, S. E. et al. (2004), “An introductory letter in advance of a telephone survey may increase response rate”, *Canada Communicable Disease Report*, 30, 121-123.
- Murata, T., Gwartney, P. A. (1999), “Question Salience, Question Difficulty and Item Nonresponse in Survey Research”, paper presented at the International Conference on Survey Nonresponse, Portland, Oregon.
- Parks, R. B.; Kennedy, J. M., and Hecht, L. F. (1994), “A Demographic Analysis of the Impact of Presurvey Letters on Cooperation Rates in Urban Neighborhoods”, paper presented at the Annual Meeting of the American Association for Public Opinion Research, Danvers, MA.
- Parsons, J., Owens, L., and Skogan, W. (2002), “Using Advance Letters in RDD Surveys: Results of Two Experiments”, *Survey Research*, 33, 1, 1-2.
- Parsons, R. J.; Medford, T. S. (1972), “The Effect of Advance Notice in Mail Surveys of Homogeneous Groups”, *Public*

- Opinion Quarterly*, 36, 258-259.
- Peytchev, A., et al. (2006), "Web survey design: paging versus scrolling", *Public Opinion Quarterly*, 70, 596-607.
- Slocum, W. L., Empey, L. T., and Swanson, H.S. (1956), "Increasing Response to Questionnaires and Structured Interviews", *American Sociological Review*, 21, 221-225.
- Stafford, J. E. (1966), "Influence of preliminary contact on mail returns", *Journal of Marketing Research*, 3, 410-411.
- Tourangeau, R., Rips, L. J., and Rasinski, K. (2000). *The Psychology of Survey Response*. Cambridge University Press.
- Vehovar, V., Manfreda, K. L., and Batagelj, Z. (2000), "Design Issues in Web Surveys", Proceedings of the Survey Research Methods Section, American Statistical Association, 983-988.
- Walker, B. J., Burdick, R. K. (1977), "Advance correspondence and error in mail surveys", *Journal of Marketing Research*, 14, 379-382.
- Woodruff, S. I., Mayer, J. A., and Clapp, E. (2006). "Effects of an introductory letter on response rates to a teen/parent telephone health survey", *Evaluation Review*, 30, 817-823.
- Wright, M. (1995), "The Effect of Pre-Notification on Mail Survey Response Rates: An Experimental Result", *Marketing Bulletin*, 6, 59-64.
- Yammarino, F. J., Skinner, S. J., and Childers, T. L. (1991), "Understanding Mail Survey Response Behavior: A Meta-Analysis", *Public Opinion Quarterly*, 55, 613-639.
- Yu, J., Cooper, H. (1983), "A Quantitative Review of Research Design Effects on Response Rates to Questionnaires", *Journal of Marketing Research*, 20, 36-44.