

Survey-Based Experimental Research via the Web: Some Observations

Mike Brennan, Nathan Rae and Mathew Parackal

This paper describes two Web-based surveys that incorporated experimental designs, and discusses some of the problems encountered. One problem related to the use of "cookies" and cgi scripts, which were needed to track respondents and assign them to treatment groups. These scripts could not be processed by older versions of browsers, causing a considerable reduction in sample size. There was also drop-off throughout the questionnaire, for reasons that are unclear. Secondly, although incentives were used, these failed to generate either a high, or a speedy, response. These findings highlight the importance of identifying appropriate incentives and developing effective Web-survey management procedures.

Keywords: web, survey research, cookie, cgi, incentives, response rate

Introduction

The phenomenal growth in Web usage for both communication and commerce, coupled with the unique technological capabilities of the medium, presents unique opportunities for Web-based survey research. Obvious advantages of the Web over traditional survey media (telephone, mail and face-to-face) are the relatively low cost of the fieldwork, potentially quick response and turn-around times, and the possibility of very large sample sizes.

A problem with the Web as a survey medium, however, is that relatively little is yet known about it. While there are tried and true methods for obtaining high response rates in conventional surveys, the procedures for managing a Web-based survey differ from those used by other methods, and it is likely that response rates and speed of response differ as well. However, there is little published information about these factors, or the problems that may be encountered conducting a survey via the Web.

The purpose of this research note is twofold: First, to describe the design and methodology used in two surveys conducted via the Web, both of which incorporated experimental designs. Second, to report the response rates and comment on some of the problems encountered. The first study was designed to test two versions of a purchase probability scale and the relative effectiveness of four versions of a banner ad, incorporating sound and/or animation. The second study was designed to extend the experiment on banner ad design. The two studies used a similar experimental design and methodology, but differed in the way respondents were recruited.

Method

Sample

The purpose of Study 1 was to estimate demand for five new billing options and two new support services being considered by a local Internet provider. Thus the sample comprised all clients on their mailing list. These clients were contacted using a letter enclosed with a mailed

newsletter. This letter explained the purpose of the research, invited participation, offered an incentive in the form of a prize draw for \$200 worth of goods or services from the Internet provider, and provided the URL for the survey Web page. The URL and details of the study were also posted on the provider's Home Page. No reminder letters or emails were used.

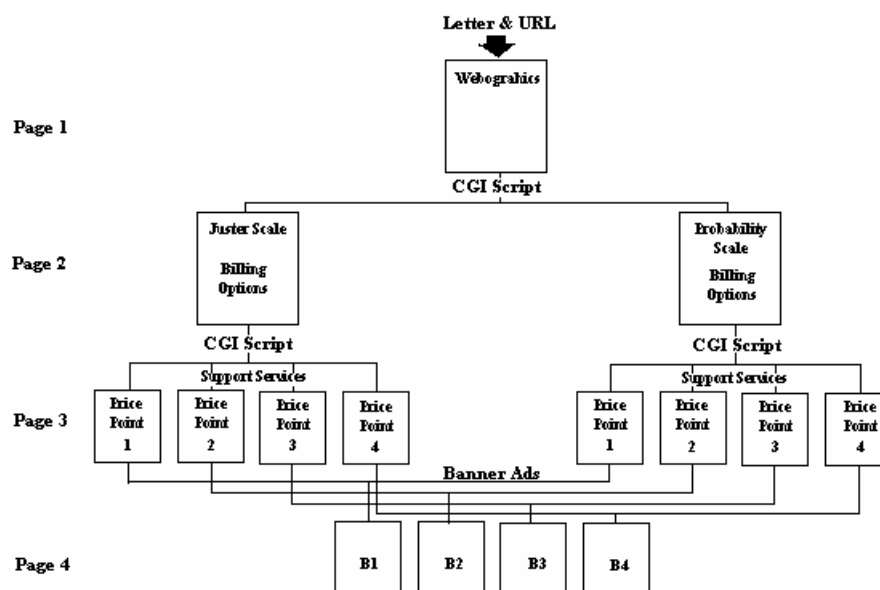
The purpose of Study 2 was to evaluate the effectiveness of sound and animation in banner ads. Respondents were recruited by placing notices on five New Zealand Web newsgroups. These notices were reposted at four day intervals. The notice briefly described the research, offered a \$50 cash prize draw incentive for participating, and gave the URL of the survey Web site.

Further details of these studies are provided elsewhere (Parackal & Brennan 1998; Rae & Brennan 1998).

Experimental Design

Both studies were conducted in two phases. Phase 1 was used to collect webographic information, assign respondents to treatment groups, and expose them to a particular banner ad. Phase 2 was a follow-up survey that examined respondent's perceptions and recall of the banner ad they were exposed to in Phase 1. In Study 1, Phase 1 used four separate pages and Phase 2 used two. In Study 2, both Phase 1 and Phase 2 used two pages. The design for Study 1 is shown in Figure 1.

Figure 1. Experimental Design for Study 1



The design for Study 2 was similar, except pages 2 and 3 were omitted from Phase 1, and eight rather than four treatment groups were used for the banner ad experiment.

Procedure

As neither study used a sampling frame, respondents had to be allocated to the different treatment groups after, rather than before, they came to the Web site to participate in the survey. Although a sampling frame existed for Study 1, access was denied. Study 2 had no

sampling frame, which is a typical situation for Web groups. The allocation of respondents to treatment groups was accomplished by means of CGI (Common Gateway Interface) scripts.

In a web-based survey, different versions of a question (i.e., different treatments) exist as separate files, one for each version. So, rather than having a series of complete questionnaires to cover all combinations of questions (treatments), as would be necessary in a mail survey, CGI scripts are used to construct each respondent's questionnaire as required.

In this study, the questionnaire was broken into separate sections. Submitting one section invoked a CGI script, which determined which section (treatment) the respondent was allocated to next. Because the first page of the Web survey was common to all respondents, it was used to collect webographic and demographic data. This webographic information was used; for example, to ensure that a respondent was not allocated to a treatment involving sound if their computer was not equipped to handle it.

Each time a respondent submitted a page, it was sent automatically via email to the researcher. As the different pages of the questionnaire were submitted separately, it was necessary to have some way of identifying these as belonging to a particular respondent, so they could be collated for analysis. This task was accomplished by using a "cookie". A cookie is the message sent to a Web browser by a Web server. A respondent's Web browser saves the message as a text file (*cookie.txt*). Every time the browser requests a page from the server, the message is sent to the server.

The cookie saved the email address when a respondent submitted the first page of the survey, and attached this to subsequent pages submitted by that respondent. The cookie also retrieved the Web page, to which the CGI script assigned a respondent.

The final step was to collate the emailed questionnaires for analysis. This was done using a Paradox script, which matched email addresses, and codes put in the subject field that identified the page number of the questionnaire, and put the extracted data into a Paradox database. This database was then imported into SPSS.

Results

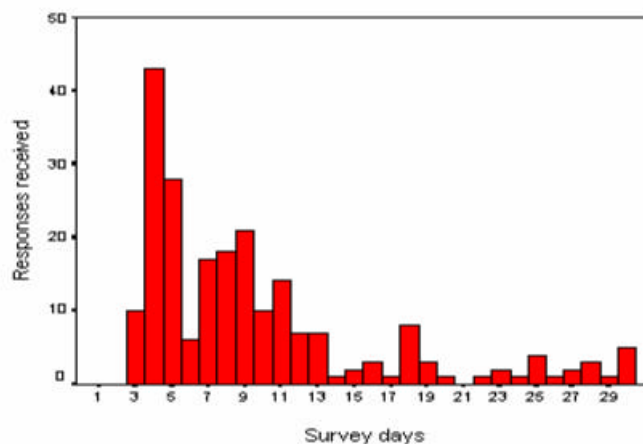
Speed of Response

The daily responses to Study 1 are shown in Figure 2 and the weekly responses are summarised in Table 1. While the responses came in quickly over the first week, they rapidly tapered off, although responses were still being received in the fifth week.

Table 1. Weekly responses to Study 1.

Week	n	%
1	107	48.6
2	35	15.9
3	34	15.5
4	24	10.9
5	20	9.1
Total	220	100.0

Figure 2. Daily responses in Study 1



The rates of response to Study 2 are shown in Table 2. As one might expect, the first listing on the newsgroups elicited the highest response, with the response to subsequent listings dropping off steadily. After almost a month, the final sample size was still quite small.

Table 2. Number of respondents acquired after each newsgroup posting.

Newsgroup Posting	Responses		Date
	n	%	
Post 1	57	40	22/10/97
Post 2	42	29	26/10/97
Post 3	28	19	30/10/97
Post 4	17	12	10/11/97
Total	144	100	

The results of these studies suggest that it is probably unreasonable to expect a speedy response to a Web based survey, at least without some sort of incentive to encourage a quick response. While the incentives used in the first study may have increased the rates of participation, this is not certain, and they certainly were not particularly effective in this regard. The sample sizes, even after a month, were still quite low. Furthermore, the incentives were not very effective at producing a speedy response. For the newsgroup sample, however, it is unclear whether the low response was due to a low response rate, or because the newsgroups did not have many visitors, as the level of participation in the newsgroups is unknown. Either way, more research is needed to identify techniques that can substantially improve both the speed and level of participation, in a situation where individually directed reminders are not an option. Otherwise, one of the potential advantages of conducting a survey via the Web is lost.

Attrition

Although 220 people responded to the first page of the questionnaire in Study 1, almost 30% failed to complete all three pages (see Table 3), leading to an unacceptably high loss of data. There are several feasible reasons for this, although it is not possible to draw firm conclusions. It may be that some respondents simply didn't wish to continue the survey. Some may have been unable or unwilling to complete the whole survey at one time, although it was not very long (about 10 minutes), and did not come back. At least seven respondents were known to have completed different parts of the survey at different times. Unfortunately, their data had to be excluded from the analysis because they had been assigned to different treatment groups on the different occasions, thereby invalidating the results. Others attempted to complete different parts of the survey using a different computer or browser, and because the cookie couldn't track them, the different sections could not be collated.

Table 3. Sample size and response rates through each phase of Study 1.

Phase	Sample Size	Response Rate
	n	%
Phase 1 - page 1	220	
Phase 1 - page 2	167	75.9%
Phase 1 - page 3	155	70.5%
Phase 2 - page 1	98	63.2%
Phase 2- page 2	85	54.8%

A likely explanation for the majority of cases was that either the respondent refused to accept the cookie, or was using a version of browser unable to do so. This would explain the over 120 emails (submitted pages) received for which the respondent could not be identified. It is also likely that some browsers could not cope with the JavaScript and cgi scripts used to format and allocate the pages, which meant that they would have been unable to view the scales or graphics used in the questionnaires. While a considerable number of potential respondents alerted us to these problems via email, it is likely that a possibly large number of those who experienced these problems simply abandoned the survey.

It is also possible that quite a high proportion of the potential respondents had their screens set at a low resolution (e.g. 640 x 480) whereas the graphics used in the questionnaire were optimised for a resolution of 800 x 600. Although this would not have stopped respondents from completing the survey, it would have altered the way they viewed the questionnaire. That is, those respondents who had their screens set to 640x480 would have had scroll bars across the page, whereas the other would not have. The effect this may have had on responses, if any, is unknown.

Another observation of note is the drop off between Phase 1 and Phase 2 in each study. In Study 1, the response rate for phase two was 63.2% (see Table 3). This is much higher than the response rate of 46% obtained in Study 2 (see Table 4), and much higher than the 47% reported by *MBinteractive* (1997) in their study conducted for the IAB. Also of note is the

low initial refusal rate. Very few respondents in either study refused permission to be contacted for Phase 2.

Table 4. Sample size and response rates through each phase of Study 2.

Phase	Sample Size n	Response Rate %
Phase 1 - page 1	144	
Phase 2 - page 1	66	45.8%
Phase 2 - page 2	62	43.1%

Discussion

The results of these two studies suggest that conducting a survey via the Web is unlikely to produce a fast rate of return unless an appropriate incentive is used, or the topic is one of particular interest to potential respondents. It is not clear, however, what form an incentive should take, but it is clear that the prize draws used in this study were not effective.

A large number of responses in Study 1 could not be identified. This meant that data was lost because these pages could not be collated. This problem of unidentifiable emails is likely to be the result of participants either refusing a cookie or not being able to accept one. Unfortunately, it is not possible to easily overcome this problem. The problem associated with respondents not being able to accept cookies will reduce over time as more Web users upgrade to newer Web browsers. However, it is difficult to predict whether the proportion of people who refuse to accept cookies will rise or fall. Some people consider cookies an invasion of privacy, and therefore refuse them. Fortunately, only a minority of Web users currently feel this way. An alternative to using a cookie would be to have respondents enter their email address, or some code of their own choosing, onto each page of the questionnaire. However, the rate of compliance one might expect needs to be investigated.

Some attrition was due to some respondents completing the different phases of the study, and sometimes even attempting to complete different pages of the survey, on different computers or using different browsers, or both. As a result, the cookie was not activated for some pages. Although the need to use the same computer and browser was stated in the email inviting respondents to participate in Phase 2, this message was clearly not forceful enough. Programming in an error message may help.

The survey in Phase 1 of Study 1 could easily have been undertaken as a mail survey. Although respondents were recruited by mail, they were required to use their computer to respond to the survey. As this probably requires more effort than simply finding a pen, or responding to a request received on-line, one might expect the initial response rate to be lower than to a comparable mail survey. The (estimated) response rate of 22% is poor, and somewhat lower than what one would expect from the first mailout of a well designed mail survey. Thus these results suggest that, as in mail surveys, Web based surveys may require the use of reminders in order to produce respectable response rates. Of course, this is not

possible unless a list is available. Further research is needed to determine the relative cost-effectiveness of using reminders, either via email or, if feasible, by mail.

In addition, it is now common practice to use some type of incentive to encourage people to respond to mail surveys. Prize draws are among the least effective types of incentive. Further research to identify suitable incentives for Web survey research is needed. Some caution is required, however, for a very attractive incentive may well prompt multiple submissions from people wishing to increase their chances of winning. Multiple submissions can be impossible to detect and their impact on results would be unknown but potentially large. Thus methods for minimising multiple and frivolous participation need to be developed.

One of the potential benefits of conducting surveys via the Web is the presumed availability of large samples. While sample size will be limited by the size of the list or newsgroup used to recruit respondents, in theory this can be overcome by increasing the number of these groups in the study. Study 2 demonstrated that this limitation might be more difficult to overcome than it appears. A large local (New Zealand) sample was required. Even the five major New Zealand newsgroups used in the study were insufficiently large to produce the desired sample size in an acceptable time period. In this case, the five New Zealand newsgroups were the only ones suitable for selection due to the nature of the advertisement used. While other newsgroups exist in New Zealand, a posting to these may have generated a hostile reaction, as the posting may have been considered spam (junk mail) which is not tolerated.

An alternative procedure for obtaining respondents is to use an interruption technique, in which the Web users are interrupted while loading a Web page at a site they are trying to access. The technique involves adding a piece of code to the homepage of the participating Web site that instructs the Web browser of potential participants to load a page containing the survey's introductory letter. The potential participant at this point has the option of either completing the survey or of refusing. The key advantage of this technique is the ability to easily record refusals at the first stage. MBinteractive (1997) reported success with the interruption technique and has shown that a response rate of 45 percent can be achieved. However, the downside of this technique is the difficulty in getting a Web host to allow this intrusion on their readers.

Conclusions

The Web has enormous potential as a survey medium, but suffers from serious limitations, at least for the immediate future. As there are no easily accessible sampling frames for the general Web population, it is not possible to draw a probability sample, and even if it were, the sample might not be representative of the population at large. (Although this is not necessarily a major problem for experimental studies, e.g., see Kingsley & Anderson 1998). However, this does mean that for experimental studies the assignment of respondents to treatment groups has to occur after an initial response, rather than before. Unfortunately, some widely-used versions of browsers cannot handle the scripts required to make these assignments. This problem will decrease in importance over time, as people upgrade their software. But the problem of recruiting respondents will remain, at least until there is a comprehensive directory of URLs.

Currently, the most common procedure for recruiting respondents via newsgroups is less than satisfactory, partly due to the relatively slow response times and the possibility of

considerable bias. Further research is required to identify effective incentives to improve this means of recruitment. A more promising approach may be to use a multi-modal recruitment strategy, involving a range of techniques such as snowballing via email, print and Web banner advertising. Work is currently in progress to assess the effectiveness of such an approach (e.g., Rae & Brennan 1998b; Smee, Brennan, Hoek & Macpherson 1998).

References

- Kingsley P & Anderson T (1998). *Market Research in an Electronic Community: Generalizing Results From Web-based Surveys*. In L Pelton and P Schnedlitz (Eds). Proceedings of the AMA Marketing Exchange Colloquium, Vienna, Austria, 322-331.
- MBInteractive (1997). *IAB Online Advertising Effectiveness Study*. Millward Brown Interactive and Internet Advertising Bureau, Draft report.
- Parackal M & Brennan M (1998). Obtaining purchase probabilities via a Web based survey: the Juster Scale and the Verbal Probability Scale. *Marketing Bulletin*, 9, 67-75.
- Rae N & Brennan M (1998). The relative effectiveness of sound and animation in Web banner advertisements. *Marketing Bulletin*, 9, 76-82.
- Rae N & Brennan M (1998b). *Advertising on the Web: Effects of Schema Congruity and the Effectiveness of a Single Adview*. Paper presented at ANZMAC98, University of Otago, November 30 – December 2.
- Smee A; Brennan M; Hoek J & Macpherson T (1998). *A Test of Procedures for Collecting Survey Data Using Electronic Mail*. Paper presented at ANZMAC98, University of Otago, November 30 – December 2.

Acknowledgement: The authors wish to acknowledge the cooperation and support of New Zealand Post Direct Marketing Services, and in particular, Matthew Pickering.

Mike Brennan is a Senior Lecturer, and Nathan Rae and Mathew Parackal are postgraduate students, in the Department of Marketing, Massey University, Palmerston North, New Zealand.