

E-Research: Ethics, Security, Design, and Control in Psychological Research on the Internet

Brian A. Nosek* and **Mahzarin R. Banaji**

Harvard University

Anthony G. Greenwald

University of Washington

Differences between traditional laboratory research and Internet-based research require a review of basic issues of research methodology. These differences have implications for research ethics (e.g., absence of researcher, potential exposure of confidential data and/or identity to a third party, guaranteed debriefing) and security (e.g., confidentiality and anonymity, security of data transmission, security of data storage, and tracking participants over time). We also review basic design issues a researcher should consider before implementing an Internet study, including the problem of participant self-selection and loss of experimental control on the Internet laboratory. An additional challenge for Internet-based research is the increased opportunity for participant misbehavior, intentional or otherwise. We discuss methods to detect and minimize these threats to the validity of Internet-based research.

The potential of the information highway to advance understanding of psychological science is immense, and it is likely that the Internet will decisively shape the nature of psychological research. Yet as any researcher who has attempted to use the Internet to obtain data will have discovered, a host of methodological issues

*Correspondence concerning this article should be addressed to Brian Nosek, Department of Psychology, Harvard University, William James Hall, 33 Kirkland Street, Cambridge, MA 02138 [e-mail: bnosek@wjh.harvard.edu]. We thank Katelyn McKenna for sharing ideas that were directly incorporated in this article and Susan Bouregy for discussions about Internet-based experiments. This research was supported by National Science Foundation grants SBR-9422241, SBR-9422242, SBR-9709924, and SBR-9710172, and National Institute of Mental Health grants MH-41328, MH-01533, and MH-57672 to M. R. Banaji and/or A. G. Greenwald.

require consideration because of differences between standard laboratory research and Internet-based research concerning research methodology. They concern the treatment of participants, the security of the data that are transmitted and obtained, and the internal and external validity of the data. In this article, we identify and review selected issues concerning (1) conducting ethical research on the Internet, (2) personal and data security, (3) experimental design, and (4) experimental control. We raise these issues only insofar as they concern Internet research with an emphasis on on-line experimental and quasi-experimental design. For each methodological issue raised, some potential solutions or strategies are offered.

Ethics

The absence of an experimenter engaged in face-to-face interaction with the participant removes the most obvious source of coercion that has been a source of concern in psychological experimentation. The removal of this concern is a nontrivial benefit of Internet-based research and should play an important role in the cost versus benefit analysis of a research program. Here, we focus on the threats to ethical treatment of participants that can potentially affect Internet-based research. Three differences between Internet and standard laboratory research have implications for ethics: absence of a researcher, uncertainty regarding adequate informed consent and debriefing, and potential loss of participant anonymity or confidentiality. These aspects of Internet research will have a direct influence on the practices of informed consent, debriefing, research with (and without) children, and the protection of privacy on the Internet.

In any study, Internet-based or otherwise, participants ought to be given enough information to judge whether they wish to participate. In Internet laboratories, the manner in which informed consent is obtained is especially important because of the inability of an experimenter to respond to adverse reactions while the experiment is underway. Consent forms should be especially clear and accompanied by FAQs (frequently asked questions) that anticipate potential questions and concerns.

Debriefing

Internet participants could involuntarily end participation in the event of a computer or server crash, a broken Internet connection, a program error, or even a power outage. Participants might also voluntarily end participation because of boredom, frustration, confusion, being late for class, not wanting to miss a favorite television show, or hearing the footsteps of an approaching supervisor. Whatever the cause, early withdrawal from a study is a threat to ensuring adequate debriefing.

There are several options available to an Internet-researcher to enable debriefing even if participants leave the experiment early:

1. Participants could be required to enter an e-mail address at the beginning of a study. Debriefing statements could later be e-mailed to the participant. (This option is not available in studies that assure full anonymity.)
2. A “leave the study” button, made available on every study page, would allow participants to leave the study early and still direct them to a debriefing page.
3. The program driving the experiment could automatically present a debriefing page if the participant prematurely closes the browser window.

In face-to-face debriefings, researchers can make idiosyncratic adjustments for participants who have found the experience unsettling. A number of strategies are available to the Internet researcher to accommodate debriefing requirements that may vary across participants:

1. A list of FAQs that address participant concerns immediately and provide feedback that those concerns are normative, even expected.
2. Debriefing and FAQs should be engaging, making it more likely that the participant will appreciate the purpose of the study and the importance of psychological research.
3. Researchers can offer participants an e-mail address to which to send their questions and concerns about the study.
4. The researcher can be available in a chat room following participation to interactively address concerns or answer questions. In this type of design, the study might be made available only when the researcher is on-line.

Protection of Children

Internet research designed to use children as participants should proceed with care similar to that in standard laboratory settings. A more difficult issue for Internet research is controlling participation in research *not* designed for children. Whereas the participation of an 8-year-old (even a 16-year-old) is not likely to pass unnoticed in a standard laboratory, “catching” such participants on the Internet can be difficult. Asking participants to report their age may be sufficient to remove minors from data analysis, but it is not sufficient to prevent them from completing the study itself. There are, however, strategies available to minimize the opportunity and/or likelihood that children will participate:

1. Design decisions for an experimental Web site should maximize its appeal to adults while simultaneously minimizing its appeal to children. This might be

as simple as avoiding cartoons or popular-culture images that attract children or as involved as gearing the site description and text toward adults.

2. Participant recruiting could target adults by advertising only in adult-dominated list-servs, chat rooms, Web sites, and other places.
3. A password for participation could be required that is available only through adult-targeted advertisements or direct invitation.
4. Experimenters could implement an “adult check” system. Such systems require individuals to register with a centralized database by providing some evidence that they are 18 or older.

U.S. regulations have been developed to protect the privacy of children under the age of 13. Internet researchers should review these instructions before implementing their own projects, particularly the Children’s Online Privacy Protection Act. Resources are available at <http://www.ftc.gov/bcp/online/pubs/online/kidzprivacy.htm>; <http://www.ftc.gov/bcp/online/edcams/kidzprivacy/resources.htm>; <http://www.ftc.gov/ogc/coppa1.htm>; and <http://www.ftc.gov/bcp/online/pubs/buspubs/coppa.htm>

Protecting Participants

The preceding discussion suggests that researchers must pay attention to a different set of issues in Internet research compared to those in standard laboratory paradigms. That does not, however, mean that protecting participants is more difficult in Internet research. In fact, there may be better guarantees for participant protection in Internet research. For instance, the physical absence of a researcher decreases situational demands to remain in an experimental situation that is uncomfortable or unrewarding. Although explicitly coercive tactics are not likely to pass review by ethics committees, implicit situational pressures in most research designs (such as politeness norms) may well discourage participants from prematurely discontinuing participation. The physical absence of a researcher in Internet study designs eliminates these social demands to continue participation, thereby allowing participants greater freedom to withdraw. This fact should not lead investigators to place greater responsibility on the participant to self-regulate participation, but rather to make good use of this added advantage in ensuring the ethical treatment of participants (for additional discussion on Internet research ethics, see Frankel & Siang, 1999).

Security: Protecting Participant Privacy

Many experimental designs do not require identifying information from the participants. In this case, a number of factors indicate that Internet research can guarantee anonymity even more effectively than standard laboratory research.

First, the data can be encrypted, rendering them meaningless to any third party that might intercept them in transit between the participant's machine and the server. Additionally, if no identifying information is collected, the only piece of information that could possibly be used to identify a participant is the Internet Protocol (IP) address. This information typically accompanies data transmitted from the participant's machine to the server, but a researcher can choose not to record this datum. Even IP addresses do not identify information for many Internet surfers. For one, IP addresses identify machines, not individuals. Also, most Internet service providers use floating IP addresses, which identify a particular machine only for the course of a single session. For users who connect to the Internet through such providers, the IP address will change each time they log on. Even so, a small proportion of users have a single machine with a fixed IP address by means of which they can reliably be identified as the "owner" of that machine. Because of this, the IP address must be considered identifying information. Even if IP addresses do not identify a single machine, they can often be used to identify the school, company, or location from which the connection occurred.

There are many types of research designs, however, that do require identifying information. The Internet does afford a perception of anonymity, but in some circumstances that perception may be false. Three elements of Internet research require consideration for protection of confidentiality: data transmission, data storage, and poststudy interaction with participants.

Data Transmission

In an Internet study, there is a small but real possibility that the data will be intercepted by a third party. The standard approach for dealing with this possibility when conducting a Web page-based study is to implement secure server line (SSL) technology. SSL is an encryption technology that encodes information from the client machine to make it meaningless if it is intercepted in transit. If this technology is out of reach (because of cost or sophistication), simpler methods of encryption may effectively protect participant confidentiality. For example, responses to survey questions could be transformed by an algorithm, leaving the researcher to reverse the transformation during analysis. Also, questionnaire items could be given identifying labels that, even if intercepted in transmission, are meaningless to anyone but the researcher. Finally, separating transmissions of identifying information from transmissions of experimental data can increase the assurance of confidentiality. For example, demographic data could be collected with one questionnaire and all experimental data collected with a second questionnaire, with only a randomly assigned identification number to associate them. In survey research conducted through e-mail, it is a simple matter to save responses separately from the e-mail addresses and other identifying information provided in the headers that contain identifying information.

Data Storage

Although only the most intrepid of data burglars may be able to gain access to a laboratory's locked filing cabinet, data stored in files on an Internet-connected server are significantly more vulnerable. The procedures for securing data vary extensively across server designs, so techniques for maintaining strong security of stored data will not be reviewed here. Even so, Internet researchers should be sure to investigate the procedures necessary for securing data on their servers.

Internet Sampling

Internet Samples: Greater Diversity

The Internet has a clear sampling advantage for populations that are difficult to access because (1) the sample is difficult to bring to a laboratory, (2) the population is small, or (3) group members are difficult to find. In addition, although participant recruitment is often constrained to a local community, without leaving the office, Internet researchers from the South can collect samples from the North, psychologists from Japan can sample Nisei and Issei in the United States, and Asians and Asian Americans can be simultaneously sampled via a Web site set up in England.

To date, efforts comparing Internet and laboratory samples find them to show comparable results for a variety of experimental effects (Birnbaum, 1999, 2000; McGraw, Tew, & Williams, 2000). Even so, potential sampling biases should be taken seriously if the researcher intends to claim generalizability from the sample to the population.

Recruiting Participants for Internet Research

Whatever the needed sample, the method of recruiting will have a strong impact on who chooses to participate. Who participants are will depend on two features of participant recruitment: the accessibility of the study and the type of advertising.

Accessibility. Internet researchers have three options for the availability of a study for participation: open, specific, and invited accessibility. Open access refers to studies open to whoever is able to find the Web site. *Open accessibility* studies have the advantage of sampling from the broadest pool possible but are disadvantageous to the extent that it is difficult to control the type of sample that ultimately chooses to participate, not to mention the nonrandom nature of the sample.

Specific accessibility studies constrain participation to individuals who meet specific selection criteria assessed before beginning the study. For example,

researchers interested in age bias among elderly people might have participants who visit their Web site complete a preselection questionnaire with age as one of the preselection questions. Participants meeting selection criteria would be assigned to the age study, whereas others would be assigned to an alternative study or dropped. The specific design also enables a reduction in self-selection biases. A Web site administering multiple studies can advertise itself at a general level (e.g., psychology studies) rather than specifically recruiting for a specific topic (e.g., age attitudes). If random assignment to a study does not occur until after the participant agrees to participate, then selection bias is reduced. However, like any study, attention should be paid to disclosing enough information so that participants can make an informed choice about whether they wish to participate.

A final accessibility option is a more private or *invited accessibility* design. To control participation, researchers can contact a randomly selected sample of participants (for example, from among those who take part in on-line forums for the elderly) and provide each individual with a unique access code and a link directing him or her to the Web site. Only those with viable access codes would then be admitted to the experimental area of the Web site, so those who simply surf onto the site would not be eligible for participation. This procedure also allows the researcher to verify that each participant is engaged in the study on only one occasion. Additional assurance can come from the use of small files called “cookies” that can be stored on user machines. These can be an effective means of tracking visits to a Web site from individual machines. However, users have the option of setting their browsers to reject cookies, rendering them ineffective.

Recruiting. Researchers interested in attracting as many people to their research Web site as possible will often use a *widespread* recruitment method. Widespread recruitment can be as easy as registering the site with some popular search engines (e.g., Excite). A more intensive (and costly) recruitment effort might involve advertising with ad banners at popular Web sites with general appeal (e.g., www.cbs.com). Despite the fact that these methods are the most popular ways to promote a Web site, they are not necessarily the most effective at netting traffic. Simply posting a Web site may not be sufficient to generate adequate samples quickly. Researchers interested in widespread recruitment of participants for an e-mail-based survey might simply post the survey in hundreds of newsgroups and list-servs.

Becoming “viral” is often considered the most effective means for generating traffic to a Web site. Becoming viral depends on having a Web site that is interesting enough that visitors are likely to forward the address to friends. This behavior results in exponential growth in Web site traffic. For example, in November 2000 an artist at www.pixelspill.com produced humorous “proposed Florida ballots” poking fun at the 2000 presidential election debacle in Florida. Showing significant viral spreading, daily traffic at his Web site went from 1,700 to 35,000 to 263,000

in 3 days. Similarly, individuals may forward an e-mail survey to friends who might also be interested in participating.

Although widespread recruiting offers the best opportunity for recruiting a large and diverse sample, it does come with a significant risk: loss of control in the content of advertising for the site. When a research Web site is widely advertised, and especially when it becomes viral, links to the Web site set up by others may present misleading information that biases the sample choosing to visit through that link. For example, the authors of this article maintain a Web site on which participants can measure their implicit attitudes and implicit knowledge of social groups, individuals (political candidates), and themselves (Nosek, Banaji, & Greenwald, in press). Significant press coverage of the site opening and a viral response to the site content led to substantial traffic such that, 3 years after opening, the site averages over 700 hits per day. However, a subset of links to the Web site suggests that the implicit measures reveal “true” preferences (i.e., the task is a lie detector), a claim not endorsed by the researchers or espoused at the Web site. Even though the researchers did not write or control these advertisements, they still operate as recruiting tools. If widespread, these types of advertisements could affect both the types of individuals who decide to participate and the invalid expectations of those participants upon arrival.

Targeted advertising is another recruitment method that can help to increase the researcher’s control over who hears about the study and how the study is described. In targeted advertising, particular groups of interest are contacted directly. For example, messages to newsgroups or chat rooms about auto repair are likely to generate samples interested in auto repair. Targeted advertising, if accompanied by password access or other screening, guarantees that a Web site will be visited only by people contacted directly.

Conclusions for Internet sampling and recruiting. Internet samples are likely to underrepresent populations that have low levels of access to the Internet (e.g., minorities, the poor). The challenges to generalizability posed by Internet samples place some limits on Internet research. Nonetheless, in this regard, the challenges for Internet research are similar to the challenges of all experimental research paradigms. The Internet can be a boon to theory testing in psychology: Large samples can increase the power of tests and allow more confident interpretation of null results, special populations are more accessible, and experimenter bias is minimized.

Web-Based Experimental Design

Creating an Effective Lab Setting

Aronson, Ellsworth, Carlsmith, and Gonzales (1990) described four essential components for the creation of an effective laboratory setting: (1) *coherence*: each

event should appear to be an integral element of the study, (2) *simplicity*: avoid unnecessary complication, (3) *involvement*: participants should never lose interest in the study tasks (also known as “impact”), and (4) *consistency*: the design should contribute to the *creation of the same basic state* in all participants (e.g., similar mood, distractions in the environment, and cognitive load). These elements are as important for Internet research as they are for any other type of design.

The absence of an experimenter introduces some challenges for the coherence and simplicity of Internet-based designs. Although such absence significantly minimizes researcher bias, it increases the likelihood of misunderstanding. Any confusion will lead the participant either to discontinue participation or to incorrectly complete the experimental materials. Also, without a researcher present, it is more difficult to communicate the cover story. Live researchers are able to adapt a script in order to maintain study integrity, but the Internet typically does not allow flexibility in this regard. Finally, live researchers can catch problems in the design quickly, especially if the problems lie in the cover story or instructions. In Internet research, feedback about a manipulation or instruction that “just isn’t working” is not readily available. As a consequence, greater effort will be needed in Internet research to ensure that the instructions are coherent, simple, and similarly interpreted by all participants.

Maximizing experimental impact. In standard laboratory research participants typically appear at a laboratory at a designated time, are greeted by an experimenter, given verbal instructions, and shown into a room expressly designed for the experiments. These features of the experimental setting assist in creating the appropriate signals for setting up the research environment. Unfortunately, none of these signals is available in Internet research. Effective Internet studies have to utilize other means to maximize experimental impact. In short, participants should be made aware of what they will gain, as well as give, by participating. There is no definitive prescription on how to increase the impact of an experimental design in Internet research. The procedural variations that will maximize experimental impact will vary significantly with the nature of the study.

Standardizing the experience for participants. The traditional laboratory environment is usually designed to minimize distracting information and to create the same basic state in all participants. Most laboratory walls are bare, chairs are generally plain and lacking cushy comfort, and colors are neutral and bland. Unfortunately, psychologists do not have keys to Internet participants’ homes and offices to reproduce that venerable, if sterile, ambiance. The variation in environments during Internet research participation is likely to lead to much greater variability in states of the participants compared to standard laboratory settings. Internet participants might participate in any variety of environments—at home, in the office during a coffee break, or in an Internet café at 3 A.M. Also, these participants might complete the study alone, with a group of friends, while holding

a baby, while talking on the phone, while listening to music, or while smoking marijuana (only a very small minority will likely be doing all five at once). The variation in environments is a challenge to researchers trying to create the same basic state in participants.

Some strategies for creating the same basic state (or at least control for variation in state) on the Internet include (1) presenting a list of “requirements” for completing the research, including items such as “be in a quiet place,” “reserve 15 min to complete the task,” and “close other programs”; (2) having all participants complete a warm-up task designed to get participants involved and focused on the study (and, perhaps, standardize participant states); and (3) having a list of questions at the end of the study asking participants to report various distractions (e.g., “were there other people in the room while you completed this study?”). Extra effort to create a similar state among participants will pay off by increasing the power of the experimental design.

Maintaining Experimental Protocols

Components of studies that are relatively straightforward to control in a traditional laboratory become remarkably difficult to control in a Web-based laboratory. For example, the following are events that are unlikely to occur during a standard laboratory study but could easily occur during an Internet study: (1) immediately following a priming induction, the participant could be distracted with a phone call and stop participating for 15 min before returning to finish the study; (2) in a study with a surprise memory task, a participant could return to the original list by hitting the back button on the browser; or (3) after finishing a study, a participant might decide to participate in the same study again. A selection of strategies for dealing with these types of threats to the procedural integrity of Internet studies is reviewed below.

Timing

Many studies depend on completion of the measures in a fixed order and within a specified period of time. Without any safeguards, Internet participants could discontinue participation in the middle of a study only to finish it minutes, hours, or even days later. Controlling the amount of time available for each section of a study is straightforward if an applet (a small program designed to run over the Internet) is used, because deadlines and time limits can be programmed directly into administration of the materials. However, the easiest and most commonly used approach for presenting and collecting information is HTML (Hypertext Markup Language) forms. These forms are easy to produce and have the advantage of being familiar to Internet users. However, some additional design considerations are needed to maximize experimental control

when using HTML. Timing details might be handled with some of the following strategies:

1. Include a specific instruction at the beginning of the study stating the minimum and maximum amount of time needed to complete the study.
2. Include a time stamp in the data collected with HTML forms. Participants can be included or excluded based on the time difference between the submission of the various forms the study.
3. Code the HTML page to automatically forward to a different page after a specified period of time. If that time expires before the participant submits his or her data, the participant can be removed from the study.

Direction and flow. Another issue for Internet studies is the back and reload buttons on Internet browsers. These buttons allow the participant to self-navigate through a study, which is generally undesirable for experimental design. The use of an applet (or similar technology) can eliminate problems of self-navigation because the opportunity to back up can be excluded. However, for users of HTML, a couple of other options exist:

1. A study can be run in a window that does not contain a back or reload button, so that participants do not have an easy means of self-navigating through the study.
2. With use of some programming to help manage presentation of HTML pages such as CGI (common gateway interface) scripts or JSPs (Java server pages), a “tilt” mechanism can be established that ejects participants from the study who attempt to go back or reload.

Multiple data points for single individuals. Experimental design requires that data assumed to be from different participants are clearly identified as such. This requirement disallows participation of a given individual more than once in a single study. Complete prevention of multiple submissions from single individuals is difficult. Even so, a variety of solutions are available to prevent or disallow multiple submissions from a single individual. Below is a selection of strategies for detection or prevention of multiple submissions:

1. During informed consent, ask participants if they have participated in the study before.
2. HTML forms typically, by default, record the IP address of the source. In data analysis, multiple submissions from the same IP address can be removed. In addition to sacrificing anonymity, this solution is not ideal for two reasons. First, acceptable data will be needlessly removed. There are many instances in which data from the same IP address is not coming from a single individual:

a proxy server may service multiple computers from a single IP, a floating IP address may connect to two different computers on two different occasions, and more than one individual may use a single computer. Second, using only one instance from each IP address does not guarantee that all multiple submissions from individuals are removed. Participants could participate in a study from more than one IP address by either using more than one machine, or more likely, using a machine that connects to the Internet with a floating IP address.

3. Require participants to register (have a log-in process). This could be as simple as requiring participants to enter a piece of identifying information (e.g., e-mail address) or as complex as setting up an individual "account" with log-in name and password. With this information, a participant ID tag could be attached to each questionnaire and only the first instance from any participant would be used, or, more impressively, a CGI script could be written to prevent participants from participating in studies that they have already completed.

Deceptive Participants: Intentional or Otherwise

In any method of research, identifying participants who are deliberately or unintentionally responding inaccurately can be challenging. More attention may need to be paid to this issue for Internet research because of the decreased accountability associated with lack of an experimenter and absence of situational cues increasing the impact of the experimental situation. Most techniques for identifying deceptive participants in the standard laboratory can be applied for use in Internet research. In this section, we discuss deceptive practices that are unique to Internet research and identify some methods for deception detection that may be particularly useful for Internet designs.

Two or more participants appearing as one. In most experimental designs, a participant is defined as a single individual, not the collective thoughts of two or more individuals. The Internet, however, can be a group activity. Families, friends, or colleagues may enter an Internet study and take turns responding to questions or discuss them as a group. This type of behavior may be rare. Even so, there is no way to guarantee that it does not occur. To minimize the likelihood that groups of individuals are counted as a single participant, a researcher can (1) state clearly in the instructions that the study is for an individual and not for a group, (2) use marker questions more than once in a study, such as birth date or height (because these are more likely to change if more than one individual is completing the study), and (3) explicitly ask participants to report how many individuals were involved (or were in the room) when the study was completed.

Detecting intentional deception. The Internet allows individuals to conceal their identity and even to adopt an alternative one (e.g., a 20-year-old might pose as

a 70-year-old in a chat room). Also, lack of an experimenter may decrease accountability to such a degree that participants “play” with the study and make up responses that do not accurately reflect their beliefs or identity. Below are some strategies that may help to detect deceptive responding:

1. Repeat questions with definitive answers (e.g., birth date) or that are not likely to change during the course of a study (e.g., occupation). Deceptive responders may be detected by changes in responses to these types of questions.
2. Ask questions in which one of the provided responses is unlikely to be true. Participants who are not taking the study seriously or who are trying to be deliberately deceptive may be more likely to select odd responses. (Excessive use of this strategy may backfire and decrease the impact of the experimental design.)
3. Require that the participant supply identifying information at the beginning of the study. (This approach should be treated with care because of the loss of anonymity and the importance of confidentiality.)
4. Give participants the opportunity to correct their responses by showing them their responses a second time before recording them.

Additional Web-Based Research Methodologies

The Internet is a burgeoning resource, not only for experimental and quasi-experimental research, but also for other research methods, including surveys, natural and archival research, and interviews or participant observations. In this section we briefly discuss some issues of particular relevance to these methods.

Surveys

There are two ways to conduct surveys on the Internet. The first is to set up a Web page containing the survey. Currently, SPSS software is available to create an on-line survey from a data file and automatically encode survey responses into that data file. The methods discussed earlier in this article apply to recruitment, security, and control for this type of survey. Surveys can also be carried out through e-mail, and different techniques are required for reaching specific participant populations. A variety of virtual communities including newsgroups, chat rooms, and MUDs (multiple-user dialogues) offer a space for people with specific interests (e.g., clam baking), traits (e.g., narcoleptics), or qualities (e.g., identical twins) to gather and converse. For psychologists, these groups offer unique opportunities to access survey data from specific participant populations. For instance, one can easily obtain a random sample of active newsgroup participants, as the e-mail address for each user is contained in the header of his or her newsgroup posting, and surveys can then be mailed out to a randomly selected subset of users. However,

with this method it is not possible to obtain a random sample of newsgroup readers who do not actively participate in discussion. Some electronic groups (e.g., Yahoo) provide e-mail addresses for all newsgroup members whether they are active participants or not, enabling random selection. In other cases, users are not identified by their e-mail addresses but rather by a chosen nickname and must be contacted individually to participate. (For a more detailed description of newsgroups, chat rooms, and MUDs and surveying techniques within them, see McKenna & Bargh, 2000).

Natural and Archival Research

The Internet presents a unique opportunity to study individuals and groups within a naturalistic setting without the presence of an intrusive researcher. One can, for instance, observe and code the verbal content of newsgroup posts or chat room comments to test hypotheses about individual or group behavior. One can easily collect data on individuals' or groups' reactions to naturally occurring historical events such as wars, elections, or public figures. One useful feature of the Internet is that one need not wait for events to happen, nor wait months while data are being collected, because archives of newsgroup posts are available through several sources. The most popular source is Google Groups at <http://groups.google.com/>. Messages dating back to 1995 are organized within a searchable database by the newsgroup, date, and time the article was originally posted, as well as by author. Archived posts for the more newly formed Yahoo electronic groups also exist. However, Yahoo does not provide a searchable database for these posts, which appear within digests specific to each group. Thus, the researcher must first identify the groups, and then individually work through the digests and articles within them.

Interviews and Participant Observation

Both in-depth interviews and participant observation can be profitably combined with the other methodologies discussed above, as well as with one another. Glaser, Dixit and Green (this issue) opted for the latter in their study of racist behavior within chat rooms associated with White supremacist groups. Interviews carried out with users in chat rooms and MUDs or conducted through e-mail exchanges can provide valuable insight and rich data about a given phenomenon. The free-response format of interviews provides information that preconfigured surveys and other data collection methods might miss. Insights gleaned through interviews can be empirically tested by including survey questions addressing such issues or through laboratory or Web page-based experimental designs. Participant observation can also yield rich data and is a method that can be particularly beneficial

to survey research. Taking part in the groups under study aids in gaining the trust of the members of those groups and can substantially increase response rates. The researcher must take care, however, to attend to principles of research ethics and not to influence or otherwise affect the responses participants make to the survey by his or her presence in the group. Researchers should participate in no more than half of the groups under study, so that the possibility of participant bias can be checked by comparing survey results from groups in which one participated with those from groups in which one did not. Similarly, as in the study by Glaser et al., when one combines interviews with participant observation, it is important that the interviewer(s) be blind to the hypotheses being tested to reduce the chance of eliciting biased responses.

Conclusion

The challenges of the Internet should not deter investigators from taking advantage of this powerful medium for discovery and education. Yet the responsibility for conducting research that meets the highest standards for ethical treatment and advances science is first and foremost the responsibility of scientists themselves. Expenditure of the requisite resources for research design at the onset of a research enterprise will have major benefits for the ultimate quality of experimental designs and scientific discoveries.

References

- Aronson, E., Ellsworth, P. C., Carlsmith, J. M., & Gonzalez, M. H. (1990). *Methods of research in social psychology*. New York: McGraw-Hill.
- Birnbaum, M. H. (1999). Testing critical properties of decision making on the Internet. *Psychological Science*, 10, 399–407.
- Birnbaum, M. H. (2000). *Psychological experiments on the Internet*. San Diego, CA: Academic.
- Frankel, M. S., & Siang, S. (1999, June). *Ethical and legal aspects of human subjects research on the Internet*. Workshop report for the American Association for the Advancement of Science, Washington, DC.
- McGraw, K. O., Tew, M. D., & Williams, J. E. (2000). The integrity of Web-delivered experiments: Can you trust the data? *Psychological Science*, 11(6), 502–506.
- McKenna, K. Y. A., & Bargh, J. A. (2000). Plan 9 from cyberspace: The implications of the Internet for personality and social psychology. *Personality and Social Psychology Review*, 4(1), 57–75.
- Nosek, B. A., Banaji, M. R., & Greenwald, A. G. (in press). Harvesting implicit group attitudes and beliefs from a demonstration Website. *Group Dynamics*.

BRIAN NOSEK is a graduate student at Yale University working under the supervision of Mahzarin Banaji and currently an exchange scholar at Harvard University. He earned his BS from California Polytechnic State University, San Luis Obispo, in 1995 with a major in psychology and minors in computer science and women's studies.

MAHZARIN R. BANAJI received her BA from Nizam College in India in 1976 and her PhD from Ohio State University in 1986. From 1986 to 2001, she was a Professor of Psychology at Yale University. She is now the Richard Clarke Cabot Professor of Psychology, at Harvard University. She has served as Associate Editor for the *Journal of Experimental Social Psychology*, *Psychological Review*, and the *Encyclopedia of Psychology*. She has received a Guggenheim Foundation Fellowship, the James McKeen Cattell Fellowship, and the Gordon Allport Prize for Intergroup Relations.

ANTHONY G. GREENWALD received his BA from Yale College in 1959 and his PhD from Harvard University in 1963. Before moving to his present position as Professor of Psychology at University of Washington, he was a faculty member at Ohio State University from 1965 to 1986. He holds a Research Scientist Award from the National Institute of Mental Health, is a fellow of the Society of Experimental Psychologists, and received the Donald T. Campbell Award from the Society for Personality and Social Psychology. He was previously Editor of the *Journal of Personality and Social Psychology* and Chair of the Society of Experimental Social Psychology.