Observing the User Experience

A Practitioner's Guide to User Research

Mike Kuniavsky

MORGAN KAUFMANN PUBLISHERS

AN IMPRINT OF ELSEVIER

SAN FRANCISCO SAN DIEGO NEW YORK BOSTON
LONDON SYDNEY TOKYO

quotations. A presentation of the results of each test will be scheduled for everyone affected. The presentation will allow the analyst to answer questions about the results and give further explanations of the proceedings of the test.

The results of all the regular focus groups will be collected into a single report, which will be sent by email as soon as it is complete. In addition to outlining the procedures used and providing a summary of the trends as they apply to the research goals, it will analyze any differences observed between various market segments. There will be a separate report from the final focus group that will compare the values and reactions of users of company Y's services to those observed with company X's.

Maintenance

It's important to revise the research plan every time new knowledge comes in. Everything is subject to change as your team and the company's understanding of the user's experience grows. The research goals, especially, should be reevaluated, refined, and rewritten to take every piece of additional knowledge into account.

Since every piece of research you do is likely to affect your understanding of a number of different research goals, all knowledge about the user experience should be consolidated whenever possible. A good way to do this is to create a minisite for your intranet that contains all the reports and goals and that links each goal to the information that applies to it, and to enter all problems into the development group's bug-tracking software.

Eventually, you should have a set of interlinked documents that, together, make up a more-or-less complete picture of your user population.

CHAPTER 6

Universal Tools: Recruiting and Interviewing

No matter what the research, there are two things you're always going to be doing: looking for people who will give you the best feedback and asking them questions. These two elements, *recruiting* and *interviewing,* make up the backbone of every successful research project, and although some of this information is covered in other chapters, they deserve a chapter of their own.

Recruiting

Even if everything else is perfect, if you get the wrong people to talk about your product, your research can be worse than useless since it gives you confidence in results that don't represent the views and behaviors of your real users. Every product has a target audience, from toasters to missile guidance systems. You need to understand the experience of the people who are actually going to want to use, understand, and buy your product. Anyone else's experience will be of marginal use, or even deceptive. So if you're making a missile guidance system and you invite the North American Toaster Enthusiasts to discuss how it should be improved, you're going to get little feedback that will help you make a product for the Army (and you'll probably end up with a missile that has a 30-second pop-up timer).

The process of finding, inviting, and scheduling the right people for your research is called *recruiting,* and it consists of three basic steps: determining the target audience, finding representative members of that audience, and convincing them to participate in your research.

Although the details of who is recruited vary from project to project, the process generally takes about two weeks with roughly the schedule shown in Table 6.1.

Warning It is useful to have someone recruiting full-time, rather than trying to do it in between other work. Recruiting is almost a full-time job in itself, and it always turns out to be more work than anticipated. To reduce costs, a temp or an intern can recruit full-time if given some training.

Timing	Activity
t − 3 weeks to t − 2 weeks	Determine target audience(s).
t − 2 weeks	Recruit initial pool or prescreen initial group from a database of qualified participants if one is available.
t − 2 weeks to t − 1 week	Screen for final candidates.
t − 2 weeks to t − 1 week	Send invitations to primary qualified candidates.
t − 1 week	Send invitations to secondary qualified candidates.
t − 3 days	Create schedule for all candidates and contact list for alternatives.
t + 1 day	Follow up with participants and researchers.

t = testing day (or days)

Budget at least two hours per participant for recruiting and scheduling—more if you've never recruited that kind of participant before.

Pick Your Audience

Before you begin doing any kind of user experience research, you should have a solid idea of who is going to be using your product. If you don't have a complete audience profile and a thought-out purpose for the product, the project is in danger of veering off into a direction that no one wants or needs.

An audience profile should have been done before the recruiting process begins, but it needs to be refined for recruiting. Start with demographic, Web use, and technological profiles of the audience, and narrow from there. As described in Chapters 4 and 11, a *demographic* profile is one that describes a person's physical and employment characteristics; a *Web use* profile describes someone's Web experience; and a *technological* profile describes their experience with computer technology in general. You want a subset of people who have qualities that make them good research participants. For example, you probably don't care if your regular users have visited competitors' sites, but if you're recruiting for a competitive analysis focus group, then you need to get people with experience using competitors' products. If you want to usability-test a new version of your site, you may want to recruit groups of people who have experience with the old interface. Or you could do the opposite and pick people who have never seen the product to see how they react to it "fresh." These are considerations that don't normally play a large role in determining the target audience for the product as a whole, but they can be important for determining what makes a good user research audience.

Before beginning the recruiting process, you need to flesh out what an ideal research audience is for the specific piece of research you're going out to do. Start by making the demographic and Web use profiles of the product's target audience explicit. If you're making a product that helps people schedule events, your primary audience may be young professionals who throw a lot of parties and have some Internet experience.

EVENT SCHEDULING PROFILE

Demographics

Ages: 25–35

Gender: male or female

College educated

Income: $60K+

Typical reasons for including groups in a target market for the product's marketing—such as the amount of disposable income they have or if they make up the largest proportion of the population—may have nothing to do with providing good feedback to your research. Ask yourself about what makes the research audience different from the target market as a whole. What kind of people will give the best feedback for your research goals?

Which market segments are most affected by the research?
Is there a single group or multiple groups?
Which factors most affect this research?
What are desirable characteristics? What are undesirable?

Explore the answers to these questions, and modify your profile accordingly. Remove factors that are not going to affect how people use or view the product, and add information. Focus on isolating the factors that will make them an ideal *research* audience.

EVENT SCHEDULING RESEARCH PROFILE

Demographics

Ages: 25–35

Gender: male or female

College educated

Web Use

Must have a personal computer at home or work

Must use the Internet

Must have one or more years of Internet experience

Must use the Internet three to five hours a week for personal tasks (shopping, researching, etc.)

Should have some experience with Web productivity applications (calendars, email, portfolio management, etc.)

Behavior

Should regularly hold social events involving at least three other people and have held a social event recently

Social events must involve explicitly inviting people

Be careful not to overdetermine your target audiences. If you find yourself trying to define multiple conditions to the point where you're not sure you can find people that meet them, or if some of the restrictions are mutually exclusive, consider breaking the research into several groups. Thus, if you're making a Web site for truck drivers and you need the perspectives of both long-haul and short-haul drivers, rather than mixing the two groups into one piece of research or finding drivers with experience in both fields, consider running the same research with two groups.

Find Your Audience

Once you have the research profile, you need to go out and look for those people. There are two steps to this: finding *some* people and finding the *right* people.

Note The decision about whether to work with a commercial recruiting firm or do it yourself largely depends on available resources. Recruiting is a time-consuming process, especially for relatively rare groups or for people with complex schedules. However, basic recruiting can be a fairly straightforward process if organized well and given sufficient time.

Commercial recruiters have databases with thousands of people from all walks of life. For them, the task of finding people is fairly straightforward: they search through their databases for basic demographic matches, and then they screen those people for the Web use, technological, and behavior criteria that you've determined. If you're going to use a commercial service, then your next task is to create an appropriate script for the recruiter to filter the general population for the criteria specific to your research. This is called a *screener*. However, if you are doing the recruiting yourself, you need to build your own database.

Your Personal Database

Building your own database is not difficult, but it is time consuming. If you want to recruit for research that's going to happen in the next month, it's probably better to go through a professional service. Once your database is built, however, it's a valuable resource that you can use for years and for all different kinds of research.

The more people you have in your database, the better, but only if they are likely to use, understand, and appreciate your product. Getting a database full of the wrong people is more hassle than it's worth.

Recruiting by email is the most efficient method, although it's possible to use all kinds of methods to fill your database. First, set up a special email account for your recruiting effort. It's important to collect all recruiting-related information in a single place, and a dedicated email account accomplishes this quickly.

Start with Friends and Family

For your first round, send out a note to everyone in the company or put a notice in the company newsletter. Invite your co-workers' friends, roommates, significant others, and family to help the company by participating in your research (in exchange for an honorarium, of course). Ask for people who have no direct connection with the product you're testing. Tell them to send email to the special address you've set up.

Reply to the email with a message containing an explanation of the process and a questionnaire that asks for their basic demographic and Web use information. The questionnaire shouldn't take more than five minutes to fill out and should be as specific as possible so that the responses can be more easily compared. Include questions that will give you insight into people's behavior relative to the kinds of products you expect to do research with, but don't ask questions that are so specific that people's answers are likely to change quickly. Thus, a question about the classes of products the respondent regularly buys is more useful in the long run than a question about the last product bought.

Thank you for expressing interest in helping ________ make the best products we can. We are continually striving to create services and products that are innovative, compelling, useful, and usable. As part of this effort, we're implementing a system by which select members of the general public evaluate our products while they are still in their formative

stages to determine where those products succeed in their design and where they need to be revised.

That's where you come in!

However, before we can set you loose on our newest and greatest ideas, we need some information to help us determine which people are best suited for which products. Thus, this questionnaire. Please fill it out and return it to the address below as soon as you can. As soon as we can, we will contact you to schedule a day and time for the evaluation. This could be from several months from now to this week, depending on which products are ready.

Email completed questionnaires to mikek@adaptivepath.com

Questions

Instructions: for multiple-choice answers, please place an X in the brackets next to the most appropriate choice, like this: [X].

How many hours do you use the Web every week for work?
[ ] None
[ ] Less than 5
[ ] 5 to 10
[ ] 10 to 20
[ ] 20 to 40
[ ] 40+

How many leisure hours do you spend on the Web per week?
[ ] None
[ ] Less than 5
[ ] 5 to 10
[ ] 10 to 20
[ ] 20 to 40
[ ] 40+

What kind of information do you most frequently look at?

What is your primary source of news and information?

(Continued)

What computer platform are you most familiar with?
[ ] Apple Macintosh
[ ] Microsoft Windows
[ ] Linux/Unix
[ ] Other: (please specify)

What software are you most familiar with?

Are you available to come into our downtown San Francisco offices to participate in these evaluations?

Personal Information
Name:

Please check [X] how you'd prefer we contact you.
[ ] Email address:
[ ] Phone number:
[ ] Fax number:

Thank you!

—Mike Kuniavsky, mikek@adaptivepath.com
Adaptive Path

Once the responses start rolling in, you need to collect them into a database. Since the database is not likely to be large by database standards, the software doesn't need to be particularly heavyweight. Microsoft Access or FileMaker Pro are good general-purpose database programs, but even a spreadsheet will work as long as it allows you to do full-text searches on all the fields (though spreadsheet-based databases start getting awkward after 100 or 200 entries, so starting with dedicated database software makes for easier maintenance in the long run). Allocate plenty of time—at least several hours a day if you're in a large company—for data entry in the first couple of days after the request goes out. If the list is large and you expect many responses, consider automating some of the data entry with a Web form. Often, just having consistent output, as provided by a script that

emails what's typed into a form, is all that's necessary to make the data entry task significantly more efficient.

Note You should try to get 10 times as many people in your first round of recruiting as you expect to invite to your first round of research.

Expand Your Horizon

Once you've recruited your first group of friends and family and you've done your first round of research, you're going to need to continue to build your database. For a research program of two to five research projects per month, you need to have a database of about 500 people. If you're planning on doing regular research, building your database should become part of your regular "infrastructure maintenance" ritual, always going on at a low level.

Having recruited a basic, friendly group of people, you should have all the tools to continue recruiting ever larger groups. Finding those groups, however, can be a challenge. There are no foolproof methods. For some audiences, it's fairly easy to find people who are willing to participate in your research, for example, "people who buy groceries." Others, such as "Fortune 500 CEOs," are incredibly hard. Here are some techniques that have worked in the past.

- *Community email mailing lists and online bulletin boards.* Some communities have their own email mailing lists. Whether it's people who live in a certain area or people who use a certain product or people who are in the same industry, there are a lot of mailing lists in the world. Lists that cover a specific city or area are probably the most useful since you don't end up sending mail to people who can't possibly go to your research. Make sure that you know the community's rules and you have the organizers' permission before sending mail. Never *spam* (send unsolicited email). Spam doesn't work, and people get really angry when they get it. Personalize the invitation to the group, making it clear why you're sending the message and what the benefits of participation are.
- *Neighbors.* Your corporate neighbors are a great research resource. They may be familiar with your company, and it's easy for them to come into your offices during a lunch break. Send a note to your neighbors' office managers and ask them to post an invitation in a public place.

- *Your users.* If your site already collects information about its users (either through sales or through marketing research), you can invite people from that list. Filter the list for people in your geographic area, and send messages to those who have agreed to accept mail from your company. Again, never spam, and clearly explain why they were chosen and the nature of the research.
- *New employees.* Although they may not be unbiased, they won't have the intimate knowledge of in-house people and projects, and they'll be motivated and available. One company I know of invites every new employee, whether the receptionist or a vice president, to participate in a research project as soon as he or she starts work. However, be sensitive to the ethical and legal implications in making employees do work that could be considered to be an evaluation of their abilities; the invitation should make it clear that such research is not a job requirement and that results of such evaluations will not be used to judge job performance.
- *Past research participants.* People who match your recruiting criteria are likely to know other people like themselves. If you're planning to do multiple rounds of research with the same target audience, tell the people you've recruited to tell their friends about your research and how they can join the participant pool (you can even make up handouts for them to take home from tests that explain the research and how others can participate).
- *Ads.* You can take out ads on your site and on other sites. Make sure that the ads are clear about the location-specific nature of your recruiting effort and that they specify that an incentive will be paid. And make them attractive (they're ads, after all, and they need to compete with all the other ads). An animated banner can, for example, read "San Francisco resident? Want 60 bucks? Click here!"
- *Traditional methods.* Since you need people who can come into your office, you can also use traditional advertising methods. Classified ads in local newspapers can work, though an ad in the Help Wanted section might get confused for a job offer. If you're in an urban area, there may be leaflet distribution services that can put your invitation on every corkboard in town (college campuses are an especially good source of such bulletin boards and are a great way to recruit college students if they're part of your audience). If you're ambitious, you can even take out a billboard (though if you do, please send me a note since I've never heard of anyone who's actually done that, and I'd like to know how well it works).

If you're at a loss for how to target a certain group of people, first find one person that matches your research participant profile. Then sit down with him or her and brainstorm about other ways that similar people can be reached. Talk about places he or she goes—online and offline—and habits and interests. Suggest some of the methods in the preceding list and ask how to tailor them specifically to his or her lifestyle.

However, don't get to the point where everyone you're recruiting has identical qualities or has been recruited in the same way. Every attribute in a profile has a range of values that could satisfy it, and recruiting should strive to get a diversity of people within the range. For example, when participant ages are defined to be between 25 and 35, it would be better if the recruits ranged throughout the whole range rather than all being 35.

Additional Recruiting Tips

- Rather than including lengthy explanations in email, you can make a Web page that explains the research program and then send email that points people to it. This can be the same page that collects their demo or Web info, but it can be a purely informative page, too. The page should explain the purpose of the recruiting, the incentive, the occasional nature of the research (so people don't think it's a job), and how they can get more information. Add frequently asked questions to it as the recruiting process matures.
- Clearly state the location of the research and the geographic area from which you'd like to recruit in all invitations. You can't predict who will read a request, and it's generally of little use to have people offer to participate from halfway around the world (or even 100 miles away) unless you're willing to pay them for their transportation costs.

- Keep track of when and how people found out about your program. This will tell you which techniques work best and, moreover, how your recruiting methods may affect the perceptions of the participants. People who have purchased something from your company will know more about your product and have stronger opinions than people recruited through your neighbors, who will in turn have stronger opinions than people who saw an ad in a newspaper.
- Update every person's entry every six to eight months to know whether his or her information has changed and whether he or she is still interested. Note the last time the entry was updated in the database and regularly send email to people whose last update is more than six months old, asking them to confirm that the information in your database is still current.
- Try to keep the number of repeat recruits down. In most cases, it's generally acceptable to recruit someone who hasn't participated in any kind of usability or marketing research in six months. In some cases, however, this requirement will have to be waived—for example, when time is short or the person has a rare profile or he or she is in a popular market that gets recruited a lot (IT managers are few, and the targets of much market research, for example).
- When someone gives you especially good feedback, make sure to note it in the database (you can have a check box for "great response"). If you ever need a couple of people on short notice, you can just pull people who you know are articulate and insightful.
- If time allows, run a small recruiting test. Recruit one or two people and run them through the research. If they provide good feedback and they're in the right audience, proceed with recruiting everyone else the same way.

Prescreening

Once you have a database of potential participants, you need to start finding the right people and inviting them to participate. Two or three weeks before the research is supposed to happen, prefilter the database for the people who are close to the group you intend to invite. Use your database searching tools to filter for the general demographic Web use segments that you're interested in. Try to get people that represent the breadth of your target audience while staying within its parameters. Use as much of the information you've collected as you can to whittle down the list to the most promising subjects. For example, if you're going to be testing a chat room component of your service, you may want to look for people who have listed sites with strong community components as some of their favorite sites.

Take the initial list of people you've selected from the database and separate it into *primary* and *secondary* lists, noting the most promising ones. You will use this list later to prioritize whom you schedule and when.

Once you have the list, send everyone on it a screener.

Note Never send out an invitation or a screener to your whole list. Otherwise, you'll be sorting through inappropriate responses and annoying a bunch of people who are not qualified or interested.

The Screener

The screener is probably the most important part of the recruiting process. It's a script that filters out the people who will give you good responses from the ones who merely match the demographic segment. The screener is read by the recruiter over the phone or sent in an email.

Getting the screener right can get you people who are likely to be interested in the product and can speak about their experiences intelligently and eloquently. Getting the screener wrong means getting people who are, at best, only marginally interested in what you have to offer and, at worst, uninterested and inarticulate.

Screeners vary from project to project and from recruiter to recruiter, but there are some general rules that apply to most.

1. *Stick to 20 questions.* There's a reason that game exists. It's possible to find out almost anything about someone in 20 questions. Most target audiences can be defined in 10 to 15 questions, and if the people are prescreened through your database, you can get away with fewer than 5.
2. *Make it short.* It should be possible to get through a whole screener in five to ten minutes.
3. *Be clear and specific.* The person responding to the question should know exactly what kinds of answers are expected.

4. *Never use jargon.* Use simple, straightforward, unambiguous language.
5. *Ask for exact dates, quantities, and times.* This eliminates the problem of one person's "occasionally" being another's "all the time."
6. *Every question should have a purpose.* Each question should help determine whether this person is in the audience or not. Don't ask questions that are incidental or "nice to know" since answers to them will not be useful in recruiting and take everyone's time. Nice-to-know questions can be asked at the beginning of the test.
7. *Order questions from general to specific,* especially in telephone screening. The earlier a question is in the screener, the more people it should eliminate from the pool of potential participants. This saves both the recruiter and the participant time since earlier questions weed out later questions that would be irrelevant otherwise. For example, if a person's age is a major factor in defining the target audience while the amount of Internet experience he or she has is fairly flexible, put the age question before the Internet experience question.
8. *Questions should not lead.* There should be no value judgments or answers implicit in the questions. "Are you bothered by the excessive lag times on the Web?" implies that the person perceives lag times and that he or she should be bothered by them. Instead, phrase questions in a more general (but not more ambiguous) way, and then look for specific responses. "Are there things on the Web that regularly bother you? If so, what?" could get at the same kinds of issues as the first question, without its bias.
9. *Clearly state the format of the research.* State what the research is for, when it's going to happen, how long it's going to take, how much the incentive is, and whether the participants should do anything ahead of time (and whether they should do nothing ahead of time).
10. *Build in flexibility.* Let the recruiter know the acceptable parameters for answering each question, so that you don't needlessly dismiss people who can provide valuable feedback.

Sample Phone Screener

This is a telephone screener for a site that offers an online calendar. The target audience is mostly existing users, but the screener has been expanded to include a couple of potential users and a power user (to get a taste for several different kinds of user experiences).

This screener consists of three sections: an introduction for the recruiter that specifies the ideal target audience, an introduction for the participants, and the main set of questions and selection criteria. It is a relatively complex screener that would probably be too involved for a first user test.

Warning Although it's often tempting to try to recruit for several target audiences, it's critical that their descriptions be very close together and that the bulk of the recruits come from a single audience. Otherwise, you risk fragmenting your data to the point that there is insufficient information to be able to identify patterns. One technique to avoid this is to prioritize your audiences and recruit enough people (usually four to six) for the primary audience before you begin recruiting for secondary audiences.

Target Audience

6–9 people total: 4–6 current users, 1 power user, 1–2 nonusers (but likely potential users)

Current eCalendar Users

People who use eCalendar on a regular basis and have used it recently

Men or women
Any age, but prefer 35–50
Personal income $60K+
Have Internet access at home or work
Use the Web five or more hours a week
Have one or more years of Internet experience
Have created at least five eCalendar events in the last month
Have used eCalendar for at least two months

eCalendar Power User

Someone who uses eCalendar frequently and regularly uses the advanced features

Men or women, but prefer men
Any age, but prefer 40+
Creates at least 20 eCalendar events per month
Has used reminders and calendar overlays
[rest of profile the same as current users]

(Continued)

Potential Users

Have never used eCalendar, but may be aware of it
Have at least five scheduled appointments per week
[rest of profile the same as current users]

Note that some of the criteria are specific, such as the amount of time on the Web and the number of events entered into the calendar, whereas others (such as the gender and age) have ranges.

Logistics

Available on January 20 or 21 between 8 AM and 6 PM
Can come into the downtown SF office
Are not immediately affiliated with eCalendar or its competitors (but can be aware of it)
Have not participated in any usability or marketing research in the last six months

Current User Screener

Hello, my name is __________ from __________. We are seeking a few people who are interested in participating in a paid evaluation of a product that you may find useful. This is not a sales call, and no sales or solicitation efforts will be made at any time.

The evaluation will consist of a one-on-one interview on January 20 or 21 in downtown San Francisco. It will be during working hours and will last about one hour. If you participate, you will receive a cash stipend of $60. The interview will be strictly for research, and all of your comments will be confidential. If you are interested in participating, I need to ask you a few questions to see if you've had the kind of experiences we're looking for.

The introduction simultaneously sets people's expectations and serves as the first round of questioning since people will immediately say if they're unavailable on the given date or if they're unwilling to participate. It's also careful to describe the general nature of the research without going into specifics that may skew people's responses.

Note Alternatively, you can use a more generic introduction ("We are conducting a study, and I would like to ask you a few questions"), but you run the risk of going through the whole process just to find that the person is not available on a given day or is uninterested in participating in this type of research.

Question	Answers	Instructions
1. Do you or any member of your household work in any of the following businesses or industries?	Market research Advertising or media sales Public relations Usability or quality assurance Web design or development	IF YES TO ANY, TERMINATE

Eliminate people who work in industries that can present a conflict of interest. People who are in advertising, usability, Web design, and market research should almost always be eliminated since they're too aware of the kinds of issues that research is aiming to uncover and are unlikely to give an unbiased perspective (even if they want to).

"Terminate" is an instruction to the recruiter that tells him or her to stop the recruiting process and wrap up the interview. There is some termination text provided at the end of the screener.

2. We're looking for people of various ages. Which of the following categories includes your age?	Less than 30	TERMINATE
	30 to 34	ASK Question 3
	35 to 39	
	40 to 45	
	46 to 50	
	More than 50	TERMINATE
3. We're also looking for people of various income levels. Roughly speaking, what's your personal yearly income?		IF LESS THAN $60K, TERMINATE; OTHERWISE, ASK Question 4
4. Do you have a personal computer at home or work?	Yes	ASK Question 5
	No	TERMINATE

(Continued)

Question	Answers	Instructions
5. Do you have Internet access at home or work?	Yes No	ASK Question 6 TERMINATE
6. When did you start using the Internet?		IF LESS THAN 1 YEAR TERMINATE; OTHERWISE, ASK Question 7
7. On average, how many hours a week do you estimate you use the Web?		IF LESS THAN 5, TERMINATE; OTHERWISE, ASK Question 8

Although many of these questions will have already been filtered for if these names were pulled out of a database, it's still useful to make sure that the information is accurate, so it's a good idea to verify the information again. Additionally, the age and income questions may be considered to be invasive by some people. In many cases, that information doesn't affect people's behavior and the questions can be eliminated; in cases where the information is of secondary importance, it's possible to move the questions to the end.

8. How many scheduled meetings or events do you have to keep track of per week?		IF LESS THAN 5, TERMINATE; OTHERWISE, ASK Question 9
9. Have you kept track of any scheduled meetings with an online service in the last month?	Yes No	ASK Question 10 TERMINATE
10. How many?		IF LESS THAN 5, TERMINATE; IF MORE THAN 20, ASK Question 11; OTHERWISE, ASK Question 13
11. Which online calendar service or services did you use?		IF eCalendar IS MENTIONED, ASK Question 12; OTHERWISE, TERMINATE

Question	Answers	Instructions
12. How long have you used eCalendar?		IF 2 MONTHS OR MORE, ASK Question 13; OTHERWISE, TERMINATE
13. Which of the following eCalendar features have you used in the past?	Reminders The address book Calendar overlays The buddy list	IF overlays AND reminders, CONSIDER FOR POWER USER SCHEDULING; ASK Question 14
14. Are you currently working on any projects with eCalendar.com or another company that makes online calendars?	Yes No	TERMINATE ASK Question 15
15. Have you ever participated in a market research interview or discussion group?	Yes No	ASK Question 16 SCHEDULE
16. When was the last time?		IF LESS THAN 6 MONTHS, TERMINATE; OTHERWISE, ASK Question 17

Note When you are short on time or when you can interview only a couple of people, it's sometimes useful to filter for people you've interviewed before rather than filtering them out. People who are known to give good, honest, articulate feedback provide a shortcut to useful information. In such situations, it's important to inform the analyst about the participant's background since it'll probably affect the way he or she interprets the results.

When people have participated in user research recently, they're more likely to give unconsciously biased responses since they'll be familiar with the format of the research and may try to anticipate the "appropriate" answers. Since you're almost always looking for unguarded, unbiased responses, it's generally a good idea to eliminate these people from the research unless you have no other choice. Moreover, some people see the incentive payments as a good way to supplement their income and may try to get into any kind of marketing or usability research project. Inviting such people should be avoided entirely since they're unlikely to provide natural or truthful responses.

17. In a couple of sentences, describe what your favorite Web sites have been lately and why you like them.		[NOTE DOWN] TERMINATE IF INARTICULATE; OTHERWISE SCHEDULE.

Open-ended questions like this serve two purposes. They give the recruiter an idea of how articulate a potential participant is, and they can collect information that's not easily formatted as a multiple-choice question. Save them for the end, and don't put more than one in any given screener since they're time consuming and don't filter many people out. (That said, some recruiters prefer to ask the open-ended question at the beginning since it's less intimidating and can catch terminally inarticulate participants early in the process.)

Note If you're using the more generic introduction, you can also replace the termination statement with something less specific, such as "That's all the questions I have. Thank you very much for participating in our study."

Question	Answers	Instructions
TERMINATE		That's all the questions I have. Thank you very much for participating. Although we're not scheduling people who fit your profile right now, we may call you again for a different research project.

NOTE: Potential user screener is identical, minus questions 12,13, and 14.

That's it for the questions. Would you be willing to come into our downtown San Francisco offices for a one-hour paid interview on January 20 or 21? You will be reimbursed $60 for your time, and your help will be greatly appreciated by a development team that's currently making a product you may be interested in.

Email Screeners

Email screeners are similar to telephone screeners in terms of the kinds of questions that are asked, except that email screeners are designed to be answered by the participant. This makes them closer to a survey questionnaire in formation. They need to clearly specify the full range of answers that are expected, and they need to be explicit in their descriptions. They should also have fewer and simpler questions. This makes it easier on the participants since they don't feel compelled to write long descriptions and are more willing to respond (since answering 20 questions with sentences is too daunting and time consuming for most people). A question such as

How many scheduled meetings or events do you have to keep track of per week?

would be written in an email as

Roughly how many scheduled meetings or events do you have to keep track of in a typical week? (please check one)
[ ] None
[ ] 1 to 5
[ ] 6 to 10
[ ] 11 to 15
[ ] 16 to 20
[ ] More than 20

Once written and formatted in such a way that it'll look good in any email program, a screener can be emailed to the invitation list (don't forget to set the "Reply to:" address to be your catchall research email!). Avoid HTML email because many email programs can't display HTML correctly. In addition, if you have a system that can quickly create online surveys, you can just put up a Web page with all the questions in a quickly accessible fill-in form (but give people the option to respond via email in case they don't have Web access or have problems with the online form).

Scheduling

Once you have your list of potential participants, you need to create a schedule. First, define a *scheduling window.* What are the appropriate times for the researchers? Are there times that are inappropriate for your target audience? (Physicians and network administrators, for example, are almost impossible to schedule during weekday working hours.) Are there key observers who need to be there at certain times? Company meetings? Holidays?

With the scheduling windows in hand, start scheduling. There are lots of ways to schedule people, and you should feel free to use whatever procedure is right for you. The following sequence was used successfully at Wired Digital and is similar to the procedures used by other recruiters:

- Write the invitation.
- Invite primary candidates.
- Receive responses and schedule primary candidates.
- Invite secondary candidates.
- Receive responses and schedule secondary candidates.
- Confirm primary candidates.
- Confirm secondary candidates.
- Send thank-you notes to unscheduled candidates.
- Make email and telephone confirmation calls to all participants the day before their scheduled time.
- Create and distribute a schedule of all participants.
- Create and put up signs guiding participants to the testing facility.

For events with fixed schedules, such as focus groups or when critical observers have limited schedules, invitations are straightforward.

For more flexible situations, schedules should be constructed around the participant's preferences. It's often as effective to ask candidates for a list of times when they're available as it is to dictate times when they should show up. Allowing participants to drive the scheduling also shows them, up front, that their input is valuable. Asking them to provide a primary time and a couple of backup times eliminates most conflicts.

The Invitation

The invitation, whether delivered by phone or email, should reiterate what the research is about, why participation is important, how much the participant will be paid, where the research will take place, and when they should show up (or, in the case of site visits, when they can expect the visitors to arrive).

It's time again for another Wired Interface Testing Group usability study, and you've been selected from our database to receive this invitation to attend. The next series will be held on WEDNESDAY, DECEMBER 8, 2003, and THURSDAY, DECEMBER 9, 2003. If you would like to participate, you MUST RESPOND TO THIS SCREENER BY THIS FRIDAY, DECEMBER 3, 2003.

<u>SESSION DETAILS:</u>
-DATES: WED, 12/8 and THURS, 12/9
-LENGTH: One hour
-LOCATION: Wired Digital offices, San Francisco

An honorarium of $60.00 will be paid to you by check within two weeks of your test session, and you will be required to sign a nondisclosure agreement prior to beginning the study.

Please respond to the following statements or questions by FRIDAY, DECEMBER 3, 2003:

*If you CANNOT participate in this series, that's no problem. But to help us keep our records up to date, please check one of the following statements:

[ ] I can't make it this time, but keep me in the WITG database for future tests.
[ ] I won't be able to partipate in the WITG program any longer. Please remove my name and address from the database.

*If you CAN participate in this session, please answer the questions below.

For each of the sites below, please indicate how many times you visit it in an average week. Put 0 if you've visited it before, but don't go there

(Continued)

often; 1 for once a week; 2 for twice a week; and so on. Leave it blank if you've never been there:
_____ New York Times On The Web
_____ Wired News
_____ CNN
_____ Yahoo
_____ ZDNet
_____ Slashdot
_____ Salon
_____ Other: ________________ (please list)

Now please indicate your FIRST and SECOND preference for an appointment by choosing a day and time between the hours of 11:00 AM and 6:00 PM ONLY:

1. App't Day and Time—FIRST PREFERENCE
   [ ] Wed, 12/8/03 at ________ (11:00 AM - 6:00 PM ONLY)
   [ ] Thur, 12/9/03 at ________
   [ ] Either day at ________

2. App't Day and Time—SECOND PREFERENCE
   [ ] Wed, 12/8/03 at ________ (11:00 AM - 6:00 PM ONLY)
   [ ] Thur, 12/9/03 at ________
   [ ] Either day at ________

OK, that's it for now. As soon as we've received everyone's answers and had a chance to coordinate a schedule, we'll get back to you with an update of our test schedule. Thanks very much for your participation!

—Mike and Kevin

Confirmation and Reconfirmation

Since the scheduling process can involve a lot of dates and times flying around, at Wired we found it useful to request a final confirmation that the times we scheduled were the times to which the

participants thought they were agreeing. This consists of a short, clear, message that says something like this.

Dear Terri,

Thank you for responding to our invitation to next week's usability study. We'd like to set up your appointment for the following day and time:

Day: Monday, Nov 15, 2003
Time: 5:00 PM

*VERY IMPORTANT—YOU MUST CONFIRM YOUR APPOINTMENT: in order to confirm this appointment time, please either REPLY TO THIS EMAIL or CALL 415-235-3468!! If you don't respond, we will assume you are not coming, and we'll have to schedule another participant in your time slot.

All confirmations should also contain a line to the effect of "Because we have tight schedules, it's critical that you show up exactly at the time you're scheduled or a few minutes early," as well as specific transportation and parking instructions.

OUR ADDRESS

We're located at 582 Market Street at Montgomery, Suite 602, in the Hobart Building. There is a BART and Muni station right in front of the building, and all the trains stop there. The number 2, 7, 9, and 14 buses run at street level. If you're driving, the closest parking is an underground pay lot at Mission at Second that is best approached from Mission Street going east.

OTHER GOOD STUFF TO KNOW

This session should last about one hour, and you'll receive an honorarium in the amount of $60.00 by check within two weeks following your

(Continued)

participation. The schedule we have set up for this test is very tight, so if you think you're going to be running late, PLEASE CALL (415) 235-3468 TO INFORM US. If at all possible, plan to arrive 10 minutes before your scheduled appointment time, and you'll really be a star!

Thanks again and we'll see you next week!

—Mike and Kevin

In addition, we found it helpful to standardize the format of all correspondence with a one- to two-sentence summary of the message at the beginning; a standard layout for all messages; and a contact person's name, email address, and phone number at the end. Encourage the participants to contact that person as soon as they know they won't be able to make an appointment or if they have any questions. Keeping the text tight, but with humor and personality, helps make the process feel less clinical, and people are more likely to read it.

Finally, the single most effective way of preventing no-shows is a combination email and telephone reminder the day before the scheduled appointment. It reminds people, subtly reinforces the importance of their input, and gives you time to schedule an alternate in case they can't make it.

Choosing an Incentive

An incentive is something that encourages and rewards people for their help. As such, it needs to do two things: it needs to convince someone that they should share their time and experience, and it needs to communicate their value to your company. Never shortchange people on an incentive.

Although the chance to give input into a favorite product is a strong motivator (and should always be emphasized), for most people the best incentive is cash. Everyone knows what to do with it, and it's immediately useful. The basic cash incentive formula in the San Francisco Bay Area in 2003 is $0.75 for every minute spent doing general consumer research and $1 per minute for one-on-one interviews and for professionals' time. So a one-hour interview should carry a $60 incentive, and a two-hour focus group can pay each participant between $70 and $120 (depending on the complexity of the profile). Similar amounts apply in Boston and Chicago.

Extra incentive should be paid to people with unique perspectives or who are willing to come in at awkward hours. How much extra will be based on the conditions. If you need someone to come in tomorrow, an additional $20 to $40 to the basic incentive may be appropriate. A busy manager may not be willing to break up his or her day for less than $200. It's often useful to budget for a range and then let the recruiter up the incentive to get good candidates.

For certain people, no amount of money will be enough to encourage them to participate. For these people, you need to come up with alternative incentives. Executives may value meeting other executives in the same industry, so presenting the research as an opportunity to hobnob with peers may be attractive to them. Charity donations are sometimes valued, as are gift certificates to luxury restaurants or theater tickets. A luxury cruise line site had retired millionaires who liked to travel as its customers. It recruited them by giving them all-expenses paid flights to New York. It cost the site several thousand dollars per participant, but they were able to get their exact target audience from all over the country. Merely offering the same amount in cash would not have showed the understanding and respect for the participants' priorities to convince them to participate.

Sometimes people are willing to participate for less than you're offering, or even for free, but it's not a good idea to take advantage of such audiences. For one thing, groups of people who are available to do research for cheap are a self-selecting bunch: either they have a lot of time on their hands or they're really eager to talk to you or they need the money. The people who fit any of those criteria are unlikely to be representative of your target audience.

Pitfalls

The Wrong People

Sometimes you recruit the wrong group. Maybe you misworded a key screener question, and everyone answered it opposite to what you had intended. Maybe you forgot a key element when determining the target audience. Maybe the recruiter always picked the minimal criteria when you only wanted a couple of people with that description. Regardless, sometimes you will mistakenly invite the wrong people. Every researcher has stories about coming into a room and realizing that it's a roomful of people who shouldn't be there at all.

Preparation, a good screener, a carefully chosen target audience, and asking some key questions at the beginning of the research will minimize this situation, but it still happens. If you find yourself in such a situation, there are two things you can do: cancel the session and re-recruit, or try to get as much information out of the group, hoping that some of their feedback will be usable. Sometimes having the wrong group can be enlightening since the comparison of answers with a group of the right people can lead to a better understanding of the whole population (maybe your target audience isn't so unique after all?). However, when time is short and accurate results are critical, cancellation is a perfectly acceptable option, with the efforts channeled toward getting a group of the right people.

Before you decide to re-recruit (or you start blaming the recruiter), you have to determine why the recruiting failed. In some cases, it will be really obvious—you wanted newbies, but not people who have never seen a computer—but sometimes it may not be clear why a certain group didn't fit your vision of an ideal audience. Identify the elements that resulted in the wrong group being brought in: is it because of their age? Is it because of their experience? It could be that you have not specified a key element ("factory workers" versus "factory workers who work on the shop floor," for example). Or maybe your expectation of how people behave may not match how your target audience thinks. In some cases, there may be *no* way to recruit the right group because it simply doesn't exist.

Warning Always pay people who show up what was promised, even if you made a mistake in recruiting them. They had no idea they were the wrong audience and showed up in good faith. Likewise, if the recruiter scheduled the wrong people and the problem lies with your screener, the recruiter should also be paid in full.

No-Shows

It seems to be human nature that about a quarter of any invited group won't show up, for one reason or another. You can compensate for this by scheduling extra people and removing unexcused no-shows from the database.

If you would like to guarantee that you don't get any no-shows, you can double-schedule every time slot. This makes for twice the recruiting and scheduling work and twice the incentive expense (since you have to pay everyone who shows up, whether or not you need them to participate in research), so it's rarely worth the trouble, but it's useful in certain situations such as when a senior stakeholder is scheduled to observe the research. You can also use the extra people for other kinds of research if you have the resources and a secondary research project. You could have a usability test as a primary research project and a questionnaire or card-sorting task as a secondary project since these don't need as much interaction with a researcher as the usability test. You could even have members of the development team who aren't observing the usability test interview the participants (provided that the team members who do this have some user research experience).

Likewise, you can schedule a "floater" for a block of time, say, three hours, to hang out in case someone doesn't show up. However, qualified people who have two to three hours in the middle of their day are difficult to find, so this isn't as attractive an option as it may seem at first. Floaters should be paid the proportional amount of incentive for the amount of time they spend waiting (two to three times what someone scheduled for a single slot would get), and they, too, can participate in secondary research if they're not needed as substitutes.

If the schedule and facility rental budget permits, you can also create "makeup" slots a day or two after the initial batch of research. People who can't make their initially scheduled time may be amenable to rescheduling during one of the additional slots.

"Snow Days"

Carolyn Snyder, a user experience consultant and principal of Snyder Consulting, suggests that in areas with snowfall that there be alternative days decided ahead of time in case the roads are bad. If there's a possibility of inclement weather, she suggests picking a school near the testing facility and telling people to check the school closure report. If that school is closed, odds are that roads are impassable, so people automatically know they're rescheduled for the alternative day.

Make sure you note down who did not show up. Unless the people contacted you or the recruiter ahead of time to warn of their absence, do not pay them and remove them from the database, or tell the recruiter to do likewise.

Bias

The methods you use to recruit can create a bias in the kinds of people you find. If you recruit 20-somethings with $30K incomes from a discussion list on day trading, you're likely to get people with different perspectives than if you get people with the same demographic makeup from a mailing list about labor relations. Those differences may not matter if you're going to be talking about music, but they probably will if you talk about politics.

Nearly all recruitment techniques bias the population. A truly random sample is impossible, so there's no getting away from some amount of bias. The question to answer is whether the bias matters for this project. Often it does not, but when you recruit, you should mind the potential for bias and how it can affect the research. Include a field in the database for how each person was recruited, and evaluate the sources that come up often for their potential to skew the results. So, for example, if half of your participants come from the same source, you may want to consider whether that recruitment method is affecting your results.

Anonymity

Your company's image and its reputation can also bias the attitudes of people coming to test with you. A famous name can be intimidating, but an unknown company or product can seem more trivial than it really is. Preconceptions may alter the way people behave and affect the results of research.

If you don't want people to know your company's name as they are recruited and scheduled, it may be useful to maintain an alternative identity. Don't deceive people, but working through an alternative email address (even one from a free service) and describing the company and research without giving any specifics can be useful. In general, people don't mind if you tell them that you'd rather not reveal the name of the company or its products.

Teenagers' Schedules

Teenagers are hard to recruit even though they're one of the most sought-after demographics. They're difficult to find, difficult to engage, they are often difficult to contact, they need parental consent to participate, they have complicated schedules, and there are state and federal laws that restrict their employment. To verify that their schedules haven't changed, you should get as much contact information as you can (portable telephone numbers, pager numbers, parents' numbers) and contact them the day before the research, but prepare for last-minute changes and overbook more than for adult groups.

Building and Space Preparation

Although not technically a recruiting pitfall, a common problem with scheduling participants is that the company isn't prepared to greet them. If they're arriving at an office, everyone they meet should be expecting them. The doorman should know where to send them, the receptionist should know they're coming. There should be clear signs pointing them in the right direction.

Working with a Professional Recruiter

A professional recruiter can take a lot of the headache out of doing the recruiting. Pros have access to a large number of people in a broad range of markets and can schedule them in short order. They can be a great asset if you are short on time, if you don't want to maintain your own database, or if you aren't sure how to recruit a certain audience.

However, working with professionals is more involved than just picking up a phone and saying that you need people that match this-and-such a profile. Recruiters do not know your business or your research requirements. Careful preparation and collaboration on the part of the research team is necessary for the recruiter to be able to get the people you need when you need them.

Where to Find a Recruiter

Since recruiters serve the market research industry, they are often associated with companies that offer a broad range of market research services, from providing special conference rooms to focus group moderation to designing and running whole research projects. If you're in a major metropolitan area, finding one should not be hard. You can generally look up "marketing research" in the phone book and find half a dozen. You can also contact the American Marketing Association *(www.ama.org)*, the Marketing Research Association *(www.mra-net.org)*, or the European Society for Opinion and Market Research *(www.esomar.nl)*, and they can point you toward other resources.

When choosing a recruiting firm, do some research. As with any service, ask for quotes from several different firms and request references to their last couple of user experience research clients. If you're recruiting for usability testing, ask them whether they do recruiting for "one-on-one interviews" (or, ideally, for usability testing) and if they recruit for offsite facilities. Follow up on the references they provide, and ask about the accuracy of the recruiting, how receptive the recruiter was to changes, how many people were recruited, and how specific the recruiting criteria were. Ask the recruiting firm whether they have ever done usability recruiting for Web sites or software. Most big firms have done some, but many have only dabbled in it and are not yet experienced in interpreting people's responses for Web usability research.

You're not likely to find out how they get their user lists and how big their databases are since that's the heart of their recruiting business, but if you can, it's useful information for comparison.

Since recruiting may only be part of their business, full-featured marketing research companies may be reluctant to just recruit, and may insist on renting out their research spaces (which can be expensive) or providing other services. In such cases, you may want to consider their services, but if you're only interested in recruiting, then it's important to find a place that's willing to specialize in it.

For small, fast user research projects, the best bet is to find an independent recruiter who specializes in recruiting and scheduling. These types of companies are rarer than the general market research services, but they exist in most metropolitan areas. If you can't find one, ask some colleagues or check with one of the usability or human factors associations such as the Usability Professionals Association or Human Computer Interaction SIG of the Association of Computing Machinery (ACM SIG-CHI).

What Recruiters Can Provide

Every professional recruiter should handle all the aspects of finding the right people, scheduling them, answering their questions, and reminding them of their appointments. In addition, an experienced recruiter can also help you focus your research by narrowing your target audience to just to the people you're interested in. If you are not sure about how to specify the right market for your research, ask your recruiter to help you put together a profile.

Some recruiters will write a screener for you, and in practice, they'll often modify your screener to better fit their style and policies. You should participate in the process by requesting a copy before the recruiting begins to ensure that appropriate emphasis is placed on the elements that are of prime importance to you.

In addition to providing you with the profiles of the scheduled participants, many recruiters are willing to take notes about the responses of people who do not get scheduled. Although this tally does not qualify as a survey, it can be interesting information since it can reveal unexpected patterns in your target audience.

Many recruiters will do incentive payment administration, too, cutting checks or providing cash as appropriate. In addition, larger recruiting companies can recruit in various geographic markets for you.

What Recruiters Need from You

The most important thing to provide a recruiter is a complete audience description. They will be happy to recruit just about any description you give them, but since they don't know your business, they can't filter based on unstated assumptions you have about your audience. If your description isn't sufficiently specific, you're likely to get people you don't want. If you say that you're looking for people with less than a year of Internet experience without saying that you want them to have at least some Internet familiarity, the recruiter may get you a bunch of people who have never seen a browser. Be as specific as possible. If you're looking for white-collar workers, define what you mean by "white-collar workers."

Include a list of whom to exclude. What are industries that present a potential conflict of interest? How much past participation in research is too much? What companies make competing products to yours?

Provide enough time (at least a full week) and avoid changing the parameters of the recruiting after it has begun. Changes are a hassle for the recruiter, and they're likely to pass the cost of that hassle on to you. If you cancel, be prepared to pay at least part of the incentive fee (all of it if you cancel on the same day) and all the recruiting costs.

Tell the recruiter where to get qualified candidates. If you already have a list of customers, such as people who have signed up for a newsletter, and you can filter it for your geographic area, offer the recruiter the list (though be aware that list members will likely know your product and have positive preconceptions about it, otherwise they wouldn't have signed up for the newsletter). If you don't know where to find candidates, you should give the recruiter as much information as you can about where to look for the target audience and how prevalent you expect them to be in the population as a whole. If you have no idea, then you're essentially asking the recruiter to do a telephone survey, which can take a long time and be quite expensive.

Provide direction for how to handle marginal candidates. Clearly note which qualities are flexible and how flexible they are. ("We prefer 25- to 35-year-olds, but will accept people between 20 and 40 if they match all the other criteria well.")

Describe the research to the recruiter. This will help the recruiter understand how to answer questions and may give him or her additional ideas for how to structure the target market description. Is it a series of groups? Is it a series of one-on-one interviews? Will it be done at a special facility? Will it be at the participant's house? Will it be focused on their attitudes, their experiences, or how well they can use a prototype? Tell the recruiter how much of this information is appropriate to tell the participants.

Finally, explain any terminology that's necessary so that the recruiter can interpret people's responses appropriately. If you're looking for IT managers who regularly buy "hot-swappable, fault-tolerant, low RF, Mil-spec, narrow-gage U-racks," you should probably tell the recruiter something about what all those words mean. The recruiter may not need to use these words in conversation, but knowing what they mean will help the recruiter understand the questions he or she is asking.

What They Cost

As I'm writing this, in 2002 in San Francisco, recruiting typically costs between $60 and $200 per participant scheduled, with $100 as the typical rate. Consumers fall on the lower end of the scale and professionals on the upper end. For common groups with few restrictions, such as grocery shoppers who have used the Web at least once, it may be even cheaper. For other groups, such as human resource vice presidents who run organizations with enterprise-wide knowledge management systems, the cost may be significantly higher.

Other services, such as screener writing or response tabulation, can be rolled into the whole cost or charged on an hourly basis, with hourly rates of between $50 and $100 an hour.

When There Are Recruiting Problems

Don't accept bad recruiting. If it's clear that the priorities and questions in the screener were not strictly followed, ask for your money back or for some better recruiting. First, however, make sure that your audience description did not allow the recruiter to interpret it in a way you hadn't intended. Most recruiters, although not elated by the prospect, will re-recruit participants who didn't fall into the target description.

Further, as a final courtesy to the recruiter, tell him or her when a participant was particularly good or bad. This will help the recruiter in future recruiting efforts.

Interviewing

Most of the research described in this book boils down to one technique: the interview. Observation is critical, but to really know the user's experience, you have to ask him or her about it, and that's an interview. The usability interview—the other tool that's a basic part of nearly all user experience research—differs from the kind of interview an investigative journalist or a prospective employer would hold. It's more formal, more standardized, and as a kind of *nondirected interview,* tries to completely remove the perspective of the person asking the questions from the interview.

The Interview Structure

Nearly every user experience interview, whether it's a one-person lunchtime chat or a ten-person focus group, has a similar underlying structure. It's an hourglass shape that begins with the most general information and then moves to more and more specific questions before stepping back for a bigger perspective and concluding with a summary and wrap-up. Here is one way of dividing a standard interview process into six phases.

1. *Introduction.* All participants introduce themselves. In groups, it's important to know that the other people in the group are somewhat like you in order to feel comfortable, so a group introduction emphasizes the similarities between all the participants, including the interviewer. In contrast, an individual interview introduction establishes the role of the interviewer as a neutral, but sympathetic, entity.
2. *Warm-up.* The process of answering questions or engaging in a discussion needs everyone to be in an appropriate frame of mind. The warm-up in any interview is designed to get people to step away from their regular lives and focus on thinking about the product and the work of answering questions.
3. *General issues.* The initial product-specific round of questions concentrates on the issues that surround the product and how people use it. The focus is on attitudes, expectation, assumptions, and experiences. Asking these kinds of questions early prevents the assumptions of the product development team from skewing people's perceptions. Often, the product isn't even named during this phase.
4. *Deep focus.* The product, or product idea, is introduced, and people concentrate on the details of what it does, how it does it, whether they can use it, and what their immediate experience of it is. For usability testing, this phase makes up the bulk of the interview, but for contextual inquiry, where the point is to uncover problems, it may never enter the discussion.
5. *Retrospective.* This phase allows people to evaluate the product or idea in a broader light. The discussion is comparable to the "General issues" phase, but the discussion is focused on how the ideas introduced in the "Deep focus" phase affect the issues discussed earlier.
6. *Wrap-up.* This is generally the shortest phase of the interview. It formally completes the interview so that the participants aren't left hanging when the last question is asked, and it brings the discussion back to the most general administrative topics.

Warning Do a dry run with every new interview script. Run through it with a colleague or a sample participant, complete with all recording devices and prototypes, and then revise it appropriately.

Nondirected Interviewing

A famous scientist once asked the following question on a survey:

> *Does your employer or his representative resort to trickery in order to defraud you of a part of your earnings?**

This is a leading question. Before you read on, think about what makes this a leading question. What in it implies a "right" answer? What is the actual information the author is trying to elicit? What would have to be different for the question not to be a leading question?

The scientist who wrote it was Karl Marx, and he clearly had an answer that he was expecting, and it wasn't "no."

Leading questions are the bane of all social research since they inject the prejudices of the person asking a question into a situation that should be completely about the perspective of the person answering it. But avoiding directed questioning is easier said than done. It requires a constant vigilance on the part of the person asking the questions and a deeply held belief in the need to know people's thoughts *unconditionally.*

Nondirected interviewing is the process of conducting interviews that do not lead or bias the answers. It's the process of getting at the user's thoughts, feelings, and experiences without filtering those thoughts through the preconceptions of the interviewer.

The Neutral Interviewer

As the person writing and asking the questions in a nondirected interview, your job is to step outside everything you know and feel about your product. Forget all the hard work and creativity. Put away all hopes for success and all fears of failure. Ignore everything you've ever heard or thought about it. See it in a completely neutral light, as if it's not yours at all. It's merely a thing you're asking questions about, a thing that you care nothing about.

This seems harsh, but it's necessary in order to be able to understand the feedback people give you, both positive and negative,

*[T.B. Bottomore and Maximilien Rubel, eds., Karl Marx: Selected Writings in Sociology and Social Philosophy (New York: McGraw-Hill, 1956), p. 208; as cited in Earl Babbie, Survey Research Methods (Belmont, California: Wadsworth, 1990), p. 37]

and relate that to the process of making the product into what *they* want and need, not what you think they want and need. Otherwise, you'll always be seeing either the silver lining or the cloud, when you need to be seeing both.

Zen aside, asking questions so as to not bias the respondent's answer involves a lot of self-imposed distance and a rigorously critical examination of your assumptions. This can be especially difficult when the product under examination is one you are intimately familiar with or one you have a lot of interest in. At first, it's going to feel like you're expending a lot of energy not to ask the obvious questions or that your questions are coming out stilted. With some experience, it becomes clearer which questions lead people and how to phrase questions so that you get the most natural responses. Eventually—when you've achieved nondirected question enlightenment—your questions will sound natural, analysis will be easier, and the unbiased answers you get will give you greater confidence in your results.

Composing Nondirected Questions

Most important, every question should be focused on the person answering it. It should focus on *experience,* not extrapolation. Our understanding of our own behavior rarely corresponds to how we really behave. When we try to put ourselves into others' shoes, we idealize and simplify. That's useful in trying to understand people's ideals, but it's rarely useful in understanding their behavior. A question such as "Is this a useful feature?" can be easily misinterpreted as "In the universe of all things, do you think that someone somewhere could find some use for this feature?" Even if most people take it at face value, the potential of misunderstanding makes all replies questionable. "Is this feature valuable to the work you do right now?" clarifies the perspective.

Similarly, questions should concentrate on *immediate experience.* People's current behavior better predicts their future behavior than do their predictions. If you ask people "Is this interesting to you?" they may imagine that at some point they could find it interesting and say yes. But the things that are interesting in theory are quite different from the things that people will remember and return to. If they find something compelling right now, they're likely to continue to find it compelling. Thus, the responses to "If it were available right now, would you use it? Why?" will be more useful.

Questions should be *nonjudgmental.* The person answering the question should not think that you're expecting a specific answer or that any answer is wrong. You can (and should) state this explicitly, but it works better if the question reinforces that view. "Don't you think that this would be better if it was also available on PDAs?" implies that the person asking the question thinks that it would be a good idea and that they will disapprove if they hear otherwise. "If this feature were available tomorrow on PDAs, would you use it?" doesn't imply that there's an expected answer (though it suffers from being a binary question, as described later). An even better approach would be to ask, "Is there any other way you'd like to use a feature like this?" and then prompt to discuss PDAs after they've stated their initial thoughts.

Questions should be *focused on a single topic.* A question that has an "and" or an "or" linking two ideas leads to ambiguity since it's often unclear which part of the question is being answered. "How would this product be useful to you in school or at work?" is actually two questions. An answer to it may insufficiently differentiate between them.

Keep questions *open-ended.* If given a limited choice, people will choose one of the options, even if their view lies outside those options or if more than one is acceptable. They'll adjust their definitions of the options and pick the one that's closest to how they feel. But that's not how they really feel. You should always provide an out from a close-ended question, unless you're absolutely sure that the options cover all the possibilities. That's rarely the case since you're most often looking for the shades of meaning. "Which feature from the following list is most important to you?" assumes that there *are* features that are important, and it assumes that there is one that's more important than any other. A better way would be to say "Rate from 1 to 5 how important each of the following features is to you, where 1 is least important and 5 is most important. Put 0 if a feature is completely unimportant. Write down any features we may have missed" or, ignoring the feature naming scheme entirely, "Does the product do anything that's particularly useful to you? If so, what is it? What makes it useful?"

Avoid *binary questions.* They're an especially insidious form of close-ended questions. Binary questions are of the form "yes/no" or "true/false" or "this/that," and they force people to make a black-and-white choice when their attitude may not lie near either extreme. "Is this a good product?" misses a lot of the subtlety in people's attitudes. Although it may be nice to get a quick sample of people's off-the-cuff opinions, it's much more valuable to know what they find good and bad about the idea, rather than just whether they think the whole thing is good or bad. "What, if anything, do you like about this product?"

Running a Nondirected Interview

A nondirected interview is conducted just as you would any other interview, except that you have to listen more closely to the meaning of your words and the words of the person you're talking to for signs of bias. There are a number of things you can do to increase the quality of the responses.

Define terms. Words are ambiguous and easily misused. "That thing" can refer to a button, a feature, or the whole site. Personal definitions of words can be different from either the dictionary definition or the development team's definition. Someone may speak of a simple function as a "module," whereas the development team may call complex clusters of functions "modules." When using a technical term, make sure that you clearly define it first. Whenever possible, use the respondent's definition of a word (even if it's not how you use it), but make sure that you understand what that definition is first (which may mean asking the respondent to define it). This is especially important in group interactions, where everyone can come in with different definitions.

Don't force opinions. There are times when we just don't have an opinion about something. We may have never thought about a given question in qualitative terms, or we may not have enough information about it in order to form an opinion. When asked for an opinion, most people will form one, but it's not going to be carefully considered or deeply held. When asking a question that requires an opinion, it's good to make sure that the people answering are likely to have an opinion already. "Would this be better if it were done automatically?" may not make any sense to someone who has no experience with "this."

Restate answers. One of the best techniques to cut through problems with questions is to bounce the respondent's answer back at him or her using different words. It clarifies a lot of the subtlety of terminology and verifies that you've understood the answer and that the respondent understood the question. Immediately after someone has finished a thought, you can say something like "So I hear you saying that . . ." and state it as you just understood it, but using different words. However, avoid substituting the "correct" terminology for the words that the person has used. Investigate his or her understanding of the terminology first. So if someone refers to the "order summary," but it's really the "confirmation page," ask the person to elaborate what he or she expects to find on an "order summary" before using the term confirmation page in restating the point.

Follow up with examples, but always wait for an undirected answer first. Sometimes people understand a question, but may not know how to start answering it. If you are precise with your wording, it shouldn't be an issue. Occasionally, though, you may want to ask a question that's intentionally broad, to see how people understand a concept or what their most general thoughts are. Prepare an example (or two) for the questions you feel may need examples. After the participants have given their initial answers, you can refocus their thoughts with an example. Say you're running a focus group that's brainstorming new features. If they're defining features too narrowly and seem to have reached an impasse, you can say, "Now what if it were to email you whenever items you liked were on sale?" and see if the participants can come up with other ideas along the same lines. Don't give more than a couple of examples since that tends to frame people's perceptions too strongly.

Use artifacts to keep people focused on the present and to trigger ideas. Artifacts are the material products of people's work: the notes, the papers, the tools, and so on. Bring participants back to their immediate environment by asking questions that have to do with the physical objects (or the software objects) that they deal with on a regular basis. When someone is talking about "shopping carts" in the abstract, ask about "this shopping cart." When you're in the field and they're talking about how a certain procedure is done, ask them to show it to you with the actual objects. The idealized situation people imagine and discuss in the abstract is often different from the practical situation in which they live, and the objects they use help remind them of the grungy details that are missing from the ideal.

Be aware of your own expectations. Watch for situations that surprise you or when you find yourself predicting the interviewees' next statement. Despite the exhortations at the beginning of this section, it's impossible to be a blank slate when coming into an interview situation. There are going to be things you assume or expect from the interaction, and these are going to affect how you run the interview. If you're aware of these assumptions, it makes avoiding them easier.

Never say the participant is wrong. Even if someone's understanding of how a product works or what it's for is completely different from what was intended, never tell the person that his or her perspective is wrong. Study the person's perspective and try to understand where it comes from and why he or she has it. It may well be that the person's understanding doesn't match others' or yours, but it's never wrong.

Listen carefully to the questions that are asked of you. Questions reveal a lot about how people understand a product or a situation, and they're important to understanding people's experience and expectations. Probe why people are asking the question. If someone asks, "Is that how it's supposed to work?" for example, answer with a question that reveals more of the person's mental model: "Is that how you think it works?" or "Is that how you expected it to work?"

Note Observers can be present during interviews. Having an observer present makes the interview less intimate, but observers can be useful as note takers or just as a second set of eyes. The extent of their participation should be determined by the moderator, but there generally shouldn't be more than one in-room observer, and he or she should always be introduced. I've found that it works well to create special times when observers are allowed to ask questions.

Keep questions simple, both in language and in intent. Use questions to uncover assumptions and perceptions, not prove points or justify actions. A good question does the minimum necessary to elicit a perspective or view, and no more. Analysis of the answers will provide the meaning that can prove and justify. Questions should focus on getting the clearest raw information.

Always review your tapes. It's easy to miss a key statement or a subtle distinction when relying on your memory and notes. Always spend some time with your tapes—whether audio or video—verifying that your views of the discussion accurately represent what happened and how future research can be conducted better.

Common Problems

- Close-ended questions that should be open-ended. "Which of these three logos do you like the most?" is not particularly useful if they don't like any of them. "Is there anything you like about any of these logos?" will tell you what underlying characteristics people find compelling, if any. That will allow you to tailor the logo to those characteristics rather than to an arbitrary choice.

- Questions with complex answers posed as binary questions. "Is the Daily Update an important feature to you?" ignores all the reasons it would or would not be. Maybe they don't plan on checking the site every day, but a weekly update would be great. Maybe there's no need for an update at all. "Is there anything about the Daily Update that you find interesting?" will tell you which parts of it are interesting.
- Loaded words or words with multiple meanings. Be precise in the words that you use. "When you're trying to find something in a site and you get hopelessly lost, what do you do?" "Hopelessly" is imprecise. It can be interpreted by one person as meaning "pretty lost" and by another as "lost without any possibility of ever finding anything." Rewriting the question as "What do you do if, in course of looking for something on a site, you realize that you don't know how to get back to an earlier point?"
- Asking people to predict the future. As mentioned earlier, when people try to project their actions into the future, they often oversimplify and idealize to the extent that their predictions have little to do with what they actually do. People are much better at explaining the reasons for their actions as they're doing them than they are at predicting their actions ahead of time. If you're interested in how someone will behave in a given situation, put him or her into that situation (or a suitable simulation).
- Invocation of authority or peer pressure. For example, "Most people say that it's pretty easy to find information with this tool. Was that your experience, too?" or "Our designers have a lot of experience making navigation tools, and they came up with this one. How well did it work for you?" These questions can almost always be simplified to the actual question being asked: "Describe your experience using this tool."
- Assuming you know the answer. I've found myself half-listening to a response to a question, assuming that it's going to be a variant on what I've already heard, only to do a double take

when someone answers in a way that I'm totally unprepared for. Sometimes people even use many of the same words as what you're expecting, but a critical negation or spin may reverse or fundamentally change the meaning of what they're saying. Listen carefully to every word.
- Assuming that they can answer the question. Not everyone knows what they know and what they don't know. If you ask someone whether something is the best in its class, you're assuming that he or she is familiar enough with all the products in the class and that he or she can make a balanced, knowledgeable evaluation of all the products.

Problems don't just arise in the formulation of questions. The interpretation of answers also depends on the way questions are asked. There are a couple of behaviors to watch out for when asking questions, so that you can catch them and follow up quickly, making later analysis less ambiguous.

- People won't always say what they believe. Sometimes they'll say yes to avoid conflict when they mean no. Watch for the clues about what they really mean. These can take the form of hesitant answers or answers that are inconsistent with previous statements. There can be even more subtle cues, such as someone shaking his or her head no while saying yes or suddenly losing articulation. Attempt to catch such situations as they're happening and ask the person to clarify. Often, just giving the person the floor gives him or her the confidence to say what he or she really means.
- People will sometimes answer a different question from the one you asked. In a situation where someone is thinking hard about a topic—maybe because he or she is in the middle of a task or trying to remember a situation—he or she may easily mishear the specifics of your question. Sometimes participants have their own agenda and really want to discuss things you're not asking about. Listen carefully for what they're really saying and whether it's directly related to what you're asking. If it's clearly off track, interrupt, and ask the question again, using slightly different wording and emphasis. Don't be afraid to be persistent.

When to Break the Rules

Clearly, following all these rules and suggestions will make for a pretty dry conversation, and that may be worse than the bias it eliminates. People should feel comfortable talking to you and answering questions honestly. You should feel comfortable talking to them.

So take all these rules as suggestions when constructing your questions and try to follow through as much as possible. However, feel free to improvise and humanize your interviews by providing examples or letting the participant "off the hook" if a question seems too difficult to answer as it was posed. An interview can be both nondirected and comfortable. Ultimately, the best interview is the one that provides the information you need when you need it. What it takes to do that will be different in every interview. These rules and guidelines will help you get the best information you can, but only you will know how to implement them appropriately.

Videotaping Interviews

Every single interview and interaction should be videotaped, if at all possible. Many people consider video documentation a fancy form of audio recording. Sometimes that's true, but it can reveal crucial moments in any interaction that just can't be captured on audio. A crucial shrug while someone is saying yes, but they really mean no, can be the crux in understanding the person's perspective correctly. A momentary pause of a mouse over one button before clicking on another can reveal the core confusion in a feature. Plus, it frees the moderator from having to simultaneously take notes and think about moderating.

Videotaping is quite inexpensive and, if introduced and placed carefully, quickly disappears into the background for most people, so it's a relatively unobtrusive technique. The video camera can be introduced in the beginning of the interview, placed on a tripod in an inconspicuous location, and the interview can continue normally. The tape then becomes a permanent record that can be mined for critical nuances and exact quotations (both verbal and physical).

Photography uses less equipment and allows you to collect a close-up record of specific items and arrangements in an interview, but it creates a disruptive process where the researcher stops the flow of conversation in order to take a picture. However, in some situations—such as on-location contextual inquiry interviews in security-conscious organizations—it's the only way to document. In those cases, it should be coupled with an audio recording of the interview.