

CHAPTER

9

Usability testing on 10 cents a day

KEEPING TESTING SIMPLE—SO YOU DO ENOUGH OF IT

About once a month, I get one of these phone calls:



As soon as I hear “launching in two weeks” (or even “two months”) and “usability testing” in the same sentence, I start to get that old fireman-headed-into-the-burning-chemical-factory feeling, because I have a pretty good idea of what’s going on.

If it’s two weeks, then it’s almost certainly a request for a disaster check. The launch is fast approaching and everyone’s getting nervous, and someone finally says, “Maybe we better do some usability testing.”

If it’s two months, then odds are that what they want is to settle some ongoing internal debates—usually about something very specific like color schemes. Opinion around the office is split between two different designs; some people like the sexy one, some like the elegant one. Finally someone with enough clout to authorize the expense gets tired of the arguing and says, “All right, let’s get some testing done to settle this.”

And while usability testing will sometimes settle these arguments, the main thing it usually ends up doing is revealing that the things they were arguing about aren't all that important. People often test to decide which color drapes are best, only to learn that they forgot to put windows in the room. For instance, they might discover that it doesn't make much difference whether you go with the horizontal navigation bar or the vertical menus if nobody understands the value proposition of your site.

Sadly, this is how most usability testing gets done: too little, too late, and for all the wrong reasons.

Repeat after me: Focus groups are not usability tests.

Sometimes that initial phone call is even scarier:



When the last-minute request is for a focus group, it's usually a sign that the request originated in Marketing. When Web sites are being designed, the folks in Marketing often feel like they don't have much clout. Even though they're the ones who spend the most time trying to figure out who the site's audience is and what they want, the designers and developers are the ones with most of the hands-on control over how the site actually gets put together.

As the launch date approaches, the Marketing people may feel that their only hope of sanity prevailing is to appeal to a higher authority: research. And the kind of research they know is focus groups.

I often have to work very hard to make clients understand that what they need is usability testing, not focus groups. Here's the difference in a nutshell:

- › In a **focus group**, a small group of people (usually 5 to 8) sit around a table and react to ideas and designs that are shown to them. It's a group process, and much of its value comes from participants reacting to each other's opinions. Focus groups are good for quickly getting a sampling of users' opinions and feelings about things.
- › In a **usability test**, one user at a time is shown something (whether it's a Web site, a prototype of a site, or some sketches of individual pages) and asked to either (a) figure out what it is, or (b) try to use it to do a typical task.

Focus groups can be great for determining what your audience wants, needs, and likes—in the abstract. They're good for testing whether the idea behind the site makes sense and your value proposition is attractive. And they can be a good way to test the names you're using for features of your site, and to find out how people feel about your competitors.

But they're *not* good for learning about whether your site works and how to improve it.

The kinds of things you can learn from focus groups are the things you need to learn early on, *before* you begin designing the site. Focus groups are for EARLY in the process. You can even run them late in the process if you want to do a reality check and fine-tune your message, but *don't* mistake them for usability testing. They *won't* tell you whether people can actually use your site.

Several true things about testing

Here are the main things I know about testing:

- › **If you want a great site, you've got to test.** After you've worked on a site for even a few weeks, you can't see it freshly anymore. You know too much. The only way to find out if it really works is to test it.

Testing reminds you that not everyone thinks the way you do, knows what you know, uses the Web the way you do.

I used to say that the best way to think about testing was that it was like travel: a broadening experience. It reminds you how different—and the same—people are, and gives you a fresh perspective on things.

But I finally realized that testing is really more like having friends visiting from out of town. Inevitably, as you make the tourist rounds with them, you see things about your home town that you usually don't notice because you're so used to them. And at the same time, you realize that a lot of things that you take for granted aren't obvious to everybody.

- **Testing one user is 100 percent better than testing none.** Testing always works, and even the worst test with the wrong user will show you important things you can do to improve your site. I make a point of always doing a live user test at my workshops so that people can see that it's very easy to do and it always produces an abundance of valuable insights. I ask for a volunteer and have him try to perform a task on a site belonging to one of the other attendees. These tests last less than ten minutes, but the person whose site is being tested usually scribbles several pages of notes. And they always ask if they can have the recording of the test to show to their team back home. (One person told me that after his team saw the recording, they made one change to their site which they later calculated had resulted in \$100,000 in savings.)
- **Testing one user early in the project is better than testing 50 near the end.** Most people assume that testing needs to be a big deal. But if you make it into a big deal, you won't do it early enough or often enough to get the most out of it. A simple test early—while you still have time to use what you learn from it—is almost always more valuable than a sophisticated test later.

Part of the conventional wisdom about Web development is that it's very easy to go in and make changes. The truth is, it turns out that it's not that easy to make changes to a site once it's in use. Some percentage of users will resist almost any kind of change, and even apparently simple changes often turn out to have far-reaching effects, so anything you can keep from building wrong in the first place is gravy.

- › **The importance of recruiting representative users is overrated.** It's good to do your testing with people who are like the people who will use your site, but it's much more important to test early and often. My motto—as you'll see—is “Recruit loosely, and grade on a curve.”
- › **The point of testing is not to prove or disprove something. It's to inform your judgment.** People like to think, for instance, that they can use testing to prove whether navigation system “a” is better than navigation system “b”, but you can't. No one has the resources to set up the kind of controlled experiment you'd need. What testing *can* do is provide you with invaluable input which, taken together with your experience, professional judgment, and common sense, will make it easier for you to choose wisely—and with greater confidence—between “a” and “b.”
- › **Testing is an iterative process.** Testing isn't something you do once. You make something, test it, fix it, and test it again.
- › **Nothing beats a live audience reaction.** One reason why the Marx Brothers' movies are so wonderful is that before they started filming they would go on tour on the vaudeville circuit and perform scenes from the movie, doing five shows a day, improvising constantly and noting which lines got the best laughs. Even after they'd settled on a line, Groucho would insist on trying slight variations to see if it could be improved.



Mrs. Teasdale (Margaret Dumont) and Rufus T. Firefly eavesdrop in *Duck Soup*.

Lost our lease, going-out-of-business-sale usability testing

Usability testing has been around for a long time, and the basic idea is pretty simple: If you want to know whether your software or your Web site or your VCR remote control is easy enough to use, watch some people while they try to use it and note where they run into trouble. Then fix it, and test it again.

In the beginning, though, usability testing was a very expensive proposition. You had to have a usability lab with an observation room behind a one-way mirror, and at least two video cameras so you could record the users' reactions *and* the thing they were using. You had to recruit a lot of people so you could get results

THE TOP FIVE PLAUSIBLE EXCUSES FOR NOT TESTING WEB SITES



We don't have the time.

It's true that most Web development schedules seem to be based on the punchline from a Dilbert cartoon. If testing is going to add to everybody's to-do list, if you have to adjust development schedules around tests and involve key people in preparing for them, then it won't get done. That's why you have to make testing as small a deal as possible. Done right, it will save time, because you won't have to (a) argue endlessly, and (b) redo things at the end.



We don't have the money.

Forget \$5,000 to 15,000. If you can convince someone to bring in a camcorder from home, you'll only need to spend about \$300 for each round of tests.



We don't have the expertise.

The least-known fact about usability testing is that it's incredibly easy to do. Yes, some people will be better at it than others, but I've never seen a usability test fail to produce useful results, no matter how poorly it was conducted.



We don't have a usability lab.

You don't need one. All you really need is a room with a desk, a computer, and two chairs where you won't be interrupted.



We wouldn't know how to interpret the results.

One of the nicest things about usability testing is that the important lessons tend to be obvious to everyone who's watching. The serious problems are hard to miss.

that were statistically significant. It was Science. It cost \$20,000 to \$50,000 a shot. It didn't happen very often.

But in 1989 Jakob Nielsen wrote a paper titled "Usability Engineering at a Discount"¹ and pointed out that it didn't have to be that way. You didn't need a

¹ *Proceedings of the Third International Conference on Human-Computer Interaction, Boston, MA, Sept. 1989.*

usability lab, and you could achieve the same results with a lot fewer users.

The idea of discount usability testing was a huge step forward. The only problem is that a decade later most people still perceive testing as a big deal, hiring someone to conduct a test still costs \$5,000 to \$15,000, and as a result it doesn't happen nearly often enough.

What I'm going to commend to you in this chapter is something even more drastic: Lost our lease, going-out-of-business-sale usability testing.

I'm going to try to explain how to do your own testing when you have *no* money and *no* time. Don't get me wrong: *If you can afford to hire a professional to do your testing, by all means do it!* But *don't* do it if it means you'll do less testing.

	TRADITIONAL TESTING	LOST-OUR-LEASE TESTING
NUMBER OF USERS PER TEST	Usually eight or more to justify the set-up costs	Three or four
RECRUITING EFFORT	Select carefully to match target audience	Grab some people. Almost anybody who uses the Web will do.
WHERE TO TEST	A usability lab, with an observation room and a one-way mirror	Any office or conference room
WHO DOES THE TESTING	An experienced usability professional	Any reasonably patient human being
ADVANCE PLANNING	Tests have to be scheduled weeks in advance to reserve a usability lab and allow time for recruiting	Tests can be done almost any time, with little advance scheduling
PREPARATION	Draft, discuss, and revise a test protocol	Decide what you're going to show
WHAT/WHEN DO YOU TEST?	Unless you have a huge budget, put all your eggs in one basket and test once when the site is nearly complete	Run small tests continually throughout the development process
COST	\$5,000 to \$15,000 (or more)	\$300 (a \$50 to \$100 stipend for each user) or less
WHAT HAPPENS AFTERWARDS	A 20-page written report appears a week later, then the development team meets to decide what changes to make	The development team (and interested stakeholders) debrief over lunch the same day

How many users should you test?

In most cases, I tend to think the ideal number of users for each round of testing is three, or at most four.

The first three users are very likely to encounter nearly all of the most significant problems,² and it's much more important to do more rounds of testing than to wring everything you can out of each round. Testing only three users helps ensure that you *will* do another round soon.³

Also, since you will have fixed the problems you uncovered in the first round, in the next round it's likely that all three users will uncover a new set of problems, since they won't be getting stuck on the first set of problems.

Testing only three or four users also makes it possible to test and debrief in the same day, so you can take advantage of what you've learned right away. Also, when you test more than four at a time, you usually end up with more notes than anyone has time to process—many of them about things that are really “nits,” which can actually make it harder to see the forest for the trees.


In fact this is one of the reasons why I've almost completely stopped generating written reports (what I refer to as the “big honking report”) for my expert reviews and for usability tests. I finally realized that for most Web teams their ability to *find* problems greatly exceeds the resources they have available to fix them, so it's important to stay focused on the most serious problems. Instead of written reports, nowadays I report my findings in a conference call with the entire Web team, which may last for an hour or two. By the end of the call, we've all agreed which problems are most important to fix, and how they're going to fix them.

² See Jakob Nielsen's March 2000 Alertbox column “Why You Only Need to Test with 5 Users” at www.useit.com for a good discussion of the topic.

³ If you're hiring someone to do the testing for you and money is no object, you might as well test six or eight users since the additional cost per user will be comparatively low. But only if it won't mean you'll do fewer rounds of testing.

ONE TEST WITH 8 USERS


8 users



Eight users may find more problems in a single test.


But the worst problems will usually keep them from getting far enough to encounter some others.

TOTAL PROBLEMS FOUND: 5




TWO TESTS WITH 3 USERS

First test: 3 users




Three users may not find as many problems in a single test.

Second test: 3 users



But in the second test, with the first set of problems fixed, they'll find problems they couldn't have seen in the first test.

TOTAL PROBLEMS FOUND: 9



Recruit loosely and grade on a curve

When people decide to test, they often spend a lot of time trying to recruit users who they think will precisely reflect their target audience—for instance, male accountants between the ages of 25 and 30 with one to three years of computer experience who have recently purchased expensive shoes.

The best-kept secret of usability testing is the extent to which *it doesn't much matter who you test*.

For most sites, all you really need are people who have used the Web enough to know the basics.

If you can afford to hire someone to recruit the participants for you *and it won't* reduce the number of rounds of testing that you do, then by all means be as specific as you want. But if finding the ideal user means you're going to do fewer tests, I recommend a different approach:

Take anyone you can get (within limits) and grade on a curve.

In other words, try to find users who reflect your audience, but don't get hung up about it. Instead, try to make allowances for the differences between the people you test and your audience. I favor this approach for three reasons:

- › **We're all beginners under the skin.** Scratch an expert and you'll often find someone who's muddling through—just at a higher level.
- › **It's usually not a good idea to design a site so that only your target audience can use it.** If you design a site for accountants using terminology that you think all accountants will understand, what you'll probably discover is that a small but not insignificant number of accountants won't know what you're talking about. And in most cases, you need to be addressing novices as well as experts anyway, and if your grandmother can use it, an expert can.
- › **Experts are rarely insulted by something that is clear enough for beginners.** Everybody appreciates clarity. (True clarity, that is, and not just something that's been “dumbed down.”)

The exceptions:

- › **If your site is going to be used almost exclusively by one type of user and it's no harder to recruit from that group,** then do it. For instance, if your audience will be almost entirely women, then by all means test just women.
- › **If your audience is split between clearly defined groups with very divergent interests and needs,** then you need to test users from each group at least once. For instance, if you're building a university site, for at least one round of testing you want to recruit two students, two professors, two high school seniors, and two administrators. But for the other rounds, you can choose any mix.


- › **If using your site requires specific domain knowledge** (e.g., a currency exchange site for money management professionals), then you need to recruit people with that domain knowledge for at least one round of tests. But don't do it for every round if it will reduce the number of tests you do.

When you're recruiting:

- › **Offer a reasonable incentive.** Typical stipends for a one-hour test session range from \$50 for "average" Web users to several hundred dollars for professionals from a specific domain, like cardiologists for instance. I like to offer people a little more than the going rate, since (a) it makes it clear that I value their opinion, and (b) people tend to show up on time, eager to participate. Remember, even if the session is only 30 minutes, people usually have to block out another hour for travel time. Also, I'd rather have people who are curious about the process than people who are desperate for the money.
- › **Keep the invitation simple.** "We need to have a few people look at our Web site and give us some feedback. It's very easy, and would take about forty-five minutes to an hour. And you'll be paid \$__ for your time."
- › **Avoid discussing the site (or the organization behind the site) beforehand.** You want their first look to tell you whether they can figure out what it is from a standing start. (Of course, if they're coming to your office, they'll have a pretty good idea whose site it is.)
- › **Don't be embarrassed to ask friends and neighbors.** You don't have to feel like you're imposing if you ask friends or neighbors to participate. Most people enjoy the experience. It's fun to have someone take your opinion seriously and get paid for it, and they often learn something useful that they didn't know about the Web or computers in general.

Where do you test?

All you really need is an office or conference room with two chairs, a PC or Mac (with an Internet connection, if you're testing a live site), a camcorder, a long video cable, and a tripod.



Test subject (A) sits in front of computer monitor (B), while facilitator (C) tells him what to do and asks questions. Camcorder (D) powered by squirrel (E) is pointed at the monitor to record what the subject sees.

Meanwhile, cable (F) carries signal from camcorder to TV (G) in a nearby room where interested team members (H) can observe.

You can use the video cable to run the signal from the camcorder to a TV in another office—or even a cubicle—nearby so everyone on the development team can watch without disturbing the user.

The camcorder needs to transmit what the user sees (the computer screen or the designs on paper, depending on what you're testing) and what the user and the facilitator say. In a usability lab, you'll often see a second camera used to show the observers the user's face, but this isn't necessary: The user's tone of voice usually conveys frustration pretty effectively.

You can buy the camcorder, TV, cable, and tripod for less than \$600. But if your budget won't stretch that far, you can probably twist somebody's arm to bring in a camcorder from home on test days.

I don't recommend using the camcorder to videotape the sessions. In fact, I used to recommend not doing any video recording at all, because the tapes were almost never used and it made the whole process more complicated and expensive.

In the past few years though, three things have changed: PCs have gotten much faster, disk drives have gotten much larger, and screen recording software has improved dramatically. Screen recorders like Camtasia⁴ run in the background on the test PC and record everything that happens on the screen and everything the user and the facilitator say in a video file you can play on the PC. It turns out that these files are very valuable because they're much easier to review quickly than videotape and they're very easy to share over a network. I recommend that you always use a screen recorder during user tests.

Who should do the testing?

Almost anyone can facilitate a usability test; all it really takes is the courage to try it. With a little practice, most people can get quite good at it.

Try to choose someone who tends to be patient, calm, empathetic, a good listener, and inherently fair. Don't choose someone whom you would describe as "definitely not a people person" or "the office crank."

Who should observe?

Anybody who wants to. It's a good idea to encourage everyone—team members, people from marketing and business development, and any other stakeholders—to attend.

When people ask me how they can convince senior management that their organization should be investing in usability, my strongest recommendation doesn't have anything to do with things like "demonstrating return on

⁴ *There are a number of screen recorders available, but I'm partial to Camtasia, made by TechSmith, the same company that makes the screen capture program SnagIt (<http://www.techsmith.com>). It's very reliable and has a number of extremely useful features, and it costs about \$300. For \$1,000 more, they have a product called Morae specifically designed for capturing usability tests—sort of like Camtasia on steroids—which allows observers to view the test live on a networked PC, eliminating the need for a camcorder.*

investment.” The tactic that I think works best is getting management to observe even one user test. Tell them that you’re going to be doing some usability testing and it would be great for the Web team’s morale if they could just poke their head in for a few minutes. In my experience, executives often become fascinated and stay longer than they’d planned, because it’s the first time they’ve seen their site in action and it’s often not nearly as pretty a picture as they’d imagined.

What do you test, and when do you test it?

The key is to start testing early (it’s really *never* too early) and test often, at each phase of Web development.

Before you even begin designing your site, you should be testing comparable sites. They may be actual competitors, or they may be sites that are similar in style, organization, or features to what you have in mind.

Use them yourself, then watch one or two other people use them and see what works and what doesn’t. Many people overlook this step, but it’s invaluable—like having someone build a working prototype for you for free.

If you’ve never conducted a test before testing comparable sites, it will give you a pressure-free chance to get the hang of it. It will also give you a chance to develop a thick skin. The first few times you test your own site, it’s hard not to take it personally when people don’t get it. Testing someone else’s site first will help you see how people react to sites and give you a chance to get used to it.

Since the comparable sites are “live,” you can do two kinds of testing: “Get it” testing and key tasks.

- **“Get it” testing** is just what it sounds like: show them the site, and see if they get it—do they understand the purpose of the site, the value proposition, how it’s organized, how it works, and so on.
- **Key task testing** means asking the user to do something, then watching how well they do.

As a rule, you'll always get more revealing results if you can find a way to observe users doing tasks that they have a hand in choosing. It's much better, for instance, to say "Find a book you want to buy, or a book you bought recently" than "Find a cookbook for under \$14." When people are doing made-up tasks, they have no emotional investment in it, and they can't use as much of their personal knowledge.

As you begin designing your own site, it's never too early to start showing your design ideas to users, beginning with your first rough sketches. Designers are often reluctant to show work in progress, but users may actually feel freer to comment on something that looks unfinished, since they know you haven't got as much invested in it and it's still subject to change. Also, since it's not a polished design, users won't be distracted by details of implementation and they can focus on the essence and the wording.

Later, as you begin building parts of the site or functioning prototypes, you can begin testing key tasks on your own site.

I also recommend doing what I call Cubicle tests: Whenever you build a new kind of page—particularly forms—you should print the page out and show it to the person in the next cubicle and see if they can make sense out of it. This kind of informal testing can be very efficient, and eliminate a lot of potential problems.

A sample test session

Here's an annotated excerpt from a typical—but imaginary—test session. The site is real, but it has since been redesigned. The participant's name is Janice, and she's about 25 years old.

INTRODUCTION

Hi, Janice. My name is Steve Krug, and I'm going to be walking you through this session.

You probably already know, but let me explain why we've asked you to come here today. We're testing a Web site that we're working on so we can see what it's like for actual people to use it.

I want to make it clear right away that we're testing the *site*, not you. You can't do anything wrong here. In fact, this is probably the one place today where you don't have to worry about making mistakes.

We want to hear exactly what you think, so please don't worry that you're going to hurt our feelings.⁶ We want to improve it, so we need to know honestly what you think.

As we go along, I'm going to ask you to think out loud, to tell me what's going through your mind. This will help us.

This whole first section is the script that I use when I conduct tests.⁵

I always have a copy in front of me, and I don't hesitate to read from it, but I find it's good to ad lib a little, even if it means making mistakes. When the users see that I'm comfortable making mistakes, it helps take the pressure off them.

⁵ A copy of the script is available on my Web site (www.sensible.com) so you can download it and edit it for your own use.

⁶ If you didn't work on the part that's being tested, you can also say, "Don't worry about hurting my feelings. I didn't create the pages you're going to look at."

If you have questions, just ask. I may not be able to answer them right away, since we're interested in how people do when they don't have someone sitting next to them, but I will try to answer any questions you still have when we're done.

We have a lot to do, and I'm going to try to keep us moving, but we'll try to make sure that it's fun, too.

You may have noticed the camera. With your permission, we're going to record the computer screen and what you have to say. The recording will be used only to help us figure out how to improve the site, and it won't be seen by anyone except the people working on the project. It also helps me, because I don't have to take as many notes. There are also some people watching the screen in another room.

If you would, I'm going to ask you to sign something for us. It simply says that we have your permission to record you, but that it will only be seen by the people working on the project. It also says that you won't talk to anybody about what we're showing you today, since it hasn't been made public yet.

Do you have any questions before we begin?

No. I don't think so.

It's important to mention this, because it will seem rude not to answer their questions as you go along. You have to make it clear before you start that (a) it's nothing personal, and (b) you'll try to answer them at the end if they still want to know.

At this point, most people will say something like, "I'm not going to end up on *America's Funniest Home Videos*, am I?"

Give them the release and non-disclosure agreement (if required) to sign. Both should be as short as possible and written in plain English.⁷

⁷ You'll find a sample recording consent form on my Web site.

BACKGROUND QUESTIONS

Before we look at the site, I'd like to ask you just a few quick questions. First, what's your occupation?

I'm a router.

I've never heard of that before. What does a router do, exactly?

Not much. I take orders as they come in, and send them to the right office.

Good. Now, roughly how many hours a week would you say you spend using the Internet, including email?

Oh, I don't know. Probably an hour a day at work, and maybe four hours a week at home. Mostly that's on the weekend. I'm too tired at night to bother. But I like playing games sometimes.

How do you spend that time? In a typical day, for instance, tell me what you do, at work and at home.

Well, at the office I spend most of my time checking email. I get *a lot* of email, and a lot of it's junk but I have to go through it anyway. And sometimes I have to research something at work.

I find it's good to start with a few questions to get a feel for who they are and how they use the Internet. It gives them a chance to loosen up a little and gives you a chance to show that you're going to be listening attentively to what they say—and that there are no wrong or right answers.

Don't hesitate to admit your ignorance about anything. Your role here is not to come across as an expert, but as a good listener.

Notice that she's not sure how much time she really spends on the Internet. Most people aren't. Don't worry. Accurate answers aren't important here. The main point here is just to get her talking and thinking about how she uses the Internet and to give you a chance to gauge what kind of user she is.

Do you have any favorite Web sites?

Yahoo, I guess. I like Yahoo, and I use it all the time. And something called Snakes.com, because I have a pet snake.

Really? What kind of snake?

A python. He's about four feet long, but he should get to be eight or nine when he's fully grown.

Wow. OK, now, finally, have you bought anything on the Internet? How do you feel about buying things on the Internet?

I've bought some things recently. I didn't do it for a long time, but only because I couldn't get things delivered. It was hard to get things delivered, because I'm not home during the day. But now one of my neighbors is home all the time, so I can.

And what have you bought?

Well, I ordered a raincoat from L.L. Bean, and it worked out *much* better than I thought it would. It was actually pretty easy.

OK, great. We're done with the questions, and we can start looking at things.

OK, I guess.

Don't be afraid to digress and find out a little more about the user, as long as you come back to the topic before long.

REACTIONS TO THE HOME PAGE

First, I'm just going to ask you to look at this page and tell me what you think it is, what strikes you about it, and what you think you would click on first.

For now, don't actually click on anything. Just tell me what you *would* click on.

And again, as much as possible, it will help us if you can try to think out loud so we know what you're thinking about.

The browser has been open, but minimized. At this point, I reach over and click to maximize it.



Well, I guess the first thing I notice is that I like the color. I like the shade of orange, and I like the little picture of the sun [at the top of the page, in the eLance logo].

Let's see. [Reads.] "The global services market." "Where the world comes to get your job done."



I don't know what that means. I have no idea.

"Animate your logo free." [Looking at the Cool Stuff section on the left.] "3D graphics marketplace." "eLance community." "eLance marketplace."



In an average test, it's just as likely that the next user will say that she hates this shade of orange and that the drawing is too simplistic. Don't get too excited by individual reactions to site aesthetics.

There's a lot going on here. But I have no idea what any of it is.

If you had to take a guess, what do you think it might be?

Well, it seems to have something to do with buying and selling...something.

[Looks around the page again.] Now that I look at the list down here [the Yahoo-style category list halfway down the page], I guess maybe it must be services. Legal, financial, creative...they all sound like services.



So I guess that's what it is. Buying and selling services. Maybe like some kind of online *Yellow Pages*.

OK. Now, if you were at home, what would you click on first?

This user is doing a good job of thinking out loud on her own. If she wasn't, this is where I'd start asking her, "What are you thinking?"

I guess I'd click on that 3D graphics thing.
I'm interested in 3D graphics.

Before you click on it, I have one more question. What about these pictures near the top of the page—the ones with the numbers? What did you make of them?



I noticed them, but I really didn't try to figure them out. I guess I thought they were telling me what the steps in the process would be.

Any reason why you didn't pay much attention to them?

No. I guess I just wasn't ready to start the process yet. I didn't know if I *wanted* to use it yet. I just wanted to look around first.

OK. Great.

I ask this question because the site's designers think most users are going to start by clicking on the pictures of the five steps, and that everyone will at least look at them.

TESTING A TASK

OK, now we're going to try something else.

Can you think of something you might want to post as a project if you were using this site?

Hmm. Let me think. I think I saw "Home Improvement" there somewhere. We're thinking of building a deck. Maybe I would post that.

So if you were going to post the deck as a project, what would you do first?

I guess I'd click on one of the categories down here. I think I saw home improvement. [Looks.] There it is, under "Family and Household."

So what would you do?

Well, I'd click....[Hesitates, looking at the two links under "Family and Household."]

Family & Household

Food & Cooking, Gardening,
Genealogy, Home Improvement,
Interior Design, Parenting, Pets, Real
Estate...

RFPs | Fixed-Price

Now I give her a task to perform so we can see whether she can use the site for its intended purpose.

Whenever possible, it's good to let the user have some say in choosing the task.

Well, now I'm not sure *what* to do. I can't click on Home Improvement, so it looks like I have to click on either "RFPs" or "Fixed-Price." But I don't know what the difference is.

Fixed price I sort of understand; they'll give me a quote, and then they have to stick to it. But I'm not sure what RFPs is.

As it turns out, she's mistaken. Fixed-price (in this case) means services available for a fixed hourly rate, while an RFP (or Request for Proposal) is actually the choice that will elicit quotes. This is the kind of misunderstanding that often surprises the people who built the site.

Well, which one do you think you'd click on?

Fixed price, I guess.

Why don't you go ahead and do it?

From here on, I just watch while she tries to post a project, letting her continue until either (a) she finishes the task, (b) she gets really frustrated, or (c) we're not learning anything new by watching her try to muddle through.

I'd give her three or four more tasks to do, which should take not more than 45 minutes altogether.

Review the results right away

After each round of tests, you should make time as soon as possible for the development team to review everyone's observations and decide what to do next. I strongly recommend that you do three or four tests in a morning and then debrief over lunch.

You're doing two things at this meeting:

- **Triage**—reviewing the problems people saw and deciding which ones need to be fixed.
- **Problem solving**—figuring out how to fix them.

It might seem that this would be a difficult process. After all, these are the same team members who've been arguing about the right way to do things all along. So what's going to make this session any different?

Just this:

The important things that you learn from usability testing usually *just make sense*. They tend to be obvious to anyone who watches the sessions.

Also, the experience of seeing your handiwork through someone else's eyes will often suggest entirely new solutions for problems, or let you see an old idea in a new light.

And remember, this is a cyclic process, so the team doesn't have to agree on the perfect solution. You just need to figure out what to try next.

Typical problems

Here are the types of problems you're going to see most often when you test:

- **Users are unclear on the concept.** They just don't get it. They look at the site or a page and they either don't know what to make of it, or they think they do but they're wrong.
- **The words they're looking for aren't there.** This usually means that either

(a) the categories you've used to organize your content aren't the ones they would use, or (b) the categories are what they expect, but you're just not using the names they expect.

- **There's too much going on.** Sometimes what they're looking for is right there on the page, but they're just not seeing it. In this case, you need to either (a) reduce the overall noise on the page, or (b) turn up the volume on the things they need to see so they "pop" out of the visual hierarchy more.

Some triage guidelines

Here's the best advice I can give you about deciding what to fix—and what not to.

- **Ignore "kayak" problems.** In any test, you're likely to see several cases where users will go astray momentarily but manage to get back on track almost immediately without any help. It's kind of like rolling over in a kayak; as long as the kayak rights itself quickly enough, it's all part of the so-called fun. In basketball terms, no harm, no foul.

As long as (a) everyone who has the problem notices that they're no longer headed in the right direction quickly, and (b) they manage to recover without help, and (c) it doesn't seem to faze them, you can ignore the problem. In general, if the user's second guess about where to find things is always right, that's good enough.

Of course, if there's an easy and obvious fix that won't break anything else, then by all means fix it. But kayak problems usually don't come as a surprise to the development team. They're usually there because of some ambiguity for which there is no simple resolution. For example, there are usually at least one or two oddball items that don't fit perfectly into any of the top-level categories of a site. So half the users may look for movie listings in Lifestyles first, and the other half will look for them in Arts first. Whatever you do, half of them are going to be wrong on their first guess, but everyone will get it on their second guess, which is fine.⁸

⁸ You may be thinking "Well, why not just put it in both categories?" In general, I think it's best for things to "live" in only one place in a hierarchy, with a prominent "see also" crosslink in any other places where people are likely to look for them.

- › **Resist the impulse to add things.** When it's obvious in testing that users aren't getting something, most people's first reaction is to add something, like an explanation or some instructions.

Very often, the right solution is to take something (or things) away that are obscuring the meaning, rather than adding yet another distraction.

- › **Take “new feature” requests with a grain of salt.** People will often say, “I'd like it better if it could do x .” It always pays to be suspicious of these requests for new features. If you probe deeper, it often turns out that they already have a perfectly fine source for x and wouldn't be likely to switch; they're just telling you what they like.
- › **Grab the low-hanging fruit.** The main thing you're looking for in each round of testing is the big, cheap wins. These fall into two categories:
 - › **Head slappers.** These are the surprises that show up during testing where the problem and the solution were obvious to everyone the moment they saw the first user try to muddle through. These are like found money, and you should fix them right away.
 - › **Cheap hits.** Also try to implement any changes that (a) require almost no effort, or (b) require a *little* effort but are highly visible.

And finally, there's one last piece of advice about “making changes” that deserves its own section:

Don't throw the baby out with the dishes

Like any good design, successful Web pages are usually a delicate balance, and it's important to keep in mind that even a minor change can have a major impact. Sometimes the real challenge isn't fixing the problems you find—it's fixing them *without* breaking the parts that already work.

Whenever you're making a change, think carefully about what else is going to be affected. In particular, when you're making something more prominent than it was, consider what else might end up being de-emphasized as a result.

One morning a month: that's all we ask

Ideally, I think every Web development team should spend one morning a month doing usability testing.

In a morning, you can test three or four users, then debrief over lunch. That's it.

When you leave lunch, the team will have decided what you're going to fix, and you'll be done with testing for the month. No reports, no endless meetings.

Doing it all in a morning also greatly increases the chances that most team members will make time to come and watch at least some of the sessions, which is highly desirable.

If you're going to try doing some testing yourself—and I hope you will—you'll find some more advice about how to do it in a chapter called “Usability testing: The Movie” that was in the first edition of this book.⁹ My next book is going to be all about do-it-yourself usability testing, but I do *not* want you to wait for it before you start testing. Start now.

⁹ You can download the chapter for free at <http://www.sensible.com/secondedition>.