


THE PSYCHOPATHOLOGY OF EVERYDAY THINGS



"Kenneth Olsen, the engineer who founded and still runs Digital Equipment Corp., confessed at the annual meeting that he can't figure out how to heat a cup of coffee in the company's microwave oven."¹

You Would Need an Engineering Degree to Figure This Out

"You would need an engineering degree from MIT to work this," someone once told me, shaking his head in puzzlement over his brand new digital watch. Well, I have an engineering degree from MIT. (Kenneth Olsen has two of them, and he can't figure out a microwave oven.) Give me a few hours and I can figure out the watch. But why should it take hours? I have talked with many people who can't use all the features of their washing machines or cameras, who can't figure out how to work a sewing machine or a video cassette recorder, who habitually turn on the wrong stove burner.

Why do we put up with the frustrations of everyday objects, with objects that we can't figure out how to use, with those neat plastic-wrapped packages that seem impossible to open, with doors that trap people, with washing machines and dryers that have become too con-



1.1 Carelman's Coffeepot for Masochists. The French artist Jacques Carelman in his series of books *Catalogue d'objets introuvables (Catalog of unfindable objects)* provides delightful examples of everyday things that are deliberately unworkable, outrageous, or otherwise ill-formed. Jacques Carelman: "Coffeepot for Masochists." Copyright © 1969-76-80 by Jacques Carelman and A. D. A. G. P. Paris. From Jacques Carelman, *Catalog of Unfindable Objects*, Balland, éditeur, Paris-France. Used by permission of the artist.

fusing to use, with audio-stereo-television-video-cassette-recorders that claim in their advertisements to do everything, but that make it almost impossible to do anything?

The human mind is exquisitely tailored to make sense of the world. Give it the slightest clue and off it goes, providing explanation, rationalization, understanding. Consider the objects—books, radios, kitchen appliances, office machines, and light switches—that make up our everyday lives. Well-designed objects are easy to interpret and understand. They contain visible clues to their operation. Poorly designed objects can be difficult and frustrating to use. They provide no clues—or sometimes false clues. They trap the user and thwart the normal process of interpretation and understanding. Alas, poor design predominates. The result is a world filled with frustration, with objects that cannot be understood, with devices that lead to error. This book is an attempt to change things.

The Frustrations of Everyday Life

If I were placed in the cockpit of a modern jet airliner, my inability to perform gracefully and smoothly would neither surprise nor bother me. But I shouldn't have trouble with doors and switches, water faucets and stoves. "Doors?" I can hear the reader saying, "you have trouble

opening doors?" Yes. I push doors that are meant to be pulled, pull doors that should be pushed, and walk into doors that should be slid. Moreover, I see others having the same troubles—unnecessary troubles. There are psychological principles that can be followed to make these things understandable and usable.

Consider the door. There is not much you can do to a door: you can open it or shut it. Suppose you are in an office building, walking down a corridor. You come to a door. In which direction does it open? Should you pull or push, on the left or the right? Maybe the door slides. If so, in which direction? I have seen doors that slide up into the ceiling. A door poses only two essential questions: In which direction does it move? On which side should one work it? The answers should be given by the design, without any need for words or symbols, certainly without any need for trial and error.

A friend told me of the time he got trapped in the doorway of a post office in a European city. The entrance was an imposing row of perhaps six glass swinging doors, followed immediately by a second, identical row. That's a standard design: it helps reduce the airflow and thus maintain the indoor temperature of the building.

My friend pushed on the side of one of the leftmost pair of outer doors. It swung inward, and he entered the building. Then, before he could get to the next row of doors, he was distracted and turned around for an instant. He didn't realize it at the time, but he had moved slightly to the right. So when he came to the next door and pushed it, nothing happened. "Hmm," he thought, "must be locked." So he pushed the side of the adjacent door. Nothing. Puzzled, my friend decided to go outside again. He turned around and pushed against the side of a door. Nothing. He pushed the adjacent door. Nothing. The door he had just entered no longer worked. He turned around once more and tried the inside doors again. Nothing. Concern, then mild panic. He was trapped! Just then, a group of people on the other side of the entranceway (to my friend's right) passed easily through both sets of doors. My friend hurried over to follow their path.

How could such a thing happen? A swinging door has two sides. One contains the supporting pillar and the hinge, the other is unsupported. To open the door, you must push on the unsupported edge. If you push on the hinge side, nothing happens. In this case, the designer aimed for beauty, not utility. No distracting lines, no visible pillars, no visible hinges. So how can the ordinary user know which side to push



1.2 A Row of Swinging Glass Doors in a Boston Hotel. A similar problem to the doors from that European post office. On which side of the door should you push? When I asked people who had just used the doors, most couldn't say. Yet only a few of the people I watched had trouble with the doors. The designers had incorporated a subtle clue into the design. Note that the horizontal bars are not centered: they are a bit closer together on the sides you should push on. The design almost works—but not entirely, for not everyone used the doors right on the first try.

on? While distracted, my friend had moved toward the (invisible) supporting pillar, so he was pushing the doors on the hinged side. No wonder nothing happened. Pretty doors. Elegant. Probably won a design prize.

The door story illustrates one of the most important principles of design: *visibility*. The correct parts must be visible, and they must convey the correct message. With doors that push, the designer must provide signals that naturally indicate where to push. These need not destroy the aesthetics. Put a vertical plate on the side to be pushed, nothing on the other. Or make the supporting pillars visible. The vertical plate and supporting pillars are *natural* signals, *naturally* interpreted, without any need to be conscious of them. I call the use of natural signals *natural design* and elaborate on the approach throughout this book.

Visibility problems come in many forms. My friend, trapped between the glass doors, suffered from a lack of clues that would indicate what part of a door should be operated. Other problems concern the *mappings* between what you want to do and what appears to be possible, another topic that will be expanded upon throughout the book. Consider one type of slide projector. This projector has a single button to control whether the slide tray moves forward or backward. One button to do two things? What is the mapping? How can you figure out how to control the slides? You can't. Nothing is visible to give the slightest hint. Here is what happened to me in one of the many unfamiliar places I've lectured in during my travels as a professor:

The Leitz slide projector illustrated in figure 1.3 has shown up several times in my travels. The first time, it led to a rather dramatic incident. A conscientious student was in charge of showing my slides. I started my talk and showed the first slide. When I finished with the first slide and asked for the next, the student carefully pushed the control button and watched in dismay as the tray backed up, slid out of the projector and plopped off the table onto the floor, spilling its entire contents. We had to delay the lecture fifteen minutes while I struggled to reorganize the slides. It wasn't the student's fault. It was the fault of the elegant projector. With only one button to control the slide advance, how could one switch from forward to reverse? Neither of us could figure out how to make the control work.

All during the lecture the slides would sometimes go forward, sometimes backward. Afterward, we found the local technician, who explained it to us. A brief push of the button and the slide would go

Taste (7) für Diawechsel am Gerät

Diawechsel vorwärts = kurz drücken,
Diawechsel rückwärtz = länger drücken.

Button (7) for changing the slides

Slide change forward = short press,
Slide change backward = longer press.

1.3 Leitz Pravidit Slide Projector. I finally tracked down the instruction manual for that projector. A photograph of the projector has its parts numbered. The button for changing slides is number 7. The button itself has no labels. Who could discover this operation without the aid of the manual? Here is the entire text related to the button, in the original German and in my English translation:


forward, a long push and it would reverse. (Pity the conscientious student who kept pushing it hard—and long—to make sure that the switch was making contact.) What an elegant design. Why, it managed to do two functions with only one button! But how was a first-time user of the projector to know this?

As another example, consider the beautiful Amphithéâtre Louis-Laird in the Paris Sorbonne, which is filled with magnificent paintings of great figures in French intellectual history. (The mural on the ceiling shows lots of naked women floating about a man who is valiantly trying to read a book. The painting is right side up only for the lecturer—it is upside down for all the people in the audience.) The room is a delight to lecture in, at least until you ask for the projection screen to be lowered. "Ah," says the professor in charge, who gestures to the technician, who runs out of the room, up a short flight of stairs, and out of sight behind a solid wall. The screen comes down and stops. "No, no," shouts the professor, "a little bit more." The screen comes down again, this time too much. "No, no, no!" the professor jumps up and down and gestures wildly. It's a lovely room, with lovely paintings. But why can't the person who is trying to lower or raise the screen see what he is doing?

New telephone systems have proven to be another excellent example of incomprehensible design. No matter where I travel, I can count upon finding a particularly bad example.

When I visited Basic Books, the publishers of this book, I noticed a new telephone system. I asked people how they liked it. The question unleashed a torrent of abuse. "It doesn't have a hold function," one woman complained bitterly—the same complaint people at my university made about their rather different system. In older days, business phones always had a button labeled "hold." You could push the button and hang up the phone without losing the call on your line. Then you could talk to a colleague, or pick up another telephone call, or even pick up the call at another phone with the same telephone number. A light on the hold button indicated when the function was in use. It was an invaluable tool for business. Why didn't the new phones at Basic Books or in my university have a hold function, if it is so essential? Well, they did, even the very instrument the woman was complaining about. But there was no easy way to discover the fact, nor to learn how to use it.

I was visiting the University of Michigan and I asked about the new



1.4 Plate Mounted Over the Dial of the Telephones at the University of Michigan. These inadequate instructions are all that most users see. (The button labeled "TAP" at the lower right is used to transfer or pick up calls—it is pressed whenever the instruction plate says "TAP." The light on the lower left comes on whenever the telephone rings.)

system there. "Yech!" was the response, "and it doesn't even have a hold function!" Here we go again. What is going on? The answer is simple: first, look at the instructions for hold. At the University of Michigan the phone company provided a little plate that fits over the keypad and reminds users of the functions and how to use them. I carefully unhooked one of the plates from the telephone and made a photocopy (figure 1.4). Can you understand how to use it? I can't. There is a "call hold" operation, but it doesn't make sense to me, not for the application that I just described.

The telephone hold situation illustrates a number of different problems. One of them is simply poor instructions, especially a failure to relate the new functions to the similarly named functions that people already know about. Second, and more serious, is the lack of visibility of the operation of the system. The new telephones, for all their added sophistication, lack both the hold button and the flashing light of the old ones. The hold is signified by an arbitrary action: dialing an arbitrary sequence of digits (*8, or *99, or what have you: it varies from one phone system to another). Third, there is no visible outcome of the operation.

Devices in the home have developed some related problems: functions and more functions, controls and more controls. I do not think that simple home appliances—stoves, washing machines, audio and television sets—should look like Hollywood's idea of a spaceship control room. They already do, much to the consternation of the consumer who, often as not, has lost (or cannot understand) the instruction

manual, so—faced with the bewildering array of controls and displays—simply memorizes one or two fixed settings to approximate what is desired. The whole purpose of the design is lost.

In England I visited a home with a fancy new Italian washer-drier combination, with super-duper multi-symbol controls, all to do everything you ever wanted to do with the washing and drying of clothes. The husband (an engineering psychologist) said he refused to go near it. The wife (a physician) said she had simply memorized one setting and tried to ignore the rest.

Someone went to a lot of trouble to create that design. I read the instruction manual. That machine took into account everything about today's wide variety of synthetic and natural fabrics. The designers worked hard; they really cared. But obviously they had never thought of trying it out, or of watching anyone use it.

If the design was so bad, if the controls were so unusable, why did the couple purchase it? If people keep buying poorly designed products, manufacturers and designers will think they are doing the right thing and continue as usual.

The user needs help. Just the right things have to be visible: to indicate what parts operate and how, to indicate how the user is to interact with the device. Visibility indicates the mapping between intended actions and actual operations. Visibility indicates crucial distinctions—so that you can tell salt and pepper shakers apart, for example. And visibility of the effects of the operations tells you if the lights have turned on properly, if the projection screen has lowered to the correct height, or if the refrigerator temperature is adjusted correctly. It is lack of visibility that makes so many computer-controlled devices so difficult to operate. And it is an excess of visibility that makes the gadget-ridden, feature-laden modern audio set or video cassette recorder (VCR) so intimidating.

The Psychology of Everyday Things

This book is about the psychology of everyday things. POET emphasizes the understanding of everyday things, things with knobs and dials, controls and switches, lights and meters. The instances we have just examined demonstrate several principles, including the importance

of visibility, appropriate clues, and feedback of one's actions. These principles constitute a form of psychology—the psychology of how people interact with things. A British designer once noted that the kinds of materials used in the construction of passenger shelters affected the way vandals responded. He suggested that there might be a psychology of materials.

AFFORDANCES

"In one case, the reinforced glass used to panel shelters (for railroad passengers) erected by British Rail was smashed by vandals as fast as it was renewed. When the reinforced glass was replaced by plywood boarding, however, little further damage occurred, although no extra force would have been required to produce it. Thus British Rail managed to elevate the desire for defacement to those who could write, albeit in somewhat limited terms. Nobody has, as yet, considered whether there is a kind of psychology of materials. But on the evidence, there could well be!"²

There already exists the start of a psychology of materials and of things, the study of affordances of objects. When used in this sense, the term *affordance* refers to the perceived and actual properties of the thing, primarily those fundamental properties that determine just how the thing could possibly be used (see figures 1.5 and 1.6). A chair affords ("is for") support and, therefore, affords sitting. A chair can also be carried. Glass is for seeing through, and for breaking. Wood is normally used for solidity, opacity, support, or carving. Flat, porous, smooth surfaces are for writing on. So wood is also for writing on. Hence the problem for British Rail: when the shelters had glass, vandals smashed it; when they had plywood, vandals wrote on and carved it. The planners were trapped by the affordances of their materials.³


Affordances provide strong clues to the operations of things. Plates are for pushing. Knobs are for turning. Slots are for inserting things into. Balls are for throwing or bouncing. When affordances are taken advantage of, the user knows what to do just by looking: no picture, label, or instruction is required. Complex things may require explanation, but simple things should not. When simple things need pictures, labels, or instructions, the design has failed.

A psychology of causality is also at work as we use everyday things.



1.5 **Affordances of Doors.** Door hardware can signal whether to push or pull without signs. The flat horizontal bar of *A* (above left) affords no operations except pushing: it is excellent hardware for a door that must be pushed to be opened. The door in *B* (above right) has a different kind of bar on each side, one relatively small and vertical to signify a pull, the other relatively large and horizontal to signify a push. Both bars support the affordance of grasping: size and position specify whether the grasp is used to push or pull—though ambiguously.

1.6 **When Affordances Fail.** I had to tie a string around my cabinet door to afford pulling.



Something that happens right after an action appears to be caused by that action. Touch a computer terminal just when it fails, and you are apt to believe that you caused the failure, even though the failure and your action were related only by coincidence. Such false causality is the basis for much superstition. Many of the peculiar behaviors of people using computer systems or complex household appliances result from such false coincidences. When an action has no apparent result, you may conclude that the action was ineffective. So you repeat it. In earlier days, when computer word processors did not always show the results of their operations, people would sometimes attempt to change their manuscript, but the lack of visible effect from each action would make them think that their commands had not been executed, so they would repeat the commands, sometimes over and over, to their later astonishment and regret. It is a poor design that allows either kind of false causality to occur.

TWENTY THOUSAND EVERYDAY THINGS

There are an amazing number of everyday things, perhaps twenty thousand of them. Are there really that many? Start by looking about you. There are light fixtures, bulbs, and sockets; wall plates and screws; clocks, watches, and watchbands. There are writing devices (I count twelve in front of me, each different in function, color, or style). There are clothes, with different functions, openings, and flaps. Notice the variety of materials and pieces. Notice the variety of fasteners—buttons, zippers, snaps, laces. Look at all the furniture and food utensils: all those details, each serving some function for manufacturability, usage, or appearance. Consider the work area: paper clips, scissors, pads of paper, magazines, books, bookmarks. In the room I'm working in, I counted more than a hundred specialized objects before I tired. Each is simple, but each requires its own method of operation, each has to be learned, each does its own specialized task, and each has to be designed separately. Furthermore, many of the objects are made of many parts. A desk stapler has sixteen parts, a household iron fifteen, the simple bathtub-shower combination twenty-three. You can't believe these simple objects have so many parts? Here are the eleven basic parts to a sink: drain, flange (around the drain), pop-up stopper, basin, soap dish, overflow vent, spout, lift rod, fittings, hot-water handle, and cold-water handle. We can count even more if we start taking the faucets, fittings, and lift rods apart.


The book *What's What: A Visual Glossary of the Physical World* has more than fifteen hundred drawings and pictures and illustrates twenty-three thousand items or parts of items.⁴ Irving Biederman, a psychologist who studies visual perception, estimates that there are probably "30,000 readily discriminable objects for the adult."⁵ Whatever the exact number, it is clear that the difficulties of everyday life are amplified by the sheer profusion of items. Suppose that each everyday thing takes only one minute to learn; learning 20,000 of them occupies 20,000 minutes—333 hours or about 8 forty-hour work weeks. Furthermore, we often encounter new objects unexpectedly, when we are really concerned with something else. We are confused and distracted, and what ought to be a simple, effortless, everyday thing interferes with the important task of the moment.

How do people cope? Part of the answer lies in the way the mind works—in the psychology of human thought and cognition. Part lies in the information available from the appearance of the objects—the psychology of everyday things. And part comes from the ability of the designer to make the operation clear, to project a good image of the operation, and to take advantage of other things people might be expected to know. Here is where the designer's knowledge of the psychology of people coupled with knowledge of how things work becomes crucial.

CONCEPTUAL MODELS

Consider the rather strange bicycle illustrated in figure 1.7. You know it won't work because you form a *conceptual model* of the device and mentally simulate its operation. You can do the simulation because the parts are visible and the implications clear.

Other clues to how things work come from their visible structure—in particular from *affordances*, *constraints*, and *mappings*. Consider a pair of scissors: even if you have never seen or used them before, you can see that the number of possible actions is limited. The holes are clearly there to put something into, and the only logical things that will fit are fingers. The holes are affordances: they allow the the fingers to be inserted. The sizes of the holes provide *constraints* to limit the possible fingers: the big hole suggests several fingers, the small hole only one. The mapping between holes and fingers—the set of possible operations—is suggested and constrained by the holes. Moreover, the operation is not sensitive to finger placement: if you use the wrong fingers,



1.7 Carelman's Tandem "Convergent Bicycle (Model for Fiancés)." Jacques Carelman: "Convergent Bicycle" Copyright © 1969-76-80 by Jacques Carelman and A. D. A. G. P. Paris. From Jacques Carelman, *Catalog of Unfindable Objects*, Balland, éditeur, Paris-France. Used by permission of the artist.

the scissors still work. You can figure out the scissors because their operating parts are visible and the implications clear. The conceptual model is made obvious, and there is effective use of affordances and constraints.

As a counterexample, consider the digital watch, one with two to four push buttons on the front or side. What are those push buttons for? How would you set the time? There is no way to tell—no evident relationship between the operating controls and the functions, no constraints, no apparent mappings. With the scissors, moving the handle makes the blades move. The watch and the Leitz slide projector provide no visible relationship between the buttons and the possible actions, no discernible relationship between the actions and the end result.

Principles of Design for Understandability and Usability

We have now encountered the fundamental principles of designing for people: (1) provide a good conceptual model and (2) make things visible.

PROVIDE A GOOD CONCEPTUAL MODEL

A good conceptual model allows us to predict the effects of our actions. Without a good model we operate by rote, blindly; we do operations as we were told to do them; we can't fully appreciate why, what effects to expect, or what to do if things go wrong. As long as things work properly, we can manage. When things go wrong, however, or when

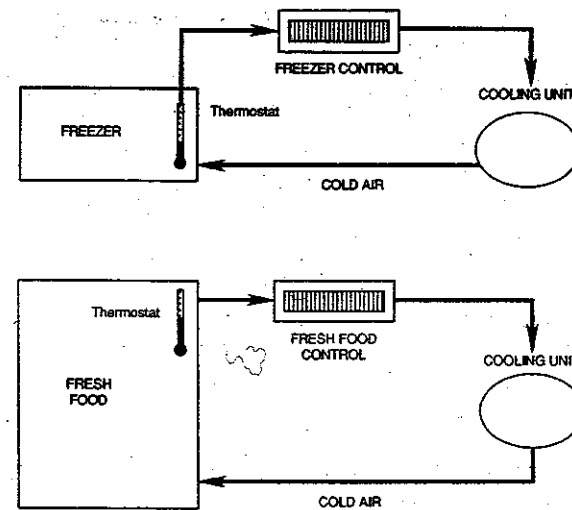
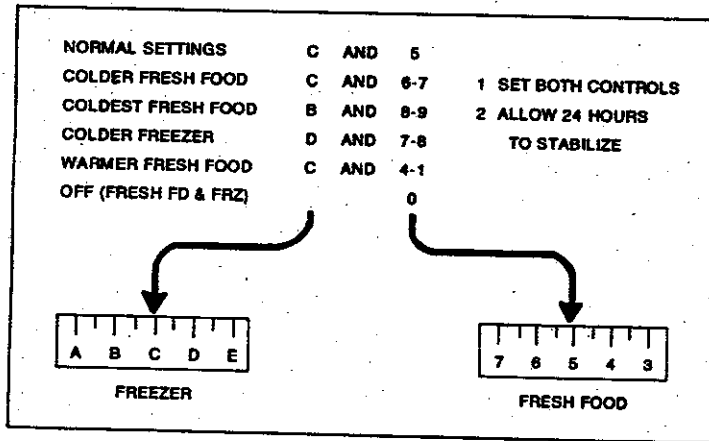
we come upon a novel situation, then we need a deeper understanding, a good model.

For everyday things, conceptual models need not be very complex. After all, scissors, pens, and light switches are pretty simple devices. There is no need to understand the underlying physics or chemistry of each device we own, simply the relationship between the controls and the outcomes. When the model presented to us is inadequate or wrong (or, worse, nonexistent), we can have difficulties. Let me tell you about my refrigerator.


My house has an ordinary, two-compartment refrigerator—nothing very fancy about it. The problem is that I can't set the temperature properly. There are only two things to do: adjust the temperature of the freezer compartment and adjust the temperature of the fresh food compartment. And there are two controls, one labeled "freezer," the other "fresh food." What's the problem?

You try it. Figure 1.8 shows the instruction plate from inside the refrigerator. Now, suppose the freezer is too cold, the fresh food section just right. You want to make the freezer warmer, keeping the fresh food constant. Go on, read the instructions, figure them out.

1.8 My Refrigerator. Two compartments—fresh food and freezer—and two controls (in the fresh food unit). The illustration shows the controls and instructions. Your task: Suppose the freezer is too cold, the fresh food section just right. How would you adjust the controls so as to make the freezer warmer and keep the fresh food the same? (From Norman, 1986.)



1.9 Two Conceptual Models for My Refrigerator. The model A (above) is provided by the system image of the refrigerator as gleaned from the controls and instructions; B (below) is the correct conceptual model. The problem is that it is impossible to tell in which compartment the thermostat is located and whether the two controls are in the freezer and fresh food compartment, or vice versa.




Oh, perhaps I'd better warn you. The two controls are not independent. The freezer control affects the fresh food temperature, and the fresh food control affects the freezer. And don't forget to wait twenty-four hours to check on whether you made the right adjustment, if you can remember what you did.

Control of the refrigerator is made difficult because the manufacturer provides a false conceptual model. There are two compartments and two controls. The setup clearly and unambiguously provides a simple model for the user: each control is responsible for the temperature of the compartment that carries its name. Wrong. In fact, there is only one thermostat and only one cooling mechanism. One control adjusts the thermostat setting, the other the relative proportion of cold air sent to each of the two compartments of the refrigerator. This is why the two controls interact. With the conceptual model provided by the manufacturer, adjusting the temperatures is almost impossible and always frustrating. Given the correct model, life would be much easier (figure 1.9).

Why did the manufacturer present the wrong conceptual model?

1.10 Conceptual Models. The *design model* is the designer's conceptual model. The *user's model* is the mental model developed through interaction with the system. The *system image* results from the physical structure that has been built (including documentation, instructions, and labels). The designer expects the user's model to be identical to the design model. But the designer doesn't talk directly with the user—all communication takes place through the system image. If the system image does not make the design model clear and consistent, then the user will end up with the wrong mental model. (From Norman, 1986.)



Perhaps the designers thought the correct model was too complex, that the model they were giving was easier to understand. But with the wrong conceptual model, it is impossible to set the controls. And even though I am convinced I now know the correct model, I still cannot accurately adjust the temperatures because the refrigerator design makes it impossible for me to discover which control is for the thermostat, which control is for the relative proportion of cold air, and in which compartment the thermostat is located. The lack of immediate feedback for the actions does not help: with a delay of twenty-four hours, who can remember what was tried?

The topic of conceptual models will reappear in the book. They are part of an important concept in design: *mental models*, the models people have of themselves, others, the environment, and the things with which they interact. People form mental models through experience, training, and instruction. The mental model of a device is formed largely by interpreting its perceived actions and its visible structure. I call the visible part of the device the *system image* (figure 1.10). When the system image is incoherent or inappropriate, as in the case of the refrigerator, then the user cannot easily use the device. If it is incomplete or contradictory, there will be trouble.

MAKE THINGS VISIBLE

The problems caused by inadequate attention to visibility are all neatly demonstrated with one simple appliance: the modern telephone.

I stand at the blackboard in my office, talking with a student, when my telephone rings. Once, twice it rings. I pause, trying to complete my sentence before answering. The ringing stops. "I'm sorry," says the student. "Not your fault," I say. "But it's no problem, the call now transfers to my secretary's phone. She'll answer it." As we listen we hear her phone start to ring. Once, twice. I look at my watch. Six o'clock: it's late, the office staff has left for the day. I rush out of my office to my secretary's phone, but as I get there, it stops ringing. "Ah," I think, "it's being transferred to another phone." Sure enough, the phone in the adjacent office now starts ringing. I rush to that office, but it is locked. Back to my office to get the key, out to the locked door, fumble with the lock, into the office, and to the now quiet phone. I hear a telephone down the hall start to ring. Could that still be my call,

making its way mysteriously, with a predetermined lurching path, through the phones of the building? Or is it just another telephone call coincidentally arriving at this time?

In fact, I could have retrieved the call from my office, had I acted quickly enough. The manual states: "Within your pre-programmed pick-up group, dial 14 to connect to incoming call. Otherwise, to answer any ringing extension, dial ringing extension number, listen for busy tone. Dial 8 to connect to incoming call." Huh? What do those instructions mean? What is a "pre-programmed pick-up group," and why do I even want to know? What is the extension number of the ringing phone? Can I remember all those instructions when I need them? No.

Telephone chase is the new game in the modern office, as the automatic features of telephones go awry—features designed without proper thought, and certainly without testing them with their intended users. There are several other games, too. One game is announced by the plea, "How do I answer this call?" The question is properly whined in front of a ringing, flashing telephone, receiver in hand. Then there is the paradoxical game entitled "This telephone doesn't have a hold function." The accusation is directed at a telephone that actually *does* have a hold function. And, finally, there is "What do you mean I called you, you called me!"

Many of the modern telephone systems have a new feature that automatically keeps trying to dial a number for you. This feature resides under names such as automatic redialing or automatic callback. I am supposed to use this feature whenever I call someone who doesn't answer or whose line is busy. When the person next hangs up the phone, my phone will dial it again. Several automatic callbacks can be active at a time. Here's how it works. I place a phone call. There's no answer, so I activate the automatic callback feature. Several hours later my telephone rings. I pick it up and say "Hello," only to hear a ringing sound and then someone else saying "Hello."

"Hello," I answer, "who is this?"

"Who is this?" I hear in reply, "you called me."


"No," I say, "you called me, my phone just rang."

Slowly I realize that perhaps this is my delayed call. Now, let me see, who was I trying to call several hours ago? Did I have several callbacks in place? Why was I making the call?

The modern telephone did not happen by accident: it was carefully designed. Someone—more likely a team of people—invented a list of features thought desirable, invented what seemed to them to be plausible ways of controlling the features, and then put it all together. My university, focusing on cost and perhaps dazzled by the features, bought the system, spending millions of dollars on a telephone installation that has proved vastly unpopular and even unworkable. Why did the university buy the system? The purchase took several years of committee work and studies and presentations by competing telephone companies, and piles of documentation and specification. I myself took part, looking at the interaction between the telephone system and the computer networks, ensuring that the two would be compatible and reasonable in price. To my knowledge, nobody ever thought of trying out the telephones in advance. Nobody suggested installing them in a sample office to see whether users' needs would be met or whether users could understand how to operate the phone. The result: disaster. The main culprit—lack of visibility—was coupled with a secondary culprit—a poor conceptual model. Any money saved on the installation and purchase is quickly disappearing in training costs, missed calls, and frustration. Yet from what I have seen, the competing phone systems would not have been any better.

I recently spent six months at the Applied Psychology Unit in Cambridge, England. Just before I arrived the British Telecom Company had installed a new telephone system. It had lots and lots of features. The telephone instrument itself was unremarkable (figure 1.11). It was the standard twelve-button, push-button phone, except that it had an extra key labeled "R" off on the side. (I never did find out what that key did.)

The telephone system was a standing joke. Nobody could use all the features. One person even started a small research project to record people's confusions. Another person wrote a small "expert systems" computer program, one of the new toys of the field of artificial intelligence; the program can reason through complex situations. If you wanted to use the phone system, perhaps to make a conference call among three people, you asked the expert system and it would explain how to do it. So, you're on the line with someone and you need to add a third person to the call. First turn on your computer. Then load the expert system. After three or four minutes (needed for loading the program), type in what you want to accomplish. Eventually the computer will tell you what to do—if you can remember why you want to




1.11 **British Telecom Telephone.** This was in my office at the Applied Psychology Unit in Cambridge, England. It certainly looks simple, doesn't it?

1.12 **Two Ways to Use Hold on Modern Telephones.** Illustration A (below left) is the instruction manual page for British Telecom. The procedure seems especially complicated, with three 3-digit codes to be learned: 681, 682, and 683. Illustration B (below right) shows the equivalent instructions for the Ericsson Single Line Analog Telephone installed at the University of California, San Diego. I find the second set of instructions easier to understand, but one must still dial an arbitrary digit: 8 in this case.

HOLD

This feature allows you to hold an existing call, then to replace the handset or to make another call. The held call may be retrieved from the holding extension or from any other extension within the system.

TO HOLD THE CALL



You may use your extension normally.

TO RETRIEVE THE CALL AT YOUR PHONE



TO RETRIEVE THE CALL AT SOMEONE ELSE'S PHONE



CALL HOLD/CALL PARK

With party on line

- Press R key
- Listen for recall dial tone (three beeps and dial tone)
- Hang up handset

TO RETRIEVE FROM SAME PHONE

- Lift handset you are connected to the call

TO RETRIEVE FROM ANOTHER PHONE

- Lift handset
- Dial extension where call was parked; listen for busy tone
- Dial 8; you are connected to the call

NOTE: Call will remain parked for 3 minutes before re-ringing

do it, and if the person on the other end of the line is still around. But, as it happens, using the expert system is a lot easier than reading and understanding the manual provided with the telephone (figure 1.12).

Why is that telephone system so hard to understand? Nothing in it is conceptually difficult. Each of the operations is actually quite simple. A few digits to dial, that's all. The telephone doesn't even look complicated. There are only fifteen controls: the usual twelve buttons—ten labeled 0 through 9, #, and *—plus the handset itself, the handset button, and the mysterious "R" button. All except the "R" are the everyday parts of a normal modern telephone. Why was the system so difficult?

A designer who works for a telephone company told me the following story:

"I was involved in designing the faceplate of some of those new multifunction phones, some of which have buttons labeled "R." The "R" button is kind of a vestigial feature. It is very hard to remove features of a newly designed product that had existed in an earlier version. It's kind of like physical evolution. If a feature is in the genome, and if that feature is not associated with any negativity (i.e., no customers gripe about it), then the feature hangs on for generations.

"It is interesting that things like the "R" button are largely determined through examples. Somebody asks, 'What is the "R" button used for?' and the answer is to give an example: 'You can push "R" to access loudspeaker paging.' If nobody can think of an example, the feature is dropped. Designers are pretty bright people, however. They can come up with a plausible-sounding example for almost anything. Hence, you get features, many many features, and these features hang on for a long time. The end result is complex interfaces for essentially simple things."

As I pondered this problem, I decided it would make sense to compare the phone system with something that was of equal or greater complexity but easier to use. So let us temporarily leave the difficult telephone system and take a look at my automobile. I bought a car in Europe. When I picked up the new car at the factory, a man from the company sat in the car with me and went over each control, explaining its function. When he had gone through the controls once, I said fine, thanked him, and drove away. That was all the instruction it took. There are 112 controls inside the car. This isn't quite as bad as it

sounds. Twenty-five of them are on the radio. Another 7 are the temperature control system, and 11 work the windows and sunroof. The trip computer has 14 buttons, each matched with a specific function. So four devices—the radio, temperature controls, windows, and trip computer—have together 57 controls, or just over 50 percent of the ones available.

Why is the automobile, with all its varied functions and numerous controls, so much easier to learn and to use than the telephone system, with its much smaller set of functions and controls? What is good about the design of the car? Things are visible. There are good mappings, natural relationships, between the controls and the things controlled. Single controls often have single functions. There is good feedback. The system is understandable. In general, the relationships among the user's intentions, the required actions, and the results are sensible, nonarbitrary, and meaningful.

What is bad about the design of the telephone? There is no visible structure. Mappings are arbitrary: there is no rhyme or reason to the relationship between the actions the user must perform and the results to be accomplished. The controls have multiple functions. There isn't good feedback, so the user is never sure whether the desired result has been obtained. The system, in general, is not understandable; its capabilities aren't apparent. In general, the relationships among the user's intentions, the required actions, and the results are completely arbitrary.

Whenever the number of possible actions exceeds the number of controls, there is apt to be difficulty. The telephone system has twenty-four functions, yet only fifteen controls—none of them labeled for specific action. In contrast, the trip computer for the car performs seventeen functions with fourteen controls. With minor exceptions, there is one control for each function. In fact, the controls with more than one function are indeed harder to remember and use. When the number of controls equals the number of functions, each control can be specialized, each can be labeled. The possible functions are visible, for each corresponds with a control. If the user forgets the functions, the controls serve as reminders. When, as on the telephone, there are more functions than controls, labeling becomes difficult or impossible. There is nothing to remind the user. Functions are invisible, hidden from sight. No wonder the operation becomes mysterious and difficult. The controls for the car are visible and, through their location and mode of operation, bear an intelligent relationship to their action. Visi-

bility acts as a good reminder of what can be done and allows the control to specify how the action is to be performed. The good relationship between the placement of the control and what it does makes it easy to find the appropriate control for a task. As a result, there is little to remember.

THE PRINCIPLE OF MAPPING

Mapping is a technical term meaning the relationship between two things, in this case between the controls and their movements and the results in the world. Consider the mapping relationships involved in steering a car. To turn the car to the right, one turns the steering wheel clockwise (so that its top moves to the right). The user must identify two mappings here: one of the 112 controls affects the steering, and the steering wheel must be turned in one of two directions. Both are somewhat arbitrary. But the wheel and the clockwise direction are natural choices: visible, closely related to the desired outcome, and providing immediate feedback. The mapping is easily learned and always remembered.

Natural mapping, by which I mean taking advantage of physical analogies and cultural standards, leads to immediate understanding. For example, a designer can use spatial analogy: to move an object up, move the control up. To control an array of lights, arrange the controls in the same pattern as the lights. Some natural mappings are cultural or biological, as in the universal standard that a rising level represents more, a diminishing level, less. Similarly, a louder sound can mean a greater amount. Amount and loudness (and weight, line length, and brightness) are additive dimensions: add more to show incremental increases. Note that the logically plausible relationship between musical pitch and amount does not work: Would a higher pitch mean less or more of something? Pitch (and taste, color, and location) are substitutive dimensions: substitute one value for another to make a change. There is no natural concept of more or less in the comparison of different pitches, or hues, or taste qualities. Other natural mappings follow from the principles of perception and allow for the natural grouping or patterning of controls and feedback (see figure 1.13).

Mapping problems are abundant, one of the fundamental causes of difficulties. Consider the telephone. Suppose you wish to activate the callback on "no reply" function. To initiate this feature on one tele-



1.15 Seat Adjustment Control from a Mercedes-Benz Automobile. This is an excellent example of natural mapping. The control is in the shape of the seat itself: the mapping is straightforward. To move the front edge of the seat higher, lift up on the front part of the button. To make the seat back recline, move the button back. Mercedes-Benz automobiles are obviously not everyday things for most people, but the principle doesn't require great expense or wealth. The same principle could be applied to much more common objects.

phone system, press and release the "recall" button (the button on the handset), then dial 60, then dial the number you called.

There are several problems here. First, the description of the function is relatively complex—yet incomplete: What if two people set up callback at the same time? What if the person does not come back until a week later? What if you have meanwhile set up three or four other functions? What if you want to cancel it? Second, the action to be performed is arbitrary. (Dial 60. Why 60? Why not 73 or 27? How does one remember an arbitrary number?) Third, the sequence ends with what appears to be a redundant, unnecessary action: dialing the number of the person to be called. If the phone system is smart enough to do all these other things, why can't it remember the number that was just attempted; why must it be told all over again? And finally, consider the lack of feedback. How do I know I did the right action? Maybe I disconnected the phone. Maybe I set up some other special feature. There is no visible or audible way to know immediately.

A device is easy to use when there is visibility to the set of possible actions, where the controls and displays exploit natural mappings. The principles are simple but rarely incorporated into design. Good design takes care, planning, thought. It takes conscious attention to the needs of the user. And sometimes the designer gets it right:

Once, when I was at a conference at Gmunden, Austria, a group of us went off to see the sights. I sat directly behind the driver of the brand new, sleek, high-technology German tour bus. I gazed in wonder at the hundreds of controls scattered all over the front of the bus.


"How can you ever learn all those controls?" I asked the driver (with the aid of a German-speaking colleague). The driver was clearly puzzled by the question.

"What do you mean?" he replied. "Each control is just where it ought to be. There is no difficulty."

A good principle, that. Controls are where they ought to be. One function, one control. Harder to do, of course, than to say, but essentially this is the principle of natural mappings: the relationship between controls and actions should be apparent to the user. I return to this topic later in the book, for the problem of determining the "naturalness" of mappings is difficult, but crucial.

I've already described how my car's controls are generally easy to use. Actually, the car has lots of problems. The approach to usability used in the car seems to be to make sure that you can reach everything and see everything. That's good, but not nearly good enough.

Here is a simple example: the controls for the loudspeakers—a simple control that determines whether the sound comes out of the front speakers, the rear, or a combination (figure 1.14). Rotate the wheel from left to right or right to left. Simple, except how do you know which way to rotate the control? Which direction moves the sound to the rear, which to the front? If you want sound to come out of the front speaker, you should be able to move the control to the front. To get it out of the back, move the control to the back. Then the form of the motion would mimic the function and make a natural mapping. But the way the control is actually mounted in the car, forward and backward get translated into left and right. Which direction is which? There is no natural relationship. What's worse, the control isn't even labeled. Even the instruction manual does not say how to use it.



1.14 The Front/Rear Speaker Selector of an Automobile Radio. Rotating the knob with the pictures of the speaker at either side makes the sound come entirely out of the front speakers (when the knob is all the way over to one side), entirely out of the rear speakers (when the knob is all the way the other way), or equally out of both (when the knob is midway). Which way is front, which rear? You can't tell by looking. While you're at it, imagine trying to manipulate the radio controls while keeping your eyes on the road.

The control should be mounted so that it moves forward and backward. If that can't be done, rotate the control 90° on the panel so that it moves vertically. Moving something up to represent forward is not as natural as moving it forward, but at least it follows a standard convention.

In fact, we see that both the car and the telephone have easy functions and difficult ones. The car seems to have more of the easy ones, the telephone more of the difficult ones. Moreover, with the car, enough of the controls are easy that I can do almost everything I need to. Not so with the telephone: it is very difficult to use even a single one of the special features.

The easy things on both telephone and car have a lot in common, as do the difficult things. When things are visible, they tend to be easier

than when they are not. In addition, there must be a close, natural relationship between the control and its function: a *natural mapping*.

THE PRINCIPLE OF FEEDBACK

Feedback—sending back to the user information about what action has actually been done, what result has been accomplished—is a well-known concept in the science of control and information theory. Imagine trying to talk to someone when you cannot even hear your own voice, or trying to draw a picture with a pencil that leaves no mark: there would be no feedback.

In the good old days of the telephone, before the American telephone system was divided among competing companies, before telephones were fancy and had so many features, telephones were designed with much more care and concern for the user. Designers at the Bell Telephone Laboratories worried a lot about feedback. The push buttons were designed to give an appropriate feel—tactile feedback. When a button was pushed, a tone was fed back into the earpiece so the user could tell that the button had been properly pushed. When the phone call was being connected, clicks, tones, and other noises gave the user feedback about the progress of the call. And the speaker's voice was always fed back to the earpiece in a carefully controlled amount, because the auditory feedback (called "sidetone") helped the person regulate how loudly to talk. All this has changed. We now have telephones that are much more powerful and often cheaper than those that existed just a few years ago—more function for less money. To be fair, these new designs are pushing hard on the paradox of technology: added functionality generally comes along at the price of added complexity. But that does not justify backward progress.

Why are the modern telephone systems so difficult to learn and to use? Basically, the problem is that the systems have more features and less feedback. Suppose all telephones had a small display screen, not unlike the ones on small, inexpensive calculators. The display could be used to present, upon the push of a button, a brief menu of all the features of the telephone, one by one. When the desired one was encountered, the user would push another button to indicate that it should be invoked. If further action was required, the display could tell the person what to do. The display could even be auditory, with speech instead of a visual display. Only two buttons need be added to the

telephone: one to change the display, one to accept the option on display. Of course, the telephone would be slightly more expensive. The tradeoff is cost versus usability.⁷

Pity the Poor Designer

Designing well is not easy. The manufacturer wants something that can be produced economically. The store wants something that will be attractive to its customers. The purchaser has several demands. In the store, the purchaser focuses on price and appearance, and perhaps on prestige value. At home, the same person will pay more attention to functionality and usability. The repair service cares about maintainability: how easy is the device to take apart, diagnose, and service? The needs of those concerned are different and often conflict. Nonetheless, the designer may be able to satisfy everyone.

A simple example of good design is the 3½-inch magnetic diskette for computers, a small circle of "floppy" magnetic material encased in hard plastic. Earlier types of floppy disks did not have this plastic case, which protects the magnetic material from abuse and damage. A sliding metal cover protects the delicate magnetic surface when the diskette is not in use and automatically opens when the diskette is inserted into the computer. The diskette has a square shape: there are apparently eight possible ways to insert it into the machine, only one of which is correct. What happens if I do it wrong? I try inserting the disk sideways. Ah, the designer thought of that. A little study shows that the case really isn't square: it's rectangular, so you can't insert a longer side. I try backward. The diskette goes in only part of the way. Small protrusions, indentations, and cutouts prevent the diskette from being inserted backward or upside down: of the eight ways one might try to insert the diskette, only one is correct, and only that one will fit. An excellent design.

Take another example of good design. My felt-tipped marking pen has ribs along only one of its sides; otherwise all sides look identical. Careful examination shows that the tip of the marker is angled and makes the best line if the marker is held with the ribbed side up, a natural result if the forefinger rests upon the ribs. No harm results if I hold the marker another way, but the marker writes less well. The ribs are a subtle design cue—functional, yet visibly and aesthetically unobtrusive.

The world is permeated with small examples of good design, with the amazing details that make important differences in our lives. Each detail was added by some person, a designer, carefully thinking through the uses of the device, the ways that people abuse things, the kinds of errors that can get made, and the functions that people wish to have performed.

Then why is it that so many good design ideas don't find their way into products in the marketplace? Or something good shows up for a short time, only to fall into oblivion? I once spoke with a designer about the frustrations of trying to get the best product out:

It usually takes five or six attempts to get a product right. This may be acceptable in an established product, but consider what it means in a new one. Suppose a company wants to make a product that will perhaps make a real difference. The problem is that if the product is truly revolutionary, it is unlikely that anyone will quite know how to design it right the first time; it will take several tries. But if a product is introduced into the marketplace and fails, well that is it. Perhaps it could be introduced a second time, or maybe even a third time; but after that it is dead: everyone believes it to be a failure.

I asked him to explain. "You mean," I said, "that it takes five or six tries to get an idea right?"

"Yes," he said, "at least that."

"But," I replied, "you also said that if a newly introduced product doesn't catch on in the first two or three times; then it is dead?"

"Yup," he said.

"Then new products are almost guaranteed to fail, no matter how good the idea."

"Now you understand," said the designer. "Consider the use of voice messages on complex devices such as cameras, soft-drink machines, and copiers. A failure. No longer even tried. Too bad. It really is a good idea, for it can be very useful when the hands or eyes are busy elsewhere. But those first few attempts were very badly done and the public scoffed—properly. Now, nobody dares try it again, even in those places where it is needed."

The Paradox of Technology

Technology offers the potential to make life easier and more enjoyable; each new technology provides increased benefits. At the same time,

added complexities arise to increase our difficulty and frustration. The development of a technology tends to follow a U-shaped curve of complexity: starting high; dropping to a low, comfortable level; then climbing again. New kinds of devices are complex and difficult to use. As technicians become more competent and an industry matures, devices become simpler, more reliable, and more powerful. But then, after the industry has stabilized, newcomers figure out how to add increased power and capability, but always at the expense of added complexity and sometimes decreased reliability. We can see the curve of complexity in the history of the watch, radio, telephone, and television set. Take the radio. In the early days, radios were quite complex. To tune in a station required several adjustments, including one for the antenna, one for the radio frequency, one for intermediate frequencies, and controls for both sensitivity and loudness. Later radios were simpler and had controls only to turn it on, tune the station, and adjust the loudness. But the latest radios are again very complex, perhaps even more so than early ones. Now the radio is called a tuner, and it is littered with numerous controls, switches, slide bars, lights, displays, and meters. The modern sets are technologically superior, offering higher quality sound, better reception, and enhanced capability. But what good is the technology if it is too complex to use?

The design problem posed by technological advances is enormous. Consider the watch. A few decades ago, watches were simple. All you had to do was set the time and keep them wound. The standard control was the stem: a knob at the side of the watch. Turning the knob wound the spring that worked the watch. Pulling the knob out and turning it made the hands move. The operations were easy to learn and easy to do. There was a reasonable relation between the turning of the knob and the resulting turning of the hands. The design even took into account human error: the normal position of the stem was for winding the spring, so that an accidental turn would not reset the time.

In the modern digital watch the spring is gone, replaced by a motor run by long-lasting batteries. All that remains is the task of setting the watch. The stem is still a sensible solution, for you can go fast or slow, forward or backward, until the exact desired time is reached. But the stem is more complex (and therefore more expensive) than simple push-button switches. If the only change in the transition from the spring-wound analog watch to the battery-run digital watch were in how the time was set, there would be little difficulty. The problem is that new technology has allowed us to add more functions to the

watch: the watch can give the day of the week, the month, and the year; it can act as a stop watch (which itself has several functions), a count-down timer, and an alarm clock (or two); it has the ability to show the time for different time zones; it can act as a counter and even as a calculator. But the added functions cause problems: How do you design a watch that has so many functions while trying to limit the size, cost, and complexity of the device? How many buttons does it take to make the watch workable and learnable, yet not too expensive? There are no easy answers. Whenever the number of functions and required operations exceeds the number of controls, the design becomes arbitrary, unnatural, and complicated. The same technology that simplifies life by providing more functions in each device also complicates life by making the device harder to learn, harder to use. This is the paradox of technology.

The paradox of technology should never be used as an excuse for poor design. It is true that as the number of options and capabilities of any device increases, so too must the number and complexity of the controls. But the principles of good design can make complexity manageable.

In one of my courses I gave as homework the assignment to design a multiple-function clock radio:

You have been employed by a manufacturing company to design their new product. The company is considering combining the following into one item:

- AM-FM radio
- Cassette player
- CD player
- Telephone
- Telephone answering machine
- Clock
- Alarm clock (the alarm can turn on a tone, radio, cassette, or CD)
- Desk or bed lamp

The company is trying to decide whether to include a small (two-inch screen) TV set and a switched electric outlet that can turn on a coffee maker or toaster.


Your job is (A) to recommend what to build, then (B) to design the control panel, and finally (C) to certify that it is actually both what customers want and easy to use.

State what you would do for the three parts of your job: A, B, and C. Explain how you would go about validating and justifying your recommendations.

Draw a rough sketch of a control panel for the items in the indented list, with a brief justification and analysis of the factors that went into the choice of design.


There are several things I looked for in the answer. (Figure 1.15 is an unacceptable solution.) First, how well did the answer address the

1.15 Possible Solution to My Homework Assignment. Completely unacceptable. (Thanks to Bill Gaver for devising and drawing this sample.)



real needs of the user? I expected my students to visit the homes of potential users to see how their current devices were being used and to determine how the combined multipurpose device would be used. Next, I evaluated whether all the controls were usable and understandable, allowing all the desired functions to be operated with minimum confusion or error. Clock radios are often used in the dark, with the user in bed and reaching overhead to grope for the desired control. Therefore the unit had to be usable in the dark by feel only. It was not supposed to be possible to make a serious mistake by accidentally hitting the wrong control. (Alas, many existing clock radios do not tolerate serious errors—for example, the user may reset the time by hitting the wrong button accidentally.) Finally, the design was expected to take into account real issues in cost, manufacturability, and aesthetics. The finished design had to pass muster with users. The point of the exercise was for the student to realize the paradox of technology: added complexity and difficulty cannot be avoided when functions are added, but with clever design, they can be minimized.

THE PSYCHOLOGY OF EVERYDAY ACTIONS



During my family's stay in England, we rented a furnished house while the owners were away. One day, our landlady returned to the house to get some personal papers. She walked over to her filing cabinet and attempted to open the top drawer. It wouldn't open. She pushed it forward and backward, right and left, up and down, without success. I offered to help. I wiggled the drawer. Then I twisted the front panel, pushed down hard, and banged the front with the palm of one hand. The cabinet drawer slid open. "Oh," she said, "I'm sorry. I am so bad at mechanical things."

Falsely Blaming Yourself

I have studied people making errors—sometimes serious ones—with mechanical devices, light switches and fuses, computer operating systems and word processors, even airplanes and nuclear power plants. Invariably people feel guilty and either try to hide the error or blame themselves for "stupidity" or "clumsiness." I often have difficulty getting permission to watch: nobody likes to be observed performing badly. I point out that the design is faulty and that others make the

same errors. Still, if the task *appears* simple or trivial, then people blame themselves.¹ It is as if they take perverse pride in thinking of themselves as mechanically incompetent.

I once was asked by a large computer company to evaluate a brand new product. I spent a day learning to use it and trying it out on various problems. In using the keyboard to enter data, it was necessary to differentiate between the the "return" key and the "enter" key. If the wrong key was typed, the last few minutes' work was irrevocably lost.

I pointed this problem out to the designer, explaining that I myself had made the error frequently and that my analyses indicated that this was very likely to be a frequent error among users. The designer's first response was: "Why did you make that error? Didn't you read the manual?" He proceeded to explain the different functions of the two keys.

"Yes, yes," I explained, "I understand the two keys, I simply confuse them. They have similar functions, are located in similar locations on the keyboard, and as a skilled typist, I often hit "return" automatically, without thought. Certainly others have had similar problems."

"Nope," said the designer. He claimed that I was the only person who had ever complained, and the company's secretaries had been using the system for many months. I was skeptical, so we went together to some of the secretaries and asked them whether they had ever hit the "return" key when they should have hit "enter." And did they ever lose their work as a result?

"Oh, yes," said the secretaries, "we do that a lot."

"Well, how come nobody ever said anything about it?" we asked the secretaries. After all, they were encouraged to report all problems with the system.

The reason was simple: when the system stopped working or did something strange, the secretaries dutifully reported it as a problem. But when they made the "return" versus "enter" error, they blamed themselves. After all, they had been told what to do. They had simply erred.

Of course, people do make errors. Complex devices will always require some instruction, and someone using them without instruction should expect to make errors and to be confused. But designers should take special pains to make errors as cost-free as possible. Here is my credo about errors:

If an error is possible, someone will make it. The designer must assume that all possible errors will occur and design so as to minimize the chance of the error in the first place, or its effects once it gets made. Errors should be easy to detect, they should have minimal consequences, and, if possible, their effects should be reversible.

Misconceptions of Everyday Life

Our lives are filled with misconceptions. This should not be surprising: we must frequently deal with unfamiliar situations. Psychologists love errors and misconceptions, for they give important clues about the organization and operation of our minds. Many everyday misunderstandings are classified as "naive" or "folk" understandings. And not just plain folk hold these misconceptions: Aristotle developed an entire theory of physics that physicists find quaint and amusing. Yet Aristotle's theories correspond much better to common-sense, everyday observations than do the highly refined and abstract theories we are taught in school. Aristotle developed what we might call naive physics. It is only when you study the esoteric world of physics that you learn what is "correct" and are able to understand why the "naive" view is wrong.

ARISTOTLE'S NAIVE PHYSICS


For example, Aristotle thought that moving objects kept moving only if something kept pushing them. Today's physicist says nonsense: a moving object continues to move unless some force is exerted to stop it. This is Newton's first law of motion, and it contributed to the development of modern physics. Yet anyone who has ever pushed a heavy box along a street or, for that matter, hiked for miles into the wilderness, knows that Aristotle was right: if you don't keep on pushing, the movement stops. Of course, Newton and his successors assume the absence of friction and air. Aristotle lived in a world where there was always friction and air resistance. Once friction is involved, then objects in motion tend to stop unless you keep pushing. Aristotle's theory may be bad physics, but it describes reasonably well what we can see in the real world. Think about how you might answer the following questions.

1. I take a pistol and, carefully aiming it on a level, horizontal line, I fire a bullet. With my other hand, I hold a bullet so that the bullet in the pistol and the one in my hand are exactly the same distance from the ground. I drop the bullet at the same instant as I fire the pistol. Which bullet hits the ground first?

2. Imagine someone running across a field carrying a ball. As you watch, the runner drops the ball. Which path (a, b, or c in figure 2.1) does the ball take as it falls to the ground?

The physicist says the answer to the bullet problem is trivial: both bullets hit the ground at the same time. The fact that one bullet is traveling horizontally very rapidly has absolutely no effect on how fast it falls downward. Why should we accept that answer? Shouldn't the speeding bullet develop some lift—sort of like an airplane—so that it will stay up a bit longer because it is kept up by the air? Who knows? The theory of physics is based upon a situation where there is no air. The popular misconception is that the pistol bullet will hit the ground long after the dropped bullet; yet this naive view doesn't seem so strange.

2.1 A Running Man Drops a Ball. Which path does the ball take as it falls to the ground, path A, B, or C? When this question was asked of sixth-grade students in Boston schools, only 3 percent answered A, the right answer; the others were evenly divided between B and C. Even high school students did not do well: of forty-one students who had just studied Newtonian mechanics for a month and a half, only 20 percent got the right answer; the others were almost equally divided between B and C. (The study was performed by White & Horwitz, 1987. The figure is reprinted from *Intuitive Physics* by McCloskey. Copyright © 1983 by Scientific American, Inc. All rights reserved.)



In the case of the falling ball, our prediction is that the ball will drop straight down. In fact, the falling ball follows trajectory A (figure 2.1). As it is carried by the runner, it is set into horizontal motion. It then maintains the same forward speed upon being released, even as it also falls to the ground.³

Naive physics—and naive views of psychology and other fields—are often sensible, even if wrong. But at times they can get us into trouble. Yet we must have a way to digest the unfamiliar, for people are explanatory creatures.

PEOPLE AS EXPLANATORY CREATURES

Mental models, our conceptual models of the way objects work, events take place, or people behave, result from our tendency to form explanations of things. These models are essential in helping us understand our experiences, predict the outcomes of our actions, and handle unexpected occurrences. We base our models on whatever knowledge we have, real or imaginary, naive or sophisticated.

Mental models are often constructed from fragmentary evidence, with but a poor understanding of what is happening, and with a kind of naive psychology that postulates causes, mechanisms, and relationships even where there are none. Some faulty models lead to the frustrations of everyday life, as in the case of my unsetting refrigerator, where my mental model of its operation (figure 1.9 A) did not correspond to reality (figure 1.9 B). Far more serious are faulty models of such complex systems as an industrial plant or passenger airplane. Misunderstanding there can lead to devastating accidents.

Consider the room thermostat. How does it work? Here is a device that offers almost no evidence of its operation except in a highly round-about manner. We walk into a room and feel too cold: so we walk over to the thermostat and set it higher. Eventually we feel warmer. Note that the same thing applies to the temperature control for a cooking oven (or a pottery kiln, or an air conditioner, or almost any device whose temperature is to be regulated). Want to bake a cake, but the oven is off? Set the oven thermostat and the oven gets to the desired temperature. Is the room too hot? Set the thermostat on the air conditioner. Fine, but how does the thermostat work?

If you are in a cold room, in a hurry to get warm, will the room heat more quickly if you turn the thermostat all the way up? Or if you want

the oven to reach its working temperature faster, should you turn the temperature dial all the way to maximum, then turn it down once the desired temperature is reached? Or to cool a room most quickly, should you set the air conditioner thermostat to its lowest temperature setting?

If you think that the room or oven will heat (or cool) faster if the thermostat is turned all the way to the maximum setting, you are wrong. You hold a folk theory of thermostats. There are two commonly held folk theories about thermostats: the timer theory and the valve theory. The timer theory proposes that the thermostat simply controls the relative proportion of time that the device stays on. Set the thermostat midway, and the device is on about half the time; set it all the way up and the device is on all the time. Hence, to heat or cool something most quickly, set the thermostat so that the device is on all the time. The valve theory proposes that the thermostat controls how much heat (or cold) comes out of the device. Turn the thermostat all the way up, and you get maximum heating or cooling.⁴

The correct story is that the thermostat is just an on-off switch. It treats the heater, oven, and air conditioner as all-or-nothing devices that can be either fully on or fully off, with no in-between states. The thermostat turns the heater, oven, or air conditioner completely on—at full power—until the temperature setting on the thermostat is reached. Then it turns the unit completely off. Setting the thermostat at one extreme cannot affect how long it takes to reach the desired temperature.⁵

The real point of the example is not that some people have erroneous theories; it is that everyone forms theories (mental models) to explain what they have observed. In the case of the thermostat, the design gives absolutely no hint as to the correct answer. In the absence of external information, people are free to let their imaginations run free as long as the mental models they develop account for the facts as they perceive them.

Blaming the Wrong Cause

"Look at this!" my colleague exclaimed to me, "My computer terminal is broken. The library did it! Every time I connect it to the library catalog I have trouble. Now I can't even use the terminal to read my computer mail anymore."

"That doesn't make sense," I replied. "You can't even turn on the

power to the terminal. How could a computer program possibly do that kind of damage?"

"All I know," he said, "is that everything was working fine until I tried to look up an author in the library catalog using that new library program, and then my terminal stopped working. I always have trouble with that program. And this is simply too much of a coincidence to be anything else."

Well, it was a coincidence. It turns out that the power supply to the terminal had burned out, a fact that had nothing to do with the computer program. Coincidence is enough to set the causal wheels rolling.

Earlier I suggested that people have a tendency to blame themselves for difficulties with technology. Actually, the point is a bit more complicated. People do tend to find causes for events, and just what they assign as the cause varies. In part people tend to assign a causal relation whenever two things occur in succession. If I do some action *A* just prior to some result *R*, then I conclude that *A* must have caused *R*, even if, as in the example above, there really was no relationship between the two. The story is more complex when we intend an action to produce a desired result and fail, and there are problems when we have done the action through some intermediate mechanism.

Just where do we put the blame for failure? The answer is not clear. The psychology of blame (or, to be more accurate, of attribution) is complex and not fully understood. In part, there seems to have to be some perceived causal relationship between the thing being blamed and the result. The word *perceived* is critical: the causal relationship does not have to exist; the person simply has to think it is there. Sometimes we attribute the cause to things that had nothing to do with the action. And sometimes we ignore the real culprit.

One major aspect of the assignment of blame is that we frequently have little information on which to make the judgment, and what little we have may be wrong. As a result, blame or credit can be assessed almost independently of reality. Here is where the apparent simplicity of everyday objects causes problems. Suppose I try to use an everyday thing, but I can't. Where is the fault, in my action or in the thing? We are apt to blame ourselves. If we believe that others are able to use the device and if we believe that it is not very complex, then we conclude that any difficulties must be our own fault. Suppose the fault really lies in the device, so that lots of people have the same problems. Because everyone perceives the fault to be his or her own, nobody wants to

admit to having trouble. This creates a conspiracy of silence, maintaining the feelings of guilt and helplessness among users.

Interestingly enough, the common tendency to blame ourselves for failures with everyday objects goes against the normal attributions people make. In general, it has been found that people attribute their own problems to the environment, those of other people to their personalities.

Here is a made-up example. Consider Tom, the office terror. Today Tom got to work late, slammed the door to his office, and yelled at his colleagues. "Ah," his colleagues and staff said, "there he goes again. He's so excitable—always gets mad at the slightest thing."

Now consider Tom's point of view. "I really had a hard day," Tom explains. "I woke up late because when my clock radio turned on, I tried to hit the snooze bar to give me five minutes' more sleep; instead I reset the time so that I overslept for a whole hour. That wasn't my fault—the radio's badly designed. I didn't even have time for my morning coffee. I couldn't find a close parking spot because I was late. And then because I was in such a rush I dropped my papers all over the street and got them dirty. Then when I went to get a cup of coffee from the office machine, it was all out. None of this was my fault—I had a run of really bad events. Yes, I was a bit curt with my colleagues, but who wouldn't be under the same circumstances? Surely they understand."

But Tom's colleagues see a different picture. They don't have access to his inner thoughts or even to his morning's activities. All they see is that Tom yelled at them simply because the office coffee machine was empty. And this reminds them of another time when the same thing happened. "He does that all the time," they conclude, "always blowing up over the most minor events." The events are the same events, but there are two different points of view and two different interpretations. The protagonist, Tom, views his actions as sensible responses to the trials of life. The onlooker views Tom's actions as a result of his explosive, irascible personality.

It seems natural for people to blame their own misfortunes on the environment. It seems equally natural to blame other people's misfortunes on their personalities. Just the opposite attribution, by the way, is made when things go well. When things go right, people credit their own forceful personalities and intelligence: "I really did a good job today; no wonder we finished the project so well." The onlookers do

the reverse. When they see things go well for someone else, they credit the environment: "Joan really was lucky today; she just happened to be standing there when the boss came by, so she got all the credit for the project work. Some people have all the luck."

In all cases, whether a person is inappropriately accepting blame for the inability to work simple objects or attributing behavior to environment or personality, a faulty mental model is at work.

LEARNED HELPLESSNESS

The phenomenon called *learned helplessness* may help explain the self-blame. It refers to the situation in which people experience failure at a task, often numerous times. As a result, they decide that the task cannot be done, at least not by them: they are helpless. They stop trying. If this feeling covers a group of tasks, the result can be severe difficulties coping with life. In the extreme case, such learned helplessness leads to depression and to a belief that the person cannot cope with everyday life at all. Sometimes all that it takes to get such a feeling of helplessness is a few experiences that accidentally turn out bad. The phenomenon has been most frequently studied as a precursor to the clinical problem of depression, but it might easily arise with a few bad experiences with everyday objects.

TAUGHT HELPLESSNESS

Do the common technology and mathematics phobias result from a kind of learned helplessness? Could a few instances of failure in what appear to be straightforward situations generalize to every technological object, every mathematics problem? Perhaps. In fact, the design of everyday things (and the design of mathematics courses) seems almost guaranteed to cause this. We could call this phenomenon *taught helplessness*.

With badly designed objects—constructed so as to lead to misunderstanding—faulty mental models, and poor feedback, no wonder people feel guilty when they have trouble using objects, especially when they perceive (even if incorrectly) that nobody else is having the same problems. Or consider the normal mathematics curriculum, which continues relentlessly on its way, each new lesson assuming full knowl-

edge and understanding of all that has passed before. Even though each point may be simple, once you fall behind it is hard to catch up. The result: mathematics phobia. Not because the material is difficult, but because it is taught so that difficulty in one stage hinders further progress. The problem is that once failure starts, it soon generalizes by self-blame to all of mathematics. Similar processes are at work with technology. The vicious cycle starts: if you fail at something, you think it is your fault. Therefore you think you can't do that task. As a result, next time you have to do the task, you believe you can't so you don't even try. The result is that you can't, just as you thought. You're trapped in a self-fulfilling prophecy.

The Nature of Human Thought and Explanation

It isn't always easy to tell just where the blame for a problem should be placed. A number of dramatic accidents have come about, in part, from the false assessment of blame in a situation. Highly skilled, well-trained people are using complex equipment when suddenly something goes wrong. They have to figure out what the problem is. Most industrial equipment is pretty reliable. When the instruments indicate that something is wrong, one has to consider the possibility that the instruments themselves are wrong. Often this is the correct assessment. But when operators mistakenly blame the instruments for an actual equipment failure, the situation is ripe for a major accident.

It is spectacularly easy to find examples of false assessment in industrial accidents. Analysts come in well after the fact, knowing what actually did happen; with hindsight, it is almost impossible to understand how the people involved could have made the mistake. But from the point of view of the person making decisions at the time, the sequence of events is quite natural.

At the Three Mile Island nuclear power plant, operators pushed a button to close a valve; the valve had been opened (properly) to allow excess water to escape from the nuclear core. In fact, the valve was deficient, so it didn't close. But a light on the control panel indicated that the valve position was closed. The light actually didn't monitor the valve, only the electrical signal to the valve, a fact known by the operators. Still, why suspect a problem? The operators did look at the temperature in the pipe leading from the valve: it was high, indicating that fluid was still flowing through the closed valve. Ah, but the opera-

tors knew that the valve had been leaky, so the leak would explain the high temperature; but the leak was known to be small, and operators assumed that it wouldn't affect the main operation. They were wrong, and the water that was able to escape from the core added significantly to the problems of that nuclear disaster. I think the operators' assessment was perfectly reasonable: the fault was in the design of the lights and in the equipment that gave false evidence of a closed valve.

Similar misinterpretations take place all the time. I have studied a number of airline accidents. Consider the flight crew of the Lockheed L-1011 flying from Miami, Florida, to Nassau, Bahamas. The plane was over the Atlantic Ocean, about 110 miles from Miami, when the low oil pressure light for one of the three engines went on. The crew turned off the engine and turned around to go back to Miami. Eight minutes later, the low pressure lights for the remaining two engines also went on, and the instruments showed zero oil pressure and quantity in all three engines. What did the crew do now? They didn't believe it! After all, the pilot correctly said later, the likelihood of simultaneous oil exhaustion in all three engines was "one in millions; I would think." At the time, sitting in the airplane, simultaneous failure did seem most unlikely. Even the National Transportation Safety Board declared, "The analysis of the situation by the flightcrew was logical, and was what most pilots probably would have done if confronted by the same situation."⁶

What happened? The second and third engines were indeed out of oil, and they failed. So there were no operating engines: one had been turned off when its gauge registered low, the other two had failed. The pilots prepared the plane for an emergency landing on the water. The pilots were too busy to instruct the flight crew properly, so the passengers were not prepared. There was semi-hysteria in the passenger cabin. At the last minute, just as the plane was about to ditch in the ocean, the pilots managed to restart the first engine and to land safely at Miami. Then that engine failed at the end of the runway.

Why did all three engines fail? Three missing O-rings, one missing from each of three oil plugs, allowed all the oil to seep out. The O-rings were put in by two different people who worked on the three engines (one for the two plugs on the wings, the other for the plug on the tail). How did both workers make the same mistake? Because the normal method by which they got the oil plugs had been changed that day. The whole tale is very instructive, for there were four major failures of

different sorts, from the omission of the O-rings, to the inadequacy of the maintenance procedures, to the false assessment of the problem, to the poor handling of the passengers. Fortunately, nobody was injured. The analysts of the National Transportation Safety Board got to write a fascinating report.

I've misinterpreted signals, as I'm sure most people have. My family was driving from San Diego to Mammoth, California, a ski area about 500 miles north: a ten- to twelve-hour drive. As we drove, we noticed more and more signs advertising the hotels and gambling casinos of Las Vegas, Nevada. "Strange," we said, "Las Vegas always did advertise a long way off—there is even a billboard in San Diego—but this seems excessive, advertising on the road to Mammoth." We stopped for gasoline and continued on our journey. Only later, when we tried to find a place to eat supper, did we discover that we had taken the wrong turn nearly two hours earlier, before we had stopped for gasoline, and that we were on the road to Las Vegas, not the road to Mammoth. We had to backtrack the entire two-hour segment, wasting four hours of driving. It's humorous now; it wasn't then.

Find an explanation, and we are happy. But our explanations are based on analogy with past experience, experience that may not apply in the current situation. In the Three Mile Island incident, past experience with the leaky valve explained away the discrepant temperature reading; on the flight from Miami to Nassau, the pilots' lack of experience with simultaneous oil pressure failure triggered their belief that the instruments must be faulty; in the driving story, the prevalence of billboards for Las Vegas seemed easily explained. Once we have an explanation—correct or incorrect—for otherwise discrepant or puzzling events, there is no more puzzle, no more discrepancy. As a result, we are complacent, at least for a while.

How People Do Things: The Seven Stages of Action

I am in Italy, at a conference. I watch the next speaker attempt to thread a film onto a projector that he has never used before. He puts the reel into place, then takes it off and reverses it. Another person comes to help. Jointly they thread the film through the projector and hold the free end, discussing how to put it on the takeup reel. Two

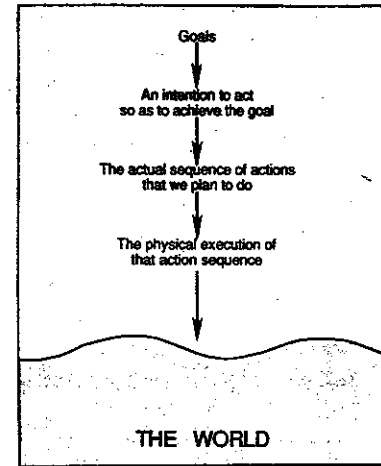

more people come over to help, and then another. The voices grow louder, in three languages: Italian, German, and English. One person investigates the controls, manipulating each and announcing the result. Confusion mounts. I can no longer observe all that is happening. The conference organizer comes over. After a few moments he turns and faces the audience, which has been waiting patiently in the auditorium. "Ahem," he says, "is anybody expert in projectors?" Finally, fourteen minutes after the speaker had started to thread the film (and eight minutes after the scheduled start of the session) a blue-coated technician appears. He scowls, then promptly takes the entire film off the projector, rethreads it, and gets it working.

What makes something—like threading the projector—difficult to do? To answer this question, the central one of this book, we need to know what happens when someone does something. We need to examine the structure of an action.

The basic idea is simple. To get something done, you have to start with some notion of what is wanted—the goal that is to be achieved. Then, you have to do something to the world, that is, take action to move yourself or manipulate someone or something. Finally, you check to see that your goal was made. So there are four different things to consider: the goal, what is done to the world, the world itself, and the check of the world. The action itself has two major aspects: doing something and checking. Call these *execution* and *evaluation* (figure 2.2).

Real tasks are not quite so simple. The original goal may be imprecisely specified—perhaps "get something to eat," "get to work," "get dressed," "watch television." Goals do not state precisely what to do—where and how to move, what to pick up. To lead to actions goals must be transformed into specific statements of what is to be done, statements that I call *intentions*. A *goal* is something to be achieved, often vaguely stated. An *intention* is a specific action taken to get to the goal. Yet even intentions are not specific enough to control actions.

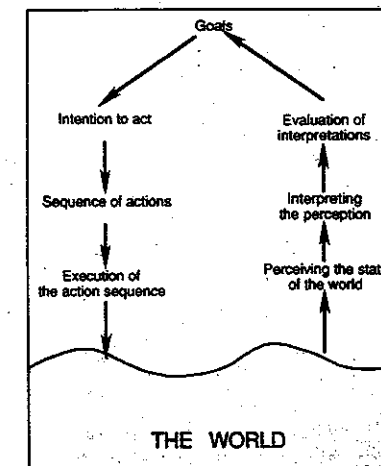
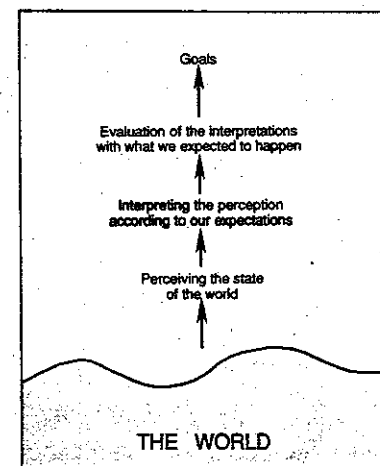
Suppose I am sitting in my armchair, reading a book. It is dusk, and the light has gotten dimmer and dimmer. I decide I need more light (that is the goal: get more light). My goal has to be translated into the intention that states the appropriate action in the world: push the switch button on the lamp. There's more: I need to specify how to move my body, how to stretch to reach the light switch, how to extend my finger to push the button (without knocking over the lamp). The goal



2.2 The Action Cycle (above left). Human action has two aspects, execution and evaluation. Execution involves doing something. Evaluation is the comparison of what happened in the world with what we wanted to happen (our goal).

2.3 Stages of Execution (above right). Start at the top with the goal, the state that is to be achieved. The goal is translated into an intention to do some action. The intention must be translated into a set of internal commands, an action sequence that can be performed to satisfy the intention. The action sequence is still a mental event: nothing happens until it is executed, performed upon the world.

2.4 Stages of Evaluation (below left). Evaluation starts with our perception of the world. This perception must then be interpreted according to our expectations and then compared (evaluated) with respect to both our intentions (from figure 2.3) and our goals.



2.5 Seven Stages of Action (below right). The stages of execution from figure 2.3 (intentions, action sequence, and execution) are coupled with the stages of evaluation from figure 2.4 (perception, interpretation, and evaluation), with goals common to both stages.

has to be translated into an intention, which in turn has to be made into a specific action sequence, one that can control my muscles. Note that I could satisfy my goal with other action sequences, other intentions. If someone walked into the room and passed by the lamp, I might alter my intention from pushing the switch button to asking the other person to do it for me. The goal hasn't changed, but the intention and resulting action sequence have.

The specific actions bridge the gap between what we would like to have done (our goals and intentions) and all possible physical actions. After we specify what actions to make, we must actually do them—the stage of execution. All in all, there are three stages that follow from the goal: intention, action sequence, and execution (figure 2.3).

The evaluation side of things, checking up on what happened, has three stages: first, perceiving what happened in the world; second, trying to make sense of it (interpreting it); and, finally, comparing what happened with what was wanted (figure 2.4).

There we have it. Seven stages of action: one for goals, three for execution, and three for evaluation.

- Forming the goal
- Forming the intention
- Specifying an action
- Executing the action
- Perceiving the state of the world
- Interpreting the state of the world
- Evaluating the outcome

The seven stages form an *approximate model*, not a complete psychological theory. In particular, the stages are almost certainly not discrete entities. Most behavior does not require going through all stages in sequence, and most activities will not be satisfied by single actions. There must be numerous sequences, and the whole activity may last hours or even days. There is a continual feedback loop, in which the results of one activity are used to direct further ones, in which goals lead to subgoals, intentions lead to subintentions. There are activities in which goals are forgotten, discarded, or reformulated.⁷

For many everyday tasks, goals and intentions are not well specified: they are opportunistic rather than planned. Opportunistic actions are

those in which the behavior takes advantage of the circumstances. Rather than engage in extensive planning and analysis, the person goes about the day's activities and performs the intended actions if the relevant opportunity arises. Thus, we may not go out of our way to go to a shop; or to the library, or to ask a question of a friend. Rather, we go through the day's activities, and if we find ourselves at the shop, near the library, or encountering the friend, then we allow the opportunity to trigger the relevant activity. Otherwise, the task remains undone. Only in the case of crucial tasks do we make special efforts to ensure that they get done. Opportunistic actions are less precise and certain than specified goals and intentions, but they result in less mental effort, less inconvenience, and perhaps more interest.

The seven-stage process of action can be started at any point. People do not always behave as full, logical, reasoning organisms, starting with high-level goals and working to achieve them. Our goals are often ill-formed and vague. We may respond to the events of the world (in what is called data-driven behavior) rather than to think out plans and goals. An event in the world may trigger an interpretation and a resulting response. Actions may be executed before they are fully developed. In fact, some of us adjust our lives so that the environment can control our behavior. For example, sometimes when I must do an important task, I make a formal, public promise to get it done by a certain date. I make sure that I will be reminded of the promise. And then, hours before the deadline, I actually get to work and do the job. This kind of behavior is fully compatible with the seven-stage analysis.


The Gulfs of Execution and Evaluation

Remember the movie projector story? People's problems threading the projector did not come from a lack of understanding of the goal or the task. It did not come from deep, subtle complexity. The difficulty lay entirely in determining the relationship between the intended actions and the mechanisms of the projector, in determining the functions of each of the controls, in determining what specific manipulation of each control enabled each function, and in deciding by the sights, sounds, lights, and movements of the projector whether the intended actions were being done successfully. The users had a problem with mappings and feedback, as they would have with the projector in figure 2.6.

The projector story is only an extreme case of the difficulties faced in the conduct of many tasks. For a surprisingly large number of every-

day tasks, the difficulty resides entirely in deriving the relationships between the mental intentions and interpretations and the physical actions and states. There are several *gulfs* that separate mental states from physical ones. Each gulf reflects one aspect of the distance between the mental representations of the person and the physical com-

2.6 Threading the Movie Projector. The dark line at the right shows the path of the film. This picture doesn't tell the whole story, for the several loops of film have to be threaded just right, neither too loose nor too taut. (From *Projectionist's manual*, Department of the Army and the Air Force, May 1966.)



ponents and states of the environment. And these gulfs present major problems for users.⁸

THE GULF OF EXECUTION

Does the system provide actions that correspond to the intentions of the person? The difference between the intentions and the allowable actions is the Gulf of Execution. One measure of this gulf is how well the system allows the person to do the intended actions directly, without extra effort: Do the actions provided by the system match those intended by the person?

Consider the movie projector example: one problem resulted from the Gulf of Execution. The person wanted to set up the projector. Ideally, this would be a simple thing to do. But no, a long, complex sequence was required. It wasn't at all clear what actions had to be done to accomplish the intentions of setting up the projector and showing the film.

Self-threading projectors do exist. These nicely bridge the gulf. Or look at VCRs. They have the same mechanical problem as film projectors: the videotape has to be threaded through their mechanism. But the solution is to hide this part of the system, to put the task on the machine, not the person. So the machinery bridges the gulf. All the user has to do is to plop in the cartridge and push the start button. It's a pity the film companies are so far behind. Well, in a while it won't matter. There won't be any film, just videotape.

THE GULF OF EVALUATION

Does the system provide a physical representation that can be directly perceived and that is directly interpretable in terms of the intentions and expectations of the person? The Gulf of Evaluation reflects the amount of effort that the person must exert to interpret the physical state of the system and to determine how well the expectations and intentions have been met. The gulf is small when the system provides information about its state in a form that is easy to get, is easy to interpret, and matches the way the person thinks of the system.

In the movie projector example there was also a problem with the Gulf of Evaluation. Even when the film was in the projector, it was

difficult to tell if it had been threaded correctly. With VCRs all you have to know is whether the cartridge is properly inserted into the machine. If it isn't, usually it won't fit right: it sticks out obviously, and you know that things are not right.

But VCRs aren't perfect, either. I remember a conference speaker who pushed the start button on the VCR and told the audience to watch the screen. No picture. She fiddled with the machine, then called for help. One, then two, then three technicians appeared on the scene. They carefully checked the power connections, the leads to the VCR, the circuits. The audience waited impatiently, giggling. Finally the problem was found: there wasn't any tape in the VCR. No tape, no picture. The problem was that once the cartridge door to that particular VCR was shut, there was no visible way to tell whether it contained a tape. Bad design. That Gulf of Evaluation sunk another user.

The gulfs are present to an amazing degree in a variety of devices. Usually the difficulties are unremarked and invisible. The users either take the blame themselves (in the case of things they believe they should be capable of using, such as water faucets, refrigerator temperature controls, stove tops, radio and television sets) or decide that they are incapable of operating the pesky devices (sewing machines, washing machines, digital watches, digital controls on household appliances, VCRs, audio sets). These are indeed the gadgets of everyday household use. None of them has a complex structure, yet many of them defeat the otherwise capable user.

The Seven Stages of Action as Design Aids

The seven-stage structure can be a valuable design aid, for it provides a basic checklist of questions to ask to ensure that the Gulfs of Evaluation and Execution are bridged (figure 2.7).

In general, each stage of action requires its own special design strategies and, in turn, provides its own opportunity for disaster. It would be fun, were it not also so frustrating, to look over the world and gleefully analyze each deficiency. On the whole, as you can see in figure 2.7, the questions for each stage are relatively simple. And these, in turn, boil down to the principles of good design introduced in chapter 1.

- *Visibility.* By looking, the user can tell the state of the device and the alternatives for action.

How Easily Can One:

Determine The Function of the Device?

Tell What Actions Are Possible?	Tell if System is in Desired State?
Determine Mapping from Intention to Physical Movement?	Determine Mapping from System State to Interpretation?
Perform the Action?	Tell What State the System is In?

2.7 Using the Seven Stages to Ask Design Questions

- *A good conceptual model.* The designer provides a good conceptual model for the user, with consistency in the presentation of operations and results and a coherent, consistent system image.
- *Good mappings.* It is possible to determine the relationships between actions and results, between the controls and their effects, and between the system state and what is visible.
- *Feedback.* The user receives full and continuous feedback about the results of actions.

Each point provides support for one or more of the seven stages of action. The next time you can't immediately figure out the shower control in a motel or work an unfamiliar television set or stove, remember that the problem is in the design. And the next time you pick up an unfamiliar object and use it smoothly and effortlessly on the first try, stop and examine it: the ease of use did not come about by accident. Someone designed the object carefully and well.