bors may be annoyed and that others can monitor your activities. The use of sound to convey information is a powerful and important idea, but still in its infancy.

Just as the presence of sound can serve a useful role in providing feedback about events, the absence of sound can lead to the same kinds of difficulties we have already encountered from a lack of feedback. The absence of sound can mean an absence of information, and if feedback from an action is expected to come from sound, silence can lead to problems.

I once stayed in the guest apartment of a technological institute in the Netherlands. The building was newly completed, with many interesting architectural features. The architect had gone to great lengths to keep the noise level low; the ventilation system could not be heard. In similar fashion, the ventilation for the room came and went through invisible slots in the ceiling (so I am told; I never did find them).

All was fine until I took a shower. The bathroom seemed to have no ventilation at all, so everything became wet, then eventually cold and clammy. There was a switch in the bathroom that I thought might be the control for an exhaust fan. When I pushed the switch, a light on it came on and stayed on. Further pushing had no effect.

I noticed that whenever I returned to the apartment after an absence, the light would be off. So each time I entered the apartment, I went into the bathroom and pushed the button. By listening closely, I could hear a slight "thump" in the distance the first time the button was depressed. I decided it was some kind of signal. Perhaps it was a call button, summoning the maid, or the janitor, or maybe even the fire department (though no one showed up). I did also consider that it might control a ventilation system, but I could hear no flow of air. I examined the inside of the entire bathroom with care, trying to find an air inlet. I even got a chair and a flashlight and examined the ceiling. Nothing.

At the end of my stay, the person driving me to the airport, explained that the button controlled the exhaust fan. The fan was on as long as the light was on, and it turned off, automatically, in about five minutes. The architect was very good at disguising the ventilation system and at keeping the noise level down.

Here is a case where the architect was too successful: the feedback was clearly lacking. The light was not enough—in fact, it was quite misleading. Noise would have been welcome. It would have signaled that there really was ventilation.

TO ERR IS HUMAN

"LONDON—An inexperienced computer-operator pressed the wrong key on a terminal in early December, causing chaos at the London Stock Exchange. The error at stockbrokers Greenwell Montagu led to systems staff working through the night in an attempt to cure the problem."[1]

People make errors routinely. Hardly a minute of a normal conversation can go by without a stumble, a repetition, a phrase stopped midway through to be discarded or redone. Human language provides special mechanisms that make corrections so automatic that the participants hardly take notice; indeed, they may be surprised when errors are pointed out. Artificial devices do not have the same tolerance. Push the wrong button, and chaos may result.

Errors come in several forms. Two fundamental categories are slips and mistakes. Slips result from automatic behavior, when subconscious actions that are intended to satisfy our goals get waylaid en route. Mistakes result from conscious deliberations. The same processes that make us creative and insightful by allowing us to see relationships between apparently unrelated things, that let us leap to correct conclusions on the basis of partial or even faulty evidence, also lead to error.

Our ability to generalize from small amounts of information helps tremendously in new situations; but sometimes we generalize too rapidly, classifying a new situation as similar to an old one when, in fact, there are significant discrepancies. False generalizations can be hard to discover, let alone eliminate.

The differences between slips and mistakes are readily apparent in the analysis of the seven stages of action. Form an appropriate goal but mess up in the performance, and you've made a slip. Slips are almost always small things: a misplaced action, the wrong thing moved, a desired action undone. Moreover, they are relatively easy to discover by simple observation and monitoring. Form the wrong goal, and you've made a mistake. Mistakes can be major events, and they are difficult or even impossible to detect—after all, the action performed is appropriate for the goal.

Slips

A colleague reported that he went to his car to drive to work. As he drove away, he realized that he had forgotten his briefcase, so he turned around and went back. He stopped the car, turned off the engine, and unbuckled his wristwatch. Yes, wristwatch, instead of his seatbelt.

Most everyday errors are slips. Intend to do one action, find yourself doing another. Have a person say something clearly and distinctly to you, but "hear" something quite different. The study of slips is the study of the psychology of everyday errors—what Freud called "the psychopathology of everyday life." Some slips may indeed have hidden, darker meanings, but most are accounted for by rather simple events in our mental mechanisms.[2]

Slips show up most frequently in skilled behavior. We don't make so many slips in things we are still learning. In part, slips result from a lack of attention. On the whole, people can consciously attend to only one primary thing at a time. But we often do many things at once. We walk while we talk; we drive cars while we talk, sing, listen to the radio, use a telephone, take notes, or read a map. We can do more than one thing at a time only if most of the actions are done automatically, subconsciously, with little or no need for conscious attention.

Doing several things at once is essential even in carrying out a single task. To play the piano, we must move the fingers properly over the keyboard while reading the music, manipulating the pedals, and listening to the resulting sounds. But to play the piano well, we should do these things automatically. Our conscious attention should be focused on the higher levels of the music, on style, and on phrasing. So it is with every skill. The low-level, physical movements should be controlled subconsciously.

TYPES OF SLIPS

Some slips result from the similarities of actions. Or an event in the world may automatically trigger an action. Sometimes our thoughts and actions may remind us of unintended actions, which we then perform. We can place slips into one of six categories: capture errors, description errors, data-driven errors, associative activation errors, loss-of-activation errors, and mode errors.

CAPTURE ERRORS

"I was using a copying machine, and I was counting the pages. I found myself counting '1, 2, 3, 4, 5, 6, 7, 8, 9, 10, Jack, Queen, King.' I have been playing cards recently."[3]

Consider the common slip called the capture error, in which a frequently done activity suddenly takes charge instead of (captures) the one intended.[4] You are playing a piece of music (without too much attention) and it is similar to another (which you know better); suddenly you are playing the more familiar piece. Or you go off to your bedroom to change your clothes for dinner and find yourself in bed. (This slip was first reported by William James in 1890.) Or you finish typing your thoughts on your word processor or text editing program, turn off the power, and go off to other things, neglecting to save any of your work. Or you get into your car on Sunday to go to the store and find yourself at the office.

The capture error appears whenever two different action sequences have their initial stages in common, with one sequence being unfamiliar and the other being well practiced. Seldom, if ever, does the unfamiliar sequence capture the familiar one.

DESCRIPTION ERRORS

A former student reported that one day he came home from jogging, took off his sweaty shirt, and rolled it up in a ball, intending to throw

2/19

it in the laundry basket. Instead he threw it in the toilet. (It wasn't poor aim: the laundry basket and toilet were in different rooms.)

In the common slip known as the description error, the intended action has much in common with others that are possible. As a result, unless the action sequence is completely and precisely specified, the intended action might fit several possibilities. Suppose that my tired student in the example formed a mental description of his intended action something like "throw the shirt into the opening at the top of the container." This description would be perfectly unambiguous and sufficient were the laundry basket the only open container in sight; but when the open toilet was visible, its characteristics matched the description and triggered the inappropriate action. This is a description error because the internal description of the intention was not sufficiently precise. Description errors usually result in performing the correct action on the wrong object. Obviously, the more the wrong and right objects have in common, the more likely the errors are to occur. Description errors, like all slips, are more likely when we are distracted, bored, involved in other activities, under extra stress, or otherwise not inclined to pay full attention to the task at hand.

Description errors occur most frequently when the wrong and right objects are physically near each other. People have reported a number of description errors to me.

Two clerks in a department store were both on the telephone to verify credit cards while simultaneously dealing with a customer and filling out a credit card form. One sales clerk had passed in back of the other to reach the charge forms. When this clerk finished preparing the sales slip, she hung up the handset on the wrong telephone, thereby terminating the other clerk's call.

A person intended to put the lid on a sugar bowl, but instead put it on a coffee cup (with the same size opening).

I had a report of someone who planned to to pour orange juice into a glass but instead poured it into a coffee cup (adjacent to the glass).

Another person told me of intending to pour rice from a storage jar into a measuring cup, but instead pouring cooking oil into the measuring cup (both the oil and the rice were kept in glass containers on the counter).

Some things seem designed to cause slips. Long rows of identical switches are perfect setups for description errors. Intend to flip one switch, instead flip a similar-looking one. It happens in industrial plants, aircraft, homes, anywhere. When different actions have similar descriptions, there is a good chance of mishap, especially when the operator is experienced and well practiced and therefore not paying full attention, and if there are more important things to do.

DATA-DRIVEN ERRORS

"I was assigning a visitor a room to use. I decided to call the department secretary to tell her the room number. I used the telephone in the alcove outside the room, with the room number in sight. Instead of dialing the secretary's phone number—which I use frequently and know very well—I dialed the room number."

Much human behavior is automatic, for example, brushing away an insect. Automatic actions are data driven—triggered by the arrival of the sensory data. But sometimes data-driven activities can intrude into an ongoing action sequence, causing behavior that was not intended.

ASSOCIATIVE ACTIVATION ERRORS

"My office phone rang. I picked up the receiver and bellowed 'Come in' at it."[5]

If external data can sometimes trigger actions, so, too, can internal thoughts and associations. The ringing of the telephone and knocking on the door both signal the need to greet someone. Other errors occur from associations among thoughts and ideas. Associative activation errors are the slips studied by Freud; you think something that ought not to be said and then, to your embarrassment, you say it.

LOSS-OF-ACTIVATION ERRORS

"I have to go to the bedroom before I start working in the dining room. I start going there and realize as I am walking that I have no idea why I am going there. Knowing myself, I keep going, hoping that something in the bedroom will remind me. . . . I get there but still cannot recall what I wanted . . . so I go back to the dining room. There I realize that my glasses are dirty. With great relief I go back to the bedroom, get my handkerchief, and wipe my glasses clean."

One of the more common slips is simply forgetting to do something. More interesting is forgetting part of the act, remembering the rest, as in the story above where the goal was forgotten, but the rest of the action continued unimpaired. One of my informants walked all the way through the house to the kitchen and opened the refrigerator door; then he wondered why he was there. Lack-of-activation errors occur because the presumed mechanism—the "activation" of the goals—has decayed. The less technical but more common term would be "forgetting."

MODE ERRORS

"I had just completed a long run from my university to my home in what I was convinced would be record time. It was dark when I got home, so I could not read the time on my stopwatch. As I walked up and down the street in front of my home, cooling off, I got more and more anxious to see how fast I had run. I then remembered that my watch had a built-in light, operated by the upper right-hand button. Elated, I depressed the button to illuminate the reading, only to read a time of zero seconds. I had forgotten that in stopwatch mode, the same button [that in the normal, time-reading mode would have turned on a light] cleared the time and reset the stopwatch."

Mode errors occur when devices have different modes of operation, and the action appropriate for one mode has different meanings in other modes. Mode errors are inevitable any time equipment is designed to have more possible actions than it has controls or displays, so the controls must do double duty. Mode errors are especially likely where the equipment does not make the mode visible, so the user is expected to remember what mode has been established, sometimes for many hours.

Mode errors are common with digital watches and computer systems (especially text editors). Several accidents in commercial aviation can be attributed to mode errors, especially in the use of the automatic pilots (which have a large number of complex modes).

DETECTING SLIPS

Although slips are relatively easy to detect because there is a clear discrepancy between goal and result, detection can only take place if

there is feedback. If the result of the action is not visible, how can a misaction be detected? Even when a mismatch is noted, the person may not believe that the error occurred. Some trail of the sequence of actions that was performed is valuable.

Even when an error has been detected, it may not be clear what the error was.

"Alice" was driving a van and noticed that the rearview mirror on the passenger side was not adjusted properly. Alice meant to say to the passenger on the right, "Please adjust the mirror," but instead said "Please adjust the window."

The passenger, "Sally," was confused and asked, "What should I do? What do you want?"

Alice repeated the request: "Adjust the window for me."

The situation continued through several frustrating cycles of conversation and attempts by the passenger to understand just what adjustments should be made to the window. The error-correction mechanism adopted by the driver was to repeat the erroneous sentence more and more loudly.

In this example, it was easy to detect that something was wrong but hard to discover what. Alice believed the problem was that she couldn't be understood or heard. She was monitoring the wrong part of the action sequence—she had a problem of level.

Actions can be specified at many different levels. Suppose I were driving my car to the bank. At any given moment, the action being performed could be described at many different levels:

- Driving to the bank
- Turning into the parking lot
- Making a right turn
- Rotating the steering wheel clockwise
- Moving my left hand upward and to the right and my right hand downward
- Increasing the tension on the sternocostal portion of the pectoralis major muscle

All these levels are active at the same time. The most global description (the one at the top of the list), is called the high-level specification. The more detailed descriptions, the ones at the bottom of the list, are called the low-level specifications. Any one of them might be in error.

4/19

It is often possible to detect that the result of an action is not as planned, but then not to know at which level of specification the error has taken place.

Problems of level commonly thwart the correction of error. My collection of slips includes several examples in which a person detects a problem but attempts to correct it at the wrong level.

One frequent example is the nonworking key, reported to me both for cars and homes. Someone goes to his or her car and the key won't work. The first response is to try again, perhaps holding the key more level or straight. Then the key is reversed, tried upside down. When that fails, the key is examined and perhaps another tried in its stead. Then the door is wiggled, shaken, hit. Finally, the person decides that the lock has broken, and walks around the car to try the other door, at which point it is suddenly clear that this is the wrong car.

In all the situations I have examined the error correction mechanism seems to start at the lowest possible level and slowly works its way higher. Whether this is universally true I do not know, but the hypothesis warrants further examination.

DESIGN LESSONS FROM THE STUDY OF SLIPS

Two different kinds of design lessons can be drawn, one for preventing slips before they occur and one for detecting and correcting them when they do occur. In general, the solutions follow directly from the preceding analyses. For example, mode errors are minimized by minimizing modes, or at least by making modes visible.

Cars provide a number of examples of how design relates to error. A variety of fluids are required in the engine compartment of an automobile: engine oil, transmission oil, brake fluid, windshield washer solution, radiator coolant, battery water. Putting the wrong fluid into a reservoir could lead to serious damage or even an accident. Automobile manufacturers try to minimize these errors (a combination of description and mode errors) by making the different compartments look different—using different shapes and different-size openings—and by adding color to the fluids so that they can be distinguished. Here design by and large prevents errors. But, unfortunately, designers seem to prefer to encourage them.

I was in a taxi in Austin, Texas, admiring the large number of new devices in front of the driver. No more simple radio. In its place was a computer display, so that messages from the dispatcher were now printed on the screen. The driver took great delight in demonstrating all the features to me. On the radio transmitter I saw four identical-looking buttons laid out in a row.

"Oh," I said, "you have four different radio channels."

"Nope," he replied, "three. The fourth button resets all the settings. Then it takes me thirty minutes to get everything all set up properly again."

"Hmm," I said, "I bet you hit that every now and then by accident."

"I certainly do," he replied (in his own unprintable words).

In computer systems, it is common to prevent errors by requiring confirmation before a command will be executed, especially when the action will destroy a file. But the request is ill timed; it comes just after the person has initiated the action and is still fully content with the choice. The standard interaction goes something like this:

USER: Remove file "My-most-important-work."

COMPUTER: Are you certain you wish to remove the file "My-most-important-work"?

USER: Yes.

COMPUTER: Are you certain?

USER: Yes, of course.

COMPUTER: The file "My-most-important-work" has been removed.

USER: Oops, damn.

The user has requested deletion of the wrong file but the computer's request for confirmation is unlikely to catch the error; the user is confirming the action, not the file name. Thus asking for confirmation cannot catch all slips. It would be more appropriate to eliminate irreversible actions: in this example, the request to remove a file would be handled by the computer's moving the file to some temporary holding place. Then the user would have time for reconsideration and recovery.

At a research laboratory I once directed, we discovered that people would frequently throw away their records and notes, only to discover the next day that they needed them again. We solved the problem by getting seven trash cans and labeling them with the days of the week.

3/19

Then the trash can labeled Wednesday would be used only on Wednesdays. At the end of the day it was safely stored away and not emptied until the next Tuesday, just before it was to be used again.

People discovered that they kept neater records and books because they no longer hesitated to throw away things that they thought would probably never be used again; they figured it was safe to throw something away, for they still had a week in which to change their minds.

But design is often a tradeoff. We had to make room for the six reserve wastebaskets, and we had a never-ending struggle with the janitorial staff, who kept trying to empty all of the wastebaskets every evening. The users of the computer center came to depend upon the "soft" nature of the wastebaskets and would discard things that they otherwise might have kept for a while longer. When there was an error—sometimes on the part of the janitorial staff, sometimes on our part in cycling the wastebaskets properly—then it was a calamity. When you build an error-tolerant mechanism, people come to rely upon it, so it had better be reliable.

Mistakes as Errors of Thought

Mistakes result from the choice of inappropriate goals. A person makes a poor decision, misclassifies a situation, or fails to take all the relevant factors into account. Many mistakes arise from the vagaries of human thought, often because people tend to rely upon remembered experiences rather than on more systematic analysis. We make decisions based upon what is in our memory; memory is biased toward overgeneralization and overregularization of the commonplace and overemphasis on the discrepant.

SOME MODELS OF HUMAN THOUGHT

Psychologists have chronicled the failures of thought, the nonrationality of real behavior. Even simple tasks can sometimes throw otherwise clever people into disarray. Even though principles of rationality seem as often violated as followed, we still cling to the notion that human thought should be rational, logical, and orderly. Much of law is based upon the concept of rational thought and behavior. Much of economic theory is based upon the model of the rational human who attempts

to optimize personal benefit, utility, or comfort. Many scientists who study artificial intelligence use the mathematics of formal logic—the predicate calculus—as their major tool to simulate thought.

But human thought—and its close relatives, problem solving and planning—seem more rooted in past experience than in logical deduction. Mental life is not neat and orderly. It does not proceed smoothly and gracefully in neat, logical form. Instead, it hops, skips, and jumps its way from idea to idea, tying together things that have no business being put together; forming new creative leaps, new insights and concepts. Human thought is not like logic; it is fundamentally different in kind and in spirit. The difference is neither worse nor better. But it is the difference that leads to creative discovery and to great robustness of behavior.

Thought and memory are closely related, for thought relies heavily upon the experiences of life. Indeed, much problem solving and decision making takes place through attempts to remember some previous experience that can serve as a guide for the present. There have been many theories of human memory. For example, every method of filing things has shown up somewhere along the line as a model for human memory. Do you file photographs neatly in a scrapbook? One theory of memory has postulated that our experiences are neatly encoded and organized, as if in a photo album. This theory is wrong. Human memory is most definitely not like a set of photographs or a tape recording. It mushes things together too much, confuses one event with another, combines different events, and leaves out parts of individual events.

Another theory is based on the filing cabinet model, wherein there are lots of cross references and pointers to other records. This theory has a good deal going for it, and it is probably a reasonable characterization of the most prominent approach today. Of course, it is not called a file cabinet theory. It goes by the names of "schema theory," "frame theory," or sometimes "semantic networks" and "propositional encoding." The individual file folders are defined in the formal structure of the schemas or frames, and the connections and associations among the individual records make the structure into a vast and complex network. The essence of the theory consists of three beliefs, all reasonable and supported by considerable evidence: (1) that there is logic and order to the individual structures (this is what the schema or frame is about); (2) that human memory is associative, with each schema pointing and referring to multiple others to which it is related or that help define the components (thus the term "network"); and (3) that much of our power for deductive thought comes from using the information in one schema to deduce the properties of another (thus the term "propositional encoding").[6] To illustrate the third concept: once I learn that all living animals breathe, I know that any live animal I will ever meet will breathe. I don't have to learn this separately for all animals. We call this the "default value." Unless told otherwise, anything I learn for a general concept applies to all of its instances by default. Default values do not have to apply to everything—I can learn exceptions, such as that all birds fly except for penguins and ostriches. But defaults hold true unless an exception shows otherwise. Deduction is a most useful and powerful property of human memory.

THE CONNECTIONIST APPROACH

We still are a long way from understanding human memory and cognition. Today, in the developing field of cognitive science, two different views are emerging. The traditional view considers thought to be rational, logical, and orderly; this approach uses mathematical logic as the scientific means to explain thought. Adherents of this approach have pioneered the development of schemas as the mechanism of human memory. A newer approach is rooted in the working of the brain itself. Those of us who follow this new approach call it "connectionism," but it also goes under the names of "neural nets," "neural models," and "parallel distributed processing." It is an attempt to model the way in which the brain itself is structured, with billions of brain cells connected into groups, many cells connected to tens of thousands of others, many all working at the same time. This approach follows the rules of thermodynamics more than it does the rules of logic. Connectionism is still tentative, still unproven. I believe that it has the potential to explain much of what puzzled us before, but part of the scientific community thinks that it is fundamentally flawed.[7]

The brain consists of billions of nerve cells—neurons—each connected to thousands of other cells. Each neuron sends simple signals to the neurons to which it is connected, each signal attempting to increase or decrease the activity of its recipient. The connectionist approach to the study of thought mimics these connections. Each connectionist unit is connected to many other units. The signals are either positive in value (called "activation" signals) or negative in value (called "inhibition"). Each unit adds up the total influence of the signals that it receives and then sends along its outward connections a signal whose value is a function of that sum. That's about all there is to it. The elements are all simple: the complexity and power come from the fact that there are a large number of interconnected units trying to influence the activities of the others. All this interconnection leads to massive interaction among the units, with the signals sometimes leading to fights and conflicts, sometimes to cooperation and stability. After a while, however, the system of interconnected units will eventually settle down to a stable state that represents a compromise among the opposing forces.

Thoughts are represented by stable patterns of activity. New thoughts are triggered whenever there is some change in the system, oftentimes because some new information arrives at the senses and changes the pattern of activation and inhibition. We can think of the interactions as the computational part of thought: when one set of units sends signals activating another, this can be interpreted as offering support for a cooperative interpretation of events; when one set of units sends signals suppressing another, it is because the two usually offer competing interpretations. The result of all this support and competition is a compromise: not the correct interpretation, simply one that is as consistent as possible with all possibilities under active consideration. This approach suggests that much of thought results from a kind of pattern matching system, one that forces its solutions to be analogous to past experiences, and one that does not necessarily follow the formal rules of logical inference.

The relaxation of interacting connectionist structures into patterns happens relatively quickly and automatically, below the surface of consciousness. We are conscious only of the end states, not of the means for getting there. As a result, in this view of the mind, our explanations of our own behavior are always suspect, for they amount to stories made up after the fact to explain the thoughts that we already have.

Much of our knowledge is hidden beneath the surface of our minds, inaccessible to conscious inspection. We discover our own knowledge primarily through our actions. We can also find out by testing ourselves, by trying to retrieve examples from memory—self-generated examples. Think of an example, then think of another example. Find a story that explains them. Then we believe that story and call it the reason or explanation for our behavior. The problem is that the story changes dramatically depending upon what examples we select. And the examples we select depend upon a large set of factors, some under our control, some not.

The connectionist approach to memory might also be called the "multiple-exposure" theory of memory.

Suppose, unbeknownst to you, your camera broke so that the film wouldn't wind. Every picture you took went right on top of all the others. If you had taken pictures of different scenes, you might still be able to make out the individual parts. But suppose you had taken a picture of a high-school graduating class, one person at a time. Each person took a turn sitting in the chair in front of the fixed camera; each smiled; each had a picture taken. Afterward, when you developed the film, you would find just one picture, a composite of all those faces. All the individual records would still be there, but on top of one another, difficult to separate out. You'd have the average high-school graduate.

Throw everything into memory on top of one another. That is a crude approximation of the connectionist approach to memory. Actually, things aren't thrown together until after a lot of processing has gone on. And memory isn't really like a multiple exposure. Still, this is not a bad characterization of the connectionist approach.

Consider what happens when two similar events are experienced: they merge together to form a kind of average, a "prototypical event." This prototype governs interpretations and actions related to any other event that seems similar. What happens when something really discrepant occurs? If it is quite different from the prototype, it still manages to maintain its identity when thrown into memory. It stands out by itself.

If there were a thousand similar events, we would tend to remember them as one composite prototype. If there were just one discrepant event, we would remember it, too, for by being discrepant it didn't get smudged up with the rest. But the resulting memory is almost as if there had been only two events: the common one and the discrepant one. The common one is a thousand times more likely, but not to the memory; in memory there are two things, and the discrepant event hardly seems less likely than the everyday one.

So it is with human memory. We mush together details of things that are similar, and give undue weight to the discrepant. We relish discrepant and unusual memories. We remember them, talk about

them, and bias behavior toward them in wholly inappropriate ways.

What has this to do with everyday thought? A lot. Everyday thought seems to be based upon past experiences, upon our ability to retrieve an event from the past and use it to model the present. This event-based reasoning is powerful, yet fundamentally flawed. Because thought is based on what can be recalled, the rare event can predominate. Think about it. Think of your experiences with computers, or VCRs, or home appliances; what probably come to mind are the unusual experiences, things that are discrepant. It doesn't matter that you may have used the device a hundred times successfully—it is the one time you got embarrassed that will come to mind.[6]

The limitations of human thought processes have important implications for everyday activities, and in fact can be called in to distinguish everyday activities from others.

The Structure of Tasks

Everyday activities are conceptually simple. We should be able to do most things without having to think about what we're doing. The simplicity lies in the nature of the structure of the tasks.

WIDE AND DEEP STRUCTURES

Consider the game of chess, an activity that is neither everyday nor simple, at least, not for most of us. When it is my turn to play, I have a number of possible moves. For each of my moves, my opponent has a number of possible responses. And for each of my opponent's responses, I have a number of possible counterresponses. The sequences can be represented on a decision tree, a diagram that in this case takes the current board position as a starting point and shows each of my possible moves, each of the possible countermoves, each possible counter-counter move, each possible counter-counter-counter move, and so on, as deep as time and energy permit. The size of the tree for chess is immense, for the number of choices increases exponentially. Suppose that at each spot there are 8 possible moves. At that spot I must consider 8 initial moves for me, $8 \times 8 = 64$ replies of my opponent, $64 \times 8 = 512$ replies I can make, $512 \times 8 = 4{,}096$ possible replies by my opponent, and then $4{,}096 \times 8 = 32{,}768$ more possibili-

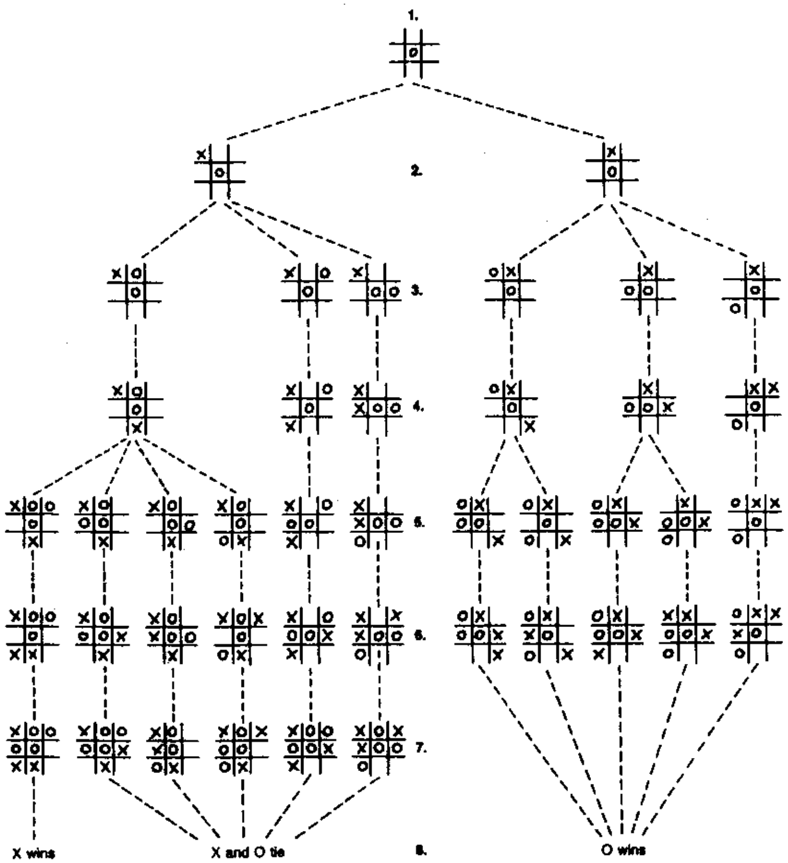

1. The first move by O.
2. Possible first moves by X.
3. Possible second moves by O.
4. Possible second moves by X.
5. Possible third moves by O.
6. Possible third moves by X.
7. Possible fourth moves by O.
8. Outcomes.

5.1 Wide and Deep Decision Tree. The game of tic-tac-toe (naughts and crosses). The tree starts at the top, with the initial state, then deepens as each successive layer considers all the alternative moves by each player. Although this diagram looks a bit complex, it is a pretty simple structure as these things go. First of all, this picture is much simplified. Only one possible first move by O is shown, and the symmetry of the board is used to reduce the number of alternatives being considered. (Only two first moves by X need be considered: the eight possibilities are really equivalent to the two shown because of the symmetry.) In the full game, there are nine possible first moves for O, eight possible replies by X, seven second moves by O, and so on, up to the third move by O, which is the first possible time for the game to be won; there are 15,120 possible sequences up to that point. Even this simple game leads to such a wide and deep decision tree that it is not possible to work out all the possibilities in the head. Expert players take advantage of simple strategies and memorized move sequences. (From *Human Information Processing,* Second Edition, by Peter H. Lindsay and Donald A. Norman, copyright © 1977 by Harcourt Brace Jovanovich, Inc. Reprinted by permission of the publisher.)

ties for me. As you can see, the decision tree gets large rapidly: looking ahead five moves means considering over 30,000 possibilities. The tree is characterized by a vast, spreading network of possibilities. There isn't space here for the decision tree for chess. But even a simple game like tic-tac-toe (or naughts and crosses) has a similar structure, shown in figure 5.1.

That decision tree for chess is even wider and deeper—wide in the sense that at each point in the tree there are many alternatives, so that the tree spreads out over a considerable area; deep in the sense that most branches of the tree go on for a considerable distance.

Everyday activities don't require the kind of complex analyses required for something like chess. In most everyday activities, we need only examine the alternatives and act. Everyday structures are either shallow or narrow.[9]

SHALLOW STRUCTURES

The menu of an ice cream store provides a good example of a shallow structure (figure 5.2). There are many alternative actions, but each is simple; there are few decisions to make after the single top-level choice. The major problem is to decide which action to do. Difficulties arise from competing alternatives, not from any prolonged search, problem solving, or trial and error. In shallow structures, there's no problem of planning or depth of analysis.

NARROW STRUCTURES

A cookbook recipe is a good example of a narrow structure (figure 5.3). A narrow structure arises when there are only a small number of alternatives, perhaps one or two. If each possibility leads to only one or two further choices, then the resulting tree structure can be said to be narrow and deep.

Just as the ice cream store menu is an example of a shallow structure, the multicourse, fixed menu meal can serve as an example of a deep structure. Although there may be many courses, for each course the diner is either automatically served the relevant dish or offered the choice of one or two dishes. The only action required is to accept one or to refuse: no deep thought is required.

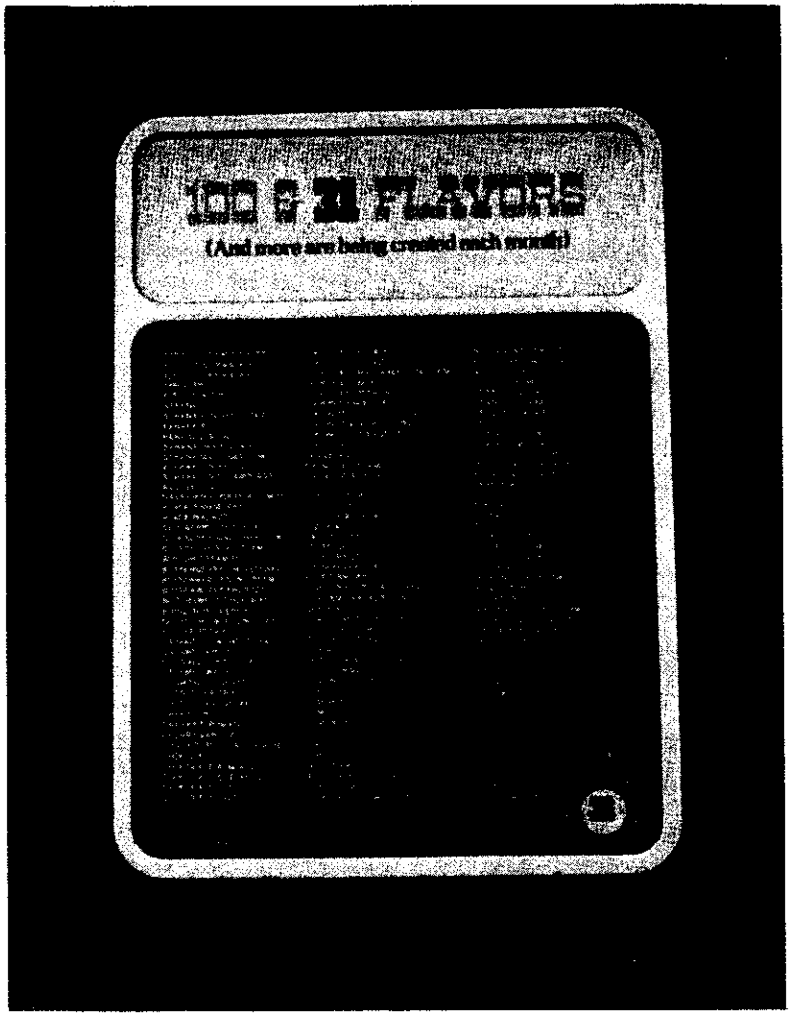

5.2 Wide and Shallow Decision Tree. A lot of alternatives, but after the first decision, few or no further decisions. In this menu from an ice cream parlor, there are many choices, but once the flavor has been chosen, the remaining decisions are simple—what kind of cone, how many scoops, and what kind of topping. (Photograph by the author of a sign at a Baskin-Robbins store.)

Another example is the sequence of steps required to start a car. You must go to the car, select the proper key, insert it in the door lock, turn the key, open the door, remove the key, get into the car, close the door, put on the seatbelt, insert the proper key into the ignition, make sure

jnd Cabrillo Sea Bass

Saute onion and garlic.
Heat 2 bottles of beer to boiling.
Place sea bass in pan.
Pour beer over fish.
Add onion, garlic, and mushrooms.
Add 4 whole garlic cloves, unpeeled.
Add cilantro.
Poach for 10 minutes (approx.)
Remove fish from pan.
Reduce stock over high heat.
Put brown rice on serving plate.
Place fish over rice.
Cover fish with stock and *jnd*Jalepeño sauce.

jnd Jalepeño Sauce

Chop 1 large onion.
Peel and chop 6 tomatillos.
Slice 2 Jalapeño chiles lengthwise.
Peel and quarter 2 tomatoes.
Place onion, tomatillos, chile, and tomatoes in sauce pan.
Cover with 1 cup red wine.
Simmer for 15 minutes to 2 hours (the longer the milder).
Add cilantro.
Serve.

5.3 Deep and Narrow Decision Tree. Few decisions need be made at any level, but to complete the task, many steps (levels) must be followed. This decision structure is characteristic of any task that has a large number of steps, each of which is relatively straightforward. An example is the steps required to follow a recipe, such as my favorite fish recipe.

the car is not in gear, start the engine, and so on. This is a deep structure, but it is narrow. There is a long series of steps, but at each point, there are few, if any, alternatives to consider. Any task that involves a sequence of activities where the action to be done at any point is determined by its place in the sequence is an example of a narrow structure.

The modern superhighway offers the driver a series of exits. The driver either starts on the road with a predetermined exit in mind or else must decide at each exit whether to stay on the road or not. In fact, road designers attempt to linearize and simplify the decision-making tasks of the driver: the relevant information is fed slowly and sequentially to the driver to minimize the mental workload and the need for overlapping processing.

Freeway design is by now a science, with a well-defined set of procedures and with societies, books, and journals devoted to it. Different countries of the world have reached different solutions to the problem of guiding the driver.

A rather complete analysis was done in Britain for the design of the M *series motorways. Each motorway exit has a carefully programmed sequence of six signs. The first precedes the exit by one mile and is* *intended to serve an alerting function, as well as to present route number information. The second precedes the exit by a half mile and gives the major towns reached by the exit (but no route number information). The third precedes the exit by a quarter mile and adds the "forward destination" (where you eventually get to if you don't exit). The fourth sign is at the exit and provides major route numbers and a few town names. The fifth sign is on the motorway beyond the exit; it is intended to play a "confirmatory" role: it displays the forward destinations and their distances. The sixth sign is on the exit ramp, in colors the reverse of all the preceding signs; it shows all the local destinations, usually on a map of the roundabout (traffic circle) found at most exits.*[10]

THE NATURE OF EVERYDAY TASKS

Most tasks of daily life are routine, requiring little thought or planning—things like bathing and dressing, brushing teeth, eating at the table, getting to work, meeting with friends, going to the theater. These are the daily activities that occupy most of our time, and there are many of them. Yet each, by itself, is relatively simple: either shallow or narrow.

What are *not* everyday activities? Those with wide and deep structures, the ones that require considerable conscious planning and thought, deliberate trial and error: trying first this approach, then that—backtracking. Unusual tasks include writing a long document or letter, making a major or complex purchase, computing income tax, planning a special meal, arranging a vacation trip. And don't forget intellectual games: bridge, chess, poker, crossword puzzles, and so on.

The tasks most frequently studied by psychologists are *not* everyday tasks. They are things like chess or algebraic puzzles, which require much thought and effort; but indeed, these pursuits have just the sort of wide and deep structures that do not characterize everyday activities.

In general, we find wide and deep structures in games and leisure activities, where the structure is devised so as to occupy the mind or to make the task deliberately (and artificially) difficult. After all, what challenge would there be if games such as chess or bridge were conceptually simple? How would interest in a mystery novel—or any novel for that matter—be sustained if the plot were straightforward and the

answers readily deducible? Recreational activities *should* be wide and deep, for we do them when we have the time and wish to expend the effort. In the everyday world, we want to get on with the important things of life, not spend our time in deep thought attempting to open a can of food or dial a telephone number.

Everyday activities must usually be done relatively quickly, often simultaneously with other activities. Neither time nor mental resources may be available. As a result, everyday activities structure themselves so as to minimize conscious mental activity, which means they must minimize planning (and especially any planning with extensive looking ahead and backing up) and mental computation. These characteristics restrict everyday tasks to those that are shallow (having no need for extensive looking ahead and backing up) and those that are narrow (having few choices at any point, and therefore requiring little planning). If the structure is shallow, width is not important. If the structure is narrow, depth is not important. In either case, the mental effort required for doing the task is minimized.

Conscious and Subconscious Behavior

Much human behavior is done subconsciously, without conscious awareness and not available to inspection. The exact relationship between conscious and subconscious thought is still under great debate. The resulting scientific puzzles are complex and not easily solved.

Subconscious thought matches patterns. It operates, I believe, by finding the best possible match of one's past experience to the current one. It proceeds rapidly and automatically, without effort. Subconscious processing is one of our strengths. It is good at detecting general trends, at recognizing the relationship between what we now experience and what has happened in the past. And it is good at generalizing, at making predictions about the general trend based on few examples. But subconscious thought can find matches that are inappropriate, or wrong, and it may not distinguish the common from the rare. Subconscious thought is biased toward regularity and structure, and it is limited in formal power. It may not be capable of symbolic manipulation, of careful reasoning through a sequence of steps.

Conscious thought is quite different. It is slow and labored. Here is where we slowly ponder decisions, think through alternatives, compare different choices. Conscious thought ponders first this approach, then that—comparing, rationalizing, finding explanations. Formal logic, mathematics, decision theory: these are the tools of conscious thought. Both conscious and subconscious modes of thought are powerful and essential aspects of human life. Both can provide insightful leaps and creative moments. And both are subject to errors, misconceptions, and failures.

Conscious thought tends to be slow and serial. Conscious processing seems to involve short-term memory and is thereby limited in the amount that can be readily available. Try consciously to solve the children's game called tic-tac-toe or naughts and crosses and you will discover that you can't, not if you try to explore all the alternatives. How can I claim that a trivial children's game cannot be done in the head? Because you don't really play by thinking it through; you play by memorizing the patterns, by transforming the game into something simpler. Try playing the following game:

Start with the nine numbers 1, 2, 3, 4, 5, 6, 7, 8, and 9. You and your opponent alternate turns, each time taking a number. Each number can be taken only once, so if your opponent has selected a number, you cannot also take it. The first person to have any three numbers that total 15 wins the game.

This is a difficult game. You will find it is very hard to play without writing it down. But this game is identical to tic-tac-toe. Why should it be hard if tic-tac-toe is easy?

To see the relationship between the game of 15 and tic-tac-toe, simply arrange the nine digits into the following pattern:

8	*1*	*6*
3	*5*	*7*
4	*9*	*2*

Now you can see the connection: any three numbers that solve the 15 problem also solve tic-tac-toe. And any tic-tac-toe solution is also a solution to 15. So why is one easy and the other hard? Because tic-tac-toe takes advantage of perceptual abilities, and because you simplify tic-tac-toe by changing it in several ways, by taking advantage of symmetries, and by memorizing ("learning") the basic opening moves and their appropriate responses. In the end, unless someone makes a slip, two players will always draw, neither one winning.

The transformations of tic-tac-toe have made a complex task into an everyday one. The everyday version doesn't require much mental effort, it does not require planning and thinking, and it is boring. Which is exactly what everyday tasks ought to be—boring, so that we can put our conscious attention on the important things of life, not the routine.

Conscious thought is severely limited by the small capacity of short-term memory. Five or six items is all that can be kept available at any one moment. But subconscious thought is one of the tools of the conscious mind, and the memory limitation can be overcome if only an appropriate organizational structure can be found. Take fifteen unrelated things and it is not possible to keep them in conscious memory at once. Organize them into a structure and it is easy, for only that one structure has to be kept in conscious memory. As a result of this power of organization to overcome the limits of working memory, explanation and understanding become essential components of conscious thought: with understanding and explanation, the number of things that can be kept consciously in mind expands enormously.

Now consider how mistakes might be made: by mismatch; by taking the current situation and falsely matching it with something in the past. Although we are really good at finding examples from the past to match the present, these examples are biased in one of two ways: toward the regularities of the past—the prototypical situation—or toward the unique, discrepant event. But suppose the current event is different from all that has been experienced before: it is neither common nor unique, it is simply rare. We won't deal well with it: we are apt to classify the rare with either the common or the unique, and either of these choices is wrong. The same powers that make us so good at dealing with the common and the unique lead to severe error with the rare.

EXPLAINING AWAY ERRORS

A reformed thief, telling of his success, put it this way: "I'm telling you . . . if I had a hundred dollars for every time I heard a dog owner tell their dog to 'shut up . . . go lie down,' while I was right outside their window, I'd be a millionaire."[11]

Mistakes, especially when they involve misinterpreting the situation, can take a very long time to be discovered. For one thing, the interpretation is quite reasonable at the time. This is a special problem in a novel situation. The situation may look very much like others we've been in; we tend to confuse the rare event with the frequent one.

How many times have you heard a strange noise while driving your car, only to dismiss it as not relevant, or unimportant? How many times does your dog bark in the night, causing you to get up and yell out, "Be quiet!" And what if the car turns out to be broken, and your mistake has increased the damage? Or there really is a burglar outside, but you've silenced the dog?

This problem is natural. There are lots of things we could pay attention to or worry about; most would be false alarms, irrelevant minor events. At the other extreme, we can ignore everything, rationally explain each apparent anomaly. Hear a noise that sounds like a pistol shot and explain it away: "Must be a car's exhaust backfiring." Hear someone yell out and think, Why can't my neighbors be quiet? Most of the time we are correct. But when we're not, our explanations seem stupid and hard to justify.

When there is a devastating accident, people's explaining away the signs of the impending disaster always seems implausible to others. Afterward, there is a tendency to read about what has taken place and to criticize: "How could those people be so stupid? Fire them. Pass a law against it. Redo the training." Look at the nuclear power accidents. Operators at Three Mile Island made numerous errors and misdiagnoses, but each one was logical and understandable at the time. The nuclear plant disaster at Chernobyl in the Soviet Union was triggered by a well-intentioned attempt to test the safety features of the plant. The actions seemed logical and sensible to the operators at the time, but now their judgments can be seen to have been erroneous.[12]

Explaining away errors is a common problem in commercial accidents. Most major accidents follow a series of breakdowns and errors, problem after problem, each making the next more likely. Seldom does a major accident occur without numerous failures: equipment malfunctions, unusual events, a series of apparently unrelated breakdowns and errors that culminate in major disaster; yet no single step has appeared to be serious. In many of these cases, the people involved noted the problem but explained it away, finding a logical explanation for the otherwise deviant observation.

The contrast in our understanding before and after an event can be dramatic. The psychologist Baruch Fischhoff has studied explanations given in hindsight, where events seem completely obvious and predictable after the fact but completely unpredictable beforehand.[13]

Fischhoff presented people with a number of situations and asked them to predict what would happen: they were correct only at the chance level. He then presented the same situation along with the actual outcome to another group of people, asking them to state how likely the outcome was: when the actual outcome was known, it appeared to be plausible and likely, whereas the others appeared unlikely. When the actual outcome was not known, the various alternatives had quite different plausibility. It is a lot easier to determine what is obvious after it has happened.

SOCIAL PRESSURE AND MISTAKES

A subtle issue that seems to figure in many accidents is social pressure. Although it may not at first seem to be relevant in design, it has strong influence on everyday behavior. In industrial settings social pressures can lead to misinterpretation, mistakes, and accidents. For understanding mistakes, social structure is every bit as essential as physical structure.

Look at airline accidents, not everyday activities for most of us, but subject to the same principles. In 1983, Korean Air flight 007 strayed over the Soviet Union and got shot down, probably because of an error in programming the flight path into the inertial navigation system (INS). Although each checkpoint was discrepant, apparently the deviations were easily explained away if the crew substituted for each point the checkpoint reading for the previous INS point. But there were significant social pressures operating as well.

The crew of flight 007 probably misprogrammed the INS, but the INS couldn't be reprogrammed in flight: if an error were detected the aircraft would have to go back to the original airport, land (jettisoning fuel to get to a safe landing weight), and then reset the INS and take off again—an expensive proposition. Three Korean Air flights had returned to their airport in the six months preceding the flight 007 incident, and the airline had told its pilots that the next pilot who returned would be punished. Was this a factor in the accident? It's hard to know, but the design of the INS sounds badly deficient. The social pressures on the crew not to find (or admit to) an error in the INS were clearly strong. But punishment for following a safety procedure is never wise. The proper approach would be to redesign either the INS's or the procedures for using them.[14]

The real culprit, almost always, is the design. Design that makes it easy to make wrong settings, or to misread an instrument, or to misclassify an event. Design of the social structure that makes false reporting of danger punishable. Turn a nuclear power plant off by mistake and you cost the company hundreds of thousands of dollars; you'll probably lose your job. Fail to turn it off when there is a real incident, and you might lose your life. If you refuse to fly a crowded airliner because the weather looks bad, the company loses lots of money and the passengers get very angry. Take off under those situations and most of the time it works out fine, which encourages risk taking. But every so often there is a disaster.

Tenerife, the Canary Islands, in 1977. A KLM Boeing 747 that was taking off crashed into a Pan American 747 that was taxiing on the runway, killing 583 people. The KLM plane should not have tried to take off then, but the weather was starting to get bad, and the crew had already been delayed for too long (even being on the Canary Islands was a diversion from the scheduled flight—they had to land there because bad weather had prevented them from landing at their scheduled destination); they had not received clearance to take off. And the Pan American flight should not have been on the runway, but there was considerable misunderstanding between the pilots and the air traffic controllers. Furthermore, the fog was coming in so neither plane could see the other.

There were time pressures and economic pressures acting together. The Pan American pilots questioned their orders to taxi on the runway, but they continued anyway. The co-pilot of the KLM flight voiced minor objections to the pilot, suggesting that they were not yet cleared for takeoff. All in all, a tragedy occurred due to a complex mixture of social pressures and logical explaining away of discrepant observations.

The Air Florida flight from National Airport, Washington, D.C., crashed at takeoff into the 14th Street bridge over the Potomac River, killing seventy-eight people, including four who were on the bridge. The plane should not have taken off because there was ice on the wings, but it had already been delayed over an hour and a half; this and other factors "may have predisposed the crew to hurry." The accident occurred despite the first officer's (the co-pilot's) concern: "Although the first officer expressed concern that something 'was not right' to the captain four times during the takeoff, the captain took no action to reject the takeoff." Again we see social pressures coupled with time and economic forces.[15]

Designing for Error

Error is often thought of as something to be avoided or something done by unskilled or unmotivated people. But everyone makes errors. Designers make the mistake of not taking error into account. Inadvertently, they can make it easy to err and difficult or impossible to discover error or to recover from it. Consider the London stock market story that opened this chapter. The system was poorly designed. It should not be possible for one person, with one simple error, to cause such widespread damage. Here is what designers should do:

1. Understand the causes of error and design to minimize those causes.
2. Make it possible to reverse actions—to "undo" them—or make it harder to do what cannot be reversed.
3. Make it easier to discover the errors that do occur, and make them easier to correct.
4. Change the attitude toward errors. Think of an object's user as attempting to do a task, getting there by imperfect approximations. Don't think of the user as making errors; think of the actions as approximations of what is desired.

When someone makes an error, there usually is good reason for it. If it was a mistake, the information available was probably incomplete or misleading. The decision was probably sensible at the time. If it was a slip, it was probably due to poor design or distraction. Errors are usually understandable and logical, once you think through their causes. Don't punish the person for making errors. Don't take offense. But most of all, don't ignore it. Try to design the system to allow for errors. Realize that normal behavior isn't always accurate. Design so that errors are easy to discover and corrections are possible.

HOW TO DEAL WITH ERROR—AND HOW NOT TO

Consider the error of locking your keys into your car. Some cars have made this error much less likely. You simply can't lock the doors (not easily, anyway) except by using the key. So you're pretty much forced to have the keys with you. I call this kind of design a *forcing function.* (More on this topic in the next section.)

In the United States, cars are required to be designed so that if the door is opened while the keys are in the ignition, a warning sound comes on. In theory, if you walk away from your car, leaving the keys in the ignition, the buzzer will call you back. Yet the signal must be ignored as often as it must be attended to. It must be ignored when you open the door of your car while the engine is running so you can hand someone something. On these occasions it is annoying; you know the door is open. And sometimes you want to or need to leave the keys in the car. There goes the buzzer—it can't distinguish deliberate actions from erroneous ones.

Warning signals are usually not the answer. Consider the control room of a nuclear power plant or the cockpit of a commercial aircraft. Thousands of instruments, each designed by someone who thought it was necessary to put in a warning signal for it. Many of the signals sound the same. Most can be ignored anyway because they tell the operator about something that is already known. And when a real emergency happens, all the warning signals seem to go on at once. Each competes with the others to be heard, preventing the person from concentrating upon the problem.[16]

Built-in warning features are bypassed for several reasons. One is that they can go off in error, disrupting perfectly sensible, proper behavior. Another is that they often conflict, and the resulting cacophony is distracting enough to hamper performance. Finally, they are often inconvenient. You can't sit in the car on a warm day, open the door to get some air, and listen to the radio. The key must be in the ignition to make the radio work, but then the door buzzes all the time. So we disconnect those warning signals, tape them over, silence the bell, unscrew the lightbulbs. Warnings and safety methods must be used with care and intelligence, taking into account the tradeoffs for the people who are affected.

FORCING FUNCTIONS

Forcing functions are a form of physical constraint: situations in which the actions are constrained so that failure at one stage prevents the next step from happening. Starting a car has a forcing function associated with it—you must put the ignition key into the ignition switch. Some

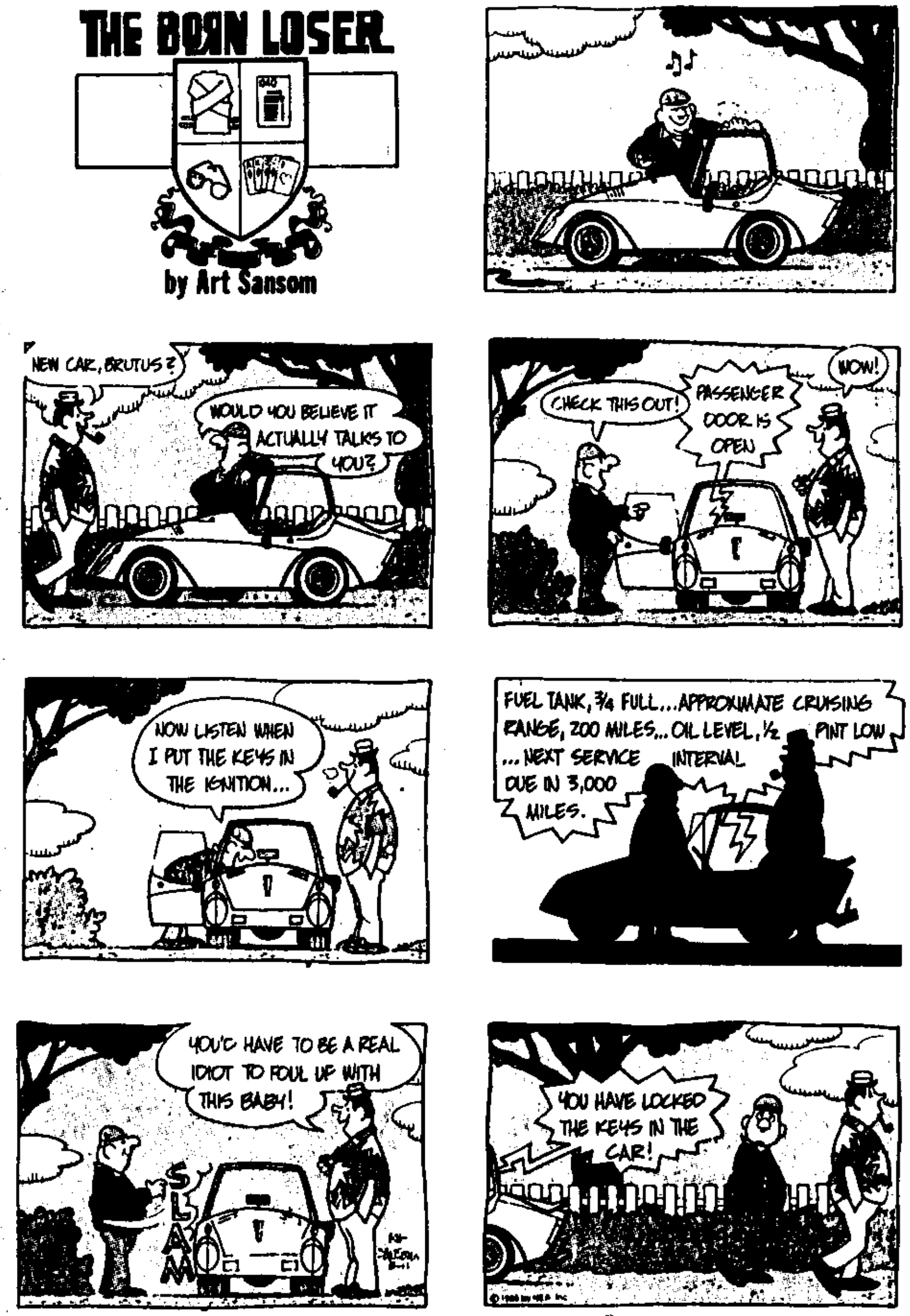

The Born Loser, May 11, 1986. Copyright © 1986 NEA Inc.

time ago, the button that activated the starter motor was separate from the ignition key, so that it was possible to attempt to start the car without the keys; the error was made frequently. In most modern automobiles, the starter switch is activated by turning the key—an effective forcing function that makes you use the key to do the operation.

There is no analogous forcing function for removing the key upon leaving the automobile. As we have already seen, those automobiles that have door locks that can be operated only by a key (from outside the vehicle) do introduce a forcing function: if you want to lock the door you can't leave the key in the car. If a forcing function is really desired, it is usually possible to find one, although at some cost for normal behavior. It is important to think through the implications of that cost—to decide whether people will deliberately disable the forcing function.

The history of seatbelts in autos provides a good example. Despite all the evidence that seatbelts are an effective means of saving lives, some people dislike them enough that they refuse to wear them, probably because the perceived risk is so much less than the actual, statistical risk. For a short period, the United States tried a forcing function on seatbelts: a special interlock was installed on each new car. If the driver's and passengers' belts were not fastened, the car would not start (and a buzzer would sound). This forcing function was so disliked that most drivers had their mechanics disconnect it. The law was quickly changed.

There seemed to be three problems. First, many people did not want to wear seatbelts, and they resented the mechanical forcing function. Second, the forcing function couldn't distinguish legitimate cases in which the seatbelt should not be buckled from illegitimate ones. Thus, if you wanted to carry a package in the passenger's seat, the weight-sensing element in the seat registered a person, so the car wouldn't start unless the passenger seat's buckle was fastened. Third, the mechanisms were not reliable, so they often failed—buzzing, stopping the engine, and being an overall nuisance. Those people who couldn't figure out how to disconnect the forcing function simply buckled the belts permanently, fastening the buckle when the seat was unoccupied and stuffing it under the seat. So if a passenger really wanted to use the belt, it couldn't be done. Moral: it isn't easy to force unwanted behavior upon people. And if you are going to use a forcing function, make sure it works right, is reliable, and distinguishes legitimate violations from illegitimate ones.

Forcing functions are the extreme case of strong constraints that make it easy to discover erroneous behavior. Not every situation allows such strong constraints to operate, but the general principle can be extended to a wide variety of situations. In the field of safety engineering, forcing functions show up under other names, in particular as specialized methods for the prevention of accidents. Three such methods are *interlocks, lockins,* and *lockouts.*

An *interlock* forces operations to take place in proper sequence (figure 5.4). Microwave ovens and television sets use interlocks as forcing functions to prevent people from opening the door of the oven or taking off the back of the television set without first turning off the electric power: the interlock disconnects the power the instant the door is opened or the back removed. The pin on a fire extinguisher or hand grenade and the safety on a rifle are other examples of interlocks; these forcing functions prevent the accidental use of the devices.

A *lockin* keeps an operation active, preventing someone from prematurely stopping it. The sad stories of those who turn off word processors without first saving their work could be avoided with the use of a lockin. Suppose the on-off switch were a "soft" switch, not really disconnecting the power, but sending a signal to the program to quit, checking that all files had been saved, and then, after all the appropriate housekeeping operations had been completed, turning off the power. (Of course, a normal power switch should also exist as an override for special situations or for when a software problem causes the soft switch to fail.)

A *lockout* device is one that prevents someone from entering a place that is dangerous, or prevents an event from occurring. A good example of a lockout occurs in stairways of public buildings, at least in the United States (figure 5.5). In cases of fire, people have a tendency to flee in panic, down the stairs, down, down, down, past the ground floor and into the basement, where they are trapped. The solution (required by the fire laws) is not to allow simple passage from the ground floor to the basement.

In the building in which I work, at the ground floor the stairs seem to end, leading directly to the building's exit door. To go down further requires finding a different door, opening it, and proceeding down the

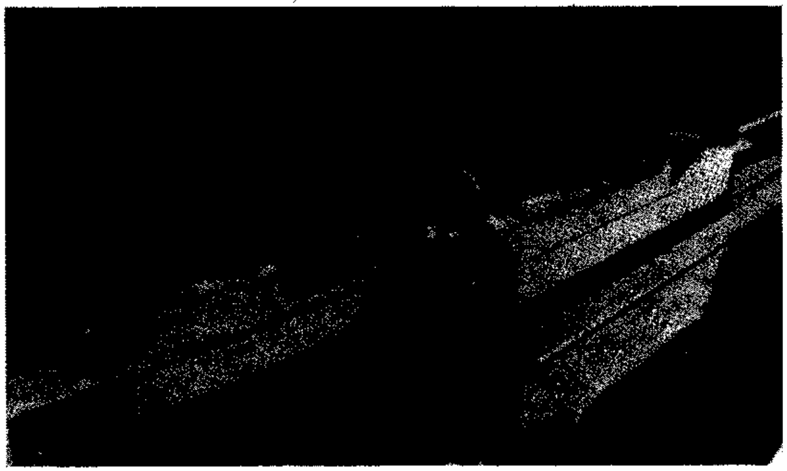

5.4 Use of an Interlock. The Nissan Stanza van was constructed with the access door for its fuel tank right in the path of the sliding passenger door (above). It could be dangerous for the door to be opened while someone was fueling the car. To overcome the problem, Nissan added a forcing function, a bar that prevented the sliding door from opening whenever the fuel tank was being filled. The bar is constructed in the form of an interlock: the cap to the fuel tank cannot be removed unless the bar is moved to its safety position (below). Furthermore, the fuel door cannot be shut again unless the bar is returned to its normal position. Finally, warning signals were added, so that if someone attempts to open the door during fueling, a buzzer sounds. All in all, a lot of effort was put into these forcing functions—which were needed only because of an unfortunate placement of the fuel tank access in the first place.

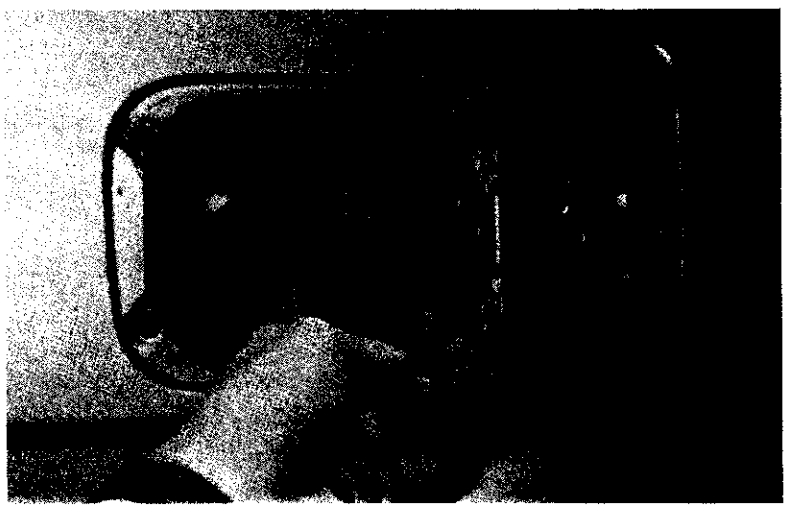

5.5 Lockout. A form of forcing function that prevents people from going down the stairs, past the ground floor, and into the basement. Although in normal times this is a nuisance, in times of fire, when people flee down the stairs in panic, the forcing function can save lives by preventing a mad dash into the basement. The bar encourages people to stop at the ground floor and leave the building.

stairs. This safety feature is usually a nuisance: we have never had a fire, yet I frequently must go from a higher floor into the basement. It's a minor nuisance, however, and it is worth the cost if it can save lives when there is a fire.

Forcing functions almost always are a nuisance in normal usage. The clever designer has to minimize the nuisance value while retaining the safety, forcing-function mechanism, to guard against the occasional tragedy.

There are other useful devices that make use of a forcing function. In some public restrooms there's a package shelf inconveniently placed on the wall just behind the cubicle door, held in a vertical position by a spring. You lower the shelf to the horizontal position, and the weight of a package keeps it there. Why not supply a permanent shelf, always horizontal, placed so that it wouldn't interfere with the opening of the door? There is room. A little thought reveals the answer: the shelf's position is a forcing function. When the shelf is lowered, it blocks the door. So to get out of the cubicle, you have to remove whatever is on the shelf and raise it out of the way. And that forces you to remember your packages. Clever design.

It is common to forget items. Examples spring readily to mind:

- Making copies of a document, but leaving the original inside the machine and walking off with only the copy.
- Using a bank or credit card to withdraw money from an automatic teller machine, then walking off without the card. This was a frequent enough error that many machines now have a forcing function: you must remove the card before the money will be delivered. Of course, you then can walk off without your money, but this is less likely than forgetting the card because money is the goal of using the machine. The possibility exists so the forcing function isn't perfect.
- Leaving a child behind at a rest stop during a car trip. I also heard about a new mother who left her infant in the dressing room of a department store.
- Losing a pen because it was taken out to write a note or a check in some public location, then put down for a moment while doing some other task—such as giving the check to the salesperson. The pen is forgotten in the activities of putting away the checkbook, picking up the goods, talking to the salesperson or friends, and so on. Or the reverse: borrowing a pen, using it, and then putting it away in pocket or purse, even though it is someone else's; this slip is an example of a capture error.

Forcing functions don't always show up where they should. Sometimes their absence causes all sorts of unnecessary confusion. Read the caution statement from the game instructions shown in figure 5.6.

All those exclamation marks! And the caution is repeated throughout the instruction manual. It won't do any good. The Nintendo Entertainment System is meant to be used by children. The instruction manual probably won't be around. Even if it is, a group of active children, anxious to try a different game, won't bother with it. I watched my own child follow the instruction faithfully for several days, then fail when asked to stop playing and come to dinner. I forgot on the few attempts I made to master the game. The only possible virtue of the warning is to protect the manufacturer: when the children repeatedly burn out the electronic circuits, the company can disclaim liability, asserting that the children violated the instructions.

4. OPERATING YOUR NES

●TO START PLAY

1. Turn your television on to Channel 3.

   Note: If Channel 3 is broadcasting in your area and interfering with the game, set the switch on the back of the Control Deck to Channel 4.

2. If your TV has an automatic fine tuning control (AFC), turn it off. (Use the manual fine tune dial to adjust the picture after inserting the game pak as described below.)

   Note: If you have a color TV that turns black and white when the AFC is turned off, leave the AFC on.

3. Make sure that the power switch on the Control Deck is off.

CAUTION !! ALWAYS MAKE SURE THAT THE POWER SWITCH ON THE CONTROL DECK IS OFF BEFORE INSERTING OR REMOVING A GAME PAK !!

4. Open the Chamber Lid on the Control Deck. | Insert a Game Pak into the Chamber (Label Facing up) and Push it all the way in. | Press Down on the Game Pak until it locks into place and close the Chamber Lid.

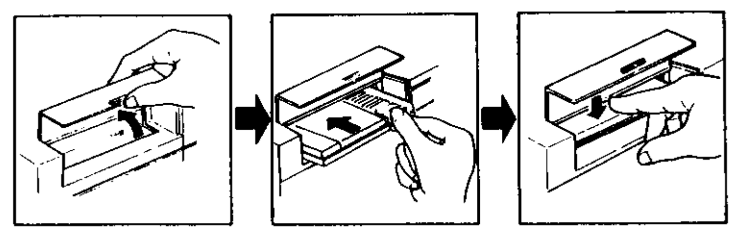

5.6 The Nintendo Children's Toy. This home video game set is intended for use by children. However, it has a complex safety instruction, one almost guaranteed to be ignored. To use the system, one inserts a "game pak" cartridge into the "chamber." The power switch should be off when inserting or removing the cartridge. In the absence of any forcing function, the instruction is almost universally disregarded (if anyone even knows about it). If order is important, there should be a forcing function. If order does not matter, the instruction should be dropped. (From the Nintendo instruction manual. Nintendo® and Nintendo Entertainment System® are trademarks of Nintendo of America Inc. © 1986 Nintendo.)

Proper design calls for a forcing function here. There are several viable schemes. The cover over the game pack compartment could control an interlock, so that it automatically turned off the power whenever it was opened. Or the power switch could move a lever blocking the top of the game pack compartment, so that the packs could not be removed or inserted unless the lever were out of the way, turning off the power. There are other possibilities. My point is, of course, that the design should have included one; without the forcing function, failure to heed the warning is almost guaranteed.

A Design Philosophy

There are lots of ways for a designer to deal with errors.[17] The critical thing, however, is to approach the topic with the proper philosophy. The designer shouldn't think of a simple dichotomy between errors and correct behavior; rather, the entire interaction should be treated as a cooperative endeavor between person and machine, one in which misconceptions can arise on either side. This philosophy is much easier to implement on something like a computer which has the ability to make decisions on its own than on things like doors and power plants, which do not have such intelligence. But the philosophy of user-centered system design still holds. Think of the user's point of view. Assume that every possible mishap will happen, so protect against it. Make actions reversible. Try to make them less costly. All the required principles have been thoroughly discussed in this book.

- Put the required knowledge in the world. Don't require all the knowledge to be in the head. Yet do allow for more efficient operation when the user has learned the operations, has gotten the knowledge in the head.
- Use the power of natural and artificial constraints: physical, logical, semantic, and cultural. Use forcing functions and natural mappings.
- Narrow the gulfs of execution and evaluation. Make things visible, both for execution and evaluation. On the execution side, make the options readily available. On the evaluation side, make the results of each action apparent. Make it possible to determine the system state readily, easily, and accurately, and in a form consistent with the person's goals, intentions, and expectations.

THE DESIGN CHALLENGE

They began work at once, and by the next September the first [typewriter] machine was finished, and letters were written with it. It worked successfully so far as to write rapidly and correctly, but trial and experience showed it to be far short of an acceptable, practicable writing machine. . . .

One device after another was conceived and developed till twenty-five or thirty experimental instruments were made, each succeeding one a little different from and a little better than the one preceding. They were put into the hands of stenographers, practical persons who were presumed to know better than anyone else what would be needed and satisfactory. Of these, James O. Clephane, of Washington, D.C., was one. He tried the instruments as no one else had tried them; he destroyed them, one after another, as fast as they could be made and sent him, till the patience of Mr. Sholes [the inventor] was exhausted. But Mr. Densmore insisted that this was the very salvation of the enterprise; that it showed the weak spots and defects, and that the machine must be made so that anybody could use it, or all efforts might as well be abandoned; that such a test was a blessing and not a misfortune, for which the enterprise should be thankful.[1]