5 Error Reduction as a Systems Problem

Neville Moray
University of Illinois at Urbana-Champaign

As other chapters in this book make clear, error is rife in medical systems, many errors are hazardous, and it is clearly desirable to reduce their occurrence and impact. In any industry, organization, or setting where error is common, there are many ways in which error reduction may be approached. One approach assumes that if people are more careful, pay more attention, and in general take more trouble over what they are doing, then errors can be reduced and their effects mitigated. This approach tends to put great emphasis on the psychology of the individual who makes the error and on training, admonition, supervision, and ever-tighter and more detailed rules, with an implication of blame attached to those who make errors. In contrast to this is an approach that sees relatively few errors as being the fault of the human who commits them, and even fewer as being blameworthy. Rather, one sees the *design* of objects, activities, procedures, and patterns of behavior as being the source of errors. This approach, epitomized by writers such as Norman (1981, 1988) and Reason (1990), emphasizes that people of good intention, skilled and experienced, may nonetheless be forced to commit errors by the way in which the design of their environment calls forth their behavior. One need not deny that people make errors because of fatigue, carelessness, or lack of training in order to espouse this approach to error. However, the fundamental claim is that *the systems of which humans are a part call forth errors from humans*, not the other way around. Only as an attribution of last resort, or because of the tendency of a legal system to be less concerned with justice than with economics, does one ascribe blame to an individual who commits an error.

This chapter begins by noting that complex systems such as health-care delivery are composed of a series of hierarchically organized subsystems. These include equipment, individuals, teams and groups, and organizations. The way in which the different levels may cause errors is described from a psychological standpoint. These causes are then discussed in the more general context of a system of constraints, and strategies to reduce errors are seen as a search for ways to alter constraints. Finally, the question is raised whether the psychology of the individual is the best point at which to try to reduce error, or whether a different philosophy is needed, based on the assumption that it is never possible to reduce the probability of error to zero. The chapter concludes with an argument for the use of an interdisciplinary approach to error management within the context of a chosen social philosophy.

There are many cases in which it is easy to observe how a system elicits human error. When North Americans drive a car for the first time in England or Japan, the habits of a lifetime tend to make drivers behave in stereotypical ways. They tend to glance to the right when looking for the image in the rear view mirror, turn to the right to move closer to the near side verge and expect to be overtaken on the left. All these behaviors are incorrect in those countries and predispose the well-intentioned driver to make errors. When people buy a kitchen range in which the relations between the positions of the burners and controls are other than those they have learned habitually to expect, the probability that they will turn on the wrong burner can be increased by a factor of 10 or more (Grandjean, 1980). Many similar examples will be found in Norman (1988).

In the context of medical error, it is easy to find similar cases. The labels on bottles may be very similar, the shapes of bottles may be confusing, the controls on equipment may violate expected stereotypes, or there may be a lack of labeling on connections in equipment that allows potentially lethal interconnections to be made. The layout of operating rooms may require wires to run across the floor in ways that make it likely that people will trip over them or accidentally pull them out of the socket. Notoriously, the handwriting of physicians may verge on the illegible, increasing the likelihood of error in interpreting orders or filling prescriptions.

Errors that arise from such sources are not caused by a lack of good intent, nor commonly from carelessness, and in many cases the nature of the psychological mechanisms that underlie them are well understood. The existence of a phenomenon such as the speed–accuracy trade-off, in which forcing people to act rapidly increases the probability of error, is well known. The very strong effect of stimulus–response stereotypes, strong expectations in a culture about the directions in which knobs turn, switches move, and so on is well documented. The relation between the architectural layout of a room, the equipment therein, and the likelihood that their combination will make people stumble, trip over cables, or otherwise behave clumsily is clear. There are strong models for the effect of display parameters on the probability of correct identification of displayed information.

In all of these cases, it is relatively easy to say how the design of the object or system should be altered to reduce error. Indeed, one can go further. Enough is now known about the relation between ergonomic design and the way that error is caused by ignoring that knowledge to make it certain that almost all errors involving these aspects of a system should be laid directly at the door of designers and manufacturers. It is trivially easy to come by the relevant information that is needed to reduce error in the design of equipment to be used by humans. This is the field of the *ergonomics* of equipment design. Anyone who manufactures or installs equipment that violates the published data on these matters is directly responsible for the vast majority of the errors that arise in its use. Errors arising from the violations of well-known design rules are the responsibility of the designer, not of the user.

On the other hand, one of the major discoveries of recent years has been that even when all that is known about ergonomics is applied to design, the probability of error cannot be reduced to zero. Moreover, further research of this type will never reduce error to zero. There are factors at work in a complex human–machine system that have far greater potency for causing errors than do ergonomic factors. Such factors lead to the notion of *systems design* rather than equipment design and into an area that has been relatively poorly studied.

The most that can be done in this chapter is to indicate some of the factors that are relevant, and what kinds of approaches will tend to produce safer systems. Few designers have been trained to take such factors into account and, more importantly, they *cannot* be taken into account by designers when design is done piecewise. This is usually the case in medical systems. The invention of a new piece of equipment, its realization in terms of physical hardware, and its purchase and installation are all done independently and separately from the design of a system, such as a hospital or other health-care delivery system, as a whole. The situation is little different in practice whether the entire hospital or health-care system is regarded as a single system, or whether some parts are regarded as subsystems of a larger whole. The particular opportunities for error will be different, but the interaction of system and component will be very similar.

TOWARD A SYSTEMS APPROACH

Figure 5.1 is a representation of the causal structure of a complex hierarchical human–machine system based on an analysis by Gaines and Moray (1985). It is quite general, so there are few systems that cannot be mapped

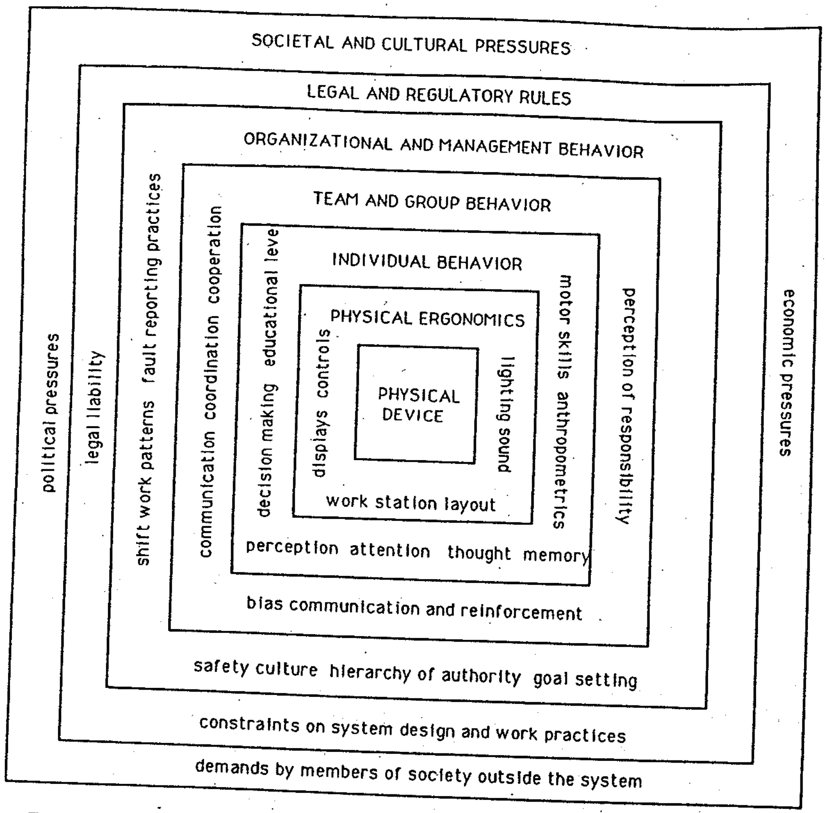

FIG. 5.1. A generic hierarchical systems oriented approach to design and analysis. The terms in upper case define levels of description. The terms in lower case describe typical variables relevant to each level of description.

onto such a diagram. Note that the term "human–machine system" is interpreted in a far wider sense than is usual. It is commonplace to use the term to talk about such things as a vehicle, an aircraft, or even a factory or power station where the relation between the human and the machine is what engineers call *tightly coupled.* A human–machine system is usually thought of as something where there is a piece of equipment that the operator controls directly or through the use of an automatic controller.

It is not usual to think of a bureaucratic organization as a human–machine system, but it is a goal of this chapter to extend the term to include all such systems. A *system* , in this sense, is any collection of components and

the relations between them, whether the components are human or not, when the components have been brought together for a well-defined goal or purpose. Thus, the management of a health-care system includes human components, such as doctors, nurses, and managers; hardware components such as computers and telephones that transmit and store information, paper and magnetic records, drugs, operating theaters, scalpels, and beds; the management *policies* that are adopted; and the financial mechanisms in place to govern the economic control of the system.

Only when the entire *system* is designed correctly will error be minimized. The components of the system must not be merely correctly designed and chosen, but the relations between the components must also be part of the design, as must the rules for its operation. If, for example, standard operating procedures are written without reference to the particular choice and layout of equipment, without reference to the training or social organization of the users, and without reference to maintenance practices and manning levels, then the system will be accident prone (as the accident at Three Mile Island proved). These aspects of system design must be integrated. Error will even then not be eliminated. It will, however, be reduced, and the effects of errors rendered more manageable.

Figure 5.1 is designed so that the causes and effects of errors and the steps needed to reduce and control them are most local in scope in the inner regions of the diagram. As we go outward, the scope of causal variables becomes increasingly global. Local causes of error are usually more readily manageable than global causes, but there is a sense in which global causes, if they can be controlled, have greater payoffs, because the effect of action at a global level will be far more pervasive throughout the system than will local intervention.

Physical Ergonomics

At the center of the system lies the physical design of equipment and the immediate work environment. In the context of health-care systems, this includes such things as:

- The legibility of labels on bottles.
- The confusability of bottles due to similar shapes and colors.
- Noise levels in working environments that may cause messages to be misunderstood.
- The position of displays and controls on equipment that may make it difficult for the health-care provider to read the former while operating the latter.

- The height of beds.
- The position of call buttons for the use of patients.

Anthropometrics, by which the dimensions of equipment are matched to the anatomical dimensions of the users, should also be included. Particularly in an ethnically heterogeneous society such as the United States, anthropometric considerations can be critical. Differing heights and arm lengths can mean a difference of as much as 30 cm in positions that can be reached by individuals sitting at a given console. Similar problems are the spacing of controls and the size of keyboards and buttons. There are also more subtle problems. It is common to make use of color coding in displays and controls, and increasingly, as computer graphics displays are built into automated or semi-automated equipment, colored displays are being used. But the naive use of color is not necessarily helpful, and a substantial proportion of people, particularly males, are color-blind to a greater or lesser degree. When color coding is used in a display, it is essential that color-blindness be considered and some other dimension, such as shape or brightness, be used to support the encoding of information in displays. It is left as an exercise to the reader to think of other such problems in ergonomics.

Almost every problem at this level of design can be solved with existing data. Books such as Grandjean (1980), Sanders and McCormick (1993), and Boff and Lincoln (1988); tools such as MANNEQUIN, available from Humancad (Melville, NY); databases such as ERGOBASE, available from Biomechanics Corporation of America (Deer Park, NY); or general texts such as that of Norman (1988) all provide suitable sources of information for the designer. It should be noted that the increasingly international nature of science and medicine can give rise to error prone systems. For example, the United States is one of a relatively small number of countries in which the position of a switch is coded so that UP = ON. The majority of countries that the writer has visited use DOWN = ON. Hence, equipment purchased from one country may have control and display conventions that violate the expectations of a user in another country—a potent source of error. Such a violation of expectation is usually easy to change by changing the physical orientation of a switch.

More difficult to change, and potentially more dangerous, is the fact that the color coding of electrical wiring differs widely between countries. It is not unknown for a technician unfamiliar with the coding conventions in one country to make a mistaken assumption when wiring a plug to carry main electrical supplies. (American readers may like to ask themselves how they would wire a cable that has one brown wire, one white wire and one yellow-and-green wire, a recent European standard.) But knowledge exists to avoid the vast majority of problems at the level of physical ergonomics of equipment design and workplace layout. There is no excuse for designers failing to make use of this knowledge. The appropriate use of existing data can solve most problems of local scope.

Individual Behavior

With the growing attention to human error in the last 20 years, there has been increasing study of the way in which errors arise at the level of individual behavior. General treatments of error in the context of the psychology of the individual will be found in Reason (1990), Norman (1981, 1988), Senders and Moray (1991), and Rasmussen and Batstone (1989). These are descriptions of the psychology of the individual, of single operators processing information and making decisions on their own.

An important distinction in the errors of individuals is between slips, mistakes, and violations (Norman, 1981; Reason, 1990). *Slips* are defined as errors where the person correctly assesses what needs to be done, and acts accordingly, but slips in carrying out the intention. A common example is typing errors, but the classical slip of a surgeon's knife is of course another, as is a slip of the pen that leads to the writing of an incorrect drug name or dose when the physician has formed a correct intention but fails to carry it out.

A *mistake*, by contrast, is an error in which the person fails to form the correct intention as to what act to perform, and so performs an action which, judged objectively, is incorrect. Note that the word *objectively* may be misleading. People may fail to decide on the *correct* action because the information they receive is inadequate or unclear. A decision based on the best information available may seem to be in error if the person taking it lacks training or experience as to what action is required. There may be a failure of memory as to what to do given the situation that has arisen. Although there may be a better solution available to someone who has perfect information and perfect training, it does not follow that the decision actually made is in error. That decision may be the best possible under the circumstances. Insofar as results fall short of what was required, it is a fault of the system for failing to support adequately the person required to make the decision.

A *violation* is usually defined as a deliberate choice to behave in a nonstandard way, such as when violating normal procedures. Violations may occur because a person believes that a nonstandard procedure provides a better chance of success than the standard procedure, or can in special cases be an example of deliberate malice. Violations can, however, also occur unconsciously if a strong pattern of behavior has been formed by habitual conscious violation so that the behavior ultimately becomes virtually automatic.

Violations and their place in the natural history of error are further discussed later in this chapter. For now, attention is focused on what slips and mistakes tell us about the origins of error. Note that even if all the design problems at the level of physical ergonomics were to be solved, slips and mistakes could still occur. Even with the best possible equipment, situations will occur where the information available is insufficient to lead to an unambiguous diagnosis of the situation, or where memory fails to guide the decision maker to choose the correct action.

Two major psychological factors affect the probability of error. The first factor is that people tend to avoid reasoning their way to solutions, and prefer to pattern match. When pattern matching, people decide that a present situation is identical to one that has occurred before and that it more or less resembles. The second mechanism is that given uncertainty as to what action to take, people will choose one that has worked before, and the more often they have successfully used a particular action, the more likely they are to choose it. These two mechanisms have been called *similarity matching* and *frequency gambling* (Reason, 1990). They appear to be pervasive, and are very strong causes of errors.

The psychology of similarity matching and frequency gambling is fairly well understood, and the implications are obvious. It is rather difficult for systems designers to find practical advice on how to counteract them, although again the standard texts mentioned earlier often contain some discussion. Appropriate counterstrategies can be based on an understanding of the psychological mechanisms involved. For example, we can think of an observer receiving information as being faced with the decision matrix shown in Fig. 5.2.

The target that is hit or missed may be an alarm, a shadow of a tumor, the name on a label, or in general anything, either perceptual or in memory, where a decision is required as to whether an event has really occurred. Usually misses and false alarms are thought of as errors, and there is a strong theory, the theory of signal detection, that describes what factors cause these errors. For example, a poor signal-to-noise ratio obviously limits detectability, reducing hits and true negatives. If an observer receives a signal that is ambiguous, whether a spoken message, a reading on an instrument, or an X-ray of a possible tumor, the ambiguity is resolved by the brain by taking into account the a priori probability of the event and the costs and payoffs associated with alternative outcomes.

All else being equal, an ambiguous event will be perceived as being the event with the highest subjective probability, and ambiguity will tend to be resolved in favor of an outcome that has the best expected payoff in terms of

JUDGEMENT BY OBSERVER	STATE OF THE WORLD: event actually occurred	STATE OF THE WORLD: event did not occur
event occurred	HIT TRUE POSITIVE CORRECT IDENTIFICATION	FALSE ALARM FALSE POSITIVE
event did not occur	MISS	TRUE NEGATIVE CORRECT REJECTION

FIG. 5.2. Decision matrix for signal detection theory of decision making.

the costs and values involved. If a person expects many targets to occur, most of their responses will be in the upper row of Fig. 5.2, whereas if targets are expected to be rare, most of the responses will be in the lower row. If false alarms have unimportant consequences, and detections are valuable, then again, most of the responses will be in the upper row, and in the lower row for the opposite situation. Note that which row the observer favors is dependent on the subjective expectations and subjective estimation of payoff, not merely on the detectability of the signal. It must be emphasized that it is not a matter of the observer failing to take enough trouble over the decision. The content of perception is itself governed by probabilities and payoffs, not just the conscious decision as to what action to take (Swets & Pickett, 1982).

There is an intimate relation between correct and erroneous behavior in situations to which the theory of signal detection applies. For a given physical situation, the percentages of hits (correct diagnosis of the state of events) can only be increased if accompanied by an increase in false alarms, whereas false alarms can only be reduced by reducing hits unless the properties of the displayed information are altered. Thus, the probability of correct behavior can only be increased at the cost of errors. Suppose we are faced with a situation in which the observer is missing too many signals. The theory of signal detection states that the probability of hits can be increased by changing the observer's expectation or by changing the payoffs associated with each of the cells in the matrix. But if hits increase for a given signal-to-noise ratio, the false alarm ratio will also *inevitably* increase. (See Swets & Pickett, 1982, for a discussion in a medical context.) If a low false alarm rate together with a high hit rate is required, then the only way to reduce misses is to redesign the system so as to increase the signal-to-noise ratio. The lighting may be altered, a new instrument that is more sensitive may be introduced, the legibility of a display or a document may be increased, and so forth. Admonition, training, and the like will not alone bring about the desired ratio of hits to false alarms when the signal is weak.

Some errors are due not to decision making but to the properties of memory. If the nature of the errors is such as to suggest that fallible memory is the cause of mistakes, then steps will have to be taken in light of what is known about the psychology of memory. This may involve an increase in the quality and frequency of training, provision of some form of decision aid whether computerized or traditional hard copy, or providing the person with access to the assistance of others. (Note that this involves a move from the psychology of the individual to the psychology of teams and groups.)

Many aspects of attention and decision making are well understood. For example, a basic fact about visual attention is that people need to look directly at a display to read highly detailed information. The area of maximum acuity and color vision is only about 1.5° diameter, and eye movements do not occur more frequently than about 2 per second in most tasks outside the laboratory (Moray, 1984, 1986). Hence, the rate at which an observer can monitor a complex visual display is extremely limited. It is also known that it is very difficult for observers to remain alert for periods of more than 10 to 20 min when keeping watch for rare events. The interaction of attention with fatigue is known to produce a kind of tunnel vision.

Among the most important aspects of decision making, particularly in situations where an observer is trying to decide what has happened in an unusual situation, is confirmation bias, a tendency to select a likely hypothesis quickly and thereafter accept only evidence supporting that hypothesis. This applies not only to medical diagnosis, but to the diagnosis of almost all situations, from equipment malfunction to understanding unexpected changes in a balance sheet or unexpected changes in the pattern of work by a team or a group of individuals. (For important discussions of decision making, see Sage, 1981, and Klein, Orasanu, Calderwood, & Zsambok, 1993).

Errors are made by individuals. The properties of a system defined by PHYSICAL DEVICE and PHYSICAL ERGONOMICS in Fig. 5.1 determine such things as the number of eye movements required to carry out a task, the signal-to-noise ratio of a display, how well memory is supported, and so on. That is, there is a clear causal path outward from the core of Fig. 5.1 toward INDIVIDUAL BEHAVIOR. We shall now see that there are equally clear causal paths inward from the outer layers of the system.

Team and Group Behavior

Although errors are made by individuals, most work situations are such that a person is a member of a group or team either directly or indirectly. The distinction between teams and groups is a matter of the degree of formal allocation of roles. Where several people carry out a single task and each has a specific role, the collection of people is a team. Where the grouping is transitory and informal, it is a group. Thus passengers in a train or aircraft are a group, whereas the crew form a team. In the medical context, the professionals in an operating theater are a team, each with well-defined roles, whereas patients in a waiting room constitute a group.

A common problem and cause of error seems to be the way in which the structure of a team dissolves into an informal group under certain circumstances, undermining the formal patterns of authority and responsibility and losing the cohesion and mutual support that are characteristic of a well-integrated team. On the other hand, situations sometimes arise where individuals are thrown together unexpectedly to operate as a team. In such cases, the advantages of a coherent team are not likely to be found, and indeed, the constraint of a team setting may itself give rise to errors if those involved are unused to working together.

Human reliability is the obverse of the tendency of humans to make errors. Research on reliability has drawn attention to several problems in teams and groups. In order to make nonhuman hardware systems more reliable, it is common to make use of redundancy. For example, if the failure rate of a component in a hardware system is higher than can be tolerated, the system can be made more reliable by incorporating several of the same components in parallel. If it is known that the probability that a particular motor will malfunction is 0.01 per hr, then by having three of them in parallel, the probability that all three will fail at the same time is 0.000001. (This assumes, of course, that the causes of malfunction are random events.) Attempts to use multiple humans in this way will, in general, fail.

The social dynamics of a group of people performing the same task are unpredictable. If people believe that another will check on whether they have performed a task correctly, the accuracy of performance may increase or decrease. It may increase if the performer believes that the second person will check to see if the work has been done poorly and will penalize the performer for bad work or reward him or her for good work. The quality of work may well decrease if the relation between two people is such that each thinks that the other will catch any imperfections in the work without there being any penalty involved. Which result is seen will depend on the nature of the hierarchy of authority in the team. In general, it is extremely difficult to predict whether having two or more people check one another's work will improve performance, degrade it, or leave it unchanged.

In strongly hierarchical teams where a piece of work must be checked by several people, there is no guarantee that accuracy will improve. Frequently, such hierarchical checking becomes merely a formality, and signing off is little more than a signature to certify that each person has seen the prior person's signature. In a case discovered by the writer, a set of operating procedures had been signed off by more than four levels of authority, and no one had discovered that the original documents were completely flawed. This is likely to happen where the work to be checked is extensive, complicated, and highly technical. The higher levels of authority are often remote from the technical expertise needed to check the details of the system, and signatures are merely used to satisfy bureaucratic guidelines. This is particularly true where management is seen as a profession in itself devoid of technical expertise in the particular domain being managed, a common philosophy today where it is assumed that someone who can manage one organization can therefore manage any other without the need for technical knowledge.

In a strongly hierarchical organization, those higher up cannot have time to check in detail work performed by those low in the hierarchy, because the quantity of information far exceeds the ability of those high in a hierarchy to process it. Generally, it overwhelms even those low in the hierarchy. Hence the technique of signing off on quality checks is often largely a superstitious behavior performed mainly to provide an audit trail of apparent authority should legal action question the system's efficiency.

It is clear that the structure of authority in teams and groups is critically important in reducing error. If a team has a very strong hierarchy, then it will be difficult for juniors to question decisions made by those at a higher level of authority even when the latter make errors. Furthermore, there will be tendency for those low in the hierarchy to be afraid to show initiative. On the other hand, if the hierarchy fails due to the absence or other unavailability of a person in authority, then the team may well dissolve into a group in which no one has clear authority to make decisions or to take action. Where intelligent or automated hardware is incorporated as part of a team (a situation becoming increasingly common with the advent of expert systems and other computerized aids), there is a similar problem. How can people decide whether and how to query the decisions of the computerized aids? Are they to be regarded as authorities of last resort? If not, when should their advice be accepted and when rejected?

It seems reasonable to expect that in a team, the collective knowledge,

skill, and wisdom will be greater than that available to any individual in the team. On the other hand, social dynamics often prevent such collective expertise from becoming available. The phenomenon of groupthink (Janis, 1972) is well established. In situations of uncertainty, members of a team or group will often tend to reinforce each other's assumptions. Where one might expect that the knowledge or skill needed to solve a problem must be more plentiful where there are more people, the ability to explore alternatives is actually reduced by the tendency of the group to come to a premature agreement that then insulates its members from further alternatives. During the accident at Three Mile Island, the team in the control room became unable to think of any explanation other than their initial hypothesis about the nature of the fault in the plant. At the end of the shift, a person who had not been involved in their discussions took a fresh look at the data and offered a completely different explanation, which was in fact the correct one (President's Commission Report, 1979). If collective behavior is to reduce error, ways must be found to reduce the tendency to groupthink and to free people from the tendency of hierarchical authority to bias the choice of action.

As discussed earlier, individuals tend to perceive and make judgments that are strongly determined by expectations about relative probabilities of events, costs and payoffs, similarities between current situations and past situations, and the success of certain actions in previous situations. All of these tendencies can be reinforced by the presence of other members of a team. There are interactions between personality styles and the degree to which people will accept authority or reject it. Social dynamics play a major role in the success with which a team can deal with unexpected events. Where the social structure of a team is well designed, where there is a fluent and free flow of information among its members, and where authoritarian hierarchical control is correctly designed (not necessarily minimized), exceptionally good and error-free performance from the group as a whole can be found. Where the social dynamics are not correctly designed, errors can go undetected, uncorrected, unobserved, and unreported. This can lead to catastrophic consequences and may hinder collective learning or improvement of performance over time. Particularly interesting is the work of people like Klein et al. (1993) and Rochlin, LaPorte, and Roberts (1987) on what makes teams effective.

The structure of authority in a team or group becomes critical when an abnormal situation occurs. This is particularly true when a beyond-design-basis event occurs whose possibility has not been foreseen. Such events frequently require a new pattern of organization and responsibility, but if the team is very tightly structured, no one will feel that they have the authority or responsibility to make radical changes in the way things are done. Teams need a balance between authoritarian coordination and flexible self-organization. It is only the latter that permits the evolution of new ways of behaving in situations for which no rules exist. On the other hand, flexibility, although supporting innovation, tends to lead to unpredictability, and the latter is often unacceptable to higher levels in an organization.

The behavior of the individual can thus be altered extensively by membership of a team or group. In some teams, individual ability may be repressed and constrained. In others, it may be nurtured and used in efficient cooperation where the strengths of one member make up for the weaknesses of another. What is clear is that how individuals will behave cannot be predicted without knowing the dynamics of the social setting within which they are working.

Organization and Management

The effects of organizational and managerial behavior affect the probability of error in yet more global ways. It is at these higher levels that policy decisions are made that indirectly but powerfully act downward to constrain the degrees of freedom in the behavior of teams and of individuals. Managerial influences can also act upward from the core of Fig. 5.1 by policies that determine the choice of equipment and the design of physical facilities.

A good example of organizational and managerial behavior that can have a major effect on the probability of error at the level of the individual is that of setting policy for shiftwork and hours of work. Just as there is a wealth of knowledge available to support the ergonomic design of equipment, so is there abundant evidence for the effect of shift work patterns, shift length, and circadian rhythms on error. It is known that errors greatly increase in the small hours of the morning, and often in the early hours of the afternoon. Errors in human information processing begin to increase significantly for shifts longer than 12 hours, and in physically demanding jobs, the errors often begin to increase at shorter intervals. Where shift work is used, certain patterns of shift rotation are less likely to cause errors than others (Tepas & Monk, 1987), and it is a managerial responsibility to decide what pattern of shift work will be implemented and how long working hours will be.

An obvious example of how managerial policies may increase or decrease error is to be seen in the long-standing problem of the hours worked by junior doctors in hospitals. The decision that junior physicians must work shifts that are frequently up to 24 hours in length is clearly a managerial policy. It is worth noting that because of the danger of increasing errors as time on shift increases, there is no other setting in hazardous industrial or military settings where people are permitted let alone required regularly to work the hours that are commonplace in hospitals.

Managerial decisions can have profound indirect effects through the choice of equipment. Decisions by management to invest in certain kinds of equipment can clearly cause major changes in the behavior of people who use the equipment. Insofar as managerial level members of an organization seldom have technical expertise or human-factors expertise, the control that they exercise over the quality of equipment, of operating instructions, of training in completely new technology, and so on is often at best remote. The opportunities for errors depend on the tasks given to members of an organization and the equipment with which they are provided. Those decisions are usually the province of management, if only because of the central role of financial constraints on what can be purchased. Although they have a direct effect by deciding what equipment to purchase, what shift schedules to implement, and so forth, they have little direct face-to- face communication with the health-care deliverers, and in that sense their effect is indirect and remote.

Just as one might expect that the collective knowledge and wisdom of a team should be greater than that of its individual members, so one might expect that the wisdom of an organization should be even greater. This will be true only if the organization is designed in such a way that it can, as whole, benefit from experience. Corporate memory seems to be extremely volatile unless very particular care is taken to enhance it. It is common to find that organizations are extremely rigid and unable to learn from past experience. Records of past errors and accidents are not kept, nor are they used to discover how to change the system. Frequently there is little significant change in practice following an error or accident, although those most closely involved are blamed and may be treated severely.

The ability of an organization to learn from the past errors of its members depends to a large extent on the attitudes and the managerial culture that are developed in the organization. If an organization is to learn, that is, to change its behavior as a result of past errors in a way that is reflected in the behavior of its members, then it must acknowledge errors that occur. If people report errors, those reports must be taken as information on the basis of which the organization can make constructive change. They must not be suppressed as undesirable "whistle blowing." The organization may need to accept the risks of making public the occurrence of errors. Errors should be seen, as they are in control theory, as signals for a needed change in practice. If, as commonly occurs, they are at all costs concealed, no learning will take place, and they will occur again. (An example of a system that permits "no fault" reporting of errors is the Aviation Safety Reporting System [ASRS] operated by the Federal Aviation Administration [FAA]. The ASRS allows pilots to report their errors anonymously so that the incidents can be recorded and analyzed without the pilot being blamed. These reports are fed back to the airline community through a publication that is available on a monthly basis, serving as a reminder of the kind of errors that even the best pilots may make.)

The whole notion of an organizational culture is linked to leadership at the managerial level. A special case is the notion of a *safety culture.* The attitude of members of an organization to safety, to corner cutting, or to violations is largely determined by managerial behavior. An open and flexible style of management can promote organizational learning so that individual violations of rules, which may occur because someone has found an objectively better way of performing some task, can become incorporated into the officially permitted or encouraged patterns of behavior. Likewise, an extreme rigidity of behavior can be imposed by authoritarian organizational structures, which will almost certainly render an organization unable to cope well with unforeseen events. Rigid rules do not necessarily make for a safer practice, especially in systems subject to many dynamic disturbances.

Legal and Societal Pressures

The outermost layers of Fig. 5.1 are causes of error that are usually remote from individuals but are still powerful. Behavioral options available to those working in a system may be tightly constrained by regulatory rules. Only certain drugs may be administered, only certain procedures undertaken. Violations of regulatory rules may have a heavy financial penalty attached. In the United States, the pattern of regulation in most industries has until recently been strongly prescriptive. Regulations prescribe what should be done under as many foreseeable situations as possible.

Except for the simplest systems, such an approach is doomed to failure. As systems and organizations become ever more complex, the number of possible events far exceeds the number of rules that can be thought up by regulators, and even where rules exist there are many problems. First, as the number of rules increases, it becomes increasingly difficult for people to learn, recall, and obey them all. Second, an individual or organization is often left in a very poor situation to deal with events unforeseen by the regulations—what is permissible in a situation for which there is no explicit regulation? Third, there is a grave danger that satisfying the regulations becomes an end in itself. The fact that the purpose of the rules is to regulate behavior with respect to some problem that affects the system is forgotten. People begin to feel that it is necessary and sufficient to carry out the rules as written, even though there may be far better ways to perform some task, and even when tasks exist for which no procedure has been specified. Treating rules as necessary and sufficient is the ultimate protection for the actor. No one can be blamed for following the rules, even when doing so causes undesirable consequences in a particular set of circumstances.

In addition to formal regulation, there is, particularly in the United

States, enormous pressure on behavior caused by the threat of legal action. In a highly litigious society, there is a fear of leaving undone what should be done and, equally, of doing things that are unnecessary. The legal system puts great pressure on an organization to make rules so that the organization will be protected by the rules. Rules allow an organization to pass blame down the hierarchy to individuals who break the rules, whether by accident or design. This in turn constrains behavior and can cause errors both of omission and commission.

Regulation and litigation are themselves driven by larger issues of social, cultural, and political pressures. Both decision making and overt behavior can be distorted by the requirements of society. A decision as to what treatment to use may be altered by knowledge of the economic situation of a patient. Pressures by shareholders for greater profits may affect decisions about work practices, manning levels, or the purchase of equipment. Cultural beliefs may render certain treatments unacceptable. At an individual level, the fears and hopes of a patient may exert strong pressure on a physician as to the choice of treatment. Union practices may predispose people to behave in certain ways, and antagonism to unions by management may be of equal influence.

ERRORS AND CONSTRAINTS

Consider again the statement that errors are made by individuals. This is often easy to see, as when people pick up a bottle next to the one that they wish to pick up, or misread a thermometer or blood pressure measuring instrument. At other levels, it is more difficult to attribute error to individuals, as when a board of management makes a collective decision about policy. But here, even if it makes sense to think of the collective decision as being an error of collective management, the decision is supported by the discussions, votes, and other activities of individuals, which in turn lead to the collective decision. The preceding sections of this chapter have discussed some ways in which influences from different levels of a system can impinge on an individual to cause or reduce error. These influences can be collectively described as the *constraints* that act on an individual to cause errors.

Constraints from Equipment

There are constraints on individuals' behavior from the equipment with which they work, acting from the center of Fig. 5.1 outwards. The equipment with which a person works is a filter that limits the information available for decisions and a filter that limits what actions can be taken. The accuracy of diagnosis, for example, depends very greatly on what measurements can be made. Computerized Axial Tomography (CAT) or Magnetic Resonance Imaging (MRI) provide data that are simply unavailable in the absence of the relevant equipment. The existence of computer databases and electronic communication alters the way in which decisions are made at all levels compared with an organization that uses only paper records.

Constraints from Physical Ergonomics

The size, shape, and legibility of displays, the positions and shapes of controls, the quality of alarms and communication subsystems, to name but a few physical properties of equipment, all constrain the way in which a person can acquire and use information and exercise control. These constraints force a person into certain patterns of thought and behavior, and those in turn can lead to a greater or lower probability of error. Even at the organizational level this is true. The quality of communication systems, the response time of computer databases, the time it takes to arrange meetings, and even the quality of air conditioning in a meeting room can have an effect on the quality of decisions made.

Constraints Within the Individual

The accuracy with which individuals acquire information and use it for decision and control is subject to the basic limitations of human psychology. These include the limits on the accuracy of perception, the volatility and accessibility of memory, the dynamics of attention, and the precision of motor skills. Other limitations arise from the complex interactions of emotion, motivation, judgment, and decision making. The human as an information processor is limited in rate and accuracy, and is biased by expectations and value systems. These limits impose powerful constraints on behavior.

Constraints in Teams and Groups

Constraints on the behavior of individuals arise from social dynamics within teams and groups. They include social pressure to conform by other members of the team, hierarchical patterns of authority and responsibility, habits of work that have grown up in a group, folklore about better ways of doing things that is passed from older members to new members, and tendencies toward groupthink. Other constraints are the quality of communication between members and whether a team is always composed of the same people or called together in an ad hoc way. Recently these factors have been emphasized by research into what is called *situation awareness*, the degree to which members of a team know from moment to moment what one another are doing and the extent to which they share a common understanding of what is happening.

Constraints from Organization and Management

Management exercises constraints in a great variety of ways. These include particularly the molding of organizational culture, which is deliberately or inadvertently developed with an emphasis on a particular combination of safety, profit and service, and the standards to which the members of the organization are expected to adhere. Patterns of shift work, available budgets, manning levels, response to the reporting of errors by oneself or others, openness to suggestions and innovation, even levels of pay and the pattern of bonuses for good work or sanctions for bad work all affect the way in which the individual members of the organization work, as well as the quality of their judgment and action.

Legal and Regulatory Rules

Legal and regulatory constraints play a very important role in determining behavior. The degrees of freedom available to an individual may be very tightly constrained by fear of litigation and by prescriptive regulatory rules. The first of these will lead to inherently conservative behavior, in which risky choices with a potential for better outcomes are rejected in favor of options that are believed to be fail-safe, not in the sense that they are necessarily best for the patient involved, but best in the sense that the person making the decision is safe from legal recourse even if things turn out badly. The second tends to promote rigid behavior, which cannot deal well with events unforeseen by the regulation and inhibits innovation.

Societal and Cultural Constraints

Society exerts pressures by its general expectations and even by pressure on individuals to adopt a particular philosophy of life. These in turn affect the options that will be considered by the individual, the choices that will be made, and the amount of risk that can be tolerated. For the health-care provider, such pressures may be in the form of demands from people outside the profession, or constraints from within the profession such as a limit on the numbers allowed to enter it, attitudes to the political structure or financial structure of the career, and so on. Questions of triage and rationing of care are intimately connected with social philosophy, whereas religious beliefs may prohibit the use of certain techniques or drugs.

DECISIONS, OUTCOMES, AND ERROR MANAGEMENT

One approach to error is to study the psychology of error. One can try to identify psychological mechanisms that cause error, but the list of constraints derived from Fig. 5.1 suggests that such an approach to error reduction is not likely to be successful.

What causes a person to record incorrectly the temperature of a patient? If the cause can be identified, then we can hope to make changes to the system so as to reduce the probability of such errors in the future. The cause of a misrecording may be due to a slip, in that the person writes down incorrectly a value that was correctly read, or it may be an error of perception when reading the thermometer, followed by correctly writing down the misperceived value. However, the cause of that misreading or miswriting may not be the perceptual or motor mechanisms of the individual. It may have arisen because the person is under great pressure to complete the task rapidly, having been urgently asked for a reading from a member of the health-care team. The person may have recently come from a meeting in which the statistical distribution of temperature in people with a certain illness was discussed, and the knowledge of the expected value of the reading from this kind of patient may predispose the observer to make the error. The display on the thermometer may have been degraded, making it difficult to read. If the data are to be recorded using a keyboard, there is a finite probability of typing errors, probably in the region of 10^{-3} to 10^{-4} per keystroke for average typists. Although the error emerges as the behavior of the individual, the cause may lie in one of the constraints listed here or from a variety of other causes. Hence, the most appropriate way to try to reduce the probability of error may not be to look at the psychology of the individual.

More generally, it is important to realize that all the constraints listed here interact. Even if the cause of *this* particular error were correctly identified and steps taken to prevent it, we can never be sure that no error will arise on the next occasion that the opportunity arises. *This* error may not recur, but in removing it, the system may have been changed in a way that makes another error more likely. (An obvious example is when computers are introduced: Errors of handwriting obviously decrease, whereas typing errors increase. Which has the higher probability, and under what circumstances?) A fundamental principle is that when we make any change in a system, that change propagates effects *throughout* the system and may cause many changes that were not foreseen. These may be at a different level in Fig. 5.1 from where the change was made.

The impact throughout the system of a change at a single level is most readily seen when changes occur at the outer levels of Fig. 5.1. Consider a change in regulatory philosophy from one where the regulations specify exact behavior to one where no specific behavior is specified, and the regulated organization is allowed to perform a task in any way it chooses, subject to satisfying the regulatory authority that the way chosen is safe. (This difference is the difference in regulatory philosophy, for example, between the United States' nuclear industry and that of Canada or the United Kingdom. The former spells out in great detail what must be done by all utilities under as many situations as possible. The latter allow utilities to do anything they can convince the regulators is safe.) The relaxation of the requirement to follow a specified pattern of behavior may have a great effect on how individuals perform their tasks, and hence one may expect a large change in the probability of error. That change may increase error or reduce it: It may increase in some parts of a system and decrease in another. The net effect on those who are served by the system may be difficult to predict. Similarly, if management changes the number of hours that a person must work, that change will have widespread effects throughout the system. These may include the level of fatigue experienced by individuals or the members of the group or team with whom they interact, the way in which handover at shift changes may alter, and changes in morale that may have important effects.

Even a change in the physical ergonomics of a piece of equipment may have more than a local effect. If a new kind of display is installed, the new symbology may be easier or harder to read. This will change the time required to carry out tasks, and may change the pattern of verbal communication between the person using the equipment and other members of the team. This in turn may change the length of time for which the patient is under the anesthetic and hence the number of operations performed each day, and so on. It is important to understand that when *any* element of a system is changed, the result is better described as a new system, not just as the old system with a change. Particularly in large complex systems, it is often very difficult, if not impossible, to foresee the results of quite minor changes because of the tendency for changes to propagate their effects throughout the system with varying time delays and unforeseen interactions (Ashby, 1956).

A lesson from this line of thought is that if one wants to reduce error and its consequences, the most productive approach may not be to ask for the particular psychological mechanism that caused the error. It may be more productive to try to change the system as a whole in such a way that the undesirable behavior or the undesirable outcome does not happen, *however it was caused.* When an error manifests itself, it will do so in the context of one of the constraints described earlier. It may seem to have been an error caused by a poorly designed piece of equipment, an inappropriate pattern of attention, a failure to communicate within a group, overlong hours of work, or an overly tight constraining regulatory rule. But it does not follow that the level of the constraint that seems to have caused the error is the same level at which one should try to prevent the recurrence of the error.

It is also important to distinguish clearly between errors and undesirable outcomes. The fact that a person's actions lead to an undesirable outcome does not mean that the person made a mistake. As we have already seen, there are many situations where, if a person follows a well-defined rule, the outcome may not be what is expected because there is some factor present that was not foreseen by the rule makers. Similarly, it can often be the case that a slip, mistake, or violation leads to a better outcome than would have the "correct" behavior as defined by a rule. In particular, violations are often ways to explore new and better solutions, and errors may also, albeit accidentally, lead to similar discoveries. Both in understanding how errors arise and in designing systems for their management, it is most important to keep the distinction between error and outcome distinct.

The systems approach to error management requires that possible solutions be considered at many levels of constraint. Suppose one were to find that infection is sometimes spread within a hospital because syringes are occasionally inadequately sterilized. One solution would be to change the rules that must be followed by personnel to ensure that sterilization is properly performed. More efficient training could be introduced. Better displays and alarms might be designed to ensure that the syringes were autoclaved for the requisite period of time. Personnel might be dismissed or fined for inadequate performance.

Those solutions are at the level of the individual, and assume that by acting to make perception, attention, memory, and skill more efficient, the errors could be reduced or eliminated. But a completely different solution is to change to disposable syringes, so that there is no need to perform sterilization at all. The first solution involves ergonomics and the psychology of individual behavior, and invokes constraints at the two innermost levels of Fig. 5.1. The change to disposable syringes is a solution at the level of organizational and managerial levels that acts indirectly through the innermost level by changing equipment. No blame or requirement to try harder is imposed on the personnel. (But note that with a change in practice the opportunity for *different* errors arises.)

Either of the suggested approaches could solve the problem. The choice between them might be constrained by economics rather than by an understanding of psychology. The important point is that each solution has its impact on the system in quite different ways. If the solution adopted still places responsibility on individual people to remember to perform tasks in a particular way, similar individual errors will probably recur in the future, even if team constraints in the form of supervisory checking were to be implemented. If the chosen solution alters the system so that there is no opportunity for the particular error to recur, then even though other errors may occur, that particular one will not.

Global or managerial solutions will not always best reduce error. There will be many cases where redesigning equipment or retraining personnel is the correct solution, and such redesign may be carried out by the user rather than the manufacturer. The point is that *no change has only local effects.* All changes propagate throughout the system. When trying to eliminate a particular error, the proposed solution and the consequences of that solution must themselves be examined in the context of the entire system. What constraints will be made stronger? What constraints will become irrelevant?

The problem of error management can be thought of as a search through the system for ways to constrain the possibility of particular errors without relaxing the constraints on others. Given an emphasis on the desirability of constraining outcomes, a detailed analysis of the psychological causes of error may not be required. Moreover, the search for solutions in the psychology of the individual may lead to solutions that will be extremely difficult to implement. How can one ensure that someone will always be alert when checking the work of another person? Almost always, it will be better to find constraints on behavior that remove the need for such checking. If the organizational culture favors safety and service rather than profit, how can that philosophy best be transmitted to newcomers? Rather than emphasizing an understanding of the deep psychological causes of error, it will be better to find ways to make safety look after itself. This does not imply that the psychological causes of error can be ignored, but that a complete systems approach to designing a health-care system is required in which the approaches from many disciplines are integrated.

CONCLUSION

Errors are made by individuals, but individuals work within systems. Systems are composed of people, things, information, and the relations and interaction among them. Systems can be analyzed and described in many ways, and the different ways emphasize different sources of constraint on the individuals who work in them. For any error that must be prevented, there will be many places in the system at which to intervene and many different ways of intervening. To use a classical metaphor, to concentrate too closely on the details of a particular error often leads to locking the stable door after the horse is stolen. To do so ensures that if there is another similar horse and a similar thief, we will be able to prevent the loss of the second horse. But in fact, in large, complex systems such as health-care delivery, there is an infinite number of horses and an infinite number of thieves. When any horse goes missing we should consider not merely locking doors, but rebuilding stables, retraining personnel—or even keeping animals other than horses.

The systems point of view emphasizes versatility in searching for solutions and shows that concentrating on ever-tighter local constraints will simply leave the system increasingly vulnerable to unforeseen events. It also emphasizes that the way to reduce error is to examine systems at all levels of constraint description. Although the process may be more time-consuming and difficult than looking for simple solutions at the level where the error is detected, it offers the hope of a greater measure of controlled adaptability, innovation, and versatility. Human error is not only a property of humans—it is a property of systems that include humans. In the end, the best way to prevent errors in giving medication may be as simple as changing the font of the print used in labels, or as complex as changing the purchasing strategies of the entire hospital. By using a systems approach with its potential for rich solutions, the number of errors can be reduced and their consequences mitigated.

Errors will always occur, and it is perhaps as well to reflect, particularly in a litigious society, that although the way in which death comes to each of us may be due to an error, death itself is not an error, but a result of life.

REFERENCES

Ashby, W. R. (1956). *An introduction to cybernetics.* London: Chapman and Hall.

Boff, K., & Lincoln, J. (1988). *Engineering data compendium.* New York: Wiley.

Gaines, B. R., & Moray, N. (1985.) *Development of performance measures for computer-based man-machine interfaces*. DCIEM-PER-FIN:JUL85, Downsview, Ontario.

Grandjean, E. (1980). *Fitting the task to the man.* London: Taylor and Francis.

Janis, I. L. (1972). *Victims of groupthink.* Boston: Houghton Mifflin.

Klein, G. A., Orasanu, J., Calderwood, R., & Zsambok, C. E. (1993). *Decision making in action: Models and methods.* Norwood, NJ: Ablex.

Moray, N. (1984). Attention to dynamic visual displays. In R. Parasuraman (Ed.), *Varieties of attention* (pp. 485–512). New York: Academic Press.

Moray, N. (1986). Monitoring behavior and supervisory control. In K. Boff, L. Kaufmann, & J. Beatty (Eds.), *Handbook of perception and human performance* (pp. 40-1–40-51). New York: Wiley.

Norman, D. A. (1981). Categorization of action slips. *Psychological Review, 88,* 1–55.

Norman, D. A. (1988). *The psychology of everyday things.* New York: Basic Books.

President's Commission Report on the Accident at Three Mile Island Unit 2. (1979). New York: Pergamon.

Rasmussen, J., Duncan, K., & Leplat, J. (1987). *New technology and human error.* Chichester, England: Wiley.

Rasmussen, J., & Batstone, R. (1989). Why do complex organizational systems fail? *Summary proceedings of a Cross-Disciplinary Workshop in "Safety Control and Risk Management."* Washington, DC: World Bank.

Reason, J. (1990). *Human error.* Cambridge, England: Cambridge University Press.

Rochlin, G. I., LaPorte, T. R., & Roberts, K. H. (1987). The self designing high-reliability organization: Aircraft carrier flight operations at sea. *Naval War College Review,* 76–90.

Sage, A. P. (1981). Behavioral and organizational considerations in the design of information systems and processes for planning and decision support. *IEEE Transactions on Systems, Man, and Cybernetics, 11,* 640–678.

Sanders, M., & McCormick, E. J. (1993). *Human factors in engineering and design* (7th ed.). New York: McGraw-Hill.

Senders, J. W., & Moray, N. (1991). *Human error: Cause, prediction and reduction.* Hillsdale, NJ: Lawrence Erlbaum Associates.

Swets, J. A., & Pickett, R. (1982). *Evaluation of diagnostic systems.* New York: Academic Press.

Tepas, D. I., & Monk, T. H. (1987). Work schedules. In G. Salvendy (Ed.), *Handbook of human factors* (pp. 819–843). New York: Wiley.