CHAPTER 3

Complexity, Coupling, and Catastrophe

To make a systematic examination of the world of high-risk systems, and to address problems of reorganization of systems, risk analysis, and public participation, we need to carefully define our terms. Not everything untoward that happens should be called an accident; to exclude many minor failures, we need an exact definition of an accident. Our key term, *system accident* or *normal accident,* needs to be defined as precisely as possible, and distinguished from more commonplace accidents. We will define it with the aid of two concepts used loosely so far, which now require definition and illustration: *complexity* and *coupling.* We also need to explain the terms *catastrophe* and *victim* as they are used in our analysis. Then we will be in a position to chart the world of organized systems, predicting which will be prone to system accidents and which will not. This chapter, drawing upon numerous examples, will construct an apparatus that will carry us through the systems discussed in the rest of the book. It constitutes a theory of systems, of their potential for failure and recovery from failure. As such, it is, I believe, unique in the literature on accidents and the literature on organizations.

Perhaps the most original aspect of the analysis is that it focuses on the properties of systems themselves, rather than on the errors that owners, designers, and operators make in running them. Conventional explanations for accidents use notions such as operator error; faulty design or equipment; lack of attention to safety features; lack of operating experience; inadequately trained personnel; failure to use the most advanced technology; systems that are too big, underfinanced, or poorly run. We have already encountered ample evidence of these problems causing accidents. But something more basic and important contributes to the failure of systems. The conventional explanations only speak of problems that are more or less inevitable, widespread, and common to all systems, and thus do not account for variations in the failure rate of different kinds of systems.

What is needed is an explanation based upon system characteristics. In Chapter 6, on marine transport, you will see how technological fixes at best make no difference, and can even make the situation worse. Chapter 8, on the space program, reveals how the best talent and organizational resources, while they certainly help, cannot overcome system accident potentials. Several chapters show how system-related production pressures defeat safety improvements. The accidents covered all chapters challenge the ready explanation of "operator error." This is not to say that the concepts and definitions of system characteristics presented here will solve all analysis problems. They are preliminary; definitional problems remain. This is a first attempt at a "structural" analysis of risky systems, but I believe it goes substantially beyond conventional "item" analyses, and yields, as we shall see in the last chapter, long-run strategies for handling risks—rather than short-run tactics of posting safety notices or levying trivial fines.

Defining Accidents

What do we mean by an *accident?* At the minimum, an accident is an unintended and untoward event. If you are driving home and take a wrong turn and lengthen your trip, it is both unintended and unfortunate. But we generally mean something more serious than this by the term *accident.* You would not, upon arriving home late, announce, "I had an accident on the way home." If you did, you would be worriedly asked, "Was anyone hurt?" or, "Was the car damaged?"

An accident, then, involves some damage to people, objects, or to both. But uncertainties remain. Suppose you scratched the paint of the car on a post while leaving a parking lot. You probably would not call it an accident (even though it was accidental) if it didn't interrupt the trip or impair the functioning of either the car or the post. We might, then, say that the damage to objects or people must be sufficient to disrupt the ongoing "task" or future tasks that will be demanded of the objects or people.

This introduces a further complication: tasks analysis. What task is involved? That depends on what we want to call the system. If I had planned to take the car to a rally the next day and show it off to other automobile buffs, I might very well call the scratch in the parking lot an accident. The system, from my point of view, involves my going to a rally to meet people, impress people, and show off my car, and it is interrupted. If the post was bent over so much that cars trying to enter the lot were endangered by what I had done, that too might qualify the event as an accident. Drivers would notify the maintenance people, "There has been an accident, the post is knocked over so far that I can hardly get into the lot." It is part of the parking system to have a post separating lanes; it is part of my recreational system to show off a beautifully maintained and polished car. The event has disturbed these systems.

But there are also degrees of disturbance to systems. The rally will not be disturbed in any perceptible way if I show up with a scratch on my '57 De Soto, or if, ashamed of the state of my car, I do not go at all. But *my* system might be greatly disturbed if I stayed home rather than meeting or impressing people. The degree of disturbance, then, is related to what we define as the system. If the rally is the system under analysis, there is no accident. If that part of my life concerned with custom cars is the system, there *is* an accident. A steam generator tube failure in a nuclear plant can hardly be anything other than an accident for that plant and for that utility. Yet it may or may not have an appreciable effect upon the nuclear power "system" in the United States.

So far we have defined an accident as unintended damage to people or objects that affects the functioning of the system we choose to analyze. But we can also affect system functioning by "damaging" symbols, communication patterns, legitimacy, or a number of factors that are not, strictly speaking, people or objects. This observation will become important if we are to consider such organizations as universities.

An accident, then, involves damage to a defined system that disrupts the ongoing or future output of that system. But not all such disruptions

should be classified as accidents; the damage must be reasonably substantial. In a nuclear plant, the momentary failure of the off-site power supply (the electricity coming into the plant to run its machinery) usually means that the reactor must be shut down. Similarly, a valve failure may result in a shutdown. But we will not call these *accidents,* even though they were not desired and have unfortunate or untoward consequences—purchasing replacement power alone may cost a few hundred thousand dollars before the reactor is restarted and the plant placed back on-line. We will call them *incidents.* Though the reactor is shut down, it does not sustain any damage. We need, then, criteria, inevitably somewhat arbitrary and rough, to distinguish between "minor" events such as these and accidents. Furthermore, we need a scheme that can apply equally to a steam generator as a system, or to a plant, or to the nuclear power industry.

I propose that the system be divided into four levels. Disruption to the third or fourth levels will be called *accidents* and disruption to the first and second will be called *incidents.* It is a scheme that need not be applied carefully in all parts of this book; sometimes we will ignore it and speak quite loosely of accidents when the meaning is obvious. But it is important for selecting accidents for analysis, for understanding safety devices, and for risk analysis; at times it will be crucial.

Consider a nuclear plant as the system. A *part* will be the first level—say a valve. This is the smallest component of the system that is likely to be identified in analyzing an accident. A functionally related collection of parts, as, for example, those that make up the steam generator, will be called a *unit,* the second level. An array of units, such as the steam generator and the water return system that includes the condensate polishers and associated motors, pumps, and piping, will make up a *subsystem,* in this case the secondary cooling system. This is the third level. A nuclear plant has around two dozen subsystems under this rough scheme. They all come together in the fourth level, the nuclear plant or system. Beyond this is the environment.

With this scheme we reserve the term *accident* for serious matters, that is, those affecting the third or fourth levels; we use the term *incident* for disruptions at the first or second level. The transition between incidents and accidents is the nexus where most of the engineered safety features (ESFs) come into play—the redundant components that may be activated; the emergency shut-offs; the emergency suppressors, such as core spray; or emergency supplies, such as emergency feedwater pumps. The scheme has its ambiguities, since one could argue interminably over the dividing line between part, unit, and subsystem, but it is flexible and

adequate for our purposes. It must be flexible because at times we might want to consider, for example, the Apollo rocket and modules as a system; at other times we might want to consider all moon shots as a system. What is an accident in the first might be a mere incident in the second.

We are now ready for a formal definition. An *accident* is a failure in a subsystem, or the system as a whole, that damages more than one unit and in doing so disrupts the ongoing or future output of the system. An *incident* involves damage that is limited to parts or a unit, whether the failure disrupts the system or not. By disrupt we mean the output ceases or decreases to the extent that prompt repairs will be required. Since we have drawn a dividing line between the unit and the subsystem, and since many of the ESFs are clustered around that dividing line, it will often mean that an ESF will be one of the components that fails.

Victims

Note that, in contrast to common usage, we have not explicitly made damage to humans a part of our definition. This is because we are primarily interested in systems and the way they work. Humans are part of all the systems considered in this book, and a group of humans (the flight crew on an airliner) or a single human (the astronaut in a space capsule) may constitute a subsystem. Damage to them will mean that we have an accident.

But it is important for analysis to treat humans in most systems as mere parts. We kill about 5,000 people a year outright in U.S. industry. The vast majority of these "accidents," however, are only "incidents" in our scheme, for no subsystem or system damage is entailed. Only a "part" has been destroyed.

I am aware that this sounds like a quite heartless scheme, but while the ultimate concern of this book is reducing actual and potential damage to humans, I believe it is the character of systems that cause that damage, and thus we need a definition of accidents that focuses upon system characteristics. The failure of what we have defined as parts and units plays a vital role in subsystem and system failures, but if the analysis is limited to these we will lose our focus upon the kinds of systems that business and government leaders decide should be built. Our analysis,

then, must be of systems, and falling off a ladder becomes a mere incident.

More important, we are concerned about those systems that have catastrophic potential—can cause damage to a great many humans. Our concern here is not the blue-collar worker who backs into a grinding machine, or the scientist who accidentally creates the proper criticality conditions for the plutonium he is experimenting with and is fatally irradiated. More or less mundane precautions and training will reduce the frequency of such accidents; fairly minimal attention to mine safety greatly reduced the number of coal mining deaths in a few years. Our concerns are more cosmic. To bring them into focus, we will turn to a victim classification scheme that will be used throughout this book.

Most of the work concerned with safety and accidents deals, rightly enough, with what I call first-party victims, and to some extent second-party victims. But in this book we are concerned with third- and fourth-party victims. Briefly, first-party victims are the operators; second-party victims are nonoperating personnel or system users such as passengers on a ship; third-party victims are innocent bystanders; fourth-party victims are fetuses and future generations. Generally, as we move from operators to future generations, the number of persons involved rises geometrically, risky activities are less well compensated, and the risks taken are increasingly unknown ones. We will take a close look at each of the four classes.

First-party victims are the operators of the system. In this book operators will include not only those actually running the system (nuclear plant operators, pilots, ship officers) but others in attendance on regular shifts, such as first-level supervisors, maintenance personnel, low-level engineering personnel, and laborers and assisting personnel. Most industrial accidents involve operators; in fact, most involve only one operator. Most industrial accidents are attributed to "operator error" or "human error" by those who study and seek to prevent such accidents. There is a growing recognition, however, that this is a great oversimplification; worse, it involves blaming the victim. It also suggests an unwitting—or perhaps conscious—class bias; many jobs, for example, require that the operator ignore safety precautions if she is to produce enough to keep her job, but when she is killed or maimed, it is considered her fault. Some operators are compensated for the risks they run by receiving higher pay than the skill level would suggest, but in industry as a whole there is no clear relationship between risk and pay.[1] There is no impersonal, fair market that rewards those that risk their lives with higher wages. Indeed,

"jumpers" or the "glow boys" in the nuclear industry, temporary help who dash into a radioactive area to make repairs, will be hired for two or three weeks' work, at only six dollars an hour, even though they can receive high doses of radiation in the few minutes they are near the core.[2] There is no evidence of compensation for the long-term effects of toxic chemicals or contaminated atmospheres. Textile workers are not compensated for brown lung disease, nor are chemical plant workers compensated for cancer showing up ten or twenty years after exposure.

Second-party victims are those associated with the system as suppliers or users, but without influence over it. They are not innocent bystanders (third-party victims), because they are aware (or could be informed) of their exposure, even though such exposure may not be entirely voluntary. The largest class of second-party victims are passengers on ships, trains, airlines, cars, and busses. To some extent they "choose" to participate in the system, and thus elect to share at least some of the risk. If I accept a ride home from a party with an inebriated driver, I accept the risk. The difference between first- and second-party victims with regard to voluntary exposure may be slight when we compare an unemployed person who has no choice but to accept risky work, and an employee who must travel extensively to keep his job. Neither has realistic options to participating in the system. But there are less ambiguous examples. Comparing second-party victims to innocent bystanders illustrates the slightly voluntary nature of the former: we feel differently about airline passengers killed in a crash than we do about innocent bystanders on the ground killed by the airliner's crash. The former accepted the risk of flying, the latter did not.

There are other types of second-party victims. For example, consider the office personnel hurt by a refinery explosion or the truck driver delivering goods for his company who just happens to be on the site when the explosion occurs. These are the voluntary actions of people who elect to participate in a system but have no influence over its operation. Nevertheless, without these second party or system-associated participants there would be no system. The refinery could not run. Like passengers, they are participants in the system and aware of some risks.

Third-party victims, innocent bystanders, have no such involvement in the system. I have heard some nuclear proponents argue, in a conference on establishing safety goals for nuclear power plants, that people can choose not to live near a nuclear plant (or choose to avoid going to events at a stadium that is in the flight path of an O'Hare Airport runway). But I think these arguments can be dismissed. Nuclear plants exist near all densely populated areas in the United States; there is simply no

practical means to avoid being within 50 miles of one. Even if one could, a severe core meltdown with breach of containment under the proper weather conditions could contaminate areas as large as the whole Northeast. The plutonium that might have crashed on Madagascar during the *Apollo 13* emergency reentry of the atmosphere would have contaminated innocent bystanders in a distant part of the world that one would have thought safe from such high-technology disasters. People are aware of risks associated with flying; but I suspect that most people living downstream from large dams do not realize the dam is near enough to endanger their lives in the case of a failure.

Fourth-party victims are, for the most part, victims of radiation and toxic chemicals. They are fetuses being carried at the time of exposure; the would-be children that damaged parents will not be able to conceive; stillborn or deformed children conceived after exposure; and all those people who will be contaminated in the future by residual substances, including those substances that will become concentrated as they move up the food chain. Note that we do not include run-of-the-mill pollution here. Nuclear plants may regularly give off radiation in doses large enough to affect fetuses or future reproductive capacity, as a few scientists argue (though most deny); but these consequences are not the result of accidents. (Although one might argue that such radiation is a design failure that should, but doesn't, interrupt the output.) Nor are the pollutants from industrial plants of concern here, unless they are released in an accident. This routine and mostly conscious contamination of the planet probably has more serious long-term consequences than any accident we shall consider in this book, other than military accidents, but it is beyond the scope of our concern.

The importance of fourth-party victims in risk analysis is growing with the increasing concern and sophistication about the long-term consequences of some systems. Yet recognition is slow. A draft version of the NRC's "Safety Goals for Nuclear Power Plants" indicated that while the NRC was aware of "inter-generational" risks from a nuclear power plant accident through genetic effects or long-term contamination, "we cannot suggest a good way to handle the issue in a safety-goal context." So they ignored it.[3] If the risks of an accident are kept low enough, they said, there will be no problem with ignoring inter-generational effects. This conclusion answers the question about consequences of accidents by saying they will be trivial because there will be so few accidents.

Fourth-party victims potentially constitute the most serious class of victims. Chemical or radioactive contamination of land areas could have far-reaching effects upon the health of future generations. Genetic defects

harm future generations in other ways, including adding the burden of lifetime care and treatment of victims. Future generations carry the burden; the present generation reaps whatever rewards there may be from the activity.

These issues are comparatively new—less than one generation old. Some influential scientists and academics, at the workshop that tried to formulate safety goals for the NRC, argued that the present generation is more important than future ones—for we need nuclear power to prevent economic and political crises—and who knows, there may be a technological fix that would mitigate the burden for the future generations of an accident in the present one. Thus at least some of the "experts" do not see the issue as a clear-cut one of our responsibility to future generations. Other experts, however, believe we have a responsibility to pass on to our offspring a world that is at least no more contaminated and degraded than the one we inherited.

Accident Definitions

We now have a definition of an accident, distinguishing it from incidents on the basis of levels in the system. The following list presents these and adds the definition of the two types of accidents: component failure accidents and system accidents, which we will now take up.

Systems are divided into four levels of increasing aggregation: units, parts, subsystems, and system.

Incidents involved damage to or failures of parts or a unit only, even though the failure may stop the output of the system or affect it to the extent that it must be stopped.

Accidents involve damage to subsystems or the system as a whole, stopping the intended output or affecting it to the extent that it must be halted promptly.

Component failure accidents involve one or more component failures (part, unit, or subsystem) that are linked in an anticipated sequence.

System accidents involve the unanticipated interaction of multiple failures.

Component failure accidents and system accidents are distinguished on the basis of whether any interaction of two or more failures is anticipated, expected, or comprehensible to the persons who designed the sys-

tem, and those who are adequately trained to operate it. A system accident, in our definition, must have multiple failures, and they are likely to be in reasonably independent units or subsystems. But system accidents, as with all accidents, start with a component failure, most commonly the failure of a part, say a valve or an operator error. It is not the *source* of the accident that distinguishes the two types, since both start with component failures; it is the presence or not of multiple failures that interact in unanticipated ways.

The vast majority of component failure accidents involve a series of failures. If a valve fails there is good chance that the pump will overheat and fail, and if this happens, the boiler is quite likely to overheat because the coolant is insufficient. Designers know this, and so do operators—though they may not be able to prevent it, or intervene in the series of failures. There will be some accidents, however, where the initial failure is so drastic that it is not worth tracing out the subsequent sequence, if there is one. If the wing comes off an airplane in flight, or an earthquake shatters a dam, we hardly need further analysis. These might be called "final accidents"; there is no possible intervention by operators, and no point in detailing sequences. We will not deal with any examples of these in this book; for one thing they are fairly rare, but more important, analysis of the system itself does little to contribute to understanding these accidents.

Incidents are overwhelmingly the most common untoward system events. Accidents are far less frequent. Among accidents, component failure accidents are far more frequent than system accidents. I have no reliable way to estimate these frequencies. For the systems analyzed in this book the richest body of data comes from the safety-related failures that nuclear plants in the United States are required to report. Roughly 3,000 Licensee Event Reports are filed each year by the 70 or so plants. Based upon the literature discussing these reports, I estimate 300 of the 3,000 events might be called accidents; 15 to 30 of these might be system accidents. So far, at least as far as we know, all the system and component failure accidents in nuclear plants have had only first-party victims, and very few of these.

Complex and Linear Interactions

What kinds of systems are most prone to system accidents? In the last chapter we kept mentioning two concepts: interactiveness, which could confuse operators, and tight coupling, which could prevent speedy recovery from an incident. With a more precise definition of these two terms, we can classify systems and be forewarned about those that are most prone to system accidents. We will take up interactions first.

The notion of baffling interactions is increasingly familiar to all of us. It characterizes our social and political world as well as our technological and industrial world. As systems grow in size and in the number of diverse functions they serve, and are built to function in ever more hostile environments, increasing their ties to other systems, they experience more and more incomprehensible or unexpected interactions. They become more vulnerable to unavoidable system accidents.

Interactiveness per se, though, is not a useful concept. Almost any organization of any size, whether public or private, will have many parts that interact once we look closely at them. The existence of many parts is no great trouble for either system designers or system operators if their interactions are expected and obvious. If a part or a unit fails in an assembly line, it is quite clear what will happen to the parts and units "downstream" of the failure, and we know that the products "upstream" will start piling up fast. We can shut down the line and fix it, or bypass that work station and assemble the missing parts later on, or temporarily shunt the product into a storage bin. These are "linear" interactions: production is carried out through a series or sequence of steps laid out in a line. It doesn't matter much whether there are 1,000 or 1,000,000 parts in the line. It is easy to spot the failure and we know what its effect will be on the adjacent stations. There will be product accumulating upstream and incomplete product going out downstream of the failure point. Most of our planned life is organized that way.

But what if parts, or units, or subsystems (that is, components) serve multiple functions? For example, a heater might both heat the gas in tank A and also be used as a heat exchanger to absorb excess heat from a chemical reactor. If the heater fails, tank A will be too cool for the recombination of gas molecules expected, and at the same time, the chemical reactor will overheat as the excess heat fails to be absorbed. This is a good design for a heater, because it saves energy. But the interactions are no longer linear. The heater has what engineers call a "common-mode"

function—it services two other components, and if it fails, both of those "modes" (heating the tank, cooling the reactor) fail. This begins to get more complex.

This source of complexity was slow to gain recognition in the nuclear power field. The first analytical model to consider common mode failures appeared only in 1967, according to a perceptive and disturbing article by one of the editors of *Nuclear Safety,* E. W. Hagen.[4]

The monumental reactor safety study issued in 1975, WASH-1400 or the Rasmussen Report, was criticized by a subsequent NRC study for giving insufficient and overly simplified attention to this problem. Hagen concludes that potential common-mode failures are "the result of adding complexity to system designs." Ironically, in many cases, the complexity is added to reduce common-mode failures. The addition of redundant components has been the main line of defense, but, as Hagen illustrates, also the main source of the failures. "To date, all proposed 'fixes' are for more of the same—more components and more complexity in system design."[5] The Rasmussen safety study relied upon a "PRA" (probabilistic risk analysis), finding that core melts and the like were virtually impossible. Hagen notes (p. 191) that PRAs, "using established techniques of reliability and statistical analysis," constitute the main source of public reassurance about such potential dangers. But, he notes, this involves a "narrow definition" of common-mode failures, and the analysts "are busily working in an area which has not been shown to be the main problem." The main problem is complexity itself, Hagen argues. With that we can agree.

Common-mode failures are just one indication of interactiveness in systems. Proximity and indirect information sources are two others. For a graphic illustration of complexity in the form of unanticipated interactions from these sources, let us go to a different system, marine transport. A tanker, the *Dauntless Colocotronis,* traveling up the Mississippi River near New Orleans, grazed the top of a submerged wreck. The wreck had been improperly located on some of the charts. Furthermore, it was closer to the surface than the charts indicated because its depth had been determined when the channel was deep, and the listing for mariners had not been corrected for those seasons when the river would be lower. The wreck, only a foot or so too high for the tanker, sliced a wound in the bottom of the tanker, and the oil began to seep out. Unfortunately, the gash occurred just at the point where the tank adjoined the pump room. Some of the oil seeped into the pump room. At first it was probably slow-moving, but the heat in the room made it less viscous,

allowing it to flow more rapidly, drawing more oil into the pump room. When enough accumulated it reached a packing "gland" around a shaft near the floor that penetrated into the engine room next to it. The oil now leaked into the engine room. In the hot engine room, evaporation was rapid, creating an explosive gas. There is always a spark in an engine room, from motors and even engine parts striking one another. (In fact, nylon rope can create sparks sufficient to cause explosions in tanker holds.) When enough gas was produced through evaporation, it ignited, causing an explosion and fire.

In this example, an unanticipated connection between two independent, unrelated subsystems that happened to be in close proximity caused an interaction that was certainly not a planned, expected, linear one. The operators of the system had no way of knowing that the very slight jar to the ship made a gash that would supply flammable or explosive substances to the pump and engine rooms; nor of knowing, until several minutes had passed, that there was a fire in the engine room; or of the extensive amount of oil involved. They tried to put it out with water hoses, perhaps not realizing it was an oil fire, but the water merely spread the oil and broke it up into finer, more flammable particles.

This accident went on for several hours, because the crew made various mistakes (e.g., a key fire door had been tied open and was not closed by an escaping crew member, allowing the fire to spread), and they did not have a clear notion of the location of the many rooms and closets and passageways in that part of the ship. Late in the accident a trained fire crew boarded the ship equipped with protective devices and proper extinguishing equipment. But when they opened one door there was a series of explosions. They immediately closed the door and made no further attempt to put out the fire in that part of the ship, believing that the oil was producing explosive gases. Subsequently, however, it was determined that there were no explosive gases involved by this time; instead, three small empty gas tanks stored just inside the door, before being exchanged for full ones, had exploded when the heat expanded the residual amounts of freon, oxygen, and acetylene in them.[6]

Thus even in the recovery stage of the accident a nonlinear interaction intervened to mislead the recovery attempt. It was a sensible place to store the tanks; who could imagine that there might be a fire, that firefighters would not be certain that there were not explosive mixtures behind the closed doors, and that these tanks would go off just as the firemen were entering the passageway? A requirement that empty tanks be stored elsewhere would hardly make the ship safer; who could know where the next fire might be?

Recall the illustration that opened this book, linking three "subsystems": breakfast, getting to the appointment, and the job interview. In the world we plan out and think through, this seems like a very linear, straightforward problem—get some food and coffee, get the car, drive to the interview. One would expect the car keys to be linked to using the car, but one would not expect the failure of the coffeepot to be linked to using the car. One would also not expect that even if the car failed, the alternative of a taxi would be linked to a contract dispute, and the neighbor's car would be unavailable just that day. These represent interactions that were not in our original design of our world, and interactions that we as "operators" could not anticipate or reasonably guard against. What distinguishes these interactions is that they were not designed into the system by anybody; no one intended them to be linked. They baffle us because we acted in terms of our own designs of a world that we expected to exist—but the world was different.

I will refer to these kinds of interactions as *complex interactions,* suggesting that there are branching paths, feedback loops, jumps from one linear sequence to another because of proximity and certain other features we will explore shortly. The connections are not only adjacent, serial ones, but can multiply as other parts or units or subsystems are reached.

The much more common interactions, the kind we intuitively try to construct because of their simplicity and comprehensibility, I will call *linear interactions.* Linear interactions overwhelmingly predominate in all systems. But even the most linear of systems will have at least one source of complex interactions, the environment, since it impinges upon many parts or units in the system. The environment alone can constitute a source of failure that is common for many components—a common-mode accident. But even the most complex systems of any size will be primarily made up of linear, planned, visible interactions.

Based upon a general knowledge of various systems, considerable experience with reading accident reports, and some site visits, I suggest that only 1 percent of all possible parts or units in a linear system are capable of producing "complex" interactions, while about 10 percent of those in a complex system will be capable of doing so. But that 10 percent represents more than a tenfold increase in the potential for system accidents. The potential interactions produced by, say, four parts or units that are interrelated in a complex, rather than linear way, is twelve. (There are twelve possible paths between the four parts.) Suppose this exists in a system where there are 400 parts or units. If 10 percent of the units had such characteristics, rather than only 1 percent, there would be forty such

parts or units. The potential complex interactions of each one of these with the remaining 399 would be in the millions, since the possible paths increase exponentially. Of course, in a large system some parts are so remote from each other that the chances of their interacting in unexpected ways can be disregarded; but many will not be so remote. Thus, increasing the proportion of possible complex interactions from 1 percent to 10 percent will have an enormous impact upon the potential for system accidents—given the fact that, since nothing is perfect, component failures are inevitable.

The possibility of many unintended interactions is recognized by designers, so they introduce buffers and other safety devices to prevent some kinds of interactions. Imagine a chemical plant where gas is expected to flow from tank A to tank B, because the pressure is kept higher in tank A than in tank B. Various things go into tank A besides the basic gas—reagents, purifiers, "inerting" gases, and so on. Then the gas flows to B, where additional substances are piped in to alter that mixture even further. This is quite straightforward and linear.

But there is a danger that it might become nonlinear. A feedback loop could occur if pressure declined in tank A because of a failure there or elsewhere, and the now-modified gas in tank B flowed back into A. This could create problems; it might even be dangerous. The engineers are aware of this, so they design a butterfly valve to be placed between the two tanks to prevent reverse flow. An engineered safety device (ESD) is installed to keep the system as linear as possible (in this case, flowing in one direction). But valves can fail, especially rarely used ones. If the pressure in tank A were to fall and the butterfly valve to fail, the feedback loop would open again, and the operators might not expect the interaction. Actually, this example is so simple that most operators would take the unplanned interaction into account as a possibility.

For a more complicated example involving quite remote and independent units, recall the case in the last chapter of the demineralized water needed to clean up the containment room in the nuclear plant. A faucet (valve) was opened to see if the water had been sent to the men, and a complex series of backflows and pressure adjustments occurred that allowed radioactive materials to enter containment. There was nothing very linear about that, to the operators in containment, the operators in the control room, or the engineers who designed the system. Nor did anyone make an error or even a mistake in design. Subsequently they could figure out what had happened and introduce measures to prevent a recurrence, but that meant altering an unexpected and unplanned inter-

action between two normally independent subsystems, core maintenance and steam generation. As I suggested, there are probably a number of other complex interactions ready to be revealed by a casual event in that very complicated piping system.

Linear interactions are overwhelmingly those of an "expected production sequence"—this is the way the system is designed to run, and anyone operating it will know that. Complex interactions will generally be those not intended in the design. No designer of the *Dauntless Colocotronis* said, "Let's put the number 5 tank next to the pump room so they can interact." There is just no way to isolate all tanks in a tanker from a room that generates sparks. Or the nonlinear interactions may be intended but rarely activated, and thus operators or designers forget about them. It could have been foreseen by a designer that demineralized water might sometimes be needed in containment, so she could have made it possible to line up various valves to provide the water. But if it is rarely used or is usually lined up before a crew enters containment, the faucet creates no problems, it is not an expected production sequence but an infrequently used system possibility (in this case, for maintenance, not production). Thus, complex interactions may be unintended ones, or intended but unfamiliar ones.

While linear interactions occur overwhelmingly in an anticipated production sequence, there is another kind of interaction that does not occur in a production sequence but is nevertheless obvious, and thus can be defended against. This is a visible interaction, even though it is outside of the normal sequence. If the operator of a crane sees that the part of the crane that holds up the load (cable, ratchet, motor, hook) has failed and the load is about to drop on a boiler on the floor, he knows what this interaction will produce. There is nothing mysterious about the connection between these events, though they are certainly not in any expected production sequence. Knowing there is a remote possibility of a large load falling, the operator usually tries to avoid passing it over the boiler, but he cannot always do so. Similarly, the designer may have considered covering the boiler or moving it, but estimated that the problems involved were too great, given the remote possibility of a large load falling at just that point.

To summarize our work so far: *Linear interactions* are those interactions of one component in the DEPOSE system (Design, Equipment, Procedures, Operators, Supplies and materials, and Environment) with one or more components that precede or follow it immediately in the sequence of production. *Complex interactions* are those in which one

component can interact with one or more other components outside of the normal production sequence, either by design or not by design.

To elaborate on the properties of these interactions as they affect the operators:

Linear interactions are those in expected and familiar production or maintenance sequence, and those that are quite visible even if unplanned.

Complex interactions are those of unfamiliar sequences, or unplanned and unexpected sequences, and either not visible or not immediately comprehensible.

A note on the terms "complex" and "linear" is in order. It is difficult to find precise terms that are also brief; I have opted for brevity. "Complex" should read "interactions in an unexpected sequence"; "linear" should read "interactions in an expected sequence." One problem with the terms *complex* and *linear* is that the opposite of complex is "simple," and the opposite of linear is "nonlinear." Linear interactions are "simple" in the sense that they are easily comprehensible, but "simple" implies unsophisticated processes and technologies, or systems with few parts and units or routine and uncomplicated operation and maintenance. But the production of, say, pharmaceuticals or F-16 fighter planes is anything but simple, though it is linear. On the other hand, "nonlinear" does not readily convey the notion of possible incomprehensibility, while the term *complex* does. I will occasionally slip in such phrases as "complexly interactive" and "linear sequences" to remind the reader that neither of the chosen terms is really satisfactory.

Another warning is in order. Since linear interactions predominate in all systems, and even the most linear systems can occasionally have complex interactions, systems must be characterized in terms of the degree of either quality. It is not a matter of dichotomies. Furthermore, systems are not linear or complex, strictly speaking, only their interactions are. Even here we must recall that linear systems have very few complex interactions, while complex ones have more than linear ones, but complex interactions are still few in number.

Finally, the reader should not equate the notion of a linear system with the physical layout of the plant or production process. It does not necessarily imply an assembly line, although these production systems tend to be linear. Nor does the designation "complex system" necessarily imply highly sophisticated technology, numerous components, or many stages of production. We will characterize universities as complex systems, but

as most large-scale organizations go, they have few of the above characteristics.

In order to understand systems, we need to go beyond the distinction between the two types of interactions. By examining, in the next section, how systems cope with hidden interactions, we will explore the basic attributes of complex systems, and be able to characterize them more systematically.

Coping with Hidden Interactions

Linear systems also have interactions that are not visible, of course, but they occur within well-defined and segregated segments of the production or maintenance sequence. Controls, such as dials, warning lights, audible alarms, and switches read the presence of these interactions and inform the operators and allow them to intervene. In systems with a fair degree of complex interactions, however, well-defined and segregated segments do not necessarily exist. Instead, jiggling unit D may well affect not only the next unit, E, but A and H also. This increases the number of controls that must be installed and monitored. The control panel of a nuclear plant is considerably denser and larger than that of an oil or coal plant, in part because so many components are linked in branching paths and feedback loops.

Attempts are continually made to reduce the number of controls by automating the subsidiary interactions and leaving only the main parameters for the operators to worry about. But this decreases the system's flexibility; the operator loses the ability to correct a minor failure in a part rather than shutting down a whole unit or subsystem. The operator cannot exit from the high-level, summary controls to the low-level specific one required to deal with a single part. Larry Hirschhorn provides an instructive example in a perceptive article on the limits of "cybernetic" (self-correcting) systems, which I shall elaborate on considerably.[7]

Consider a piston engine and a fuel pump. If particles get into a cylinder, perhaps during a maintenance procedure, or from wear from some other part, they can cause the piston to move more sluggishly. The result is that there is less output from this part of the engine. A control device monitors the output (but not the many causes of reduced output, such as foreign particles, piston wear, low octane fuel, over- or under-heating).

Noting the reduced power, the automatic control calls for more fuel from the fuel pump. This does compensate for the new source of friction, and the machine functions adequately, though less efficiently. But suppose that a surge in power is called for from the engine. The fuel pump is asked to increase its output. Soon its limit is reached. A typical cybernetic monitor will now decide that the fuel pump has failed; it will not know that the requirements expected of a (now degraded) piston/cylinder part cannot be met. An engineered safety device comes on, sending fuel from another source to the cylinder.

Any number of things might happen at this point. If the first pump is left running (since it has not, in fact, failed at all) there will be too much fuel supplied, which might burst some fuel carriers, stall the engine with fuel flooding in, overpressurize the fuel supply so that it tries to flow backwards (an eventuality that was not conceived by designers, so no protection is provided), or cause a runaway engine that rapidly overheats and explodes.

If, instead of the above, power to the fuel pump is automatically shut off because it is believed to have failed, the pump may be shut off with the valves in the open position, rather than the closed position they would be in if the pump had indeed failed. This could create back-flow or other problems since the other systems have not been shut off. Designers might have anticipated that if an undamaged pump were shut off, it would be because of a power failure that would affect other units of the system, and thus no back-flow problems would occur. In this case, however, an undamaged pump loses its power but keeps its valve setting, while other related parts or units do not lose their power and thus keep operating.

After the failure, the pump would probably be replaced, since the monitor indicated that this is what had failed. But the problem would remain, since it resided in the piston; only the automatic cybernetic control system thought the problem was the pump. This elaborated example is fabricated, but it illustrates the difficulty control systems have in reading the nature and source of failures. In a system without this added complexity, an operator presented with this sequence would probably cut back on the power demands, reasoning that either the engine or the fuel supply was malfunctioning, check the pump by varying engine speed, and conclude that there was sluggishness in the engine, which in fact is the correct conclusion.

One limitation of this example is that there is no pressing reason to have automatic controls in such a simple system (though if the motor exists in connection with many other units or subsystems, it may be

advisable). But in some systems automatic controls are necessary because there is simply not sufficient time for operator reaction. The nuclear reactors designed by Babcock and Wilcox are of this type. The reactor is extremely responsive, and this has important economic benefits. But some of its subsystems may be too responsive. With a loss of coolant, the operator may have only thirty to sixty seconds before the steam generator boils dry. Consider this accident at the Crystal River, Florida, plant in February 1980.

For an unknown reason a short circuit occurred in some of the controls in the control room. The utility said it could have been due to a bent connecting pin in the control panel, so sensitive are these devices, or the malfunction may have been caused by some maintenance work being done on an adjacent panel. The short circuit distorted some of the readings in the control system, in particular the important and very sensitive one of coolant temperature. The computer "thought" the coolant was growing too cold, so it speeded up the reaction in the core. (The Babcock and Wilcox reactor operates within a very efficient, but quite narrow temperature band.) The reactor overheated, the pressure in the core went up to the danger level, and then the reactor automatically shut down. The computer was apparently now puzzled, and correctly ordered the pressure relief valve (PORV) to open, but incorrectly ordered it to remain open until things settled down. This was an "error" on its part, because pressure dropped so quickly that it automatically caused high pressure injection to come on, and stay on, flooding the primary coolant loop—including the core, steam tubes, and pressurizer. A valve stuck and 43,000 gallons of radioactive water were dumped on the floor of the reactor building. Fortunately it was not worse; after several minutes an operator noticed the computer's error in keeping the relief valve open and closed the valves manually. Had the frequent injunction been followed that the computer knows best, and that the (dumb) operators should keep their hands off the system until its routines are carried out, the sump would have overflowed and we would indeed have had a wet reactor.[8]

Operators tend to resist the introduction of more general, high-level controls such as cathode ray tubes (CRTs), which produce a television screen display of the state of a number of units or subsystems, because they feel they cannot make selective interventions as easily. Only these high-level controls can be manipulated; the selective ones are more difficult to get to because it is assumed they will not be needed. They may be out of the way, or accessible only by a complicated series of steps that inactivate the more general controls. On the other hand, operators also

complain about the arrangement of less automated systems that allow selective intervention, because they are confronted with 15-foot banks of identically designed switches with tiny numbers above them. These are not even grouped in any way that reflects operation, but only in the way that reflects ease of installation. One of the most common "operator errors" in nuclear plants is, understandably, operating the wrong switch. Surely the choice need not be reduced discretion versus endless, identical displays.[9]

Complicating the dilemma of too few vs. too many controls is the problem of uncertain "default" status. Default status is the normal status of a control; for example, you must choose to change a switch to "on"; by default it is "off." The position of a relief valve is normally closed. The temperature reading on a drain pipe is normally cool. The position of most block valves is normally open. One might decide to have the default status represented by a green light on the control panel—go ahead, there is no danger. But one also might decide to have green represent open valves or circuits that are open to current flow—they are "on," and on is the default status. One can readily see the problem: if something is normally closed, this status cannot be represented by the color green, since green also means it is on, or open. Nor can red be reserved for "danger"; if the reactor is supposed to be on, you may want a red signal to warn you that it *is* on, even though it is not dangerous, but normal under these circumstances. But then what would green mean? (I learned the hard way that on my personal computer closing a "dip switch" means to open it!)

Some switches and valves have to be on at certain times and off at others, so there is no default position. Some systems use colored lights to handle this. If operating mode A is in effect, then switch 1 should be on, and an amber light next to the switch indicates this. If operating mode B is in effect, the switch should be off. If the switch is in the wrong position for the operating mode, a red or blinking light appears. If that part of the system is not in operation, a green light may be on, regardless of the position of the switch, or no light may be on. But modes and switches are not always clearly linked; there may well be a mode A and a mode A1, and the correct switch positions for the two would be different. Operators sometimes disconnect these sophisticated indicators (or much more commonly, ignore them) because of such complications.

These problems exist in all industrial and transportation systems, but they are greatly magnified in systems with many complex interactions. This is because interactions, caused by proximity, common mode connections, or unfamiliar or unintended feedback loops, require many

more probes of system conditions, and many more alterations of the conditions. Much more is simply invisible to the controller. The events go on inside vessels, or inside airplane wings, or in the space craft's service module, or inside computers. Complex systems tend to have elaborate control centers not because they make life easier for the operators, saving steps or time, nor because there is necessarily more machinery to control, but because components must interact in more than linear, sequential ways, and therefore may interact in unexpected ways.

In addition to there being many interactions to control, the information about the state of components or processes is more indirect and inferential in complex systems. The reactor core at Three Mile Island had no direct measure of coolant level. Given the high pressures and the flow of coolant in the core, such a measure would be very hard to provide, and would introduce more penetrations of the vessel, something to be avoided. So, operators had to estimate coolant level from indirect indicators. The presence of steam voids and hydrogen bubbles made even these unreliable.

To cite other examples, pilots in ships or airplanes can fix on the wrong star. Ships may find navigation beacons obscured by shore lights, distorted by refractions, or simply out. Routine fluctuations in pressure and temperature in chemical plants can mislead operators. Operators discount an abnormal reading, thinking it is part of the routine fluctuation. Generally they are correct; occasionally they are wrong. At TMI the operators knew the temperature on the hot leg of a drain pipe was abnormal because of a leak. During the accident, when the temperature exceeded even this abnormal reading, they treated it as a particularly wide fluctuation, whereas it really indicated an open relief valve.

In 1977 New York City experienced a massive and very costly blackout. One key contribution to the accident was an operator's expectation about the default reading for current flowing over a particular line. Normally that line carried little or no current. The operator did not know there had been two relay failures—one that would automatically lead to a high flow of current over that line; and a second that blocked the flow over the line. The operator treated the zero current reading as normal. In fact, it was abnormal, but only in this particular set of circumstances. This ambiguity led to a systematic, by-the-book sequence of actions to handle the problems that were showing up in other parts of the system, ending in the system being brought to a halt.[10] The only evidence the operator could really "see," in terms of sensory confirmation, was the lights going out in his control room.

In a typical (linear) manufacturing plant both errors and interactions

are more visible. If an operator misunderstands an instruction the supervisor will probably know it soon enough because she will see the operator set up for the wrong task. I am told that one of the advantages of old-style steam valves, where the stem of the valve will rise when opened, is that a quick glance by an operator or supervisor over a huge room will show which valves are open, and which are shut—they just stick up when open. If an employee heads for the wrong valve after misunderstanding an order, that will be quite visible too. In complex systems, where not even a tip of an iceberg is visible, the communication must be exact, the dial correct, the switch position obvious, the reading direct and "on-line."

The problem of indirect or inferential information sources is compounded by the lack of redundancy available to complex systems. If we stopped to notice, we would observe that our daily life is full of missed or misunderstood signals and faulty information. A great deal of our speech is devoted to redundancy—saying the same thing over and over, or repeating it in a slightly different way. We know from experience that the person we are talking to may be in a different cognitive framework, framing our remarks to "hear" that which he expects to hear, not what he is being told. The listener suppresses such words as "not" or "no" because he doesn't expect to hear them. Indeed, he does not "hear" them, in the literal sense of processing in his brain the sounds that enter the ear. All sorts of trivial misunderstandings, and some quite serious ones, occur in normal conversation. We should not be surprised, then, if ambiguous or indirect information sources in complex systems are subject to misinterpretation.

Transformation Processes

As we gain more experience with systems, and design them more effectively, the high degree of interactiveness may be reduced. In Chapter 5 we shall examine the case of air traffic control, where exactly this appears to have happened. It is also true that a poorly trained or inexperienced operator may see a system as replete with unsuspected interactions or "traps," but after gaining experience may find it to be more linear. Finally, new technological breakthroughs may bring linearity into a once complexly interactive system, as when the jet engine replaced the piston engine, or the transistor replaced the vacuum tube. Though there is a

tendency to then demand more of the redesigned system, and thus build in new interactions, as has been dramatically recounted in the case of a jet fighter,[11] we should emphasize that new or improved designs and more operator experience may reduce the possibilities of unexpected interactions. Some we shall examine in this book are aspects of marine transport, the air traffic control system, dams, and mining. But our major concern is with systems that appear to be irretrievably complex, or at least will remain so for the next few decades. Generally these are systems that transform their raw materials, rather than fabricate or assemble them.

Transformation processes exist in recombinant DNA technology, chemical plants, nuclear power production, nuclear weapons, and some aspects of space missions. Most of these are quite new, but it is significant that chemical processing is not. While experience has helped reduce accidents, accidents continue to plague transformation processes that are fifty years old. These are processes that can be described, but not really understood. They were often discovered through trial and error, and what passes for understanding is really only a description of something that works. Some of industrial chemical production is of this nature; much of iron and steel production was of this nature, although it has been greatly altered through scientific research.

The existence of transformation processes without full understanding certainly characterizes nuclear power; recall that the nuclear scientist who advised Governor Thornburg of Pennsylvania during the accident had asserted three years before in a scientific journal that there could be no problem with a zirconium-water reaction—the process was well understood. Yet precisely this problem produced the hydrogen bubble. Each space mission introduces new systems that may be unreliable because of lack of knowledge. Experience will increase understanding of some of these problems (our rockets do not blow up on the launch pad as readily as they once did), but the performance of space vehicles is still subject to much uncertainty. Recombinant DNA research is fraught with gaps in knowledge. Limited knowledge, then, allows unsuspected interactions, and requires many control parameters and indirect sources of information.

This completes our discussion of the attributes of systems with complex interactions. To summarize, complex systems are characterized by:

- Proximity of parts or units that are not in a production sequence;
- many common mode connections between components (parts, units, or subsystems) not in a production sequence;

- unfamiliar or unintended feedback loops;
- many control parameters with potential interactions;
- indirect or inferential information sources; and
- limited understanding of some processes.

Complex systems are not necessarily high-risk systems with catastrophic potential; universities, research and development firms, and some government bureaucracies are complex systems, as I will argue shortly. We can now briefly contrast complex and linear systems.

Linear Systems

The parts or units of linear systems that are not in a direct production sequence tend to be spatially spread out. Fabrication or assembly allow this, while transformation processes often have to be compact. Linear systems lack the common-mode connections that require proximity. It is also a design criterion to separate various stages of production for sheer ease of maintenance access or replacement of equipment. Linear systems not only have spatial segregation of separate phases of production, but within production sequences the links are few and sequential, allowing damaged components to be pulled out with minimal disturbance to the rest of the system. In complex systems, removing a component or shutting it down means temporarily severing numerous ties with consequent readjustments, capping, product storage, removal to get access, and reconfigurations because parts and units tend to be multiply linked. Linear systems also favor serial production—a series of linked but semi-independent production steps—rather than what organizational theorist James Thompson calls "pooled interdependence," where all components (including operators) must coordinate their input if the system is to function at all.[12]

In contrast to complex systems, there is minimal specialization of labor, materials, and pools of supply in linear systems. Equipment is specialized ("dedicated," engineers call it), but the people who run it tend to be generalists. Operators are trained on several tasks because they tend to rotate, bid on various jobs, or fill in for people. Maintenance people can operate equipment in an emergency, operators perform emergency maintenance. There are, of course, limits to such substitutions, including the demands set by unions. Later we shall see that substitutability is important for recovery from an accident—people can fill in and know some-

thing about the other person's job. Here, in discussing complexity and linearity, the emphasis is upon awareness of interdependencies. In complex systems, not only are unanticipated interdependencies more likely to emerge because of a failure of a part or unit, but those operating the system (or managing it) are less likely, because of specialized roles and knowledge, to predict, note, or be able to diagnose the interdependency before the incident escalates to an accident.

Consider some key jobs in complex systems and note how little substitutability there is between them: fighter pilots and maintenance; bombardiers and pilots and navigators; nuclear plant engineers and operators; operators and welders and other maintenance people rated to do specialized jobs; lab technicians and biochemists in a recombinant DNA firm; chemists with advanced degrees and lab technicians, in a chemical plant, or chemical engineers and control room operators; astronauts and ground control managers; navigators, radio operators, captain, and helmsmen on a ship.

Though I don't want to claim a vast difference between employees in complex and linear systems, the latter appear to have fewer specialized and esoteric skills, allowing more awareness of interdependencies if they appear. The welder in a nuclear plant is more specialized (and specially rated), and presumably more isolated from other personnel, than the welder in a fabrication plant. Specialized personnel tend not to bridge the wide range of possible interactions; generalists, rather than specialists, are perhaps more likely to see unexpected connections and cope with them.

What is true of labor is also true of materials and supplies. If materials and supplies are substitutable, more latitude of response is available, limiting failures to incidents and preventing failures in the first place. But complex systems appear to have more exacting requirements for materials and supplies; the fuel cannot be off-standard, nor one fuel substituted for another, whether we are dealing with nuclear plants, aircraft, spacecraft, or chemical production. Substitutions are more likely in linear systems.

Finally, linear systems have minimal feedback loops, and thus less opportunity to baffle designers or operators. There are fewer interactions of control parameters, because controls are more decentralized, and attached to special-purpose equipment. And the information used to run the system is more likely to be directly received, and to reflect actual operations.

The two systems are summarized in Table 3.1, with some summary terms listed that will be used in the rest of the book.

TABLE 3.1
Complex vs. Linear Systems

Complex Systems	*Linear Systems*
Tight spacing of equipment	Equipment spread out
Proximate production steps	Segregated production steps
Many common-mode connections of components not in production sequence	Common-mode connections limited to power supply and environment
Limited isolation of failed components	Easy isolation of failed components
Personnel specialization limits awareness of interdependencies	Less personnel specialization
Limited substitution of supplies and materials	Extensive substitution of supplies and materials
Unfamiliar or unintended feedback loops	Few unfamiliar or unintended feedback loops
Many control parameters with potential interactions	Control parameters few, direct, and segregated
Indirect or inferential information sources	Direct, on-line information sources
Limited understanding of some processes (associated with transformation processes)	Extensive understanding of all processes (typically fabrication or assembly processes)

Summary Terms	
Complex Systems	*Linear Systems*
Proximity	Spacial segregation
Common-mode connections	Dedicated connections
Interconnected subsystems	Segregated subsystems
Limited substitutions	Easy substitutions
Feedback loops	Few feedback loops
Multiple and interacting controls	Single purpose, segregated controls
Indirect information	Direct information
Limited understanding	Extensive understanding

Which Is Best?

This litany of problems with complex systems and the advantages of linear systems might suggest the latter are much preferable and complex systems should be made linear. Unfortunately, this is not the case. Complex systems are more efficient (in the narrow terms of production efficiency, which neglects accident hazards) than linear systems. There is less slack, less underutilized space, less tolerance of low-quality performance, and more multifunction components. From this point of view, for design and hardware efficiency, complexity is desirable.

It also appears that we have comparatively little choice in the design of some of our systems. Some complex systems can be redesigned to be

more linear, such as in air traffic control or with the substitution of jet engines for the very interactive piston engines. Nuclear plants could be made marginally less complex if the spent storage pool were removed from the premises. (Should a plant with a full load of fresh rods in its pool have to be evacuated and the pool left unattended or without power, the water could boil off in a few days. Without the water, they would commence to fission like sparklers on the Fourth of July.) This would eliminate some proximity and common-mode problems. Using one control room to run two different reactors also seems like an unnecessary source of common-mode problems. The steam generation and turbine systems could be separated from the nuclear system to reduce unintended feedback loops, though at considerable expense. And perhaps, as we noted in the last chapter, more "forgiving" designs exist. But by and large no extensive reductions in complexity seem possible in the nuclear power industry. The transformation system simply requires many nonlinear interactions. The same is true of chemical plants such as refineries; there is probably no efficient way to crack crude oil except in a highly interactive system.

On the whole, we have complex systems because we don't know how to produce the output through linear systems. If these complex systems also have catastrophic potential, then we had better consider alternative ways of getting the product, or abandoning the product entirely. Short of that there does not seem to be a choice possible between complexity and linearity for high-risk systems. Complexity is inherent in some forms of production. It is not intrinsically undesirable; we welcome complexity in some bureaucracies and resist the "rationalization" of our disorderly life because the unexpected interactions lead to innovations, amuse or interest us, or provide variety. If the system has catastrophic potential, however, and we cannot prevent the propagation of incidents and intervene before an accident has occurred, the issue is much graver. Let us see what the idea of tight and loose coupling has to say about this. It is our second major dimension in analyzing systems.

Tight and Loose Coupling

Engineers speak of tight and loose coupling in the normal course of their work, and in that context the terms are quite clear. It's not as simple as coupling a garden hose to the faucet, and doing it tightly if we don't want to get sprayed, but that is a fairly close analogy; *tight coupling* is a me-

chanical term meaning there is no slack or buffer or give between two items. What happens in one directly affects what happens in the other.

Sociologists and social psychologists took up the term in the mid-1970s to conceptualize a particular phenomenon.[13] Some public service organizations, public schools in particular, seemed to be characterized by an unusually large gap between official programs and actual behavior. In public schools researchers investigated the gap between new techniques and actual changes in the behavior of teachers, or between new programs and what the students actually learned. One might expect hierarchically organized bureaucracies to be capable of altering the behavior of subordinate personnel such as teachers and capable of producing outcomes in students that bear some relationship to the official goals. The explanation for the school's failure to achieve its goals was that while the programs and goals were real enough, they were only loosely connected to other matters with which the organization had to be preoccupied, such as political demands from the environment, the autonomy of professional teachers, and the ineffective mobilization of parental demands.

For example, there could be a tight connection (responsiveness) between a remedial program required by the school district and the school in that it is funded, staffed, and given space and a place in the curriculum. But the program might only be loosely coupled to other parts of the school. For example, two new teachers, presumably hired for the program, are actually the ones the school has long wanted to hire—a respected art teacher and a person who can teach a course in computers. They are used for these positions, and the two teachers who are the least respected by parents and students are assigned to the new program. The decisions might almost seem to have been made by themselves, so loosely coupled is the demand for the new program with the existing problems of the school.

The program budget is a concrete, discrete item, but its appearance does not result in purchasing the proper supplies and doing the expected remodeling. Instead, the supply orders are delayed, and the money used for gym equipment and refinishing the gym floor. There is no intent of fraud; the latter have long been pressing items, subject to an informal agreement with the district that the expenditures can be made just as soon as there is some slack in the system. The school principal reasons that the remedial program will be taken care of when the new budget passes. Space for the new program might turn out to be the most inconvenient space, on the incontrovertible grounds that all the convenient space is utilized, and the program should not disrupt successful existing programs. Existing programs are buffered, in this way, from the impact

of the new program. However, the district's intent was to alter the school's existing priorities, which would require that the remedial program should have an impact on—be tightly coupled to, and change—the existing priorities.

Finally, the school could find it easy—almost inadvertent—to assign students to the program not on the basis of their reading difficulties, but because they are disruptive, impolite, or racially stereotyped. They all look as if they need "extra help." And the classroom materials, carefully designed by educators in elite universities, may bear little relationship to the needs of the stigmatized students.

In all this a number of connections are being made—student social status and program assignment; supply shortages and the power of certain departments; low performing teachers and periphery assignments. The system is not lacking in connections, but they are "loose" ones. Some are invoked for this program but perhaps not for the next. Some were invoked by accident—a chance discovery that the way the budget was drawn up it could cover costs of the gym floor. Others might have depended upon fleeting events, perhaps a recent controversy regarding the effectiveness of two teachers. Certainly the connection between the official mandate for a remedial program and the actual practice is quite loose. The characteristics of the system—its loosely coupled nature—make it possible for it to respond to the outside demand in such a loose fashion.

Loosely coupled systems tend to have ambiguous or perhaps flexible performance standards, and they may, as in the case of the school, have little consumer monitoring, so the absence of the intended connection can remain unobserved.

You might be tempted to call it a very inefficient system, or even a ripoff of state or federal funds, and call for "tight coupling." This would make the remedial program work as everyone outside the system thought it should work. But it would be a mistake to call it inefficient. The system is quite efficient for accomplishing many things that many participants desire; they are just not what the school district or the federal government, which supplies the money, had in mind. A bit of authoritarian leadership designed to put the official program in effect might well disable parts of the system that others value. If the poor teachers remained in the regular curriculum, parents would continue to object to them; if the college preparatory subjects were put in the temporary building (temporary since the end of World War II) they might also object; everyone uses the gym and it needs fixing; and what is one to do with disruptive students who reportedly interfere with the learning of the docile ones?

Loose coupling, then, allows certain parts of the system to express themselves according to their own logic or interests. Tight coupling restricts this. Loose coupling, however, is not the same as disorganization, unless we mean lack of centralized control by that term. Informally, the school is well organized, with a variety of coherent (though occasional shifting or episodic) interests, mechanisms for accommodative arrangements, slack resources to meet unexpected challenges, and stable interaction patterns. The degree of organization is independent of the degree of coupling.

We have not strayed as far from the engineer's usage of tight and loose coupling as it may seem. Elaborating the concept as used by organizational theorists will allow us to examine the responsiveness of systems to failures, or to shocks. Loosely coupled systems, whether for good or ill, can incorporate shocks and failures and pressures for change without destabilization. Tightly coupled systems will respond more quickly to these perturbations, but the response may be disastrous. Both types of systems have their virtues and vices.

For example, a continuous processing plant requires tight coupling. In some continuous processing plants where the technology is well understood and only linear interactions take place, such as pharmaceutical plants, bakeries, confectionaries, or ball-bearing plants, the characteristics of the product are altered frequently in response to market demands. Processes are also altered frequently to reflect changes in raw materials or operating conditions. Decisions to alter the process or shift to a slightly different product must proceed quickly through the organization, producing nearly unreflective changes in operator behavior. Resources are strictly allocated, schedules strictly followed, and reporting systems must be exact. Surveillance, both of processes and people, is continuous, though it is usually through remote and unobtrusive indicators—at least when higher-level personnel are monitored. Deviations from standards are quickly noted or reported, since a long product stream is affected. Responses to deviations must be standardized and immediate. The result is high operating efficiency. To loosely couple such a production process would be to invite disasters, as well as inefficiencies. If the system is linear, rather than complexly interactive, tight coupling appears to be the optimum mode of organization.

In contrast, consider a loosely coupled linear aircraft manufacturing and assembly plant. Fabrication of the tail section will be separate from the fabrication of the fuselage section it will be attached to. Different metals are involved, different tolerances, and different heat-treating processes. The constraints on the two sections are different, and their

interdependence is minimal; the interdependence is taken care of in the design stage. Personnel practices may be different too, since particularly skilled personnel are needed for the tail construction, whereas fewer skills are needed for the fuselage, which involves more routine tasks than tail construction. This can lead to personnel decisions that resemble the allocation practices we discussed in the school, reflecting unit power and politics. Quality control, instead of being "built into" the system as in continuous processing, may be a "stand-alone" function, performing its own tests and inspections, physically isolated in the plant, and deliberately not integrated into the other functions because its independence must be maintained. But this buffering of quality control also means loose coupling, with the possibility of parochial interests developing, or unauthorized bargains with units.[14]

Characteristics of Coupling In Systems

We are now ready to more systematically lay out what is meant by tight and loose coupling in systems. These characteristics are reasonably independent of our other major dimensions—complex interactions and linearity.

1. Tightly coupled systems have more time-dependent processes: they cannot wait or stand by until attended to. Sometimes this is expressly for efficiency reasons, but generally it is because the production process does not allow for cooling and reheating, for forgetting and then relearning. Storage room may not be available, so products must move through continuously. Reactions, as in chemical plants, are almost instantaneous and cannot be delayed or extended. In loosely coupled systems, delays are possible; processes can remain in a standby mode; partially finished products (tail assemblies or students) will not change much while waiting.

2. The sequences in tightly coupled systems are more invariant. B must follow A, because that is the only way to make the product. Partially finished products cannot be rerouted to have Y done to them before X; Y depends upon X's having been performed. Though it may be expensive, in an aircraft assembly line a door or seat or the radio controls could be added later if there were a disruption. This is not the case for a nuclear or chemical plant, or a pharmaceutical production line. In a trade school the sequence of courses is carefully prescribed and more

tightly coupled than a university where it does not matter much when the science or math or language requirement is met.

3. In tightly coupled systems, not only are the specific sequences invariant, but the overall design of the process allows only one way to reach the production goal. A nuclear plant cannot produce electricity by shifting to oil or coal as a fuel; but oil plants can shift to coal and vice versa with little or no reconversion. While continuous processing plants (tightly coupled) can vary product mixes and production volume within limits, they cannot make significant changes in the way the product is made. However, manufacturing plants (generally loosely coupled) may, even on a temporary basis, eliminate heat-treating by substituting another metal; may shift from unit to batch production; contract out for subassemblies; install or remove robotic devices; substitute plastic for metal, and so on. In this sense, there are many ways to produce the item. But for tightly coupled systems such as dams, chemical plants, power grids, bakeries, and recombinant DNA technologies, there is little flexibility. Loosely coupled systems are said to have "equifinality"—many ways to skin the cat; tightly coupled ones have "unifinality."

4. Tightly coupled systems have little slack. Quantities must be precise; resources cannot be substituted for one another; wasted supplies may overload the process; failed equipment entails a shutdown because the temporary substitution of other equipment is not possible. No organization makes a virtue out of wasting supplies or equipment, but some can do so without bringing the system down or damaging it. In loosely coupled systems, supplies and equipment and humanpower can be wasted without great cost to the system. Something can be done twice if it is not correct the first time; one can temporarily get by with lower quality in supplies or products in the production line. The lower quality goods may have to be rejected in the end, but the technical system is not damaged in the meantime.

Recovery from Failure

Coupling is particularly germane to recovery from the inevitable component failures that occur. One important difference between tightly and loosely coupled systems deserves a more extended comment in this connection. In tightly coupled systems the buffers and redundancies and substitutions must be designed in; they must be thought of in advance.

In loosely coupled systems there is a better chance that expedient, spur-of-the-moment buffers and redundancies and substitutions can be found, even though they were not planned ahead of time.

Since failures occur in all systems, means to recovery are critical. One should be able to prevent an incident, a failure of a part or unit, from spreading. All systems design-in safety devices to this end. But in tightly coupled systems, the recovery aids are largely limited to deliberate, designed-in aids, such as engineered safety devices (in a nuclear plant, emergency coolant pumps and an emergency supply coolant) or engineered safety features (a more general category, which would include a buffering wall between the core and the source of coolant). While some jury-rigging is possible, such possibilities are limited, because of time-dependent sequences, invariant sequences, unifinality, and the absence of slack. In loosely coupled systems, in addition to ESDs and ESFs, fortuitous recovery aids are often possible. Failures can be patched more easily, a temporary rig can be set up, a crane moved over, a pancake landing made without power. There may, fortuitously, be plenty of space separating a burning subsystem from other systems; it may be possible and relatively harmless to flood an area to put out a fire, though no designer planned for that. Tightly coupled systems offer few such opportunities. Whether the interactions are complex or linear, they cannot be temporarily altered.

This does not mean that loosely coupled systems necessarily have sufficient designed-in safety devices; typically, designers perceive they have a safety margin in the form of fortuitous safety devices, and neglect to install even quite obvious ones. Most of the safety devices required after inspections by the Occupational Safety and Health Administration (OSHA) are fairly obvious items, such as railings, rough treads on stairs or passageways, emergency switches to shut off equipment, and requirements that the coupling for hoses be arranged so that an explosive gas such as hydrogen cannot be inadvertently supplied to an area that needs an inert gas. Even in loosely coupled systems there is still enormous room for improvement in safety features.

Tightly coupled systems are not completely devoid of unplanned safety devices. In two of the most famous nuclear plant accidents, Browns Ferry and TMI, imaginative jury-rigging was possible and operators were able to save the systems through fortuitous means. At TMI two pumps were put into service to keep the coolant circulating, even though neither was designed for core cooling. Subjected to intense radiation they were not designed to survive; one of them failed rather quickly, but the other kept going for days, until natural circulation could be established. Some-

TABLE 3.2
Tight and Loose Coupling Tendencies

Tight Coupling	*Loose Coupling*
Delays in processing not possible	Processing delays possible
Invariant sequences	Order of sequences can be changed
Only one method to achieve goal	Alternative methods available
Little slack possible in supplies, equipment, personnel	Slack in resources possible
Buffers and redundancies are designed-in, deliberate	Buffers and redundancies fortuitously available
Substitutions of supplies, equipment, personnel limited and designed-in	Substitutions fortuitously available

thing more complex but similar took place at Browns Ferry. The industry claimed that the recovery proved that the safety features worked; but the designed-in ones did not work. In both accidents the critical safety features were disabled.

What is true for buffers and redundancies is also true for substitutions of equipment, processes, and personnel. Tightly coupled systems offer few occasions for such fortuitous substitutions; loosely coupled ones offer many.

The characteristics of the two systems are summarized in Table 3.2. These are tendencies; no system has, for example, absolutely invariant sequences. Nor is a system likely to have all of the various characteristics.

The Organizational World According to Complexity and Coupling

Figure 3.1, the Interaction/Coupling Chart (I/C chart) puts interaction and coupling together in a two-variable array. The placement of systems is based entirely on subjective judgments on my part; at present there is no reliable way to measure these two variables, interaction and coupling. One may well quarrel with the placement in the figure. One good reason for a quarrel is that there is no precise specification of just what constitutes the system. Consider marine transport. On the high seas the system includes the ship, radio communication, the weather, and perhaps one other ship. But as the ship enters a crowded channel, we have not only the weather, but bank effects (the suction created as the ship passes close

FIGURE 3.1
Interaction/Coupling Chart

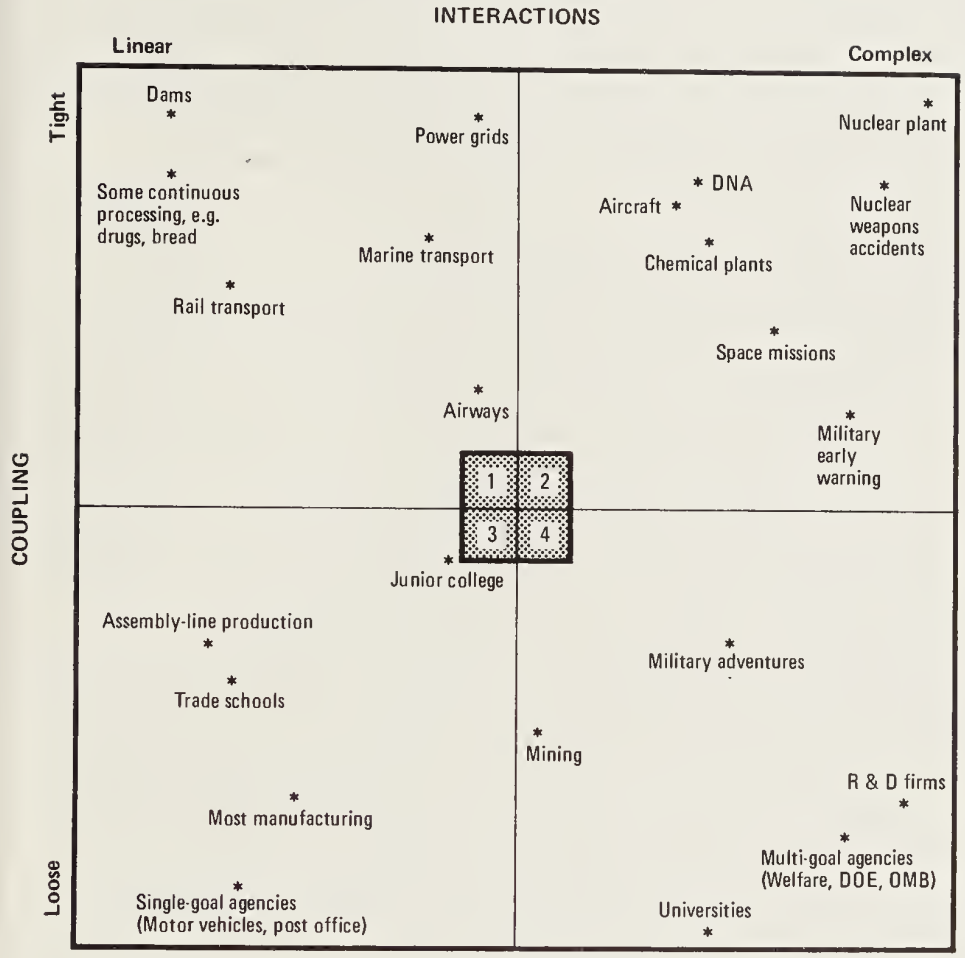

to the underwater bank of a channel), tides and current flow, wrecks and rocks, bridges, tows and other ships, a crowded radio channel, and navigation lights mixed in with the lights of highways, industrial plants, and distant towers. The high seas and the channel systems are very different. With "flying" we mix commercial airlines with recreational or sport flying. Mining includes strip mining, which should be closer to manufacturing, and underground mining. These ambiguities are a problem, but for the sake of a general argument I think the rough placement of vaguely defined systems is still useful.*

*One serious problem cannot be avoided, but should be mentioned: To some unknown extent it is quite possible that the degree of coupling and type of interactions have been

By combining our two variables in this way, a number of conclusions can be made. First, it is clear that the two variables are largely independent. Examine the top of the chart from left to right. Dams, power grids, and nuclear plants are all roughly on the same line, indicating a similar degree of tight coupling. But they differ greatly on the interaction variable. While there are few unexpected interactions possible in dams, and not that many in power grids, there are many in nuclear plants.

Or, looking across the bottom, universities and post offices are quite loosely coupled. If something goes wrong in either of these, there is plenty of time for recovery, nor do things have to be in a precise order. In the post office, mail can pile up for a while in a buffer stack without undue alarm; people tolerate the Christmas rush, just as students tolerate lines at fall registration. But in contrast to universities, post offices do not have many unexpected interactions—it is a fairly well laid out (linear) production sequence without a lot of branching paths or feedback loops. This is not true of universities. Here there are multiple functions—teaching, research, and public service, for example—and they can interact in unexpected ways. Indeed, they are expected to; synergistic interactions are desired, though "negative synergy" is also possible.

For example, let us suppose that a university reviews the record of an assistant professor of sociology, and decides that not enough has been published to warrant promotion to associate professor and the tenure or job security that goes with it. It is all quite straightforward, since research and publication are clearly specified as performance criteria for faculty and are outputs of the system. But imagine that the faculty member is a very good teacher, and the students protest the action. Imagine also that dismissal will threaten a public service program that she ran. Not only do the students protest, but some influential community members protest. To convince the students that the department is not indifferent to teaching quality, it has to institute a teacher evaluation program, and this now threatens the status of the senior faculty (associate or full professors with tenure) who have not had to pay much attention to this criterion before. Furthermore, the local churches send a delegation of church officials to protest the demise of the service program, and in the process,

---

inferred from a rough idea of the frequency of system accidents in the various systems, rather than derived from analysis of the properties of the systems independent of the nature of their failures. That is, if there are few accidents caused by air traffic control, that "must" mean it is not highly complex and tightly coupled, and then evidence for that conclusion is sought. Since the analytical scheme evolved from examination of many systems, there is no way to avoid this possible circularity. The scheme would have to be tested by examining systems not included here, as well as by collecting data based upon more rigorous concepts and verifying the placement of the systems that are included here.

demand more say in university personnel practices. The dean now has three problems on her hands: the department has to publish more (to justify firing the assistant professor), spend more resources on teaching to mollify the students, who have now made teaching a part of the formal personnel evaluation, and face a mobilized external group that demands some say in personnel decisions. This is an example of complex interactions; in contrast to the linear post office, the university is high on complexity.

But note that we are not likely to have a "system accident." Because universities are loosely coupled there is ample slack to limit the impact of this one personnel decision on other areas. Senior faculty members can be quietly reassured that student evaluations will be processed slowly, and evidence of the unreliability of evaluations will be discreetly passed about. (In one case that I witnessed, the university couldn't find the $2,000 needed to keypunch and process a student-initiated teacher evaluation program, the university's first.) Community members will be promised more access to the Student Activities Council, but realistically they will not be expected to attend often. The university research foundation is quietly told to favor sociology applications in the next cycle of grants, or joint appointments are handed out to make the sociology program look more productive. No interruption of output is experienced. The event remains an incident, with damage limited to the "part"—the teacher. While interactive, recovery is easy because of loose coupling.

The post office and the university, then, are similar with regard to coupling; both can recover from upsets because sequences are not that inevitable, there is slack in resources, substitutions are possible. But the post office is largely limited to linear interactions, while the university is full of many potentially complex ones that can reach unexpectedly into other parts of the system.

Educational systems are represented in three places in the I/C chart: universities are loosely coupled and complex, as we have seen; trade schools are somewhat loosely coupled but very linear; and junior colleges are near the middle of both dimensions. Trade schools have a well-defined curriculum and fairly explicit set of subjects students are expected to learn. This permits linear interactions to dominate and there are few surprises. However, the cafeteria style of offerings allows students to retake courses, making it somewhat more loosely coupled than junior colleges, and trade schools are also staffed by part-time, often temporary instructors who come and go with the changing student demands. Junior colleges are a bit more tightly coupled because of sequencing within programs (blocks of courses) and because of personnel inflexibility. They are

CHAPTER 9

Living with High-Risk Systems

A crucial question may have been at the back of your mind as you have read this book: What is to be done? After having looked at all these systems, what do I propose as a solution? I have a most modest proposal, but even though modest and, I think, realistic, it is not likely to be followed. I propose using our analysis to partition the high-risk systems into three categories. The first would be systems that are hopeless and should be abandoned because the inevitable risks outweigh any reasonable benefits (nuclear weapons and nuclear power); the second, systems that we are either unlikely to be able to do without but which could be made less risky with considerable effort (some marine transport), or where the expected benefits are so substantial that some risks should be run, but not as many as we are now running (DNA research and production). Finally, the third group includes those systems which, while hardly self-correcting in all respects, are self-correcting to some degree and could be further improved with quite modest efforts (chemical plants, airliners and air traffic control, and a number of systems we have not examined carefully but should mention here, such as mining, fossil fuel

power plants, highway and automobile safety). The basis of these recommendations rests not only on the system accident potential for catastrophic accidents, but also on the potential for component failure accidents. I think the recommendations are consistent with public opinions and public values.

But though we have come a long way from the broken coffee pot that started this book, before we can conclude with recommendations we must still confront three substantial objections to any recommendations that involve abandoning systems or making drastic and costly modifications. (1) My recommendations must be judged wrong if the science of risk assessment as currently practiced is correct. Current risk assessment theory suggests that what I worry about most (nuclear power and weapons) has done almost no harm to people, while what I would leave to minor corrections (such as fossil fuel plants, auto safety, and mining) has done a great deal of harm.[1] Risk assessment is a new science, but it has clearly occupied the high ground in government and many intellectual circles and will expand in the decades to come. Thus, this science deserves scrutiny.

(2) My recommendations could also be wrong if it can be shown that they are contrary to public opinions and values, or, if they are not contrary, that those public opinions and values are ill informed and should be corrected, rather than respected. It turns out that there is a fair bit of quite interesting work in cognitive psychology suggesting that the public is ill informed, and ill equipped, because of the way it reasons, to properly make important decisions about very complex matters. I think this work is in error, and will briefly outline a critique of it, though a full critique would require more detail, space, and arcane discussion than we can afford here.

(3) A third objection to my recommendations is more basic to the theory of the book. It says there is a way to run these systems safely; it simply requires authoritarian, rigidly disciplined, error-free organizations, such as, for example, the naval nuclear submarines appear to have. There is an organizational solution, this view argues; it is just that we have not been willing to put such organization in place. We will have quite a bit to say about this objection, since our whole book has been, in a sense, an organizational analysis.

We have four tasks, then: to examine the new field of risk assessment, since it counsels taking risks that I think are unacceptable and improperly evaluated; to examine the field of decision making, since it argues that the public is poorly equipped to play a role in decisions on risk; to examine organizational dilemmas inherent in high-risk systems; and, fi-

nally, to show how the analysis of these three problems, plus our analysis of system accidents and system characteristics, would lead to some modest recommendations for reducing the risks we have been told we must run. Above all, I will argue, sensible living with risky systems means keeping the controversies alive, listening to the public, and recognizing the essentially political nature of risk assessment. Ultimately, the issue is not risk, but power; the power to impose risks on the many for the benefit of the few.

Risk Assessment

Not surprisingly, the appearance of so many catastrophic man-made processes (the sexist pronoun is, for once, fitting) has occasioned public concern and, in turn, a response from the vendors of these processes and a number of social scientists. A whole new field of inquiry has resulted—risk-benefit analysis, or risk assessment. While not as dangerous as the systems it analyzes, risk assessment carries its own risks, and so we will look at both it and at the "ill-informed" public concern. As we have seen through our running commentary on the subject of "operator error," the realities are not always as the experts see them.

The activity of risk assessment is not new. People with power have always commissioned risk assessments. No important decision is likely to be made without some crude calculation of the probable benefits and the probable costs. Shamans, priests, court advisors, astrologers, lawyers, and so on have been the handmaidens of rulers and property owners throughout human history. But as power has expanded and become more centralized, so have the consequences of decisions grown, making assessment even more important. As the risks increasingly came from technological activities, scientists and engineers replaced the shamans as advisors.

Many of our potentially catastrophic systems today are not new, and already have been subject to a primitive form of risk assessment. Mining, chemical, and munitions disasters have been with us for two centuries; the failure of bridges and ships goes back to antiquity; even railroad and aircraft disasters are not all that ancient as the pace of change mounts geometrically. But third- and fourth-party victims were not present in catastrophic numbers for the older high-risk systems. Odysseus' vessel neither polluted the Mediterranean shoreline nor could destroy much of

Texas City; the World War II bombers could not crash into a building holding nuclear weapons, as happened at an unidentified overseas base in 1956;[2] chemical plants were not as large, as close to communities, or processing such explosive and toxic chemicals; airliners were not as big, numerous, or proximate to such large communities; and it is only recently that the risk of radiation from a nuclear plant accident has been visited upon almost every densely populated section of our country. The older systems now have more catastrophic potential because they are bigger and closer to us, and we have new systems that are inherently even worse. With nuclear power, nuclear weapons, and arguably, recombinant DNA, we have entirely new systems with catastrophic potential for third- and fourth-party victims—the innocent bystanders and future generations. There are few boundaries to these catastrophes, in space or in future time.

When societies confront a new or explosively growing evil, the number of risk assessors probably grows—whether they are shamans or scientists. I do not think it is an exaggeration to say that their function is not only to inform and advise the masters of these systems about the risks and benefits, but also, should the risk be taken, to legitimate it and reassure the subjects. With the increase in risk and public concern today, a new field of risk assessment has grown up, giving advice and (usually) legitimating the decisions of elites in the private and public sectors. At the behest of Congress, regulatory agencies have appeared in large numbers, and another function of risk assessors is to second-guess these agencies' awkward attempts to do a very difficult job. Risk assessors, interestingly enough, usually call for less regulation and are severe in their criticism of the agencies.[3]

The professionals in this field are generally engineers, scientists, and social scientists; they are based in universities, research organizations, government regulatory agencies, military establishments, and industry trade groups. Private, profit-making research or consulting groups, such as management consulting firms, undertake this profitable business for government and industry. Trade associations such as the Electric Power Research Institute conduct or sponsor risk-assessment studies. General Motors recently sponsored a conference on risk, and published the proceedings under the title, "How Safe Is Safe Enough."[4] The leading experts in the field were there, and a primary concern was not the risks of military or industrial activity, but the risks of regulation.

In addition to these groups, the large government and private foundations have sponsored risk-assessment studies. The National Science Foundation established programs and even a division devoted to the

topic, funding the work of management consulting firms, the Brookings Institute in Washington D.C., the Rand Corporation, and university research teams. The National Academy of Sciences (a semi-governmental body) established a committee for risk assessment some years ago headed by a prestigious authority, Howard Raiffa. The Russell Sage Foundation, arguably one of the most cautious and conservative of the large private social science foundations, recently declared this a high priority for its funding. Finally, most of the large universities now have study centers or research programs dealing with risk assessment, drawing their funding from government and industry grants, and professional journals exist to market the research results. The need is great, and the response has appeared.

This is a very sophisticated field. Mathematical models predominate; extensive research is conducted; and the esoteric matters of Bayesian probabilities, ALARA principles (as low as reasonably achievable), "discounted future probabilities," and so on are debated in courtrooms as well as academic conferences. Some of the best scientific and social science minds are at work on the problem of "how safe is safe enough."

Yet it is a narrow field, cramped by the monetarization of social good. Everything can be bought; if it cannot be bought it does not enter the sophisticated calculations. A life is worth roughly $300,000, one study concluded;[5] less if you are over sixty, even less if you are otherwise enfeebled. After taking into account age and potential earning power, a life is a life.[6] Death by diabetes should have the equivalent impact on people as death by murder, is the implication of a study that deplores the public's unawareness that the former is a cause of many more deaths than the latter. "Unfortunately," they say, "there is evidence that peoples' perception of risks are subject to large, systematic biases. . . . Such biases may misdirect the actions of public interest groups and government agencies resulting in less than optimal control of risk."[7] This bias in reasoning is due largely, they imply, to sensationalism of the media.[8] But consider how a murder death affronts human values such as dignity, and the desire for security and predictability; the researchers themselves note it is not to be equated with a diabetes death, and public estimations of death rates reflect that, but the public is still held to be "biased." To take another case, for some economists and risk assessors (often the same people) there is no difference between the death of fifty unrelated people from many communities and the death of fifty from a community of one hundred. Social ties, family continuity, a distinctive culture, and valued human traditions are unquantified and unacknowledged. Fifty thousand highway deaths a year are equivalent to a single catastrophe

with fifty thousand casualties for these experts, and they deplore the fact that the public protests nuclear plants and estimates highway deaths to be only half of what they are.[9]

The field acknowledges the difference between voluntary risks such as skiing and hang-gliding, and involuntary ones such as leaching of chemical wastes.[10] But it does not acknowledge the difference between the *imposition* of risks by profit-making firms who could reduce that risk, and the *acceptance* of risk by the public where private pleasures are involved (skiing) or some control can be exercised (driving). All are bundled up in a vague reference to market principles, as if we would not have heat and light without X number of dead miners or irradiated nuclear glow boys. The literature reflects a rational, calculative marketplace theory of cost-benefit analysis. The technical literature is fond of pointing out that we spend millions of dollars in safety devices to save one nuclear power worker, but refuse to spend $80,000 to save the driver of an automobile. (That is, the benefits from, say, an emergency core cooling system and an automatic seat belt are figured on the basis of how many lives one expects to save, and of course the cost of the ECCS and the seat belt are vastly different, as is the number of lives to be saved.)[11] It is thus irrational to spend that much money on the nuclear plant;[12] we should spend it on seat belts, highway guardrails or anti-smoking literature. It is as if there were a fixed budget category for safety, regardless of whether corporation profits or private needs are involved, and the budget, being fixed, cannot be enlarged when new risks come along.

Reading this discussion can lead one to imagine a scenario such as this: At the board meeting of a large corporation, the vice president for finance has received advice from the risk assessors. He announces that by not installing some safety devices we will kill one more worker per year. This will not affect the supply of candidates for this death, since workers will still take jobs at the plant because the labor market is depressed, and each worker decides there is a very good chance that someone else will be killed. On the benefit side, killing that worker will mean that the corporation will avoid a cost of $50 million for the safety devices. This avoids a price rise or a cut in the dividends or management bonuses. The vice president might estimate that they can avoid a one-dollar price rise on 20 million items, and avoid a cut in dividends of $30 million. By killing the worker, the public and the stockholders will obviously greatly benefit. What is a life worth? Well, he figures, 50 million is pretty high for a random, anonymous worker, so let's do it. The vice president for finance is correct; in risk analysis terms, it is a good bargain. Something similar took place at the Ford Motor Company when it

decided not to buffer the fuel tank in the Pinto, and at the General Motors Company when it rejected warnings from engineers that the Corvair would flip over for the lack of a $15 stabilizing bar.[13]

Risk-benefit analysis, with its monetarization of cultural goods and values, has been succeeded by cost-benefit analysis, with its more open concern with the dollar as the ultimate solvent for all things social. Baruch Fischhoff, in a thoughtful examination of cost-benefit analysis (the article has the engaging title, "Cost-Benefit Analysis and the Art of Motorcycle Maintenance"), notes another consequence of the monetarization of social good by economists.[14] Cost-benefit analysis is "mute with regard to the distribution of wealth in society," he notes. "Therefore, a project designed solely to redistribute a society's resources would, if analyzed, be found to be all costs (those involved in the transfer) and no benefits (since total wealth remains unchanged)." Risks from risky technologies are not borne equally by the different social classes; risk assessments ignore the social class distribution of risk.

Cost-benefit analysis also relies heavily upon current market prices for evaluating costs and benefits. Yet these reflect current economic arrangements that many might question and wish to change. For example, people with low earning power can receive lower prices on their lives. The current market price for temporary nuclear workers is quite low, given a long recession; when calculated, the costs of replacing steam generator tubes reflects this.[15] This can mean that the cost of an accident is low only because the economic system places a low value on some people. Property values near a chemical plant are likely to be low because of odors, fumes, and fire and explosion risks. When an accident takes place, the damage to the environment is calculated in terms of values already depressed because of the accident potential, rather than what the land would be worth if an electronics plant were there, or a nice park.

Another consequence of the prevailing assumptions is the argument that new risks should not be any higher than existing ones we have "already accepted" (Have we really had much choice?), and the corollary that if other industries get much riskier, the safety levels of nuclear plants or chemical plants can be lowered. This argument recently influenced the decision of the Nuclear Regulatory Commission when it set safety goals for its nuclear plants. As a colleague, Kim Scheppele, points out, as society gets more dangerous, the NRC can allow plants to get more dangerous.

Another argument that we hear much of these days is that we must push ahead with risky endeavors or other companies or nations will beat us in the competitive race to the marketplace. This line of reasoning

notes that our country was founded in risk and grew powerful by taking risks, and the social benefits have been enormous. We should, for example, push ahead with genetic engineering or the Japanese will beat us. But why push ahead if they can do it with more safeguards? If there are vast public benefits to come from genetic engineering, does it matter all that much if we buy these benefits from Japanese firms rather than from U.S. firms? It will certainly matter to the owners and investors in U.S. firms, because the oil companies and the pharmaceutical and chemical companies will lose private profits and dividends to stockholders. (Some jobs will be lost through such innovations, but those are not labor-intensive industries.) But if the Japanese will reduce the risk of a catastrophe by even a tiny amount with their better controls, the private profit losses and the indirect losses to our economy will be well worth it. What risks should we run to promote the private profit of a few? This cannot be factored into an economist's model. There, a dollar saved is a dollar saved, no matter who gets the dollar or who lives a riskier life to save it for them.

There are those who argue that we are losing our moral fiber because we no longer want to take risks with technologies.[16] But it is striking that those who feel we have abandoned risk in our search for security are speaking only of technological risks associated with large corporations and private profits, or aggressive military postures. The corporate and military risk-takers often turn out to be surprisingly risk-averse (to use the jargon of the field) when it comes to risky social experiments which might reduce poverty, dependency, and crime. Though the following list of proposals may strike one as fanciful, they all involve substantial risks that liberals and leftists have suggested, but the corporate and military risk-takers would not want to try because of the consequences for the class structure, their power, and their values. The proposals include guaranteed income maintenance plans (as found in Europe); truly progressive taxation; investment in poor and declining areas (U.S. banks pay one of the lowest income tax rates in the nation—about 3 or 4 percent—but refuse to risk investing in poor inner city areas); heroin maintenance programs to reduce crime; unilateral nuclear disarmament (a risky venture that could not only reduce the risk of nuclear accidents but promote economic prospects); withdrawal from Central America; and so on. The risks that made our country great were not industrial risks such as unsafe coal mines or chemical pollution, but social and political risks associated with democratic institutions, decentralized political structures, religious freedom and plurality, and universal suffrage.

Nor do the risk-assessment and risk-benefit studies distinguish be-

tween addiction and free choice in activities (the equivalent of our distinction between forced and unforced operator error). Along with highway fatalities, lung cancer from smoking is the favorite referent of the new body-counters. It is treated as a voluntary activity, like hang-gliding. But most of us who smoke today do so because we were barraged with advertisements and inducements that soon addicted us. In World War II every packet of field rations held its five cigarettes per meal, and the sale of cigarettes to the armed forces was heavily subsidized and untaxed. Airliners used to pass them out gratis, presumably to calm one's nerves after takeoff. No Hollywood hero was without them. The promotion was intense, and so were the private profits. So addicted was the whole economy that the government subsidies to tobacco farmers today far exceed all that is spent on warnings and research by the government. Young people see a large number of adults who have not been able to break the habit and all are still barraged by advertising and smoking TV stars. Ironically, the health of an important sector of our economy (tobacco growing, cigarette sales, and advertising) depends upon the illness of the victims; the costs of stopping smoking are not only individual (because of addiction) but corporate. This is not a matter of free marketplace decisions made by informed consumers, to be ridiculed and compared to these same people's "irrational" attacks on nuclear weapons, nuclear power, or Love Canals. Smoking is a government-supported program of addiction, for immense private profit. An individual's addiction to smoking should not be compared to the costs industry must be forced to incur to reduce brown lung disease or make safer Christmas toys.

We could say the same of alcoholism and other forms of drug abuse. These also are the referents of the risk assessors, cited to show that the public is incapable of making sensible choices between the cost of more airline safety and the cost of "substance abuse."[17] Liquor advertising is substantial; the willingness of physicians to prescribe tranquilizers and other drugs is well known, and their abuse is said to far outdistance the use of illegal substances. There is little free market choice at work here.

Finally, the risk-assessing field only infrequently distinguishes between those activities over which the person has some control, however illusory it may be, and those where she or he does not. Driving is a key one. We appear to accept risks more readily when we think our skill will play some part in avoiding the hazard. We fear and reject risks where we are passive recipients of harm. The plant, we feel, not unreasonably, should not blow up, the dam break, the air controller goof, the Ford executives fail to protect a gas tank from exploding; over these risks we have little control. But we are willing to take our risks with driving, skiing, and

parachuting. Risk assessors treat the difference as one of voluntary or involuntary risk, but I think that misses a key point. Driving to work for many of us is about as involuntary an activity as there is, but at least we have some control over it. On the other hand, although we voluntarily fly in an airliner to a distant vacation, we have no control over the aircraft or the airways. We voluntarily attend large events at stadiums that sometimes burn or collapse, but we have no control over the architects or the construction firms, or the owners who always seem to lock the safety exits. Furthermore, "active risks," as we might call them, are generally not pursued for someone else's private profit; "passive risks" generally are.

For active risks, those that the individual performing the activity has some control over, the marketplace provides at least a rudimentary, though imperfect, way of addressing safety issues. Safer skis sell better; people stopped buying Corvairs and Pintos. Though there are exceptions, people make sensible choices when they have meaningful choices, and in time manufacturers respond. This is not nearly as true for those activities where we are passive recipients of risks in the control of organizational leaders. Though the airlines are interested in safety, and it is to their economic advantage to have safe travel or usage will decline, we still require a Federal Aviation Administration to encourage, require, and police them. For such activities as nuclear waste disposal, brown lung disease, toxic contamination of Times Beach, Missouri, or the Midland area of Michigan, or the Teton dam, we cannot count on any "market" to automatically incur the costs of more safety. These activities are beyond our control. For these, the government must step in.

In more and more areas of our life the government must step in. This is not the result of the cancerous growth of government, but rather is essential because our personal control over our environment and our activities is being steadily eroded by systems that we participate in, or are passively affected by. In some cases Congress recognizes this danger. There is a huge regulatory apparatus in place trying to control the generation of nuclear power, the NRC; but the regulation of chemicals and chemical plants is quite modest. It is even more modest in the case of mining, and almost absent in the case of DNA production. Nevertheless, one frequent refrain in the risk assessment literature is lamenting overregulation. Economist Ron Howard of Stanford would do away with all regulation;[18] Starr and Whipple, pioneers in the field (from the pioneering trade association, the Electric Power Research Institute), rather than being concerned about the risks industry imposes needlessly, pity industry. The costs of regulation "stem from the litigation, misplaced invest-

ment, retrofit, and costly delays." And what causes these high costs? They "result from industry's inability to predict the acceptance of risk by the public.[19] Perhaps risk assessors might advise industry that it is easy to predict an aversion to such passive risks as mercury poison, dioxin, DES, asbestos, and brown lung disease.

Two dangers of "active risks" should be noted. Consumers will not always voluntarily pay for safer products, and often are not attentive to the risks even if they are well known. I assume it will always be thus. Second, active risks are attractive; we like to take some risks, if we feel we have personal control over them. This means that as the risk declines with better equipment, more people may be brought into the activity, those who now feel the risk is reduced to the level they will tolerate. The end result is that the accident level may not change with the new safety devices. For example, as better ski equipment appeared, and the slopes were better groomed and more safely designed, the ski resort industry began to advertise heavily to attract novice skiers. More novices meant more accidents, including more accidents for the advanced skiers they ran into. While the play was more safe, the risk was increased because of more inexperienced players. The safety record of any active risk activity must include the number of participants and the proportion of new, unskilled ones.

The risk assessors, then, have a narrow focus that all too frequently (but not always) conveniently supports the activities elites in the public and private sectors think we should engage in. For most, the focus is upon dollars and bodies, ignoring cultural and social criteria. The assessors do not distinguish risks taken for private profits from those taken for private pleasures or needs, though the one is imposed, the other to some degree chosen; they ignore the question of addiction, and the distinction between active risks, where one has some control, and passive risks; they argue for the importance of risk but limit their endorsement of the approved risks to the corporate and military ones, ignoring risks in social and political matters. As I indicated earlier, risk assessment is not as risky as the systems being assessed, but it has its unfortunate consequences for our society nevertheless.

One unfortunate implication of quantitative risk assessment is that the public should be excluded from discussions that affect them.[20] Few of the risk assessors call for this outright, most imply it; some state that the public must be involved, but only on the risk assessor's terms, and a few reject the implication and genuinely think the public has something to contribute (principally the Decision Research group, an important private company in Eugene, Oregon, that has done the best work in the field

I believe, though they hedge on the value of the contribution). Thus the range of opinion on this is broad. Most seem to take the middle position: bring the public in but control them. This is to be done by "closing the gap between the expert and the public" (that is, them and us); but the gap almost always is to be closed in one direction only—by bringing the public over to the experts' side through education.

Why the gap? Ignorance on the public's part is the main reason offered. "Public perceptions and reality dramatically differ," says Howard Raffia, a leading expert at Harvard, and the experts should be concerned even if public perceptions are based on "infirm and dubious facts"; the public should be "informed."[21] Or, as William Clark asserts, "Society's attitudes towards risks such as cancer and nuclear reactors are not readily distinguishable from its earlier fears of the evil eye."[22] By definition, the experts should know more than nonexperts. I am sure the gap exists, as the experts define it, but that could be remedied in time. There is another assumption made, however, and to that we will turn in the next section: Even if they are given the facts, the public is deficient in proper reasoning powers. Here the risk assessors have powerful support from the cognitive psychologists, and cite them.[23] Humans in general do not reason well (even experts can be found to make simple mistakes in probabilities and interpretation of evidence); heroic efforts would be needed to educate the general public in the skills needed to decide complex issues of risk. At the basis of this is a quarrel about forms of rationality in human affairs.

Three Rationalities

Why would the public puff away on cigarettes while voting against nuclear power or marching toward disarmament? One possible reason is emerging from some clever and striking work in the area of decision making and cognition. We do not reason well, the psychologists tell us: we minimize some dangers and maximize others, and do not calculate the odds as a statistician would recommend. Some of their data is convincing, but is of little relevance, I will argue; much of their data is probably spurious. While rationality in humans is certainly limited, or "bounded," as it is termed, it is possible that precisely when confronted with disorderly data and discordant goals, this limitation is its greatest strength.

It is convenient to think of three forms of rationality: absolute rationality, which is enjoyed primarily by economists and engineers; "bounded" or limited rationality, which a growing wing of risk assessors emphasize; and what I will call social and cultural rationality, which is

what most of us live by, although without thinking that much about it. Absolute rationality we have encountered in the description of risk assessors' wherein calculations can be made about risks and benefits, clearly showing which activities we should prefer—such as nuclear power over coal-fired power plants. Even including deaths from the whole fuel cycle from mining to production, nuclear power is close to risk-free (the probabilities of a meltdown and breach of containment are minuscule), while coal power kills an estimated 10,000 people a year (through mining, transportation, and pollutants from old plants without scrubbers that remove particulates from the emissions that cause acid rain). The choice is obvious in terms of absolute rationality. Why then do we have from 20 to 40 percent of the people worried about nuclear power? The answer of some is that this is irrational thinking, and the implication—though it would be unwise to do more than imply it—is that the irrational public is incapable of participating in decisions about risk. The public is "hypercritical" about nuclear power, which is one of society's "worry beads."[24] If the public is unfamiliar with a suspected carcinogen, this "can lead to irrational concern and unnecessary alarm."[25] But if it is an irrationality, the public response is a burden the experts and the elites must bear. In their view, the social harm of such response is extensive; it includes protests, demonstrations, and a Congress that refuses to follow the experts.

Cognitive psychologists, those who study the process of thinking or cognition, have generally shared this view of absolute rationality, yet conducted hundreds of experiments that showed that people were considerably less than absolute in their rationality. Gradually, in trying to find out how people think, they began to abandon notions of ignorance and irrationality, and instead began to talk of the limits on rationality, or in Herbert Simon's terms, used in another context, bounded rationality. The limits on our ability to consistently and easily make fully rational decisions might be due in part to neurological limitation, to limits on memory and attention, to lack of education, and to lack of training in probabilities and statistics. But it also seems to be due to some very practical problems and concrete experiences in daily life. Hunches and rules of thumb and rough estimates and guesses appear to be patterned and widespread. Cognitive psychologists call these guesses "heuristics," from the word for discovery, and are now identifying several specific ones. For example, the "availability heuristic"[26] suggests that rather than examining all existing cases of some phenomenon, and then basing their judgment on all this experience, people tend to judge a situation in terms of the most readily available case, the one most easily remembered. If

there has recently been an airline crash, we focus on that event and ignore all the successful flights when we think about the probability of a crash while deciding whether to take a flight or not. Or, if asked whether the letter *r* occurs more frequently as the first or third letter in thousands of common English words, we tend to say "the first" because words are more "available" to us as we recall more words beginning with *r* than containing *r* as the third letter.

It is universally granted that heuristics are useful, time-saving devices, even if they sometimes or even often get us into trouble. While some cognitive psychologists are busily giving names to rules of thumb that people use (and pointing out how fallacious they are, and how they impede rational decision making), a few have begun to search further.[27] Why might these studies be useful? The inquiry is only beginning, but I think we can make some useful observations. First, heuristics prevent a paralysis of decision making; they prevent agonizing over every possible contingency that might occur. Second, they drastically cut down on the "costs of search," the time and effort to examine all possible choices and then to try to rank them precisely in terms of their costs and benefits. Third, they undergo revision, perhaps slowly, as repeated trials led to corrections of hunches and rules of thumb, and do so without expensive conscious effort. Finally, I think, they facilitate social life by giving others a good estimate of what we are likely to do, since we appear to share these heuristics widely. We may do something an expert would disagree with, but at least joint action with other nonexperts (the vast majority of people, by definition) is possible, even if that action is not the one best line of action.

Heuristics appear to work because our world is really quite loosely coupled, and has a lot of slack and buffers in it that allow for approximations rather than complete accuracy. Because our social, daily life is like this, it takes special training to perform with precision in technological worlds that are tightly coupled. These violate our common sense experience that "things will work out," or that they are not precisely and tightly coupled. Furthermore, not everything we attend to is all that important (certainly this is true of the famous test of estimating the frequency of the letter *r* as the first and as the third letter of words—much of the work subjects are required to do in the experiments is of this doubtful relevance to real life). It is possible that for decisions we realize are crucial, we set aside the convenient but possibly faulty rules of thumb and make much more careful decisions. Unfortunately, the psychological literature is very weak on what people actually do in real life. Even the striking work on medical decision making by Eddy actually concerns the recom-

mendations that medical text and reference books make with regard to diagnosis and treatment, rather than actual practices. The books turn out to ignore or misuse probabilities and underlying rates and thus the confident recommendations are dramatically faulty, and could have severely unfortunate medical consequences.[28] Laboratory experiments simply fail us here. They can tell us how unmotivated subjects (largely undergraduates in college psychology courses) take shortcuts to solve unrealistic and often extremely complicated problems, and while suggestive, the experiments may be misleading in ways we shall shortly consider.

One important and unintended conclusion that does come from this work is the overriding importance of the context into which the subject puts the problem. Recall our nuclear power operators, or the crew of the New Zealand DC-10 on its sightseeing trip, or the mariners interpreting ambiguous signals. The decisions made in these cases were perfectly rational; it was just that the operators were using the wrong contexts. Selecting a context ("this can happen only with a small pipe break in the secondary system") is a pre-decision act, made without reflection, almost effortlessly, as a part of a stream of experience and mental processing. We start "thinking" or making "decisions" based upon conscious, rational effort only after the context has become defined. And defining the context is a much more subtle, self-steering process, influenced by long experience with trials and errors (much as the automatic adjustments made in driving or walking on a busy street). If a situation is ambiguous, without thinking about it or deciding upon it, we sometimes pick what *seems* to be the most familiar context, and only then do we begin to consciously reason. This is what appears to happen in a great many of the psychological experiments. Without conscious thought (the kind that can be easily and fairly accurately recalled), the subject says, "This is like *x*: I will do what I usually do then." The results of these experiments strongly suggest that the context supplied by the subjects is not the context the experimenter expected them to supply. With an ill-defined context, the subject of the experiment may say, "Oh, this is like situation A in real life, and this is what I generally do," while the experimenter thinks that the subject is assuming it is like situation B in real life, and is surprised by what the subject does.

For example, take the supposedly widespread failure to take the "base rate" into account (all the past events, such as the number of flights where there were no accidents, or the proportion of single people in a community). If the problem is presented in one way, the subject will use the frequency and probabilistic reasoning to estimate a rate (since we do use them in many situations in life); if the problem is presented in an-

other way, with some vague but misleading cues, the subject will rule out probabilistic reasoning. This would explain several well-known experiments, by Kahneman and Tversky, such as the one where subjects are to estimate the probability that "Jack" is an engineer or lawyer, where 70 percent of the population are lawyers and 30 percent are engineers. When no description of Jack is given, the subjects estimate the probability on the basis of the base rate. When a supposedly "irrelevant" and "useless" description is given, they ignore the base rate. But the subjects probably figure that it would be silly for the experimenters to give a description that has no meaning, so they search for the meaning, and find something. In fact, there *is* a cue in what Kahneman and Tversky assume is a "useless" description. Jack has no interest in political or social issues, and therefore is not likely to be a lawyer, but he has many hobbies such as carpentry and mathematical puzzles, which suggests that Jack is likely to be an engineer.[29] Given these cues, the subjects may ignore the base rate because they assume the cues are meant to be used as guides, in spite of the base rate.

Finally, heuristics are akin to intuitions. Indeed, they might be considered to be regularized, checked-out intuitions. An intuition is a reason, hidden from our consciousness, for certain apparently unrelated things to be connected in a causal way. Experts might be defined as people who abjure intuitions; it is their virtue to have flushed out the hidden causal connections and subjected them to scrutiny and testing, and thus discarded them or verified them. Intuitions, then, are especially unfortunate forms of heuristics, because they are not amenable to inspection. This is why they are so fiercely held even in the face of contrary evidence; the person insists the evidence is irrelevant to their "insight."

This happens in the psychological experiments; subjects insist, even after their error is explained to them, that if a fair coin has come up heads twenty times in a row, it is slightly more likely to come up tails on the twenty-first toss. Subjects insist tails is the best bet because in the long run the proportion of heads can only be .5, or 50 percent, so this unusual run of solid heads has to end. They are wrong, the experts insist, because each toss is independent of every other toss; the coin has no memory and on each toss the odds are .5 that it will be tails, and .5 it will be heads. I must admit that I am a victim of this "gambler's fallacy" too. I reason that in 1,000 tosses there will be about 50 percent heads. Twenty-one tosses is more likely to reproduce this pattern of 50 percent heads in 1,000 tosses than, say, three or four tosses, so the odds have to be made slightly better than .5 after 20 heads, and this encourages me to choose tails.

The implication of the psychologists' work is that the public is not qualified to participate in decisions about the risks they will have to endure. The public pursues an informal, probably messy "logic" that the experts do not share. At least one cognitive psychologist has recently defended intuitions. Baruch Fischhoff wonders if our intuitive judgments, or logic, might not be utilized even if it is denigrated by the experts. "It is worth asking," he writes, "whether there is not a method in people's apparent madness. Are there not decision-making criteria overlooked by formal analysis yet essential for human welfare or psychological well-being?"[30] As we shall shortly see, there are such criteria, if we examine the work of Slovic, Fischhoff, and Lichtenstein carefully.

The difference between the absolute rationalists and the cognitive psychologists who emphasize bounded rationality can be illustrated by examining the public reaction to the accident at Three Mile Island. For the rationalists, TMI was merely the occurrence of a rare event, that would be expected to occur, say, once in three hundred years for that reactor. That it occurred a few months after start up in 1979 rather than in the year 2079 is insignificant. Estimates indicate it is likely to occur sometimes, but rarely; this was simply that rare time. For the bounded rationality group this is true, but not the point. This was an unfamiliar problem, and appropriate heuristics have not been developed for it. A significant event such as this is an indication to the public of what is possible; it is a signal that these plants can have serious troubles even though the experts say they will do so only very rarely.[31] Since some experts appeared to think it could almost *never* happen, expert predictions might reasonably be questioned by the public. If experts are wrong, then this may not have been the one time in three hundred years, but the first one of many, many times over three hundred years. (Note the bounded rationality theorist is not siding with the public on the risks of nuclear power, only saying that it is not irrational that they could come to their conclusion.)

Furthermore, the public might say, if there is even a remote chance of a catastrophe, why risk it? For the bounded rationality theorist, a quite efficient and understandable logic is at work here, even though it is technically faulty. It is an efficient logic, given the fact that experts, like everyone else, are fallible and have been proven wrong in the past. It is efficient to question them. Such logic is also efficient because it motivates the public to demand: "Remove that threat; I don't want to live with a lot of threats; you are not counting in my psychic costs; find another energy source."

Finally, bounded rationality is efficient because it avoids an extensive

amount of effort. For the citizen, think of the work that could be required to decide just what the TMI accident signified. In the experts' view, the public should make the effort to answer the following questions, and if they can't, should accept the experts' answers. Did the accident fit in the technical fault tree analysis the experts constructed in WASH-1400, the Rasmussen Report (a matter of several volumes of technical writing)? How many times in the past have we come close to an accident of this type? Can we correct the system and thus learn from the accident? Was it accurately reported? Do the experts agree on what happened? Did it fit the base rate of events that led to the prediction that it was a very rare event? And so on. The experts do not have answers to some of these questions, so the public, even were they to devote some months of study to the problem, could not be assured that the answer could be known.

We should note that the public is far from hysterical about nuclear power. Even after TMI from 60 to 70 percent of the public indicate more nuclear plants should be built.[32] But this does not mean that those favoring nuclear power make this judgment on the basis of a rational calculation using extensive knowledge any more than does the small minority—from 15 to 20 percent—who favor shutting down all plants. The lay proponents of nuclear power may not be using the correct decision procedures recommended by the experts any more than the lay opponents. But note that even the more enlightened of the bounded rationality proponents, while defending the efficiency of the public's reasoning and understanding the "natural" bases of its errors, might still conclude that public opposition to nuclear power is wrong. The public's fears must be treated with respect, and a way found to bring them into policy considerations, the bounded rationalists would say; but the gap is still to be closed by bringing the public over to the experts side.

Social Rationality

The third view of rationality, social and cultural rationality (or *social rationality* for short), departs from the absolute rationality of the risk assessors and economists even more significantly than bounded rationality does. It recognizes the cognitive limits on rational choice, but holds that such limits are less consequential in accounting for poor choices than cognitive psychologists believe and are, in fact, quite beneficial in other respects. Our cognitive limits may make us human in ways we treasure.

There are at least two general reasons why we might be thankful for our limited cognitive abilities. People vary in their cognitive abilities in

absolute terms, but they also seem to vary with respect to different thinking abilities for different tasks. You and I may be equally intelligent, when measured over a number of areas, but you are good at counting while I (as I tell my quantitative colleagues) don't count. Yet I have learned how to visualize, or model, things in three-dimensional space, or perhaps have an innate capacity for it. Because of my limitation in counting, I need you, and vice versa. Our limitations bring about social bonding. Bonding by diversity in skills (which is related to limitations in cognition, incidentally) is more stable and perhaps more satisfying than bonding by addition of equal talents. That is, the standard illustration of two people moving a rock that neither could move alone as the basis for social life is a very minimal one; any partner would do, and once the rock is moved, we can part. But bonding because sometimes we need to count and sometimes we need to visualize, so we had better have each other around when these tasks appear, is a strong basis for social life. If everyone were equally rational, we would not need economists. Since we are not, we need both economists who try to see where rational, quantitative solutions will work and sociologists who try to see how social bonding can be utilized and maximized.

A second cheer for our limitations stems from your propensity, for example, to see all problems as one of measurement and counting, and my propensity to see all problems as one of social interactions. If we have a common problem and it seems to have a lot of numbers, rates, proportions, and so on in it, you are likely to move quickly to a mathematical solution. Your "heuristics" are better than mine if numbers are included. But because of your expertise, you are very likely to end up deciding that the problem should be seen in a manner that allows a quantitative analysis. Your "framing" of the problem prejudges the problem and prejudices the answer. So does mine. For you, the choice of nuclear or coal can be measured by toting up the deaths per megawatt of power produced to date by each activity. The risks of DNA research can be measured by seeing how many experiments have gone on without any accidents. But I might define the power generation problem in terms of potential deaths in a rare but conceivable catastrophe, the fact that the deaths would involve related people (communities), and potential contamination of large land areas for generations to come.

A working definition of an expert is a person who can solve a problem faster or better than others, but who runs a higher risk than others of posing the wrong problem. By virtue of his or her expert methods, the problem is redefined to suit the methods. Because you can count, and because we have data on deaths, the choice is defined as a problem of

toting up known figures. Because I look for social relations, symbolic values, and human progeny, I define the problem as one of potential consequences, not observed ones. Your concern with power lawnmowers or automobiles, where we have good accident statistics, may save more people from injury and death, than my concern with nuclear war, which happened only once and has a very low chance of happening again. We both have bounded rationalities, but our world is immeasurably enriched because our limits are not identical, that is, because we emphasize different skills or cognitions. (The enrichment will be reduced if each of us refuses to recognize the advantages of the alternative view.) The rationalist would have all be completely rational; that does not seem possible. The bounded rationality theorist would point to the unfortunate limitations on most or even all people's cognitive ability. The social rationalist would say the limits are far from unfortunate; they necessitate interdependence, and differences promote new perspectives and solutions that no one person is likely to have. Two cheers, then, for bounded rationality; it promotes the human qualities of interdependence, and it allows legitimately different values to come into play and conflict. The third cheer, not relevant here, is that the limits on every human make domination of the many by the few more difficult.

For those holding the third perspective, emphasizing social rationality, the fears of the citizens of Middletown, Pennsylvania, are a part of the cost of nuclear power. A technology that raises even unreasonable, mistaken fears is to be avoided because unreasonable fears are nevertheless real fears. A technology that produces confusion, deception, uncertainty, and incomprehensible events (as the crisis over the several days did) is to be avoided. The technology affects social bonding and social interaction, and individual psyches and peace of mind. A worker's death is not the only measure of dread; the absence of death is not the only criterion of social benefit.

Thus we can think in terms of three types of rationality: economic or absolute rationality, which requires narrow, quantitative and precise goals; bounded rationality, which emphasizes the limits on our thinking capacities and our inability to often achieve or even seek absolute rationality; and social and cultural rationality, which emphasizes diversity and social bonding. Cost-benefit analysis emphasizes the first; risk assessors are increasingly moving from the first to the second; our discussion has emphasized the last. Next we will try to show that the public also emphasizes social rationality. The public is uninformed in many respects, and certainly can make errors in reasoning, but for matters of catastrophic risk these errors seem less disabling than the alternative of neglecting the

rationality embedded in social and cultural values. The bounded rationality viewpoint has had a tremendous impact upon organizational theory. It was tentatively set forth in an influential book by James March and Herbert Simon in 1958, and then developed by March and his colleagues and students and now challenges all mainstream organizational theory.[33] This theory holds that organizational problems and solutions are tossed haphazardly into a "can" and people pick and choose as opportunities arise and move off as interest flags or new cans appear. Problems get intertwined, solutions are looking for problems to give them life, and participation is unpredictable. I think that the notion of social rationality, sketched above, is an extension of garbage can theory, indicating some reasons why the apparent inefficiency of organizations, when judged from the rational model, may prove to be quite efficient from a social rationality point of view.

The Discovery of Dread

Actually, some careful public opinion polling by Decision Research and members of a Clark University group supports the social rationality view.[34] The researchers were exploring the basis of the presumably irrational view of the public about some technologies, such as nuclear power, and compared the views of experts in various fields with the views of some members of the public, in this case, college students, members of a local business and professional association, and members of the League of Women Voters.

The experts and the lay members of the public agreed on the riskiness of several of thirty activities. Both groups rated as highly risky motor vehicles, handguns, smoking, drinking, and motorcycles. Rated as low in risk were vaccinations, power mowers, food coloring, and home appliances. But the experts and the public disagreed on others, especially nuclear power. Where 1 equals most risky and 30 the least, nuclear power was ranked as 1 by both students and League of Women Voters members. The business and professional club members ranked it 8, but the experts ranked it a very low 20 out of the 30 activities. Why the discrepancies?

The researchers had noted in other studies that the public's estimates of fatalities were greatly "biased," in part by sensational newspaper coverage and the ease of remembering recent events (the "availability heuristic"), so they checked on this. They asked the groups about the number of people likely to die next year in the United States from each of the activities, if it were a normal year as deaths go. (If you think this might be a difficult exercise you are correct; what do you think the figure for

scuba diving might be, for example, or power mowers?) The experts, being experts, gave figures closely resembling the actual figures for past years. The lay people's estimates were not very accurate; they might greatly overestimate the number of, say, power mower deaths. But more important, even though they might think that many people were killed this way, when asked to judge the riskiness of the activity by ranking the 30 activities, they might judge it as low in risk. Another activity that they felt (correctly or incorrectly) killed fewer people might be judged a more risky activity. Here we would seem to have more evidence of the inability of the public to judge risk: the public's estimates of risk do not conform with their own (often inaccurate) estimates of harm.

Fortunately, interested in this discrepancy, the Oregon researchers had included another question: What if it were a particularly disastrous year? For most of the activities the estimates of deaths did not change much, for some it changed substantially, and for nuclear power the estimate of fatalities in a bad year shot way up for the lay persons. Thus, disaster potential apparently explained the discrepancy between the perceived risk and the actual annual fatalities for nuclear power and a few other hazards.

That made sense. The public judged some risks on the possibility of a disaster, not the historical record; that would explain why the public would worry about nuclear power regardless of the number of people killed on the highways. But there were still many discrepancies between the public and the experts. Probing further, the researchers then asked the respondents to rate each of the 30 activities on the following dimensions: the degree to which the activity's risks were voluntary, controllable, known to science, known to those exposed, familiar, dreaded, certain to be fatal, catastrophic, and immediately manifested. Now the study began to pay off. Here the difference between the experts and the public all but disappeared: all of the groups gave similar ratings to each of the activities on each of the dimensions. Most strikingly, nuclear power scored at or near the extreme on all of the undesirable characteristics, for both experts and the public. "Its risks were seen as involuntary, delayed, unknown, uncontrollable, unfamiliar, catastrophic, dreaded and fatal."[35]

Note that the experts agreed with the lay people in this characterization of nuclear power, but in the same questionnaire still ranked it as only 20 in riskiness out of the 30 activities, while the three lay groups gave it ranks of 1, 1, and 8. Dread and the unknown, uncontrollable aspects were recognized by the experts, but not thought relevant in judging riskiness. But not so for the public. In fact, for the lay groups, one could predict almost exactly their assessment of risk, based upon their

assessment of how much dread was involved in the activity and the likelihood of a mishap being fatal. The ratings of dread and lethality also closely predicted their estimates of the number of fatalities that could be expected in a bad year. But this was not true of the experts. The degree to which a risk was a "dread risk" and likely to be fatal did not influence their judgment of the overall risk. Apparently, to the experts a death is a death, whether from scuba diving or irradiation.

The researchers then conducted a larger study with ninety hazards and eighteen instead of nine risk characteristics. The results were consistent with the earlier study, but more elaborate. They could be represented by three "factors" (clusters of interrelated judgments, where the three clusters are fairly independent of each other). The most important factor, which they labeled "dread risk," was associated with:

- lack of control over the activity;
- fatal consequences if there were a mishap of some sort;
- high catastrophic potential;
- reactions of dread;
- inequitable distribution of risks and benefits (including the transfer of risks to future generations), and
- the belief that risks are increasing and not easily reducible.

High on this "dread risk" factor were nuclear weapons, nuclear power, warfare, nerve gas, terrorism, crime, and "national defense." Note that nuclear weapons, nuclear power, and military activities in general are systems we classified as complex and tightly coupled—cell 2 in our I/C chart (see Figure 9.1).

It is striking that there is a parallel between the "dread risk" factor and interactively complex and tightly coupled systems. The correlation is not perfect, but still, two independent and totally different systems of classification converge, one based upon a theory of system characteristics that is independent of catastrophic potential, and another that includes catastrophic potential but much else besides, such as lack of perceived control, inequitable distribution of risks, and the belief that risks are increasing and not easily reducible. Clearly, the public's classification scheme is broader than ours in this book; it combines an intuition about systems with a perception of social characteristics and consequences that go beyond catastrophic consequences.

Factor 2, labeled "unknown risk," included risks that are

- unknown,
- unobservable,
- new, and
- delayed in their manifestation.

FIGURE 9.1
Interaction/Coupling Chart

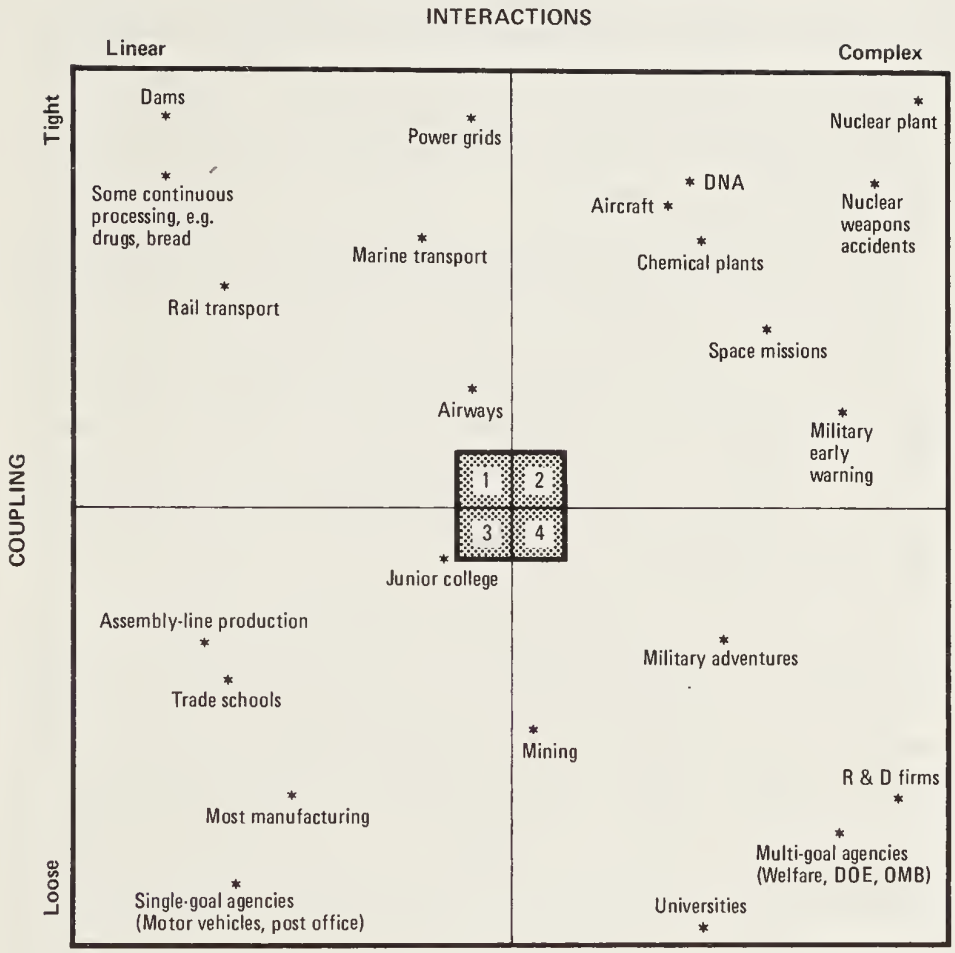

High on this factor were solar electric power, DNA research, earth orbiting satellites, space exploration, food irradiation, lasers, and nuclear power. Nuclear power is both high on dread and fairly high on the unknown dimension.

The factor of unknown risk is not conceptually related to our I/C chart as well as the dread factor. Some are processes rather than systems (solar cells, food irradiation, and lasers) and cannot be placed on our chart. There is some association of "unknown risk" with interactive complexity and tight coupling, but not for all the activities. DNA research, earth orbiting satellites and space exploration, and nuclear power are systems that we placed in the high complexity and tight coupling cell and the respondents associated with unknown risks.

Less information is given about the third cluster, "societal and personal exposure." This involved the number of people exposed and the rater's personal exposure. Hazards at the high end of this dimension were motor vehicle accidents, caffeine, alcoholic beverages, smoking, food preservatives, herbicides, and pesticides. At the low end were lasers, solar electricity, space exploration, laetrile, scuba diving, and open heart surgery. It was not as important in influencing perceived risk as the other two factors, but still made a contribution.

Unfortunately, the rated risks do not include most industrial activities we are interested in, such as chemical plants, mining, or factories. Thus, we cannot lay the I/C chart over the array produced by the factors of dread and unknown risk. But the rough comparisons that can be made suggest that our classification is more consistent with the public perception of hazards than it would be with the assessment by experts, including many risk assessors in government, industry, and the universities, who, while recognizing dreadedness and unknown risk, did not utilize it in their hazard ratings. They used simple body counts or theoretical estimates of such counts.

The dimension of dread—lack of control, high fatalities and catastrophic potential, inequitable distribution of risks and benefits, and the sense that these risks are increasing and cannot be easily reduced by technological fixes—clearly was the best predictor of perceived risk. This is what we might call, after Clifford Geertz, a "thick description" of hazards rather than a "thin description."[36] A thin one is quantitative, precise, logically consistent, economical, and value-free. It embraces many of the virtues of engineering and the physical sciences, and is consistent with what we have called component failure accidents—failures that are predictable and understandable and in an expected production sequence. A thick description recognizes subjective dimensions and cultural values and, in this example, shows a skepticism about human-made systems and institutions, and emphasizes social bonding and the tentative, ambiguous nature of experience. A thick description reflects the nature of system accidents, where unanticipated, unrecognizable interactions of failures occur, and the system does not allow for recovery.

Running Risks, or Trials Without Errors

For most of the ills that come with advanced industrialization, an increased measure of surveillance, penalties, and prohibitions would suffice to limit them. Smoking can be taxed more heavily, the subsidies to the tobacco industry eliminated, and smoking in all public places and places of work forbidden. Mine safety can be improved; smokestack scrubbers used; fluid sandbed furnaces required (they are more efficient and drastically reduce pollutants). Toxic chemicals can be much more closely regulated and illegal disposal eliminated (though organized crime might have to be eliminated to do so, if newspaper stories of their involvement in illegal disposal are correct).[37] Automobiles and highways could be made safer, and so on. We have made some of these improvements already, and could do more—if elites were willing to run the risk of alienating the offending industries. (One enormous risk which the industrialized nations may be facing is not considered in this book on normal accidents; eliminating this ill would require much more drastic measures than any of the above: This is the problem of carbon dioxide produced from deforestation primarily, but also from burning fossil fuels such as coal, oil, and wood. This threatens to create a greenhouse effect, warming the temperature of the planet, melting the ice caps, and probably causing an incredible number of other changes, most of them disastrous. If it is significant—the experts do not agree—we may have a few decades to handle this; but it may be too late. It is one of the strongest cards the nuclear addicts can play, though the enormity of the problem, by some accounts, would dwarf the capacities of nuclear industry. We would have to divert our energy and natural resources from much of industry and use it to build nuclear plants for the next generation to meet some estimates. Battalions of scientists, engineers, and operators would have to be recruited and trained, and so on.)

If the problem is not that great, it is possible that a massive and nearly impossible shift to conservation and solar power could make the difference, but recall that solar cells and collectors also require significant amounts of energy to build and distribute, though hardly as much as a thousand nuclear plants in the U.S. alone. Producing the safe systems cannot be done without further pollution, for several decades, by the major culprits, the industrialized nations. The international planning and cooperation required would be way beyond any precedents.

But it is fair to ask whether we have progressed enough as a species to handle the more immediate, short-term problems of DNA, chemical

plants, nuclear plants, and nuclear weapons. Recall the major thesis of this book: systems that transform potentially explosive or toxic raw materials or that exist in hostile environments appear to require designs that entail a great many interactions which are not visible and in expected production sequence. Since nothing is perfect—neither designs, equipment, operating procedures, operators, materials, and supplies, nor the environment—there will be failures. If the complex interactions defeat designed-in safety devices or go around them, there will be failures that are unexpected and incomprehensible. If the system is also tightly coupled, leaving little time for recovery from failure, little slack in resources or fortuitous safety devices, then the failure cannot be limited to parts or units, but will bring down subsystems or systems. These accidents then are caused initially by component failures, but become accidents rather than incidents because of the nature of the system itself; they are system accidents, and are inevitable, or "normal" for these systems.

Much can be done to make these systems somewhat safer, but accidents cannot be entirely avoided. Quality control, operator training, design experience, and environmental controls will help, but will not be sufficient. These are benign steps to reduce the frequency of system accidents. But there is one other solution that may not be so benign: instituting highly centralized, authoritarian organizational structures.

The discussion we had of thinking, or cognition, is particularly relevant to the inescapable fact that our risky enterprises are organizational enterprises. Just as some hope our risk analysis can be made more rational and scientific, so do some envision more tightly controlled, authoritarian organizations to run our risky enterprises. In their view, we simply have to eliminate "operator error." We have been dealing with organizations throughout the book. It is now time to confront explicitly this context for catastrophe. What do we know about organizations and does what we know reassure us that risks can be managed? We know quite a bit, intuitively, by simply living in the world. We also learn a lot by reading books on almost any important topic, since most important topics have organizations at least in the background. Finally, those of us who are familiar with decades of research on organizations may know an additional bit more that is useful. The topic of organizations is quite relevant to high-risk systems, at least as important as the possibilities of technological fixes.

This book started with a commonplace system accident without catastrophic potential. Better organization of our affairs might have made a difference in the coffee pot and job interview story, but not much; better

luck would have been more to the point, so that one of the several safety devices was not defeated. We then moved on to the nuclear industry, and I argued that we should not expect nuclear power to be above the industry average in organizational competence or honesty; the lack of both did not bring down TMI, though they contributed to the problem. Well-run plants have system accidents too. So do the much better-run (I suspect) chemical plants. Our review of some aircraft manufacturers suggested undue indifference to safety, but by and large, both the airways and the aircraft are not plagued with organizational problems. The marine industry is; it was analyzed as an error-inducing system, and all the variables used to make that analysis were organizational ones, including the suggestion that rather than give more technology to support the authoritarian decision-making system aboard ship, we should reorganize the decision-making system. The space missions dramatically support the notion that the operator must be left in the loop because the designers' loops are far from seamless, and faulty information and reality construction by managers, along with pressures for production, can lead to problems. Building dams, in our two case studies of the Teton failure and the radioactive tailings dam, suggested run-of-the-mill organizational failures similar to the nuclear power industry, and mining also reflects poor organizational behavior, especially in the attribution of operator error and the failure to take elementary protection in the case of potential explosive mixtures. With DNA we were especially critical; elementary precautions may be absent in the rush to publish and profit.

But we must come now to a more systematic limit on some organizations, a kind of Pushmepullyou out of the Doctor Dolittle stories (a beast with heads at both ends that wanted to go in both directions at once). The organizations at risk are the complexly interactive, tightly coupled ones in cell 2 of our Interactive/Coupling chart (see Figure 9.1). It should be considered with Figure 9.2 in order to highlight the dilemma. I will argue that complex but loosely coupled systems (cell 4, such as universities) are best decentralized; linear and tightly coupled systems (cell 1, such as pharmaceutical plants) are best centralized; linear and loosely coupled systems (cell 3, most manufacturing) can be either; but complex and tightly coupled systems (cell 2, including nuclear power) can be neither—the requirements for handling failures in these systems are contradictory.

Systems with interactive complexity (cells 2 and 4) will produce unexpected interactions among multiple failures. But while these are troublesome and unwanted, they need not bring about accidents—that is, dam-

FIGURE 9.2
Centralization/Decentralization of Authority Relevant to Crises

COUPLING	INTERACTIONS: *Linear*	INTERACTIONS: *Complex*
Tight	CENTRALIZATION for tight coupling. CENTRALIZATION compatible with linear interactions (expected, visible). Dams, power grids, some continuous processing, rail and marine transport. 1	CENTRALIZATION to cope with tight coupling (unquestioned obedience, immediate response). DECENTRALIZATION to cope with unplanned interactions of failures (careful slow search by those closest to subsystems). Demands are incompatible. Nuclear plants, weapons; DNA, chemical plants, aircraft, space missions. 2
Loose	CENTRALIZATION or DECENTRALIZATION possible. Few complex interactions; component failure accidents can be handled from above or below. Tastes of elites and tradition determine structure. Most manufacturing, trade schools, single-goal agencies (motor vehicles, post office). 3	DECENTRALIZATION for complex interactions desirable. DECENTRALIZATION for loose coupling desirable (allows people to devise indigenous substitutions and alternative paths), since system accidents possible. Mining, R&D firms, multi-goal agencies (welfare, DOE, OMB), universities. 4

age to a subsystem or the system as a whole. Accidents will be avoided if the system is also loosely coupled, (cell 4, universities and R&D units) because loose coupling gives time, resources, and alternative paths to cope with the disturbance and limits its impact. But in order to make use of these advantages of loose coupling, those at the point of disturbance must be free to interpret the situation and take corrective action. Since the disturbances are generally (not always) likely to be experienced first by operators (which include first-line supervisors and other on-duty personnel such as technicians and maintenance), this means the system should be decentralized. These personnel have two tasks: analyzing the situation and acting so as to prevent the propagation of errors. Unexpected and incomprehensible interactions will not allow immediate analysis of the cause of the accident, but given the slack in loosely coupled systems, this is not essential. It is enough that personnel perceive an unwanted system state (even though it is also unexpected and its causes are mysterious), and do so before it interacts with other units and subsystems. To

do this they must be able to "move about," and peek and poke into the system, try out parts, reflect on past curious events, ask questions and check with others. In doing this diagnostic work ("Is something amiss? What might happen next if it is?"), personnel must have the discretion to stop their ordinary work and to cross department lines and make changes that would normally need authorization. This is a part of decentralization. This delay and experimentation is possible because of loose coupling, but required by the interactive complexity.

Once an untoward state is identified, and some notion of possible consequences formed, recovery actions must take place to prevent the failure from spreading. In a loosely coupled system there is sufficient slack, resources, alternative paths, fortuitous substitutes, and safety devices to effect recovery. But the persons best able to bring these into play are those at the point of the disturbance. Thus, where systems are both complexly interactive and loosely coupled, decentralization is efficient both for diagnosing and recovering from errors. (It is also appropriate for these systems for other reasons than dealing with failures, since these systems can benefit from unexpected interactions of nonfailure events, called synergy, flexibly cope with unstandardized "raw materials," and can utilize professionals that have been socialized through professional training and thus can be left unsupervised because they have absorbed the appropriate norms.)

At the other extreme (cell 1, linear and tightly coupled), consider a continuous processing operation with a well-established technology, standardized raw materials, and linear production system. Here tight coupling is required for efficiency, and can be tolerated because the technology is well understood and the materials well controlled. When failures occur, as they inevitably will, they will not interact in unexpected and incomprehensible ways, but in expected and visible ways. The system programs responses for these infrequent but expected failures; the responses are determined at the top or in the design, and employees at all levels are expected to carry them out without question. They must be carried out immediately and precisely because of the tight coupling present, otherwise the failure could expand from part to unit to subsystem. This does not represent a problem; employees expect it, and desire it. There is no time for rumination, no fortuitous substitutions or safety devices available to bring into play, no alternative sequences to use other than those programmed in. Repeated drills will insure fast and appropriate reactions, but the reactions are prescribed ahead of time, by the central authority. The system operates much the same way in the production

mode as it does in the failure mode, so the centralization is appropriate.*

For organizations that are both linear and loosely coupled (cell 3; most manufacturing, and single goal agencies such as a state liquor authority, vehicle registration, or licensing bureau) centralization is feasible because of the linearity, but decentralization is feasible because of the loose coupling. Thus, these organizations have a choice, insofar as organizational structure affects recovery from inevitable failures. The fact that most have opted for centralized structures says a good bit about the norms of elites who design these systems, and perhaps subtle matters such as the "reproduction of the class system" (keeping people in their place). Organizational theorists generally find such organizations overcentralized and recommend various forms of decentralization for both productivity and social rationality reasons.

For the interactively complex and tightly coupled system (cell 2, including nuclear plants, nuclear weapons systems, chemical plants, space missions, and DNA) the demands are inconsistent. Because of the complexity, they are best decentralized; because of the tight coupling, they are best centralized. While some mix might be possible, and is sometimes tried (handle small duties on your own, but execute orders from on high for serious matters), this appears to be difficult for systems that are reasonably complex and tightly coupled, and perhaps impossible for those that are highly complex and tightly coupled. We saw the space missions move from a highly centralized mode in the first missions to a more decentralized one in the moon shots, and the somewhat less complex and tightly coupled space shuttle may allow still more decentralization. But I predict that the tensions between the two modes will remain, and consume a good deal of organizational energy. Chemical plants are also not extremely high on complexity and coupling, but we saw safety engineers ruminating about the problem of bringing in operators more as the process got more complex, and we saw operators overriding the centralized, automatic system to decouple parts of the plant and to manually shut down affected parts on their own. I do not know enough about this industry to say that the problem of being both centralized and decentralized is expensive and continuous, but I would expect that it is.

In the case of nuclear power I am confident that it is. I once participated in some discussions with Nuclear Regulatory Commission and

*Organizational theorists, at least since the pioneering work of Burns and Stalker, 1961 and Joan Woodward, 1965 and others in what came to be called the contingency school, have recognized that centralization is appropriate for organizations with routine tasks, and decentralization for those with nonroutine tasks. For an early statement see Perrow, 1967, and Lawrence and Lorsch, 1967. The present formulation suggests an extension of that view.

nuclear industry executives about the optimal organizational structure for nuclear power plants. We cycled endlessly through the problem of insuring rapid, unquestioning response to orders from on high (or orders in the procedures manual), and at the same time allowing discretion to operators. Regarding discretion, the operators would have the latitude to make unique diagnoses of the problem and disregard the manual, and be free of orders from remote authorities who did not have hands-on daily experience with the system. We could recognize the need for both; we could not find a way to have both. Predictably, the NRC and industry personnel, faced with the incompatibility, chose the centralized model.

The centralized design for nuclear plants recommended by most (but not all) regulatory and industry personnel goes beyond that of most manufacturing or continuous processing plants because of the catastrophic potential present. It thus has implications for society. It calls for a wartime, military model in a peacetime civilian operation. A military model reflects strict discipline, unquestioning obedience, intense socialization, and isolation from normal civilian life styles. (I am drawing upon an idealized model here; no suggestion is made that our current military organizations resemble this.) Designed to "defend a free society," it must violate that freedom. Because the military is concerned with protection from violent attacks upon our existence, it is given a latitude no other group has. Its structure is not something we generally endorse, but rather is something society tolerates because of the special problems and circumstances. Efforts to extend this model to industry in the nineteenth and early twentieth centuries failed; it was too incompatible with American social values and culture. The question arises, then, of whether we are not designing more systems that are at once both complex and tightly coupled, and that will require the extension of a military model of organization to more and more activities in society—in the name of progress, defense, competition, or whatever.

The issue came up in the discussions with industry representatives I mentioned above, but it came up far more forcefully in the inquiry into the accident at Three Mile Island.

These Ordinary Men

Immediately after the accident at Three Mile Island, President Carter appointed a prestigious commission to inquire into it, the Kemeny Commission as it came to be called. After the commission had taken testimony for some months, they met in closed sessions to draft their final report. The sessions were transcribed, and the following summary and quotes from just one meeting (September 15, 1979; the transcriptions are avail-

able in the NRC's Three Mile Island Reading Room), illustrate some key questions of this book: Do high-risk systems have to operate any differently than low-risk systems? What is the price of trying to manage both complexity and tight coupling at the same time? These questions were intensely discussed; these were immensely practical issues for a commission that would grab headlines and could influence the president and the Congress.

First, the commission had to confront the reality of how nuclear plants actually ran when they were running normally. They were surprised. Taken on tour by the utility, Metropolitan Edison, of the undamaged Unit 1 plant, they found an alarming lack of basic housekeeping and indifference to that fact. Stalagtites and stalagmites 3 feet long were growing out of leaky valves; pools of radioactive water lay about; and piles of radioactive tools, materials, and protective clothing were scattered around with slips of paper on them laconically saying "hot." Loose wiring hung from the ceilings. On another occasion one of the consultant groups working with the commission was taken on a tour of Unit 1. It reported to the commissioners at the September meeting that the engineers from the utility conducting the tour appeared not to understand the basic design of the system or the importance of such problems as loose wires hanging about.

One of the commissioners had vast industrial experience—Patrick E. Haggerty, Honorary Chairman and General Director of Texas Instruments, and executive committee member of the extremely influential Tri-Lateral Commission (which is decidedly pro-nuclear, advising the key, advanced industrial nations in the West and Japan to go nuclear). He said that there were always loose wires hanging about in a Texas Instruments plant and engineers often don't know how the whole system works. (That may be fine for TI, someone on the commission might have said, but not for a nuclear plant.) He and some other commissioners also argued that a nuclear plant can't shut down for every leaky PORV, because that creates other problems. "We all know," he said, "that the best thing you can do is to keep something running. . . ." And if there are leaky valves or plant shutdowns, the plant should not tell the local public about the problems they are having. The sociologist on the commission, Professor Cora Marrett, said "it would be ridiculous" for the utility to provide information on everything that happened that might be counter to safety, thus defending the utility's right to secrecy. The influential Washington lawyer and reputed power-broker in the Lyndon Johnson administration, Harry C. McPherson, echoed the sentiment. They should do no more than General Motors plants; there is no difference.

Commissioner Peterson, President of the National Audubon Society and former research director of duPont, was not convinced. "They've got some damn dangerous stuff out there in that building" he mused. "There is a little sign on the desk as we were in the plant," he continued, "saying right there, 'Nuclear Power Is Safe.'" Commissioner Pigford (professor of nuclear engineering at the University of California, Berkeley, and former employee of a nuclear vendor) insisted that "they thought it was safe." (One wonders why they needed the sign to remind them if that were the case.) Peterson retorted, "They knew damn well it was dangerous," and went on to say, "It's easy to work safely when you're making a suit of clothes, but not nuclear energy." Later he brought up the difference again: "There's one big helluva difference" between nuclear plants and garment factories or General Motors plants, "and the community knows that now." Commissioner Pigford stubbornly disagreed with Peterson about the difference between industries, and thus disagreed with the thesis of this book. "I don't know that; I'll examine each one," he said.

The controversy was important. I have argued all along that nuclear plants are, unfortunately, run no different from most of industry, and that we cannot expect much out of most industrial activity in the way of high reliability and safety. But they should be different because of the complexity, coupling, and catastrophic potential. The pro-industry commissioners were presented with dramatic evidence that TMI was, indeed, run like Texas Instruments and General Motors, so they argued it did not matter. I think they lost their point, though they were the majority of the commission, because McPhearson finally conceded there was a difference. But then he proposed how the difference should be handled: "There is a model of a nuclear system different from the one we have in our country, in our commercial power system," he said. "That's the naval reactor program, run with an iron fist, every decision made at the top, nobody budging down below, intense training, intense discipline on the operators." It merits our consideration, he went on, because so many of the current plants have undisciplined, untrained and, unmotivated people. "We ought seriously to consider the question of nationalization."

The discussion that followed generally approved of McPhearson's iron fist criteria, but did not like the idea of nationalization. Perhaps, in private hands, some thought, the iron fist might work. Then Commissioner Lewis, Dean of Journalism at Columbia University, broke in and finally focused the discussion squarely on risk assessment:

> I really don't believe what I hear. We are going to say in order to continue with

> nuclear power, we are up-front going to say, a certain number of lives we're willing to risk? We are going to centralize and militarize this very dangerous source of power? I mean, I don't like the Rickover thing applied to a peacetime operation. I just think in social terms it bothers me. And we're deciding to do all of these, for what? We haven't explored whether there are alternatives to getting ourselves electricity. I mean, we're really saying we're willing to risk health and safety, a serious potential accident, to centralize a source of power with a military type of group?

Acrimonious discussion followed, as one might imagine, for some time. It was not what those with ties to the industry wished to hear. Then, still frustrated, Lewis spoke again on the issue of a special kind of a system:

> You have seen the NRC and I guess the question I am asking is do you really want to trust your family's life to these people? This is the question I keep asking all through our hearings. I looked at Mr. Galena, and I looked at Mr. Denton, and I looked at Mr. Hendrie [top NRC officials active in the recovery attempt and frequently interrogated by the Commission] and the whole bunch of them; and I am saying, this is a terrible thing I am putting in the hands of these ordinary men.

But Professor Pigford would not let up. "There is no such thing as no accidents," he said. "So we have got to bite the bullet and realize that we are not going to be able to determine what is acceptable."[38]

This became the majority view. From the assertion that "plants are all the same, no special precautions are needed," they moved to a need for military management because they are special, to finally saying, "to hell with the risk assessments, let's forge ahead." The commission's report, while excoriating everyone, merely said, "Now let's all do a better job."[39]

Pigford was correct; there is no such thing as "no accidents." Normally, even space shots do not pose a catastrophic potential; but when we send up caskets of plutonium, we create one. The answer is to not send up plutonium, and we and the Soviets have stopped. But some things cannot be stopped. There have been catastrophic accidents in Admiral Rickover's nuclear naval program too, "run with an iron fist, every decision made at the top, nobody budging down below." So Professor Lewis is correct too; since there will be accidents, we should consider whether any peacetime activity is worth having if it requires that kind of a social order. Even it cannot prevent accidents. Two nuclear submarines have gone to the bottom of the ocean with all their crew.

Organizational theorists have long since given up hope of finding perfect or even exceedingly well-run organizations, even where there is no catastrophic potential. It is an enduring limitation—if it is a limitation—

of our human condition. It means that humans do not exist to give their all to organizations run by someone else, and that organizations inevitably will be run, to some degree, contrary to their interests. This is why it is not a problem of "capitalism"; socialist countries, and even the ideal communist system, cannot escape the dilemmas of cooperative, organized effort on any substantial scale and with any substantial complexity and uncertainty. At some point the cost of extracting obedience exceeds the benefits of organized activity.

What Is to Be Done

The question of what is to be done about our high-risk systems depends upon what we think the problem is. By now, you might correctly assume that I would not say it was dumb operators, but there are three other candidates I would like to dismiss fairly summarily, namely technology, capitalism, and greed, before we discuss a better candidate, "externalities." The technology argument I am concerned with here is not the one that ran through the book, arguing that elites have decided on highly risky technologies that will inevitably have system accidents. The more conventional one I wish to dismiss states simply that we are in the grip of a technological imperative that threatens to wipe out cultural values, nature, and so on. We could cite the failure of collision avoidance systems to work, and the invention of multiple warheads, in support of this. Systems are too complex, and our reach has surpassed our ability to grasp. There is a bit of this philosophy in this book, particularly in the argument for social rationality. But, first, there is no imperative inherent in the social body of society that forces technologies upon us. People—elites—*decide* that certain technological possibilities are to be financed and put into place. Second, most of our technologies do not threaten our values, nature, or our lives. One can no more be antitechnological than they can be anticulture or antinature; the hoe and the wheel and cooking meat and baking bread are technologies.

The next candidate for possible blame is the group of people making decisions—in our society, these are reputed to be capitalists, or government agents of capitalists. Critics argue that organizing an economy around private profits leads to short-run concerns and neglects the longrun consequences. We have encountered the role of production pressures (for profit) repeatedly. Furthermore, capitalists can ignore the

"externalities" or social costs of their activities (pollution, large accidents, enfeeblement of the working class) because they are borne by the public as a whole or by segments of it. The social costs to each capitalist or corporation are very small, since they are spread over society, or nonexistent if they are borne only by particular groups (workers, residents of Times Beach, Missouri, or Seveso, Italy). Government does attempt to redress the balance through regulation, but since government is overwhelmingly business oriented, or made up of businesspersons (including the lawyers who serve business interests), it does not succeed. Furthermore, government requires large-scale organizations to deal with the large-scale organizations capitalism finds it needs, and these are inherently limited and inefficient. Even if the regulatory groups were made up of impoverished environmentalists, there would be substantial inefficiencies and ineffectiveness. So runs the capitalist organization account of our problem with high-risk systems.

While I believe that capitalism accounts for a great deal of what goes on in the world, once the system was established the world was changed so substantially that invoking capitalism as a cause of these specific problems is pointless in the context of this book. Socialist countries (in part because they must compete with capitalist ones, and in part because of the limits of organized activity) behave in much the same way. They pollute, ignore the long-run costs, and in at least the Soviet sphere, enfeeble workers much more than capitalist societies do (because the workers cannot fight back). Production pressures appear to be as high or higher in some socialist countries, and vary substantially in capitalist countries (high in mines and on ships, low in air transport). Capitalism, per se, then, is not a useful or meaningful explanation.

What, then, of a less ambitious explanation than capitalism, namely greed (whether due to human nature or structural conditions)—or to put it more analytically, private gain versus the public good? This is useful for highlighting differences within capitalist and within socialist societies. Some social systems have higher ceilings on private gain than others, and some kinds of activity allow the more rampant expression of private gain than others. Presumably state bureaucrats pursue private gain as assiduously in socialist countries as they do in capitalist ones, though the private benefits are considerably limited (vacation homes and elite schools for children in socialist societies, in contrast to enormous fortunes and relatively unchecked power in capitalist ones). I don't have the impression that ceilings (such as taxation or some other form of redistribution of wealth) have much of an impact upon the creation and operation of high-risk systems, though I think it is essential for a just

society. But within each society we can identify activities where private gain is easily realized (chemical plants) and where it is much harder to realize (space missions). It would be hard, however, to classify our systems by only this criterion; some systems with the most catastrophic potential are government activities (dams, nuclear weapons) and systems with the least catastrophic potential may be private (air transport). Private gain, then, does not seem to be the overriding problem.

The question of what the problem really is has a lot to do with the role of "externalities." These are the social costs of an activity (pollution, injuries, anxieties) that are not reflected in the price of the activity. These social costs are often borne by those who do not even benefit from the activity, or, if they do, are unaware of the externalities. Externalities are important in the case of high-risk systems because of, for example, the costs of cleanup from toxic substances, or of rebuilding after a dam failure. The price of electricity from nuclear power plants does not reflect the very large government subsidies, nor the costs of the unsolved problem of long-term waste storage, nor even the unknown costs of dismantling reactors after their forty allotted years, if they run that long. Had all these been properly considered in the 1950s and included in the cost, this book would have not been written because no utility would have ordered a plant. The externalities of coal-fired power plants without proper scrubbers are enormous, and the externalities drift over several states and the Canadian border. We are now beginning to acknowledge this. Externalities are found in both profit-making and governmental activities. The Corps of Engineers is not a profit-making organization, but the possible externalities of dam failures are not included in the budget they request. The published figures on our weapons systems do not include money set aside for broken arrow accidents. Were externalities built into the price of the product, the consumer of electricity, defense, or motorcycles could make a better choice.

Outputs that are sold to the final consumer directly are able to hide externalities less easily than those that are indirect. Systems, whether public or private, that have identifiable and predictable victims are more likely to consider externalities than those where victims are random, anonymous, and probable rather than certain. Systems that have high-status, articulate, and resource-endowed operators are more likely to have the externalities brought to public attention (thus pressuring the system elites) than those with low-status, inarticulate, and impoverished operator groups. To protect their own interests, the operators in the first category can argue the public interest, since both operators and the public will suffer from externalities. The contrast here runs from airline pi-

lots at the favorable extreme, through chemical plant operators, then nuclear plant operators, to miners—the weakest group.

This suggests that in designing or redesigning high-risk systems, we need to consider not only the technology but the role of a variety of groups, including types of victims, and the true long-run social and economic costs, or externalities. Private gain still has to figure heavily in the analysis, but "structural" variables are also important. Such an analysis would consider whether or not there were direct sales to the final consumer, identifiable victims, and the degree of political independency of the operators. This would be combined with probability of loss, structure of the insurance industry, and the possibilities of lawsuits (as opposed to the much more ineffective system of workman's compensation), the amount of federal presence, and the costs that are externalized. Rather than attempt this, we will conduct a much more primitive and highly impressionistic inquiry into what can be done.

Let us state the problem as follows: How risky, in terms of only the catastrophic potential, are the high-risk systems we have been considering, and how costly would be the alternative ways of producing the same outputs, if there are any alternatives? This is essentially a risk-benefit question, but it now includes a variety of concepts that are normally not even considered.

Catastrophic Potential

Our Interaction/Coupling Chart has served us well, but now is inadequate. It presents only the theoretical placement of systems in terms of their interactive complexity and coupling, and asserts that those in cell 2 are more likely to have system accidents than the others, assuming only that there will inevitably be failures in the DEPOSE components. Establishing that argument has been the major task of this book. It says that these failures are inevitable. But by itself, this argument is without practical application. It has nothing to say about the matter of risk for society.

Now we must venture into far less categorical statements and estimates. One of these estimates is of the catastrophic potential of these systems from either system or component failure accidents. This is independent of system accident potential; once the plutonium caskets are removed from space missions there seems to be almost no catastrophic potential, despite the probability of system accidents; dams do not have system accidents, but can kill 3,000 people at one blow. Furthermore, to complicate things, system accidents even in systems with catastrophic potential may be limited to subsystem failures with no serious damage to

any humans, or even a total system failure (TMI) with no damage to humans. The marine transport system has system accidents, but catastrophic potential exists primarily where toxic and explosive cargoes are involved.

Estimating catastrophic potential from a large accident is difficult. I have ignored first-party victims entirely, assumed that for second-party victims (mostly passengers) that over 100 deaths would count as a catastrophe, but the following analysis would remain the same if the deaths totalled over 200. For third- and fourth-party victims the most catastrophic systems are estimated to be nuclear power plants, weapons systems, and DNA accidents; all of these could be very, very large indeed. Somewhat further behind are chemical plants (largely vapor cloud explosions and release of such toxins as chlorine gas) and marine accidents with toxic chemicals at sea or in port, or explosions in port. Both chemical plants and marine accidents would involve third- and fourth-party victims, but in the hundreds, normally, rather than the thousands or millions.

In column 1 of Table 9.1 I have very roughly ordered the systems of concern in this final chapter by their intrinsic potential for system accidents plus their catastrophic potential if such an accident occurred. This is the catastrophic potential from inherent system accident potential. Space missions have a fair propensity for system accidents, but almost no catastrophic potential, so they are at the bottom; so are dams, which have no system accident potential, and so no catastrophic potential from this form of accident. Nuclear power and DNA are high on both system accident and catastrophic accident potential. (They are not placed at the top, because there are other considerations. The numbers in the figure are just for ease of comparison; they have only relative significance, not absolute significance.)

But we have seen that the potential for a system accident can increase in a poorly-run organization. If there is poor regulation, poor quality control, or poor training, there is an increased chance of failures in the DEPOSE components, and these can make the unexpected interaction of failures more likely, because there are more failures to interact. This should be considered also, as it is in Column 2 of Table 9.1. This column constitutes not just the catastrophic potential from inherent system accident potential, but the catastrophic potential from this source plus what I take to be current rate of DEPOSE component failures that is in excess to what could normally be maintained. This is, how badly run is the facility? There will always be DEPOSE failures, but if there is no effective regulation (as in DNA research) there will be more than necessary. Column 1 is

TABLE 9.1
Catastrophic Potential (CP)

	1 *Inherent System Accident with CP*	*2* *Actual System Accident with CP*	*3* *Component Failure Accident with CP*		*4* *Net Catastrophic Potential (2+3)*	*5* *Cost of Alternatives*	
(High)							*(Low)*
10				20		Space	1
		DNA	Weap		Weap	Weap	
9		Weap/NucP		18		NucP/Mar	2
8	NucP			16	NucP		3
	DNA				DNA		
7	Weap		NucP	14			4
6	Fly	Fly	DNA	12		DNA	5
					Fly		
		Chem			Chem		
5	Chem	Mar	Mar/Fly/Chem	10	Mar	Dam	6
4	Mar			8			7
		Airw	Dam				
3	Airw			6			8
					Airw		
		Mine	Mine		Dam/Mine		
2	Mine		Airw	4		Chem/Mine	9
	Space	Space			Space		
1	Dam	Dam	Space	2		Air Transp	10
(Low)							(High)

the best we can expect if these industries make all the necessary efforts; Column 2 is what we are likely to expect, since they do not seem to be capable of making those efforts.

But accidents can occur simply from component failures. I think some systems are more likely to have catastrophic accidents from system failures than component failures (nuclear power plants), but with others (dams), it is the reverse. Column 3 presents some rough estimates of the potential for catastrophic accidents from component failures, given current experience.

We now have two sources of catastrophe: system accidents and component failure accidents. The simplest thing is to add them together. (There are problems with that—and indeed with all the columns—but we are conducting only a very crude analysis here.) This is done in Column 4, the net catastrophic potential of each system (or, in the case of marine transport, that part of the system that has catastrophic potential).

The result is not all that different from Column 1, but there are reorderings and some clumps have formed.

The first four columns in Table 9.1 indicate a number of conclusions. Dams have few or no system accidents, and even mismanagement is not likely to produce them. They do have component failure accidents, however, and catastrophic potential (the large ones and the small ones with radioactive wastes). If space missions have accidents, they are likely to be system accidents rather than component failure accidents; this is because it is a well-run program and the system has many redundancies. Space missions have little catastrophic potential, though. Mining has little system accident potential, and little catastrophic potential even from component failure accidents. Therefore, while it could be made much safer for first-party victims, its net catastrophic potential remains small. Marine transportation of toxic and explosive materials is inherently but modestly prone to system accidents; this potentiality is increased by poor management and regulation. We must add to that the moderate potential for component failure accidents, and the net result is a strong showing on risk.

The petrochemical industry is moderate on all counts; as with marine transport, there is room to improve the system, and reduce the risk. The airways (largely air traffic control) have managed to reduce coupling and complexity, and because of extensive backups and possibilities for decoupling, component failure accidents with catastrophic consequences are less frequent than system accidents. The airways system, then, is doing well. Flying, however, has substantial system accident potential, even though it is not increased by poor management. Component failure accidents are not common because of the extensive redundancy and safety-consciousness of aircraft builders and the skill of pilots. It is in the middle with regard to the net catastrophic potential, and probably always will have about that much potential because new technological innovations only seem to push the speeds higher and increase the flying in bad weather and crowded skies.

DNA, nuclear power, and nuclear weapons remain at the high risk end throughout. DNA's net potential, however, is substantially due to the lack of regulation, and lack of safeguards in production that we hypothesize exist. Genetic engineering could be made substantially safer. Note that it is not expected to have major component failure accidents, nor is the nuclear power industry. The reasons were different. With recombinant DNA, the failure of any step along the way will signal trouble and result in termination; it is the unexpected result of a nearly successful

series of steps that is most likely to be dangerous. With nuclear power, though there are many component failures, there already are extensive backup devices and defense in depth, such as the possibility of scram. While DNA could be made safer, it is unlikely that nuclear power could be. I suspect that this is even more true of nuclear weapons accidents. Unfortunately, component failure accidents with nuclear weapons could be catastrophic. Unlike the military early warning system (not included in this analysis), which has some degree of loose coupling that protects it from component failures, a Titan missile can literally go off with the drop of a workman's wrench and possibly release plutonium. Though unverified, I was told of a slip in a maintenance step that sent an armed missile into Canada. As described, it was not a complicated accident at all, but rather, akin to what one might expect in industry in general. Thus, this system appears to have the ultimate catastrophic potential, with a high probability of both types of accidents and of the release of lethal substances.

Do we need these systems? Given the risks, what are the benefits? What are the costs of alternative means of getting the output of these systems? Here even more subjective estimates are made, presented in Column 5. I assume that air travel cannot be substantially reduced and high-speed ground travel is not practical over distances of roughly 300 miles. Enormous investments in rail transport would be needed to make it safe and sufficient to compete with air transport even for short distances. Yet in marine transport, there is no overwhelming need to transport deadly poisons or explosives such as LNG long distances through winter storms in unsafe and poorly designed ships operated for private profit. The world economy and the economy of individual countries would not suffer if the poisons and explosives were not transported at all, or in small quantities in super-safe containers aboard specially designed and run ships. The cost of transport might double, but that is a trivial matter compared to the cost to the ecosystem. Such reductions in risk are not feasible for the chemical and mining industries. Our economy and our lifestyle is built around these industries; while some substitutions are possible, not many are likely.

We could move housing out of the flood plains and use the plains for agricultural activities only, avoiding the need for dams. This could be done in many cases, but not all, and we have to consider the "sunk costs" already in place. Moving Los Angeles (a giant flood plain) would be impossible. And dams generate electricity as well as protect flood plains—another reason we are unlikely to give up dams.

The question of nuclear weapons cannot be sensibly discussed in a

paragraph or even a chapter. I have simply followed those who call for extensive, unilateral disarmament of ballistic nuclear weapons and nuclear missiles. I do not see the Soviet threat as all the administrations have since World War II, and thus I think the benefits from abandoning strategic nuclear weapons (and tactical ones too, though that is not at issue here) would be enormous in terms of reducing international tension, improving the economy, and halting the drain of a quarter or even half of our scientific talent from the civilian economy.

The case for abandoning nuclear power strikes me as very strong. There are two pitfalls. First, the government has allowed utilities building nuclear power plants to count as profits during the building years the expected profits that will come after the plant opens. This is called Allowance for Funds Used During Construction, and shows up in the revenue side of the annual financial report. At times, this figure is as much as half or more of the profit in that year, and is paid out in the form of dividends to stockholders and bonuses to management. Were we to close down all construction and abandon these projects, not only would billions go down the drain, but some utilities might be literally and technically bankrupt. The effect upon those rich enough to invest in the stock-market, and the effect upon union and other pension funds would be extensive. The effect upon the stock market as a whole and thus our economy and the free world economy could be serious. If it occurred during a shaky time of defaulted loans to poor countries or of high unemployment, the interaction of failures could produce a subsystem or system accident in this and even other economies. Still, that might be better than a nuclear accident that contaminates a populated area of the earth.

The second problem is that of capacity and distribution. Nuclear energy produces only around 11 to 12 percent of our electrical energy. Were it to disappear, some industries and other heavy users of electricity could convert to nonelectrical energy (such as fluidized beds using coal and scrubbers), but most could not. The problem is transmission over long distances. Chicago, the Carolinas, and parts of the Northeast would need electricity from the West, Southwest, and the northeastern Canadian provinces. Recent hydroelectric capacity in Canada and the exploitation of vast natural gas fields offshore would substantially reduce the shortage in some northeastern states, but would not eliminate it. Efficient transmission to Chicago from the Pacific Northwest is not in place, though it is technically possible. Though we have excess electrical generating capacity (one reason for the cancellation of nuclear plant construction), the excess is not always in the time and place that makes its use feasible. I am sure there are some who argue the conversion back to non-nuclear power

can be done without much cost or effort, and others who deny this. I will just leave it as one of the two major problems with closing all nuclear plants. Should that possibility suddenly become attractive (if, for example, another nuclear plant has all the safety systems fail as the Salem plant did during 1983, and it is also running full-tilt and has a LOCA), I suspect we could cope with the capacity and transmission problem. We might begin planning for it now; I would expect a worse accident than TMI in ten years—one that will kill and contaminate.

Even considering those two problems, neither of which would have results as disastrous as a major release of radioactive materials, the case for shutting down all nuclear plants in the United States seems to be clear. There will be more system accidents; according to my analysis, there have to be. One or more will include a release of radioactive substances to the environment in quantities sufficient to kill many people, irradiate others, and poison some acres of land. There is no organizational structure that we would or should tolerate that could prevent it. None of our existing reactors has a design capable of preventing system accidents. Perhaps a safe one will be discovered—loosely coupled and linear—but I am doubtful. Were nuclear energy all that stood between us and starvation, that could be another matter. It is not.

I place DNA near the middle because, on the one hand, we have done well enough without it, but on the other, the potential benefits are said to be enormous. I will give it a very considerable benefit of the doubt by accepting some of the potential for good it is said to have. The benefits advertised exceed the benefits from air transport, nuclear power, and nuclear weapons. One hesitates to be optimistic, given the overly optimistic predictions about nuclear power (eliminate want from the face of the earth, for example), but the opportunities seem to be incredible for food production, for energy saving, for disease control, for repairing genetic defects, even for toxic waste disposal. The rapid abandonment of minimal safeguards probably did not occur because of these potentials for human life, but because of the prospect of immense profits and the excitement of research and knowledge accumulation. Were the prospect for human betterment the only driving force, we could go very slow and very carefully; on the scale of humankind a couple of generations is trivial. But the economic and scientific impulses are much more uncontrollable. Too many rewards for individual scientists and economic elites are at stake. We will not go slow.

Should DNA be combined with biological warfare, the immense dangers of the latter would be multiplied. There is a fair prospect that the

FIGURE 9.3
Policy Recommendations

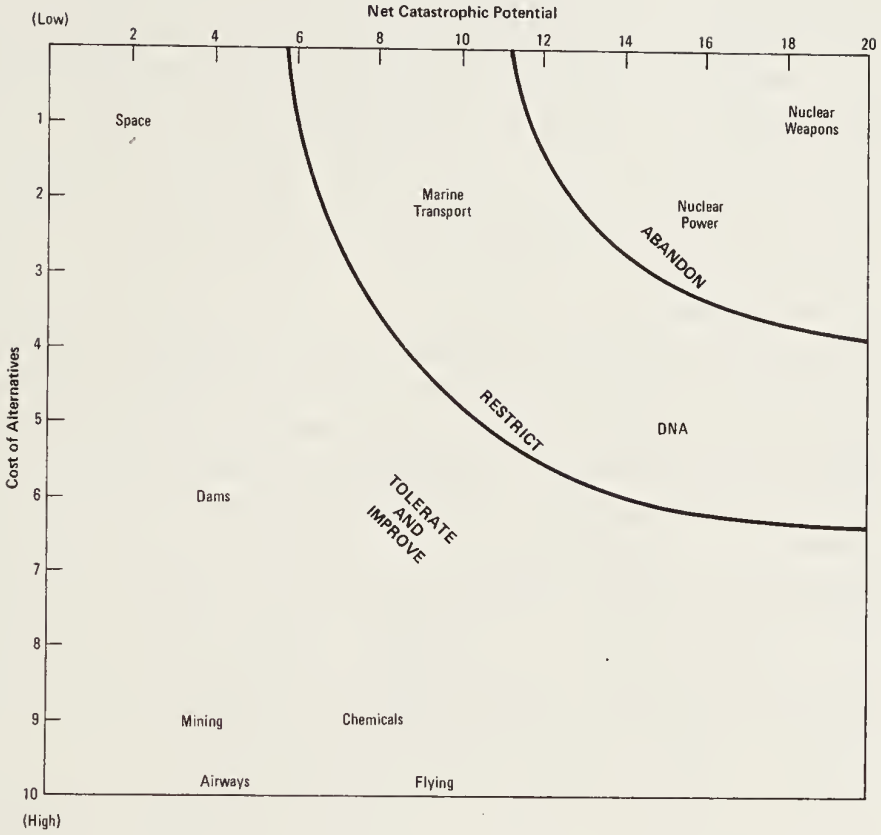

next world war will not be fought with nuclear weapons, but with biological ones. Both the USSR and the United States are known to be developing these weapons. I am ignorant of the matter, but were Genetech or Cetus (two leading DNA firms) to be found doing classified military research, I would fear that we have put the second of the two ultimate dangers to humankind into the same hands we have put the first—the military. On a long string of worry beads this one has to be near the top.

When we take Column 4, the net catastrophic potential, and combine it with Column 5, an estimate of the cost of alternatives for these systems, we produce the array given in Figure 9.3. This constitutes the

policy recommendation generated by our efforts in this book. It indicates which systems are both highly risky and not essential, and thus could be abandoned, and which that we would find it hard to do without, but that have, fortunately, less catastrophic potential. Of course, everyone is privileged to dispute the placement of systems in Figure 9.3. DNA may be such an incredible opportunity that it should be placed at the bottom left-hand corner: very risky, but very tempting. Some nuclear addicts argue that dams are much more dangerous than nuclear power plants, and this chart clearly disagrees with that (as does the public). Still, the chart in Figure 9.3 is at least explicit, and a convenient starting point for public-expert debate.

It is not particularly surprising to call for an end to nuclear power and nuclear weapons; many others are doing so. In fact, the nuclear power industry is so threatened by its failures that its public relations campaign has been stepped up even as I write this. The *New York Times* reported on May 23, 1983, that the Federal Department of Energy has been spending about $2.5 million a year on "nuclear information" and will increase its support to private groups such as the "Scientists and Engineers for Secure Energy" which organizes campus forums, and in the interests of "fairness" will not allow critics of nuclear energy to share the podium with them. They are receiving $100,000 in 1983. The Virginia Electric Power and Light utility, which we met in Chapter 2, contributed $600,000 to a pro-nuclear information committee and is asking the state regulatory committee to allow them to charge the ratepayers for it. Many other utilities are allowed to pass on this expense. Suppliers, such as G.E. and the Bechtel Corporation give handsomely, and the total industry contributions will amount to $25 or $30 million. This suggests there is something to our argument in this book; were nuclear power as safe as the decal indicated at TMI Unit 1, this campaigning would all be unnecessary.

The nuclear freeze movement seems to be spreading at the time of this writing. However, in the field of genetic engineering, the biological laboratories on campuses and in private industry have kept a very low profile; we might hope this is because the extraordinary precautions that are needed are being taken, but I have my doubts. There is nothing at all that I know of going on with regard to improving or protesting the transportation of toxic and explosive materials by sea; indeed, the same problem on land, one that we did not consider here, appears to be growing in intensity.

Much can be done in all these systems to increase safety, but our

record in each one has not been encouraging. I hope this book will remind some of that record; it has been, after all, a dismal and dismaying travelogue through the world of high-risk systems. At each turn, even in the best of the industries, we found rampant attribution of operator error to the neglect of errors by the Great Designers and the Centralized Managers. We found organizations that could not carry the burden of error-free operation, and sometimes seemed insensitive to the damage they did or could do. We may be thankful for the regulatory agencies, but too often they were shown to be ineffective, sometimes natteringly so, sometimes even criminally indifferent or co-conspirators.

But, important as these problems are, they were not the main point. The main point of the book is to see these human constructions as *systems,* not as collections of individuals or representatives of ideologies. From our opening accident with the coffeepot and job interview through the exotics of space, weapons, and microbiology, the theme has been that it is the way the parts fit together, interact, that is important. The dangerous accidents lie in the system, not in the components. The nature of the transformation processes elude the capacities of any human system we can tolerate in the case of nuclear power and weapons; the air transport system works well—diverse interests and technological changes support one another; we may worry much about the DNA system with its unregulated reward structure, less about chemical plants; and though the processes are less difficult and dangerous in mining and marine transport, we find the system of each is an unfortunate concatenation of diverse interests at cross-purposes.

These systems are human constructions, whether designed by engineers and corporate presidents, or the result of unplanned, unwitting, crescive, slowly evolving human attempts to cope. Either way they are very resistant to change. Private privileges and profits make the planned constructions resistant to change; layers upon layers of accommodations and bargains that go by the name of tradition make the unplanned ones unyielding. But they are human constructions, and humans can destruct them or reconstruct them.

The catastrophes send us warning signals. This book has attempted to decode these signals: abandon this, it is beyond your capabilities; redesign this, regardless of short-run costs; regulate this, regardless of the imperfections of regulation. But like the operators of TMI who could not conceive of the worst—and thus could not see the disasters facing them—we have misread these signals too often, reinterpreting them to fit our preconceptions. Better training alone will not solve the problem, or

more gadgets, or promises that it won't happen gain. Worse yet, we may accept the preconception that military superiority and private profits are worth the risks. This book's decoding asserts that the problems are not with individual motives, individual errors, or even political ideologies. The signals come from systems, technological, and economic. They are systems that elites have constructed, and thus can be changed or abandoned.