Nuclear Power Regulation

Seeking Safety, Doing Harm?

Elizabeth Nichols and Aaron Wildavsky

WHEN YOU SEE an opportunity to increase safety, take it. Add a new safety device and safety will actually be increased. Or will it? If nuclear power safety could be guaranteed by simply adding more safety devices, the only debate would be over how much safety is "safe enough" and what price we are willing to pay to achieve that safety. But operational safety is not merely additive. Each new device in a nuclear power plant interacts with other parts of the plant, sometimes in ways that constitute new threats to safe operation.

For many years nuclear power regulators have applied a philosophy of add-safety-wherever-possible. How safe has that left nuclear power plants? No one knows. But one of the chief threats to nuclear power safety today is the failure to recognize that individual safety systems may interfere with one another. Dealing with dangers by simply piling on safety measures is not necessarily—indeed is not often—an effective means of improving safety.

Aaron Wildavsky is professor of political science and public policy at the University of California, Berkeley, and a member of the university's Survey Research Center. Elizabeth Nichols is a doctoral candidate in sociology at the University of California, Berkeley. A longer version of this article is forthcoming in Aaron Wildavsky, Searching for Safety, *to be published later this year by Transaction Press.*

Performance Standards vs. Detailed Specification

The Nuclear Regulatory Commission (NRC) is responsible for licensing and inspecting users of nuclear materials. This includes the nuclear power industry, which now provides about 17 percent of U.S. electrical production, as well as several dozen research and medical facilities.

The way the NRC organizes its licensing and safety inspection activities may be considered either ordinary or extraordinary, depending on one's views on how complex and dangerous activities should be managed. The NRC could set performance standards for nuclear power plants and check along the way to see that they were being met. Alternatively, it could specify in detail how each step in construction, maintenance and operation of plants is to be conducted, checking each part of the process for conformance with these detailed prescriptions. The NRC uses both strategies of regulation, although the latter, regulation by detailed prescription, is dominant.

The NRC inherited much of its regulatory apparatus from its predecessor, the Division of Regulation of the Atomic Energy Commission. At the time this apparatus was constructed, there were no feasible means of calculating an overall level of nuclear power plant safety. Early regulatory measures, therefore, tended to focus on specifiable subunits (either of plant hardware,

procedures or organization) rather than on safety. Analytic techniques that are now available, while they provide a great deal of information, still cannot give complete and certain answers concerning the safety contribution of individual parts of power plants or the safety level achieved by the whole. For this reason, many still favor detailed specification both for the parts the licensee will install and the procedures the regulators will follow. Such specifications at least have the advantage of being politically defensible. Even where regulators cannot provide accurate and convincing estimates of reactor risk levels, they can point to requirements imposed and actions taken, concrete efforts to achieve public safety.

Design elements compete with one another. In the same way, safety systems may interfere with one another and potentially pose a threat to reactor safety.

Today reactor owners and operators must meet a lengthy list of technical specifications and produce detailed operating procedures for regulatory approval. The detail is overwhelming. For analogies one might think of the carving up of the whale in Melville's *Moby Dick* or the construction of the tabernacle in the Old Testament, the exact ingredients and their usage being specified precisely. The amount and character of flux in solders, the quality and form of steel or cement, and the nature of storage and handling are all specified. The composition of a screw—even the exact direction and torque and the number of times it is to be turned—may be determined. The task of inspectors is to make sure that the specifications are met and the work is accomplished as prescribed. Even the regulators are directed to go about their work according to a detailed plan.

The technical specifications for each plant are negotiated with Washington during design, construction, and start up, and may be further modified during operation. These "tech specs," of course, directly reflect the regulations formulated by the NRC. Just how each regulation is to be met, however, is based in large degree on codes established by professional groups. These codes are consensus documents produced by association committees. They tell users what minimal actions or materials are needed in the opinion of the committee to assure safe operation. The codes are frequently revised (sometimes as often as every six months), with the technical specifications for each plant identifying which version of any given code is to be considered the valid one for any particular plant. Industry codes tend to become stricter over time: if it is learned that stronger or better quality construction is possible, it is eventually adopted.

Basic codes for electrical and mechanical work in nuclear power plants actually have been incorporated as part of the code of federal regulations and have the force of law. Many other industry codes are endorsed and made part of regulatory guides on how NRC requirements might be satisfied. Some of these codes are endorsed only as modified, with the modifications almost always being more restrictive. As a result, individual elements in construction must often be built to be as safe as existing knowledge permits.

The Parts vs. the Whole

Any engineering manager knows that it is rarely possible for all the relevant aspects of a project to be optimized at the same time. In an important way, design elements compete with one another. In the same way, safety systems may interfere with one another and threaten reactor safety. Thermal insulation, for example, which is used to protect workers from the heat of large pipes, prevents proper inspection of piping details and thus may preclude the early discovery of cracks or corrosion. Nuclear security systems, which ensure careful scrutiny of personnel allowed to enter, can delay the entry of needed expert personnel during accidents. (This was the case at Three Mile Island.)

Corrective actions can themselves do damage. Nuclear welding codes, for example, require the reworking of welds in which even tiny voids occur. Rework, however, can weaken the materials used (e.g., stainless steel pipe). Given the differing skills of those who do the rework, it is not always clear that correction of minor defects increases safety. Similarly, pipe restraints are often installed to prevent damage to nearby equipment from the whipping motion generated when, and if, a pipe ruptures. These restraints, however, may produce a binding of the system that decreases safety.

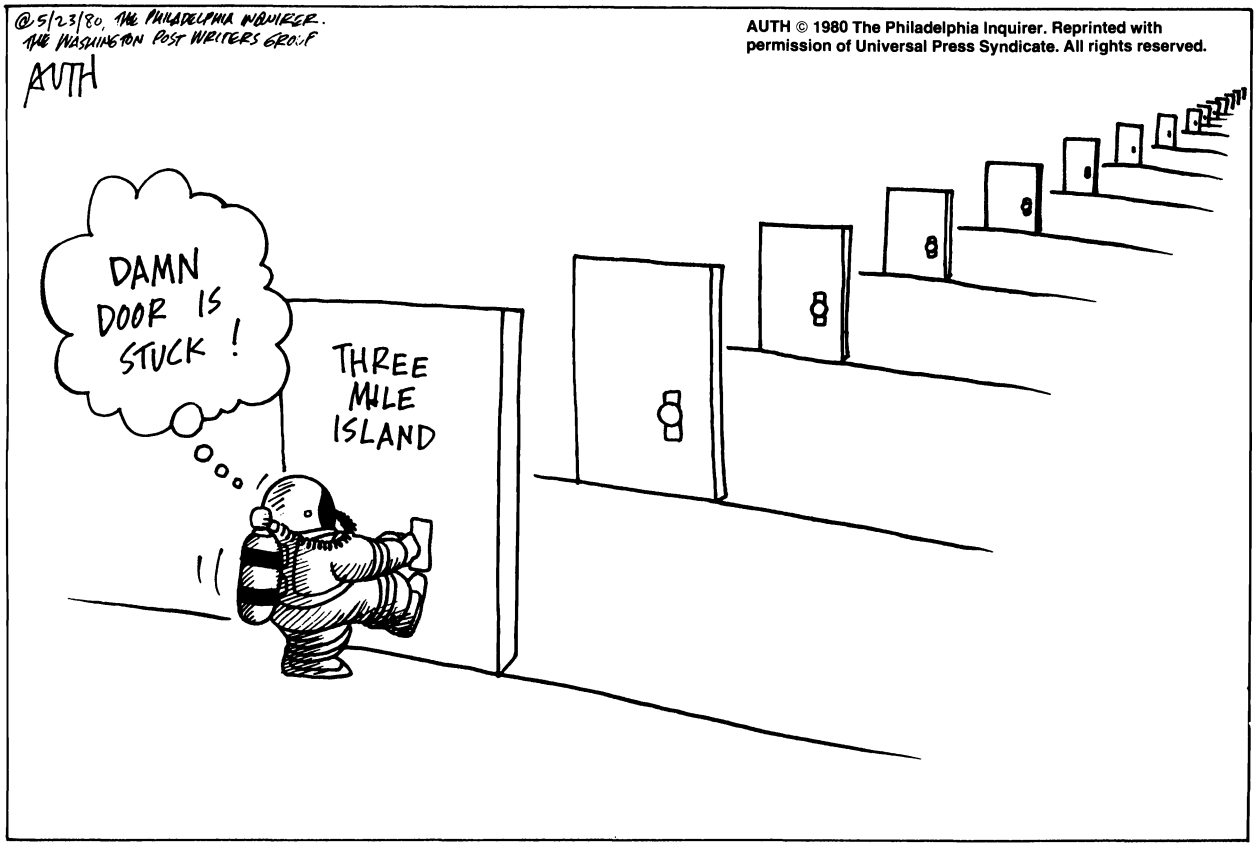

Testing can also be counterproductive. Controlling a plant during normal operation and keeping it shut down safely during emergencies depends on maintaining a reliable system of alternative electrical sources. If the probability of an offsite power failure is relatively high, the ability to prevent a series of events that would lead to reactor core damage depends on the reliability of onsite diesel power. Diesel generators must be installed, therefore, and, because they may fail to operate on demand, there are usually two or more at a given site. Given the importance of reliable onsite diesel power, it might seem reasonable to require regular testing. Testing, however, is not always advisable, since the tests themselves may make the generators less dependable.

Deciding when and if testing is worth the added risk is often a close judgment. When the Florida Power and Light Company decided to build a second unit at their St. Lucie Nuclear Power Plant in southern Florida, for example, a controversy arose concerning the sufficiency and reliability of onsite power. St. Lucie Unit 2 had two generators designed and located to be physically and electrically independent; diesel oil delivered to the site was tested and stored in separate tanks for each system; the generators were housed in a building designed to withstand hurricanes and other adverse weather conditions; a sequencer was installed to prevent rapid loading of electrical demand that might cause both generators to fail. As a means of further improving system reliability, the company suggested that during times when its power distribution grid was on "alert status" nuclear power plant personnel should "idle start" the diesel engines and run them for a short period of time to verify their availability. The NRC staff agreed that this would probably be the simplest way to determine availability, but pointed out that idle starting diesel generators and running them unloaded "could unnecessarily hamper their performance in a real emergency" and might lead to equipment failure. The NRC staff eventually concluded that such testing should not be required. The negative side-effects of testing, they believed, outweighed the benefits.

The NRC is rarely on this side of such issues. Routinely, reactor operators and other utility personnel are required to perform dozens of tests to determine the condition and operabil-

ity of their plants. Such tests often require disengaging certain electrical circuits or disabling some safety systems. Where safety systems have been installed in redundant "trains," reactors are often switched from one train to the other during tests. Valves must be properly realigned and circuits properly reconnected or safety is degraded. Tests may also require direct intervention in everyday operations. Unfortunately, testing is never unequivocally a good idea; it necessarily increases exposure to risk.

Unfortunately, testing is never unequivocally a good idea; it necessarily increases exposure to risk.

The most vivid example of this point is the accident that occurred at the Chernobyl 4 nuclear power plant in the Soviet Union. The accident killed a minimum of 31 people, and quite possibly some 200 will eventually die as a result of the direct effects of exposure to large amounts of radiation. It is the only true nuclear power disaster in history. And testing, in an effort to impose safety, was a major contributing factor.

The possibility of a "station blackout," which was the reason for the Chernobyl tests, is a very real concern since, even after a reactor is shut down, the fuel continues to produce large amounts of "decay heat" which must be carried away by the cooling system to prevent a meltdown. Cooling pumps, instrument and control panels, and even the light needed to work all require emergency power. The reactor at Chernobyl, like U.S. reactors, was equipped with diesel generators and large storage batteries. The purpose of the tests was to try to squeeze out an additional hour of electricity using the steam already present in the system and the momentum of the turbine. Initiating the test required defeating several interlocking automatic shutdown and emergency cooling systems. As a result, the operators had few available means to control the reactor once the accident began.

Strengthening the Parts

Safety hardware can be a source of danger as well. Consider the seismic design standards for nuclear power plant piping. Precautions must be taken to deal with the risk of earthquakes, of course, but it is difficult to know where to stop. Pipe supports to protect piping systems from earthquake damage have been increased in number and size. Various types of restraints have been installed. Mechanical snubbers to dampen pipe vibration have been added. Designs have changed significantly. Recently, questions have arisen concerning the safety consequences of adopting extra precautions to guard against a relatively remote worst case.

There are any number of problems with trying to respond to earthquake risks by strengthening and adding more individual supports. Such measures limit access for routine inspection and maintenance of equipment and piping. They also create a more rigid system that may be less able to withstand stress in everyday operation. Of particular concern is the heavy reliance on so-called snubbers, which anchor piping to the reactor building. Snubbers are hydraulic or mechanical devices with failure rates of their own, and whatever protection they afford must be discounted to the extent that they further complicate the system and decrease reliability. Rigid systems also require more careful alignment. Snubbers must be removed for inspection and then reinstalled, which increases the likelihood that alignment will be poor. Even when the snubbers are maintained successfully and work properly, tightly bound systems are subject to much greater stress during normal operation than more flexible systems. It was with frustration therefore that an NRC inspector told us, "Four snubbers have become 4,000 or 5,000 snubbers. Even to list them all is perhaps one-third of the technical specifications for some plants."

Many older nuclear power plants have been subject to backfitting requirements to bring them up to newly imposed seismic standards. Since these plants have relatively small containment structures, the ability to maintain and inspect piping and equipment has been compromised. Workers who must perform these services are more likely to be exposed to higher radiation doses since they must spend more time to do the same job. Even in newer construction, changing requirements may present serious difficulties when major structural features are already in place. The order in which requirements are added, therefore, must also be considered in evaluating the relationship of the parts to the safety of the whole.

The degree to which it is recognized that in-

creasing safety in one area may lead to decreased safety in another—that there is a tradeoff between the good and bad effects of safety measures—is a major determinant of the level of nuclear power safety finally achieved. It has been pointed out, for example, that thermal stresses on piping from rapid heating and cooling are far more common than dynamic stresses from earthquakes. Each time a reactor is started up or shut down, the entire piping system is subject to

The degree to which it is recognized that increasing safety in one area may lead to decreased safety in another . . . is a major determinant of the level of nuclear power safety finally achieved.

temperature changes of hundreds of degrees centigrade. While this thermal stress is a matter of concern for engineers and utility managers, it has not received much public attention. As a result, we are not likely to see changes in earthquake safety requirements which take these stresses into account. But ensuring the safety of the nuclear power plant as a whole will require the examination of such tradeoffs.

The Hydrogen Danger

Conflict between increasing safety in one area and maintaining safety in others is clearly illustrated by the long-standing controversy over how to control the hydrogen produced within the containment building during an accident. To insure that a meltdown does not rupture containment, reactors must include an emergency core cooling system to provide water to the core in cases where the primary coolant has been lost through a leak or major break. The fuel used in nuclear reactors comes in the form of small pellets of uranium sealed inside thin tubes of a noncorroding zirconium alloy. These tubes are called cladding. If the emergency core cooling system fails and heat in the reactor builds up sufficiently, the cladding will react with any water present and give off large amounts of hydrogen, a highly flammable and potentially explosive gas. Either a burn or an explosion of a large amount of such gas would place severe pressure on the containment structure.

Under federal regulations, the reactor containment building must withstand pressure and temperature conditions resulting from accidents involving loss-of-coolant without exceeding a (very low) specified leakage rate. The regulatory criterion, formalized in the federal code, recognizes explicitly that experience and experimental data (and therefore knowledge) are limited. Designers are required to calculate the peak pressures and temperature that can be expected, and then to add a safety margin. This addition is meant to insure safety even in the event of unanticipated situations.

In the accident at Three Mile Island hydrogen accumulated but did not explode, despite fears that it might. Instead there was a hydrogen burn and a resulting "pressure spike." Following the accident, the NRC ruled that new plants, such as the Tennessee Valley Authority's Sequoyah I, would have to be able to withstand a hydrogen burn at least as large as that at Three Mile Island. After careful calculation, TVA reported that the containment for their new plant would withstand three times the pressure previously required. However, the new plant was expected to survive a burn the size of the one at Three Mile Island by only a small margin.

To allay concerns, TVA proposed incorporating one of a number of alternative control methods. The alternative chosen by the TVA was to build in glow-plug igniters. Since the danger lies in allowing large amounts of hydrogen to accumulate, deliberately burning off the hydrogen as it is generated should prevent larger burns or explosions. Among the options rejected by TVA was the use of an inert atmosphere, replacing the air in the containment with nitrogen gas.

Since hydrogen requires oxygen in order to burn, the inerting of containment is a potent line of defense against the consequences of a hydrogen build-up. But this solution itself presents some very real difficulties and has remained controversial. In regard to the Vermont Nuclear Power Station some years earlier, for example, the NRC regulatory staff concluded that it would have some difficulty sustaining a satisfactory margin of safety and insisted that the utility "inert" the atmosphere of the containment building whenever operating at 80 percent power or above. The utility objected on the grounds that filling the containment with nitrogen or other inert gas would have safety disadvantages of its own. Entry into an inert containment would have to be made with self-contained breathing

apparatus. Plant personnel required to make such entries would be placed at significant risk since they would be working in close quarters with high temperatures and humidity, surrounded by projecting equipment that could easily snag air lines, and burdened with the weight and bulk of the breathing apparatus. Entering a de-inerted containment for routine inspections during scheduled outages would also be more hazardous, due to possible gas pockets. The dangers involved in entry into inert containments would greatly reduce inspection capability and, therefore, increase the risk of a serious accident.

Preserving Shutdown Capability

Even managers who favor prescriptive forms of regulation admit that the part-by-part approach has had negative consequences. Nuclear power plants are becoming more complicated and more difficult to understand, and the risk of unanticipated interactions among the parts is increasing. Nuclear regulators often emphasize that they are responsible for public safety, not for the profitability of the industry or even the operability of the plant. But safety, complexity, reliability and operability are all interrelated.

It is frequently argued that the solution to such regulatory dilemmas is for nuclear power plants to adopt reinforcing shutdown capability. All reactors would be designed to shutdown automatically at the first sign of possible trouble.

Even managers who tend to favor prescriptive forms of regulation admit that the part-by-part approach has had negative consequences.

In itself a shutdown or a "scram" does not release any radioactivity and poses no threat to public health and safety. But over the lifetime of the reactor, frequent scrams can do serious damage. Every time there is a scram, the reactor is put through a "transient," meaning it is subjected to rapid changes in temperature and pressure which produce significantly greater stress than ordinary operation. "That's something the reactor is designed to do only a few times," one regulator familiar with energy management told us. "It's like the brakes on your car. You shouldn't be using the emergency breaks every time you want to stop." The aim is to preserve shutdown capability; the question is how to do this without making the reactor so sensitive to slight operational deviations that it wears out.

Regulatory Overload

An important but hitherto neglected aspect of nuclear regulation is the workload placed on regulatory personnel under different safety regimes. An important example is the development over the past 10 to 15 years of a modular approach to nuclear power plant inspections. Under this approach there is a set of detailed specifications for the licensee to follow and a parallel set of detailed instructions by which these specifications are to be monitored.

The major drawback of all this detail is that the regulatory workload quickly outgrows the agency's resources. There are thousands of workers at a construction site and thousands of parts used in construction. In-depth examination of even a small portion of a plant soon runs into more hours than the agency's budget can cover.

As one NRC section chief, who was once assigned to inspect two large construction sites many miles apart, remarked to us:

> There were 7,500 workers on those two sites. The head of Inspection and Enforcement said some place that we inspect 1 percent of all construction. No way could I have looked at 1 percent of everything done! People can write requirements forever. But it's a case of the alligator mouth and the hummingbird stomach. Even in an operating reactor you have 250 people; you can't do a comprehensive check of everything they do.

Adding more inspection personnel might boost the parts inspected into the 1 percent range, but the discrepancy between tasks prescribed and tasks performed would still be very large.

This time constraint operates even where the agency has been careful to define its task in terms of an audit function. The NRC's Office of Inspection and Enforcement stands atop a large pyramid of inspections and reviews. At the bottom, each element is subject to inspection through the quality control program maintained by the architect-engineering firm, subcontractor, or utility subdivision responsible for building or

operating the plant. The adequacy of such programs is monitored by the quality assurance component of the licensee's organization, which is assessed in turn through NRC inspection. Just as the licensee's quality assurance unit samples the quality control work of others, so the NRC samples the work of the quality assurance unit. Only if errors begin to slip through the quality control system does the NRC undertake full-scale inspection.

Since difficulties in one plant (e.g., pipe cracks in boiling water reactors, or faulty seismic analysis) might be present in other similar plants, concern over the safety of one plant can lead to large and unexpected increases in the agency workload. A case in point is the tremendous expansion of regulatory requirements imposed after the accident at Three Mile Island. Workload increases such as this are unpredictable, making it difficult for regulators to maintain prescribed schedules while doing an adequate job of reviewing and inspecting. Often the routines officially imposed on regulatory personnel are disregarded in an effort to satisfy the most politically pressing demands on agency resources. The leak at Indian Point (New York) in 1982, for example, may have resulted, in part, from a diversion of manpower from routine inspections following the accident at Three Mile Island. In testimony before the House Committee on Government Operations, NRC representatives admitted to completing only about 30 percent of the required inspections. In a study prepared by the NRC Office of Inspection and Enforcement on the lessons learned from Three Mile Island, it was acknowledged that the diversion of manpower from routine inspections of equipment and facilities might contribute to decreased safety.

"No way I could have looked at 1 percent of everything done! People can write requirements forever. But it's a case of the alligator mouth and the hummingbird stomach."

At the same time, the conscientious effort of regulators to follow the dictates of a prescriptive safety regime comes into conflict with their need to be responsive to current safety concerns. Actual operating incidents, from which lessons about safety might be drawn, often do not fit into the neat categories that the inspectors use in their day-to-day work. Moreover, as an NRC engineer explained to us, once the difficulties observed in the field are translated into regulatory language, the official complaint against a licensee is reduced to a series of minor deviations. The real issue—the potentially grave consequences of a combination of deviations—is not addressed. The translation of the technical problem into narrow legal categories often trivializes it and, as the engineer said, "The licensee winds up thinking we're just picking on him." More concern with the consequences of system interaction would help redirect attention away from the detailed specifications and toward how the plant is actually operating.

The Nuclear Island

Why has nuclear regulation taken on these characteristics? How did the enormous growth in detailed prescriptive regulation take place? Over the past 30 years there have been some important shifts in the way the "whole" has been defined for purposes of regulation. First of all, the scope of what is to be regulated has expanded steadily. The relevant whole was once thought to be simply the nuclear steam supply system, a "nuclear island" located in the middle of an otherwise rather unremarkable power plant. Today the relevant whole includes not only a major portion of the non-nuclear hardware surrounding the nuclear island, but also many of the human and organizational aspects of the utility operation.

In the early years of nuclear power regulation, it was assumed that the nuclear aspects of plants could be separated from the rest. The nuclear steam supply system, the reactor itself, was distinguished from the turbines, generators and other systems that directly produced electric power. Nuclear regulators were to be responsible only for the nuclear parts of the plant.

Things did not turn out to be so simple, however. While the regulators' mandate to promote nuclear power encouraged a reliance on the utilities (after all the utilities were the ones with the know-how), their parallel mandate to look after nuclear safety eventually led in a different direction. Conflict emerged as the industry became more familiar with nuclear issues, and as the more technically trained regulators became con-

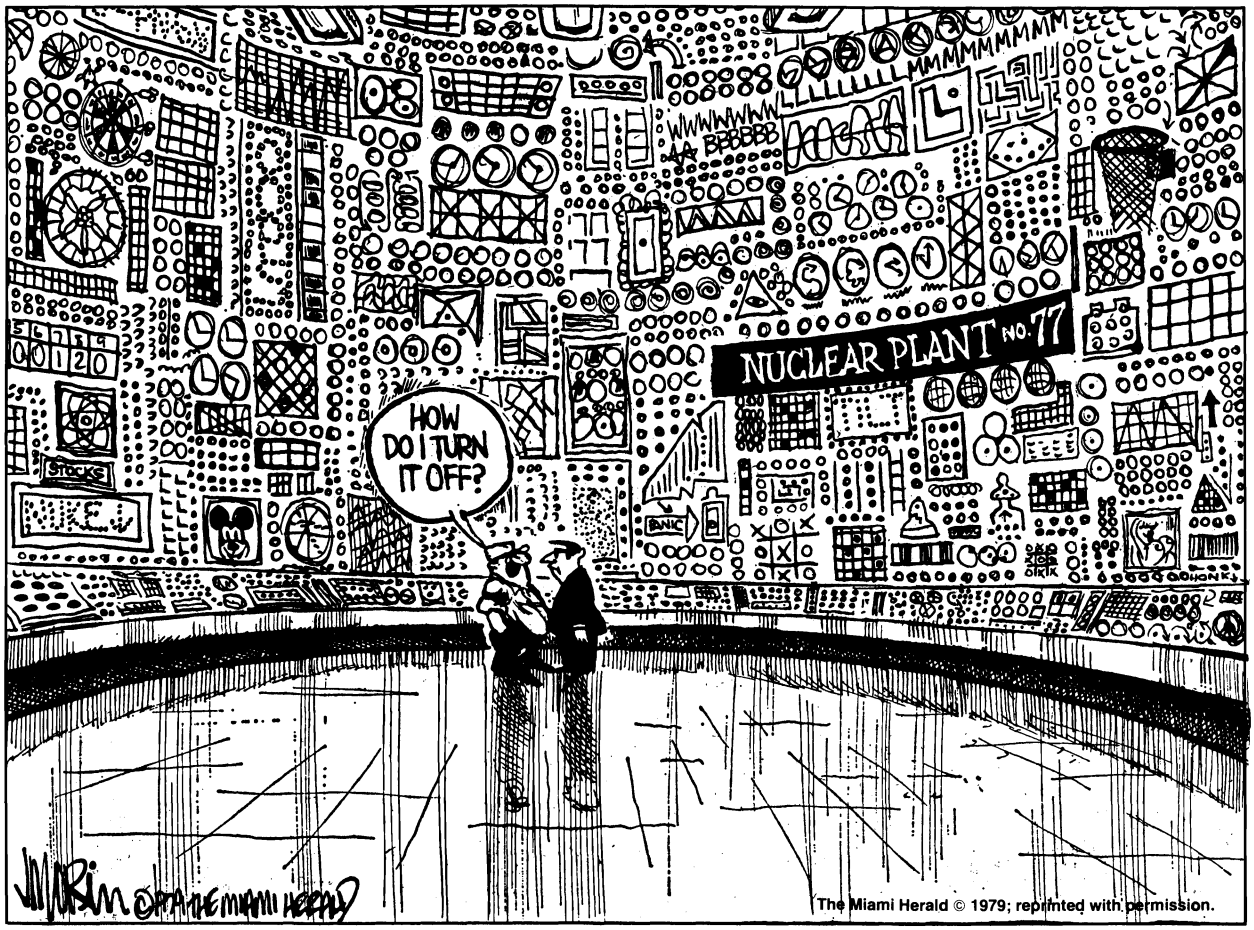

cerned with the industry's ability to perform up to its promises.

Regulations require that a utility must provide a preliminary safety analysis report at the time it applies for a construction permit. In the beginning there were few specific requirements that had to be met. This had its advantages, but at the same time the utility found it difficult to "prove" it had done an adequate job and could assure safe operation. Before the utility could receive an operating license for its new plant it had to get approval of its final safety analysis report (FSAR). As we were told by a long-time licensing reviewer, "We got into the 'bring me a rock' syndrome with approving the FSAR. We kept saying, 'give me more on this and on that.' The utilities said, 'What do you want? Sharpen up your questions!' So we began to develop a standard format: What do we really want to know?" As the standard format developed, the variation across licensing cases grew less, but the number of requests and the amount of detail required in the answer increased. As problems arose in each area, and as new solutions to old problems were discovered, these were incorporated into the regulatory process by the specialists assigned to those areas.

Whatever seemed to be an important new safety practice adopted by any one licensee or vendor was urged on the rest of the industry. "Good ideas" and "good practices" were formalized. Even when they were not made specific requirements, their use in regulatory guides meant the agency actively encouraged their adoption. The guides provided the industry with an understanding of what the regulators wanted to see included in their plans.

The utilities that had chosen to purchase nuclear power plants were surprised by the increasing scope and detail of the regulatory reviews. An NRC supervisor points out:

> [the utilities] never thought [regulators] would say what type of construction materials to use, stainless steel, or whatever, let alone be saying what types of training or management programs they should have. Their reaction was, "What the hell do you know about it?" Even today you'll get utility

executives saying, "We generate millions and millions of kilowatt hours per month and we've been doing it for 60 years!"

As regulation became more detailed and covered more territory, the number of potential points of conflict increased proportionately. Utility management, accustomed to being responsible for the safe operation of its plants as well as the production of power, found itself increasingly constrained.

The new steam supply systems were clearly not the independent nuclear islands the utilities had been led to believe they were. The division between the nuclear and the non-nuclear aspects of the plant turned out to be hard to maintain. The notion of the nuclear island was modified. A new distinction arose between systems that were "safety-related" and those that were not. This classification helped to reinforce the line between what was to be regulated and what was the sole concern of the licensee.

Well before the accident at Three Mile Island, however, there was already disagreement over which systems could be defined as "safety-related" and exactly what components were included in each system. After the accident, the report of the the President's Commission on the Accident at Three Mile Island (the Kemeny Commission) discredited the whole notion that adequate safety could be assured by oversight of only certain parts or systems. The accident was not due to the kind of large pipe break or other catastrophic failure that was anticipated and guarded against by the NRC policy of strengthening the parts. Instead it was due to a number of lesser failure conditions combined with the operators' misunderstanding of the situation confronting them, an error easier to make as a plant grows more complex. Following the accident, we were told by an NRC engineering manager, "the concept of a 'nuclear island' and a limited regulatory purview was dead. You have to look at the whole plant. There was still some debate about it until Three Mile Island, but there was no debate afterwards."

While the definition of the relevant "whole" has expanded to include all systems needed for safe shutdown, the question of what these systems and functions are and how they interact with one another is still problematic. Should the regulatory purview, for example, include the oversight of management and personnel systems as well as the hardware? Should only offsite releases and worker exposures be monitored, or the safety of the physical plant itself? And how exactly might one estimate the importance of each part for the whole?

System Interaction

An approach that emphasizes adding safety devices to achieve safety is not, of course, always unsafe. We are not suggesting that attention to the parts is inappropriate. Some individual parts used in nuclear power plants have revealed surprising frailties. (Indicators on instrument panels, for example, have shown a startling tendency to short out or give false readings when workers attempt to change the tiny indicator bulbs.) Something further is required, however: an understanding of the significance of each part for the safety of the whole, the ways in which various safety measures may reinforce or counteract each other, and the time dependencies of these interactions. Without this understanding, the real consequences of regulatory actions cannot be known.

The "regulation added—safety achieved" relationship is a contingent one. It depends on the way the regulated parts interact with one another, and on the actual physical and organizational processes involved. The selection of parts to regulate has not taken these contingencies into account. The selection is, to some extent, the result of methodical learning about what makes reactors safe, but there are many random and arbitrary elements in the selection process as well, reflecting a desire to play it safe politically by pointing to the measures one has taken, whether or not they achieve safety. Hence the relationship between adding yet another safety measure and actually achieving safety is apt to be precarious. Though we cannot deal with this matter here, we think it would be better to combine less detailed specification with more general performance standards in order to learn better how to relate parts to wholes in securing greater safety in nuclear power plants.

Conclusions

Our purpose is not to attack or defend nuclear power. Our purpose is to make a point of general interest: It is possible to do harm in the name of safety. Every act and actor is potentially dangerous; merely labeling a measure as designed to secure safety by no means guarantees that result. □