



HOKKAIDO UNIVERSITY

Title	Three Nuclear Disasters and a Hurricane : Some Reflections on Engineering Ethics
Author(s)	Davis, Michael
Citation	Journal of applied ethics and philosophy, 4, 1-10
Issue Date	2012-08
DOI	https://doi.org/10.14943/jaep.4.1
Doc URL	https://hdl.handle.net/2115/50468
Type	departmental bulletin paper
File Information	jaep4-1_micael davis.pdf



Editorial Note

The Journal of Applied Ethics and Philosophy is an interdisciplinary periodical covering diverse areas of applied ethics. It is the official journal of the Center for Applied Ethics and Philosophy (CAEP), Hokkaido University. The aim of the *Journal of Applied Ethics and Philosophy* is to contribute to a better understanding of ethical issues by promoting research into various areas of applied ethics and philosophy, and by providing researchers, scholars and students with a forum for dialogue and discussion on ethical issues raised in contemporary society.

The journal welcomes papers from scholars and disciplines traditionally and newly associated with the study of applied ethics and philosophy, as well as papers from those in related disciplines or fields of inquiry.

Shunzo Majima
Editor-in-Chief

Three Nuclear Disasters and a Hurricane:

Some Reflections on Engineering Ethics

Michael Davis

Illinois Institute of Technology, USA

Abstract

The nuclear disaster that Japan suffered at Fukushima in the months following March 11, 2011 has been compared with other major nuclear disasters, especially, Three Mile Island (1979) and Chernobyl (1986). It is more like Chernobyl in severity, the only other 7 on the International Nuclear Event Scale; more like Three Mile Island in long-term effects. Yet Fukushima is not just another nuclear disaster. In ways important to engineering ethics, it is much more like Katrina’s destruction of New Orleans than like any nuclear disaster. It is (primarily) a consequence of a natural disaster, the enormous earthquake and tsunami that wrecked much of northeast Japan. One lesson of Fukushima, one shared with Katrina, concerns the different roles engineers have at different stages in an engineering project (planning, designing, management, and operations). In the planning stage, engineers seem to have relatively little power to affect certain early large-scale trade-offs between public safety and public welfare. Another lesson may be the importance of not leaving complex technical systems untended. The events that made the disasters at Three Mile Island and Chernobyl inevitable lasted only a few minutes or hours; the events that made the disasters in New Orleans and Fukushima inevitable were spread over several days. Fukushima avoided a more serious disaster because the plants were not abandoned in the way New Orleans was. A third lesson concerns our ideas of heroism, especially our sense that heroism is sometimes one’s duty. An engineer’s duty sometimes includes protecting others from harm even at the risk of the engineer’s life.

Keywords: Chernobyl, Fukushima, Katrina, Three Mile Island, precautionary principle

This paper began with an invitation from Japan a month after the disaster at Fukushima I Nuclear Power Plant began on March 11, 2011. The paper was to be presented six months later (October 30)—at a conference where I was already scheduled to present a paper on the use of imaginary cases in ethics (Davis, 2012). The disaster was still very much the province of journalism. Its outlines certainly lacked the stability of history. Of course, even history, though it generally seems stable, is not entirely so, being subject to dispute here and there and to radical revision every now and then. At first, Fukushima’s facts changed almost daily—if by “facts” we mean those descriptive propositions about which there is general agreement. After a while, the changes were less frequent and more a matter of addition than correction. Today,

a year after I began work on the paper, the outline of the disaster seems settled. Dispute now concerns only details, such as how much land, if any, will have to be abandoned for some years, how many pre-mature deaths are to be expected because of radiation released during the disaster, and so on. At some point, I had to stop worrying about the facts and report my reflections. I stopped worrying about the facts on October 15, 2011. Since then, I have changed “a fact” only when a reader or auditor pointed out that it was no longer fact.

I am an oddity in science and technology studies (STS) because I focus not on science and technology but on scientists and technologists. Indeed, I do not write about “scientists” in general or “technologists” in general but about specific professions, for example, chemists or engineers. While most STS scholars seem

to be interested in what scientists or technologists have in common, I have focused on what distinguishes one discipline from another, for example, what distinguishes chemists from engineers (Davis 2002). I have found that the most fruitful way to study professions, especially the profession of engineering. I am, in short, not a philosopher, historian, or sociologist of technology (though scholars in those fields sometimes find my work useful). The credential that justified the invitation from Japan was a quarter-century of thinking (and writing) about engineering. (For those unfamiliar with my work on engineering, the place to start is Davis 1998.)

That invitation from Japan presented me with a practical problem. The newspapers, websites, and other sources available (at least in English) seldom identified anyone as an engineer. The stories focused on “workers”, “managers”, and machinery. I had to use what I knew about nuclear power plants in the United States to interpret the facts thus given. I had similar problems, though less severe, when interpreting the other disasters to which I chose to compare Fukushima. Interpretations are, of course, open to objection but, without interpretation, facts merely pile up, becoming in time an unmanageable heap. There is no understanding without interpretation. But interpretation relying on changing facts is necessarily the sort of time-stamped enterprise philosophers are inclined to avoid—and I would have avoided it but for that invitation from Japan. There is not much that a philosopher can do about a disaster such as that at Fukushima—except help those seeking to understand it and thereby help to prevent similar disasters. I felt I owed the Japanese that much.

This paper’s title promises “reflections” on Fukushima, not systematic or definitive understanding. Reflections are what one gets when, focusing thought on certain facts, one captures connections one happens to see, connections that seem to jump out of the dark. Reflection is a source of hypothesis rather than proof, the beginning of a discussion rather than the end. We do not need reflection when we can derive a conclusion from what we know. Reflection is useful when we want to discover a conclusion that, though far from provable given the facts we have, invites investigation. There is no algorithm for reflection, no test of success beyond useful surprise.

Why Compare These Four Disasters?

The nuclear disaster that Japan suffered at Fukushima has been compared with other major *nuclear* disasters, especially, Three Mile Island (1979) and Chernobyl (1986). It is more like Chernobyl in immediate destructiveness, the only other 7 on the International Nuclear Event Scale (the upper limit of which is 7). It

is more like Three Mile Island in probable long-term effects (though Fukushima’s long-term effects are likely to be substantially worse than Three Mile Island’s). To date, Chernobyl seems to have directly killed thirty-one reactor staff and workers, to have caused between 200,000 and 1,000,000 premature deaths worldwide, to have forced the permanent abandonment of a city of about 50,000 (Pripyat), and to have ruined perhaps a 100,000 square km of farmland. Over 300,000 people lost their homes to contamination. (All information about Chernobyl here and below is drawn from Wiki, “Chernobyl”, a source valuable both because it is easily accessed and regularly updated.)

In contrast, the radiation released from the Fukushima plant, though significant, will, it seems, leave little long-term contamination, except at the plant itself and in a plume perhaps fifty km beyond. At least six workers have exceeded lifetime legal limits for radiation and more than three hundred have received significant radiation doses. Estimates of future cancer deaths due to accumulated radiation exposures in the population living near Fukushima have ranged from none to a non-peer-reviewed “guesstimate” of a thousand. No one died in the explosions at the plant or from subsequent radiation exposure (though the tsunami killed two workers and evacuation of hospitals in the exclusion zone may have caused as many as forty-five more deaths). The earthquake or tsunami, rather than the nuclear accident, seems to be responsible for the few employees severely injured or killed at the plant. (Wiki, “Fukushima Daiichi”.)

The discussion of Fukushima below relies not only on this source but also on Wiki, “Fukushima I”. Though I shall hereafter refer to “Fukushima”, it is in fact Fukushima I (Fukushima Dai-ichi) that I shall be referring to. There is also a Fukushima II (Fukushima Dai-ni). For details, see Wiki, “Fukushima II”.

Though certainly a nuclear disaster, Fukushima is not just another nuclear disaster. In ways important to engineering, it is much more like Katrina’s destruction of New Orleans than like any other nuclear disaster. It is (primarily) a consequence of a natural—or, at least, much larger—disaster, the enormous earthquake and tsunami that wrecked much of northeast Japan on March 11, 2011, killing about 28,000 people. Fukushima has many lessons to teach, especially if we compare it with these other disasters. Here I shall focus on four lessons: The first concerns the different roles engineers have at different stages in an engineering project, especially the relative powerless of engineers to affect certain early large-scale trade-offs between public safety and public welfare. A second lesson may be the need to evaluate risk in ways beyond ordinary cost-benefit analysis when the risks are improbable but catastrophic. A third lesson is the importance of not leaving complex technical

systems untended. Engineering systems do not work long without engineers. A fourth lesson may concern the way engineers should respond, and typically do respond, to engineering disasters. They should take responsibility for limiting the harm as well as for fixing the underlying problem, even if limiting the harm involves risking their lives. To see what I mean, let us consider these four disasters in greater detail, beginning with the first.

Three Mile Island

Three Mile Island was a “normal accident”, that is, it began with ordinary failures of equipment and practice within a plant itself operating normally. Perrow 1984 also describes Three Mile Island as a “normal accident”. While I agree that it was a “normal accident” in his sense, my use of that term is somewhat different. I mean simply that the accident was a product of what engineers normally do rather than a product of incompetence, negligence, corruption, or other unusual conduct (such as experimentation).

During the night of March 27-28, 1979, workers were engaged in routine cleaning of a blockage in one of Reactor 2’s eight condensate polishers (filters for the secondary cooling loop). At 4 am, the pumps feeding the polishers stopped. We still do not know the cause of the stoppage. When a bypass valve failed to open, water ceased flowing to the secondary loop’s main feed-water pumps. These also shut down. No longer receiving water, the steam-driven generators stopped and the reactor automatically carried out an emergency shutdown. Within eight seconds, control rods were inserted into the core to halt the nuclear chain reaction. The reactor nonetheless continued to generate heat (a byproduct of normal decay). Because steam was no longer being used by the turbine, heat was no longer being removed from the reactor’s primary water loop. (Except where otherwise indicated, the discussion of Three Mile Island here and below relies for its facts on Wiki, “Three Mile Island”.)

Once the secondary’s feed-water pumps stopped, three auxiliary pumps started up automatically; but because some valves were closed for routine maintenance, the system could not pump water. So, the secondary loop was no longer working. Without the secondary loop removing heat, pressure in the primary loop began to increase, automatically triggering a relief valve. The relief valve should have closed again when the excess pressure had been released; instead, it stayed open. That open valve permitted coolant water to escape from the primary system. It was the principal mechanical cause of the coolant-loss meltdown that followed.

The mechanical failures were compounded by the failure of plant operators to recognize the situation as a

loss-of-coolant accident for more than two hours. (One cause of their failure seems to have been an indicator light blocked from view.) That initial failure led an operator to override the reactor’s automatic emergency cooling system manually. With the release valve still open, the quench tank that collected the discharge from the release valve overfilled, causing the containment building’s sump to fill and sound an alarm at 4:11 am (eleven minutes after the first pumps failed). That alarm, along with higher than normal temperatures on the discharge line and unusually high temperatures and pressures in the containment building, clearly indicated that there was a loss-of-coolant accident, but the operators did not respond to these indications. At 4:15, the quench-tank relief diaphragm ruptured and radioactive coolant began to leak out into the general containment building. This coolant was pumped from the containment building sump to an auxiliary building, outside the main containment, until the sump pumps were stopped at 4:39 am.

After almost eighty minutes of slow temperature rise, the primary loop’s four main pumps began to suffer damage as a mixture of steam and water passed through them. The operators then shut down the pumps, believing that natural circulation would continue the water movement, but steam in the system (itself the product of rising temperature) prevented coolant flow through the core. As the coolant stopped circulating, it increasingly turned to steam. Just over two hours after the first sign of trouble, the coolant level fell so low that the top of the reactor core was exposed to the steam. Intense heat then caused a reaction between the steam in the reactor core and the nuclear fuel-rod cladding. That reaction burned off the cladding and damaged the fuel pellets. The pellets then released more radioactivity into the reactor coolant, producing hydrogen gas that probably caused a small explosion in the containment building in the afternoon.

At 6 am (two hours after the incident began), there was a change of shift in the control room. A new arrival noticed that temperatures in the relief valve tailpipe and holding tanks were too high and used a backup valve to shut off the coolant venting through the relief valve. But, by then, about 120,000 liters of coolant had already leaked from the primary loop. Not until almost 7 am (almost three hours after the incident began) did contaminated water reach radiation-activated alarms. By then, the radiation in the primary coolant water was around three-hundred times higher than usual. The plant was seriously contaminated and the reactor’s core had suffered a partial meltdown.

The Nuclear Regulatory Commission (NRC) made an extensive investigation of the disaster, a typical engineering response. Its report ended with recommendations for changes in controls, quality assurance, maintenance, operator training, management,

and communication of important safety information. There was no finding of negligence or more serious wrongdoing having caused the disaster, no suggestion that major redesign of nuclear plants was needed, and no proposal to rethink the place of nuclear energy in the generation of electricity. (Rogovin 1980, pp. 89-93, focused mainly on changes in emphasis and procedures at the NRC; Kemeny 1979, pp. 61-73, focused on “attitudes and practices”). These reports do, however, contain much criticism of other aspects of how Three Mile Island operated.

Chernobyl

Chernobyl was not a normal accident. Its cause was an engineering *experiment* which, though successful, lacked proper approval. That is not to say that the experiment was unjustified, fundamentally improper, or indeed abnormal.

Even when not actively generating power, nuclear reactors require cooling to remove heat produced by the natural decay of nuclear fuel. Chernobyl’s pressurized water reactors (different in design from Three Mile Island’s) used water flowing at high pressure to remove waste heat (about 28,000 liters of water an hour). After an emergency shutdown, the core could still generate a significant amount of residual heat. If not removed, the heat could cause core damage (as it did at Three Mile Island). If the power grid failed, power to run the plant’s cooling system might be unavailable from outside for far too long.

Chernobyl’s reactors had three backup diesel generators. Each generator required fifteen seconds to start up but took over a minute to attain the speed required to run one of the main coolant pumps. Chernobyl’s engineers judged this one-minute power gap unacceptable. Too much can happen in a nuclear reactor in a minute when the cooling system is not working. Analysis indicated that one way to bridge the one-minute gap was to use the mechanical energy of the steam turbine and residual steam pressure to generate electricity to run the main coolant pumps while the generator was reaching the correct RPM, frequency, and voltage. But, of course, the analysis had to be confirmed experimentally. The engineers had to work out and then prove a specific procedure for effectively employing residual momentum and steam pressure.

Previous experiments—in 1982, 1984, and 1985—had ended in failure. The 1986 experiment was scheduled to take place at Reactor 4 during a maintenance shutdown. The experiment focused on refinements in the switching sequences of the electrical supplies for the reactor. The experiment was to begin with an automatic emergency shutdown. Because no

danger to the reactor was anticipated, the engineers did not formally coordinate the experiment with either the reactor’s chief designer or scientific manager. Indeed, the experiment did not even have the approval of the onsite representative of the Soviet nuclear oversight agency. Only the director of the plant approved it (and even his approval did not follow standard procedures).

The experiment began just after 1:23 am on April 26, 1986. The diesel generator started and sequentially picked up loads. The turbine generator supplied the power for the four main circulating pumps as it coasted down. The experiment was all but complete forty seconds later. But, as the momentum of the turbine generator that powered the water pumps decreased, the water flow decreased, producing more and more steam bubbles in the core. The reactor was now ready to begin a destructive feedback loop: The production of steam would reduce the ability of the coolant to absorb neutrons, increasing the reactor’s output of heat. The increased heat would cause yet more water to become steam, further increasing heat. During almost the entire period of the experiment, the automatic control system successfully counteracted this destructive feedback, inserting control rods into the reactor core to keep the temperature down.

If conditions had been as planned, the experiment would almost certainly have been carried out safely. The Chernobyl disaster resulted from attempts to boost the reactor power—and, therefore, temperature—once the experiment had started (something inconsistent with approved procedure). The approved procedure called for Reactor 4’s power output to be gradually reduced to 700–1000 MW. The minimum level established in the procedure (700 MW) was achieved about an hour before the experiment began. However, because of the natural dampening effect of the core’s neutron absorber, reactor power continued to decrease, even without further operator action.

As the power dropped to approximately 500 MW during the experiment, one of the engineers conducting the experiment mistakenly inserted the control rods too far, nearly shutting down the reactor. Control-room personnel soon decided to restore the power and extracted the reactor control rods, but several minutes elapsed between the extraction and the time that the power output began to increase and stabilize at 160–200 MW. The extraction withdrew the majority of control rods to the rods’ upper limit, but the rapid reduction in the power during the initial shutdown and subsequent operation at less than 200 MW led to increased dampening of the reactor core by the accumulation of xenon-135 (an unstable fission product of uranium that absorbs neutrons at a high rate). To counteract this unwanted high-absorption, the operators withdrew additional control rods from the reactor core.

Then, about the time the experiment ended, there was an emergency shutdown of the reactor. The shutdown started when someone pressed the button of the reactor’s emergency protection system. (We do not know whether the button was pressed as an emergency measure, by mistake, or simply as a routine method of shutting down the reactor upon completion of the experiment.) Because of a flaw in the design of the graphite-tip control rods, the dampening rods displaced coolant before inserting neutron-absorbing material to slow the reaction. The emergency shutdown therefore *briefly* increased the reaction rate in the lower half of the core. A few seconds after the start of the emergency shutdown, there was a massive power increase, the core overheated, and seconds later this overheating produced the first explosion. Some of the fuel rods fractured, blocking the control-rod columns and causing the control rods to become stuck at one-third insertion. Several more explosions followed, exposing the reactor’s graphite moderator to air, causing it to ignite. Since the reactor lacked a containment (a thick concrete shell), the fire in the reactor sent a plume of highly radioactive smoke into the atmosphere, causing dangerous fallout over a huge area (as much as five-hundred km away)—and, eventually, less dangerous fallout over much of the world.

The effort to halt the nuclear contamination and avert a much greater disaster soon involved over 500,000 workers and cost an estimated eighteen billion rubles, crippling the Soviet economy.

Because most of those directly involved in the Chernobyl disaster soon died of radiation poisoning, there are many uncertainties about the exact sequence of events. Nonetheless, we can be sure that the actual disaster would not have occurred had the experiment not been carried out. The Chernobyl disaster combines the “normal failures” of operators and equipment we saw at Three Mile Island with an experiment of the sort engineers often perform, though an experiment necessarily introduces the unexpected. Chernobyl was as much an engineering disaster as Three Mile Island: both the immediate and underlying causes were ordinary engineering decisions, whether in operation or design.

Fukushima

The disaster at Fukushima fits neither of these patterns. The accident was not normal or the result of an engineering experiment. It was also not the result of operator negligence, incompetence, or misconduct. The disaster began with a large earthquake, one larger than any Japan had experienced in 1400-years of recorded history (http://en.wikipedia.org/wiki/List_of_earthquakes_in_Japan, accessed April 25, 2011). The quake was

followed by an enormous tsunami. That double disaster would have happened even if the Fukushima nuclear power plant, one of the twenty-five largest in the world, had never existed. The nuclear disaster is a byproduct of that larger natural disaster.

At the time of the quake, 2:46 pm, Reactor 4 had been de-fueled while 5 and 6 were in cold shutdown for planned maintenance. The remaining three reactors shut down automatically in response to the quake. After the reactors shut down, the plant’s own generation of electricity ceased, eliminating one source of electricity used to run cooling and control systems. One of two connections to the national electrical grid also failed. That loss of power started up thirteen on-site emergency diesel generators. These would ordinarily have provided enough power to operate the reactors’ control and cooling systems until the lost connection to the national grid could be restored. Had the earthquake been the only disaster to hit the Fukushima plant on March 11, there would have been little to discuss here. The tsunami changed that.

The plant was protected by a seawall designed to withstand any tsunami up to 5.7 meters, but the great wave that struck forty-one minutes after the quake was fifteen-meters high. It flooded the entire plant, including generators and electrical switchgear in reactor basements. It also broke the remaining connection with the national electrical grid. All conventional power for cooling was lost. Only one backup remained: emergency batteries, able to run some of the monitoring and control systems for up to eight hours. Replacement batteries and mobile generators were soon dispatched to Fukushima, but collapsed bridges, debris-strewn roads, and similar obstacles delayed them. The first replacements did not arrive until 9:00 pm (six hours after the first call went in).

The arrival of the replacement batteries and mobile generators did not end the crisis, however. They had to be installed. The normal connection points were in flooded basements. There was also difficulty finding suitable cables. Work to connect batteries and generators was still continuing twenty-four hours after the quake when there was an explosion in Reactor 1’s building. The side walls of the upper level were blown away, the roof collapsed, and debris covered much of the floor and machinery.

The roof of the building was designed to provide ordinary weather protection, not to withstand an explosion or to act as containment for the reactor. In the Fukushima reactors, the primary containment surrounded the reactor’s pressure vessel. The top floor had no reactors, only water filled pools for storing new fuel ready to be craned into the reactor and used fuel ready for disposal.

This first explosion was probably caused when

hydrogen collected under the roof. Exposed fuel rods became very hot and reacted with steam, oxidizing the cladding and releasing hydrogen. The hydrogen would have leaked upward. Safety devices normally burn such hydrogen before it reaches explosive concentrations. These systems seem to have failed when the electrical power did.

Reactor 1’s containment survived the explosion. There were no large leaks of radioactive material, although there was an increase in radiation following the explosion. The explosion at Reactor 1 injured four workers. But this was only the beginning. Hydrogen gas was also collecting at the other five reactors. Over the next few days, hydrogen explosions destroyed the upper cladding of the buildings for Reactor 3 and 4 and the containment inside Reactor 2. Several fires broke out at Reactor 4. In addition, spent-fuel rods stored in the spent-fuel pools of Reactors 1–4 began to overheat as the water level dropped. Fear of radiation leaks led to evacuation of all non-essential persons within a twenty-kilometer radius of the plant.

In short, the Fukushima plant was overwhelmed by forces from outside well beyond what it was designed for. Without heroic efforts by plant staff, some of whom may die over the next few years because of exposure to radiation, the Fukushima disaster might have become at least as devastating as Chernobyl. Even with those heroic efforts, several weeks passed before the plant could be said to be more or less under control. One generator at Reactor 6 was restarted on March 17 (six days after the quake) allowing some cooling at Reactor 5 and 6, the least damaged. Connection to the power grid was restored to parts of the plant on March 20, but machinery for Reactors 1-4—damaged by flooding, fires, and explosions—could not be restarted for several months. Only in early October 2011 did coolant in all the reactors reach safe temperatures.

The Fukushima plant could have been designed to withstand the natural disaster that occurred. A breakwater three times higher than the actual breakwater could have protected the plant against the tsunami (assuming it survived the quake); the plant might have been located far enough away from the ocean to be safe from even so large a tsunami; generator-building basements might have been made waterproof; and so on. Even some less expensive arrangements might have improved what happened considerably. For example, storing more batteries on site would have allowed the cooling and control systems to function longer without repair or resupply, weeks instead of hours. But all of these changes would have been (more or less) expensive, raising the price of the electricity the plant produced. Typically, engineers, though consulted, do not make such decisions. Government regulators, senior management, or public opinion typically decide, for example, whether to protect

against a 500-year, 1,000-year, or 10,000-year quake.

Katrina

When it struck New Orleans on August 29, 2005, Katrina was a category 3 hurricane, a large storm but no larger than storms that strike the Gulf Coast almost every year. (The top of the hurricane scale is 5.) Katrina was nonetheless unusually destructive because it moved so slowly that anything in its path was subject to heavy rains and high winds for many hours. The rain and high winds were, however, only part of what caused so much destruction in New Orleans. (Except as otherwise indicated, all information in this section comes from Davis 2007.)

Even on an ordinary day, New Orleans is a city that must work to prevent flooding. One of the world’s largest rivers, the Mississippi, flows through it. From Jackson Park, the jewel of the tourist-drawing French Quarter, one of the *highest* points in the city, one can see the mighty river rushing by about two meters *above* the street. On any day of the year, the Mississippi would flood the city were it not for the levees that hold it back. Nor is the Mississippi the only watery threat. Though the oldest parts of the city are as much as ten meters above sea level, a majority of the city is below, and the sea, the Gulf of Mexico, reaches New Orleans at its back, through Lake Pontchartrain, and underground, through the water table. (While the water under New Orleans is fresh, it is as high as it is in part because the Gulf’s salt water is not lower.)

Mostly developed since 1900, the newer parts of the city are, like much of the Netherlands, dry only because water is constantly pumped out. Every year, there is more for the pumps to do. Sea level is rising about a third of a meter a century; some parts of the city have subsided by half a meter or so because the weight of buildings is compressing the soil or because pumping water from the ground allows the soil to compress. Were it not for huge screw pumps working day and night, New Orleans would today be a version of what it was when the French first settled there in 1718, a crescent-shaped string of small islands in a huge swamp. Like Venice in Italy and St. Petersburg in Russia, New Orleans is much more artificial, and therefore much more vulnerable to natural forces, than most cities. Engineers did not found New Orleans, but the city has long survived only because of engineering. The floods the city suffers from time to time are due in part to the engineering not being good enough. That is as true of Katrina as of earlier disasters, for example, the one in 1965 named for hurricane Betsy (nearly as destructive as Katrina).

Katrina flooded New Orleans because the levee system failed catastrophically. Much of the disaster,

however, occurred hours *after* the storm had moved inland as water poured through holes in levees and filled much of the city. There was no attempt to repair the levees immediately. Indeed, for many days, there were no officials in New Orleans even to report damage. Everyone who could be evacuated had been. By August 31 (two days after Katrina struck), 80% of New Orleans, a city almost emptied of inhabitants, was under water, with some parts under water almost five meters deep. The water lingered for weeks.

On March 26, 2007, a year and a half after Katrina passed through New Orleans, the Interagency Performance Evaluation Task Force (IPET) issued its (draft) *Final Report*. IPET was an independent team of more than one-hundred-fifty international and national experts from more than fifty different government organizations, universities, and private companies. The U.S. Army Corps of Engineers commissioned IPET a few weeks after Katrina hit New Orleans. It was to analyze how the levee system performed. Though many questions of detail remain unsettled, this nine-volume report, is (more or less) the last word on both the causes of the Katrina disaster and means of preventing similar disasters.

IPET reports a “system” that grew up piecemeal, only in part under the control of the Corps of Engineers, the government agency officially in charge of waterways. In some places, the system failed because a levee or other barrier to water was not high enough, often because of unanticipated subsidence rather than original design error. In other places, the system failed because, though high enough, the barriers were not designed for the forces to which they were in fact subject (an unusually slow-moving storm). Design of floodwalls along three canals was “particularly inadequate”. A series of incremental decisions between the original plan and the structures actually constructed “systematically increased the inherent risk in the system without recognition or acknowledgment” (IPET 2006, I-2). Many of the failures in the system would not have occurred had implementation of plans for reconstruction not been delayed for almost twenty-five years by inadequate funding, new laws governing the environment, and similar difficulties well beyond the control of engineers. For some important “decisions”, there was no decision-maker at all. The decisions were a mere byproduct of poor communication, poor information, poor coordination, or some combination of these.

The most important lesson IPET drew from its analysis is unsurprising: The way to avoid similar disasters is to use larger safety factors (“conservative design assumptions”) and good materials (“higher quality, less erodible”). (IPET 2006, I-3).

The flood control system now replacing the one Katrina overwhelmed is considerably more expensive

than the old one. For example, the Corps has been replacing the five-meter pilings holding canal walls in place with pilings that would go down *fifteen and a half* meters (three times as deep). The Corps agreed that the use of I-walls along the canals (without or even with the support of a simple earthen levee) was a mistake. It is replacing the canal’s I-walls with heavily-braced T-walls locked down by twenty-one meter H-piles angled out in two directions. The use of simple sand or gravel levees was also judged a mistake. The Corps is now “armoring” all levees where they seem vulnerable to overtopping, that is, covering them with something water will not soak through or quickly wear away. These are expensive changes in design that government was unwilling to pay for without a major disaster and may yet lose interest in paying for before the work is complete.

Conclusions

We can, I think, distinguish four sorts of engineering decision in these four case studies: planning, designing, management, and operations. (In a different context, I would include “disposal” in this list. I do not include it here only because none of these disasters concerns disposal as such, though Fukushima’s problems were due, in part, to fuel rods waiting disposal.)

By planning, I mean such decisions as whether to build a nuclear power plant at all, where to put it, and the upper limit of its budget. For such decisions, engineers are most important for vetoing certain options, for example, a location because the risk of earthquake makes safe construction too expensive. They are also important for suggesting alternatives, for example, conservation or a gas-fired plant rather than a nuclear plant. Engineers are not (or, at least, should not be) mere “problem solvers”. One important function they have is helping to define problems—or re-define them when it becomes clear that the client or employer has not asked the right question.

But, for any large undertaking such as a nuclear plant or flood control system, engineers are generally only one party in a complex social decision in which the other parties include employer, government officials, experts of various sorts (such as geologists), bankers, and civil society (or “the public”). Perhaps the most important contribution engineers can make to planning is developing minimum standards for evaluating and responding to specific risks and benefits of the technology in question.

By designing, I mean the actual drafting of specifications, floor plans, and so on necessary to construct or modify the technological artifact in question. Once planning has set limits, engineers are generally free to work within those limits, for example, to design

a nuclear plant that will fail slowly rather than quickly or cool rather than heat up if left alone. Only when a planning limit is too strict do engineers have a reason to restart the planning process, for example, by suggesting that the budget be raised to provide an adequate margin of safety.

By management, I mean overseeing the operations of a plant, including choosing, training, and directing operators. Much management is not technical—and is therefore not the domain of engineers. But, for nuclear plants or flood control systems, the managers will typically be engineers. For engineers, part of technical management is remaining alert to possible improvements in staff, procedures, and equipment. So, for example, a manager who noticed that operators at Three Mile Island often missed readings on an important gauge because equipment blocked their view of it should recommend, or order, that the control board or control room be redesigned to improve the view.

By operations, I mean actually doing what is necessary for the plant or other technical artifact to work. While engineers do not, in general, operate plants, they do constitute most of the operators in a nuclear plant. So, for example, at Chernobyl, they pushed the buttons that moved dampening rods into the core. While operators can be reprimanded, and their acts reversed, they are, while acting as operators, completely in control of their machines. One of the features we noted in our discussion of the three nuclear disasters is how quickly things can go wrong. What goes wrong in a nuclear plant does not, of course, go seriously wrong for just one reason. Because engineers typically design nuclear plants with a large safety factor, several systems must fail before anything goes seriously wrong. But, given the complexity of a nuclear plant, it is reasonable to expect at least one system to fail now and then because, even with proper maintenance and inspection, technical systems sometimes fail unexpectedly. That being so, it is also reasonable to expect (given the laws of statistics) that all of the independent systems will fail together sooner or later. One of the “systems” that may fail at any given time is the human operator—whether because of distraction, fatigue, poor training, misjudgment, interference, or the like.

How likely is a catastrophic failure at any moment? Not very. Perhaps only 10^{-5} at any time. But over many years and many reactors even such small risks add up. One author recently calculated that there are:

450 nuclear power plants in the world. There have been 4 meltdowns in history, one each at Chernobyl and Three Mile Island and two so far at Fukushima, as partial meltdowns count as meltdowns. That is a ~1% failure rate. (Lindsay, 2011)

This calculation means nothing unless the four meltdowns are statistically significant, that is, a good predictor of what will happen over, say, the next hundred years (rather than a chance concurrence of events—like winning the lottery three days in a row). Still, it is an empirical reminder that even a low-probability event will, given a large enough population, become highly probable.

If we look at our four disasters, two—Three Mile Island and Chernobyl—seem unrelated to any ordinary planning or design failure. Of course, with a higher budget, the Three Mile Island plant might have had more working backups for its cooling system; Chernobyl might have had a concrete containment for its reactor or a better way of controlling core temperature. But that will always, or at least almost always, be true. Engineering is about making things “safe enough” rather than “absolutely safe”.

How safe is “safe enough” is at least as much a social decision as an engineering decision. But it is an engineering decision in part. For small risks, engineers may well make the final decision. Even concerning the largest risks, engineers will be consulted and their opinion given considerable weight. No decision-maker wants to overrule the engineers on a matter of safety only to have the decision (more or less figuratively) blow up in her face.

Engineers generally evaluate risk by multiplying the harm’s (net) disvalue by the harm’s probability. This method of risk analysis works reasonably well for small harms. The method does not, I think, work at all well for the largest harms—those that, even if highly improbable, would be intolerable if realized—such as destruction of the earth or even the sort of devastation Chernobyl produced. For such intolerable harms, engineers should, I think, adopt something like the following principle of prudence in planning: *If we (society at its rational best) would reject any plausible benefit in exchange for suffering that harm, we (that part of society making the decision) should, all else equal, rule out any design that risks that harm* (however small the probability—so long as it is finite). Since this principle applies when we know both the harm in question and its probability, it is (technically) not a precautionary principle (though its spirit is much the same). It is, in this respect, more like advice frequently given to gamblers betting in games of chance with known odds (“Don’t bet more than you can afford to lose”). Precautionary principles are about dealing with uncertainty. (See, for example, Andorno, 2004.)

The principle I am proposing is only about dealing with known probabilities. Yet it is, or at least should be, an important principle in engineering. Failure is part of engineering. While engineers have a very low tolerance for failure of any kind, even in subsystems that are not

“safety sensitive”, I have yet to hear of any complicated system (even one as simple as a mechanical pencil) for which engineers have not calculated a failure rate (often, to be sure, a tiny failure rate, such as 3.4 defects per million—the famous Six Sigma). No product of engineering is (strictly speaking) “failure proof” (all things considered).

Most, perhaps all, nuclear power plants now in operation seem to have been built in violation of the planning principle suggested above (at least when the calculation of probability takes into account that human beings will operate the plant). The analogy with gambling may not be altogether fair, however. For, we always have the option of doing something much safer, such as going to the theater or buying government bonds. For nuclear energy, our choices today are more difficult. Fossil-fuel plants together (though not individually) threaten us with a world too hot to live in. Hydro-electric dams flood lowlands when they fail and are often not available as an alternative to nuclear power. Failing hydro-electric dams may have killed many more people than nuclear power-plant accidents have (depending on how deaths are calculated). Just one dam failure, that of the dam at Banqiao, China, in 1975, seems to have killed at least 26,000 people directly—and another 145,000 through resulting disease and famine (Wiki, “Banqiao”). Three Mile Island itself is only a hundred miles or so from Johnstown, Pennsylvania, the site of the “Johnstown Flood”, which killed more than 2,200 people, the result of a dam failing in 1889 (Wiki, “Johnston Flood”). In contrast, no one died at Three Mile Island and statistical deaths worldwide to be expected from the radiation that escaped is much smaller.

Nowhere has wind and geothermal met the demand for electricity in an industrial country. And so on. Even with the sort of conservation Japan has undertaken since Fukushima, there is, it seems, still a demand for electricity beyond what is available without some method of generating power that violates the principle of prudence in planning. For the time at least, we may face a choice among dangerous friends. We can only minimize the risk of disaster, not avoid it.

Two features that neither Fukushima nor Katrina share with Three Mile Island and Chernobyl is operator error and normal equipment failure. Equipment did fail at Fukushima and New Orleans—the diesel generators failed at Fukushima as did the screw pumps at New Orleans—but both these failed because of flooding, itself produced by a natural disaster (or, at least, overwhelming external events). Insofar as there were managers or operators involved in the Fukushima or Katrina disaster, they seem to have prevented an even worse outcome.

What all four disasters have in common are failures of engineering design, that is, designs that could have been better. So, for example, the canals in New Orleans

could have been designed with T-walls rather than I-walls; Fukushima could have had a higher breakwater; Chernobyl could have had a better design for its dampening rods; and Three Mile Island could have had a control board that took more account of human factors such as sight lines. And, of course, after these disasters, engineering designs made—or, in the case of Fukushima will make—such improvements. Engineers generally learn from their failures. But such failures are, all else equal, present at every disaster. They do not help us to see what, if anything, is special about Fukushima.

For me, what is special about Fukushima compared with New Orleans is precisely what makes Fukushima like Three Mile Island and Chernobyl. The engineers, and their supporting staff, stayed with the machinery—monitoring, trying to prevent things from going further wrong, and even making repairs.

How many of (what the media called) “workers” at Fukushima were engineers? I have been unable to determine that either from news sources or from contacts in Japan. My visits to nuclear plants in the United States suggest that most of those working at Fukushima would have been engineers (say, 90%)—with the remainder divided about evenly between scientists and technicians. My guess (or, as scientists like to put it, “my hypothesis”) is that most were engineers. I hope someone will find out.

The engineers at Fukushima were not as successful as the engineers at Three Mile Island and Chernobyl. Both those disaster were limited to one reactor. At Fukushima, the disaster spread to four of the six reactors—and might have spread to the other two as well but for the restarting of a diesel generator at Reactor 6 to provide power for cooling the fuel in the holding pools of Reactors 5 and 6. Workers also removed roofing from Reactors 5 and 6 to allow hydrogen to escape, thus preventing explosions similar to those that had damaged the other four units.

This aspect of what happened at Fukushima is a reminder that part of what makes engineering so reliable is that engineers design with the (usually) justified expectation that other engineers will be present to look after what they design. The works of engineering, even of civil or mechanical engineering, do not last long without continual maintenance, including continual adjustments as experience identifies unanticipated problems or unanticipated opportunities for positive improvement. The engineering experiment at Chernobyl, despite its disastrous outcome, was part of normal engineering. The engineers were trying to reduce the risks arising from the backup system’s slow startup. Even nuclear plants that are identical when commissioned, slowly differentiate as they operate, because the engineers managing a plant will continually make improvements. Those engineers should, of course, let engineers at similar plants know about the changes, thus advancing the state of the art, but

other engineers may not be able to make the necessary changes immediately because of budget or schedule, or at all because changes that they have already made bar the improvement in question. Engineers may also find an alternative way to achieve the same end. For these reasons (and perhaps others), nuclear plants, however alike at birth, tend to grow into noticeably different individuals, much as biological plants do.

Some people, especially philosophers, seem to think of those who stayed on at Fukushima—those who, for example, worked in the dark in cold waist-high radioactive water to restart the generators—as engaged in “supererogatory” conduct, that is, as engaged in conduct above and beyond what morality requires. The engineers I have discussed this with seem to view the conduct as heroic but required (supposing the “workers” in question to be engineers). An engineer who left when needed would have acted unprofessionally; he would have failed as an engineer even if he left to save his life or look after his family. Engineering sometimes requires heroism (a significantly higher standard than proposed in Alpern 1983)—or so the engineers I have talked with about this seem to think.

Acknowledgments

This article has benefited from discussion of it at: a workshop for philosophy graduate students at the Technical University-Delft, The Netherlands, May 11, 2011 (“The engineer, public safety, and economic constraints”); a seminar for the Department of Philosophy and Ethics, the Technical University-Eindhoven, The Netherlands, May 13, 2011 (“The Fukushima Nuclear Disaster: Reflections”); a talk for the Department of Philosophy and Religion, University of North Texas, Denton, October 13, 2011(“The Fukushima Nuclear Disaster: Some Issues of Engineering Ethics”); a plenary session of Sixth International Conference on Applied Ethics, Hokkaido University, Sapporo, Japan, October 30, 2011; and the Annual Meeting of the Association for Practical and Professional Ethics, Cincinnati, Ohio, March 3, 2012, as well as from comments of several reviewers for this journal.

References

Alpern, Ken (1983) “Moral Responsibility for Engineers”, *Business and Professional Ethics Journal* 2 (Winter): 39-48.

Andorno, Roberto (2004) “The Precautionary Principle: A New Legal Standard for a Technological Age”, *Journal of International Biotechnology Law* 1: 11–19.

Davis, Michael (2012) “Imaginary Cases in Ethics: A Critique”, *International Journal of Applied Philosophy* (Spring), forthcoming.

Davis, Michael (2007) “Perils of Katrina: Using that Current Event to Teach Engineering Ethics”, *IEEE Technology and Society Magazine* 26 (December): 16-22.

Davis, Michael (2002) “Do the Professional Ethics of Chemists and Engineers Differ?” *HYLE* 8 (Spring): 21-34.

Davis, Michael (1998) *Thinking like an Engineer* (Oxford University Press: New York).

Kemeny, John G. (2012) *Report of The President’s Commission on the Accident at Three Mile Island: The Need for Change: The Legacy of TMI (Washington, D.C.: The Commission, October 1979)* <http://www.threemileisland.org/downloads/188.pdf> (accessed April 18).

Lindsay, Robert (2011) “1% Failure Rate For Nuclear Power”, <http://robertlindsay.wordpress.com/2011/03/23/1-failure-rate-for-nuclear-power> (accessed April 19, 2012).

Interagency Performance Evaluation Task Force (2006) *Performance Evaluation of the New Orleans and Southeast Louisiana Hurricane Protection System: Draft Final Report of the Interagency Performance Evaluation Task Force* (1 June), <http://permanent.access.gpo.gov/lps71007/> (accessed December 5, 2011).

Perrow, Charles (1984) *Normal Accidents: Living with High-Risk Technologies* (Basic Books, NY).

Rogovin, Mitchell (1980), *Three Mile Island: A report to the Commissioners and to the Public, Volume I. Nuclear Regulatory Commission, Special Inquiry Group*, <http://www.threemileisland.org/downloads/354.pdf> (accessed April 18, 2012).

Wikipedia, “Chernobyl disaster” http://en.wikipedia.org/wiki/Chernobyl_disaster (accessed December 16, 2011).

Wiki, “Banqiao Dam” http://en.wikipedia.org/wiki/Banqiao_Dam (accessed May 5, 2012).

Wiki, “Fukushima Daiichi Nuclear Disaster” http://en.wikipedia.org/wiki/Fukushima_Daiichi_nuclear_disaster (accessed April 19, 2012).

Wiki, “Fukushima Daiichi Nuclear Power Plant” http://en.wikipedia.org/wiki/Fukushima_I_Nuclear_Power_Plant (accessed April 19, 2012)

Wiki, “Fukushima II Nuclear Power Plant” http://en.wikipedia.org/wiki/Fukushima_II_Nuclear_Power_Plant (accessed April 18, 2012).

Wiki, “Johnstown Flood” http://en.wikipedia.org/wiki/Johnstown_Flood (accessed May 5, 2012)

Wiki, “Three Mile Island Accident” http://en.wikipedia.org/wiki/Three_Mile_Island_accident (accessed April 22, 2011).

A Pluralist Ethical Decision-making Procedure

Valentin Muresan

University of Bucharest, Romania

Abstract

This paper claims that the use of *several* moral tests to assess the ethics of a new policy is unavoidable. All the efforts to make credible a methodological monism – by critical or reductionist strategies – have been unsuccessful; moreover, it must be acknowledged that even if there were a single test, when applied successively or by different people it would usually give divergent results. The main aim of the paper is to propose a pluralist procedure of ethical decision-making, using a set of proper ethical tests (such as utilitarian, Kantian, Christian, principlist and casuist) in the frame of an “ethical Delphi” procedure intended to make convergent the supposed variety of verdicts. This pluralist testing process, made by moral experts, is only a fraction of a more complex procedure intended to deliver social sanction for a new moral policy. This longer procedure also shows that the adoption of a new moral policy, rule or law is not only a question of passing a strict ethical test, but also a political (i.e. multi-criteria) decision. In general, the adoption of a new moral rule does not rely solely on an ethical test, but is essentially the outcome of a complicated social agreement. That is why in academic applications of the usual moral tests we do not *take* a moral decision on a new case, but merely *simulate* it.

Key words: ethical decision-making, ethical pluralism, ethical Delphi, pluralist model

Decisions about the moral value of an action, rule or public policy cannot be reduced to a verdict resulting from the application of traditional tests based on the major ethical theories, despite the fact that handbooks still unanimously support this view. The history of ethical test results is more one of surprises than one of predictability. You would expect, for instance, that people who adopt the same moral doctrine do this in order to approach issues in the same way, including the moral assessment of actions. We all believe that this is the main reason it is useful to embrace the same moral creed. Therefore it seems strange to find that several members of the Romanian Parliament, all active supporters of Christian morality, assessed the legalization of prostitution in opposing ways. On the other hand, it is also strange that two people who adopt different ethical theories – precisely because they offer distinct explanations of moral phenomena – can frequently

assess actions in the same manner. When a utilitarian and a Kantian – or a follower of Christian ethics and one of Muslim ethics – debate issues, it is somehow surprising to see them judging situations in the same way *in most cases*, despite the fact they declare themselves to be supporters of *opposing* ethical beliefs. Are these beliefs really opposing? In general, it appears that use of tests based on distinct or even opposing theories, such as utilitarianism and Kantianism, *can* result in different verdicts, but in most cases it results in convergent ones (Kantian and utilitarian moral duties are, ultimately, the same). On the other hand, if we dogmatically adopt a single theory and apply the *same* test repeatedly to the same action we usually get similar results, but some divergent ones also appear (see the cases of divergent utilitarian assessments of the same case given as examples in the textbooks).