Garrick Foundation
for the Advancement of the Risk Sciences
$R = \{<S_i, L_i, X_i>\}_c$

ROOTS OF QUANTITATIVE RISK ASSESSMENT[1]

B.J. Garrick
2008

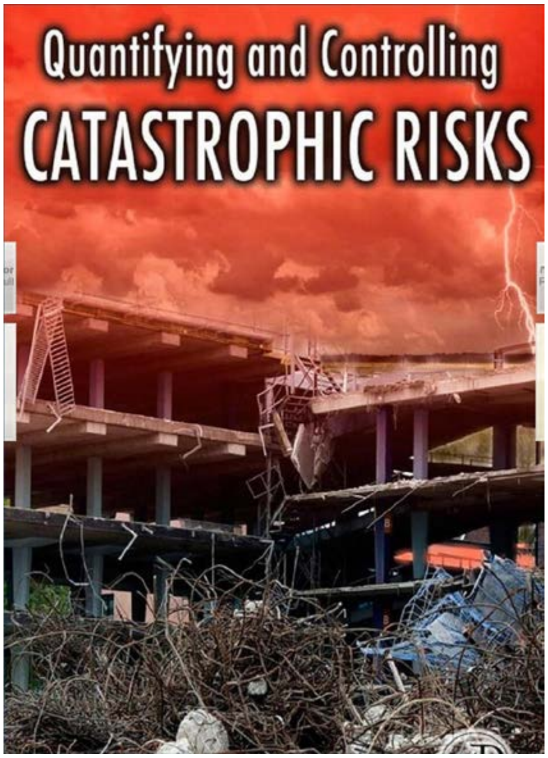

The general framework of quantitative risk assessment presented in this book is based on the "set of triplets" definition of risk. The *triplet* refers to scenarios, consequences, and likelihoods. In Chapters 1 and 2 we are explicit on the meaning given to "likelihood." The concept is illustrated in the case studies (Chapters 3 through 6). From a methodology standpoint, calculating likelihoods, or probabilities, is the central issue in "quantifying" risk. From an application perspective, the most important exercise is the structuring of the scenarios. Thus, when we consider the roots of quantitative risk assessment most of the attention is on calculating likelihoods and structuring scenarios. Calculating likelihoods primarily evolved in the fields of mathematics and mathematical physics and has a history of several hundred years. On the other hand, the formal process of structuring scenarios is a much more recent discipline developed primarily by engineers involved in designing and analyzing complex systems evolving from 20th century technology. For example, the discipline of reliability analysis and engineering has made a significant contribution to the development of integrated models of engineered systems, including the development of a variety of graphical methods for displaying interdependencies of components, subsystems, and systems.

A.1 Calculating Likelihoods

Likelihood is interpreted in a Bayesian sense. Probability, as defined in Chapters 1 and 2, is the credibility of a hypothesis based on the totality of the supporting evidence. The definition of probability has been argued for well over 200 years among the so-called subjectivists and the frequentists, sometimes referred to as the Bayesians and the frequentists, or the Bayesians and the classical statisticians. This debate is not only between Bayesians and classical statisticians, but there are different interpretations of probability within each of these groups. For example, some groups interpret Bayesian probability as a "degree of belief." Such an interpretation has a connotation of "faith" and is not the interpretation of Bayesian probability used in this book, nor is it the interpretation of E.T. Jaynes, a noted contemporary on Bayesian inference. It is not a matter of "belief" or "faith," but rather "the credibility of the hypothesis based on the totality of available evidence." Thus, the probabilities are objective in the sense that they completely depend on the "totality of available evidence." Having established our interpretation of probability, it is important to back up and make a few observations about how the principles behind probability and risk assessment evolved.

Probability theory, the foundation of contemporary risk assessment, was principally developed during the 100-year period between the mid-1600s and 1700s. Of course, there were many contributing events prior to the 1600s. For example, probability involves numbers, ranges of numbers, curves, and families of curves. Thus, the Arabic numbering system known to western civilizations for some 700 to 800 years has to be a profound event in the history of all science, including probability and risk assessment. But it was during the Renaissance period (14th to 17th century), a time of the separation of new thoughts from the constraints of medieval cultures, that the contemporary thoughts about probability and risk were essentially formulated.

Cardano and Galileo made important contributions in the 1500s on how to express probabilities and frequencies of past events, Pascal contributed to concepts of decision theory and statistical inference in the mid-1600s, and Fermet and de Méré made major contributions to the theory of numbers about the same time. Other major contributors during the 17th century were Christen Huygens, who published a popular textbook on probability theory; Gottfried Wilhelm von Leibniz, who suggested applying probability methods to legal problems; and members of a Paris monastery named Port Royal. The Port Royal group produced a pioneering work of philosophy and, probably, the first definition of risk: "Fear of harm ought to be proportional not merely to the gravity of the harm, but also to the probability of the event." Jacob

[1] *This is a replication of Appendix A of the book "Quantifying and Controlling Catastrophic Risks," Elsevier, Amsterdam, 2008, by B.J. Garrick, et al.*

Bernoulli produced the Law of Large Numbers and methods of statistical sampling forming the basis of many methods of product testing and quality control. Abraham de Moivre developed the concept of the normal distribution and standard deviation in the early 1700s.

Many modern day probability practitioners, decision analysts, and risk assessors consider Thomas Bayes, an English minister, the real father of contemporary risk assessment. In the mid-1700s, he developed a theorem rooted in fundamental logic for combining old information with new information for the assignment of probabilities. Bayesian inference reduces to the simple product and sum rules of probability theory developed by Bernoulli and Laplace. Bayes theorem provided the foundation for a unified theory of probability not bound to such properties as “randomness” and “large numbers.” The doors were opened for employing probability to address problems involving limited information. Bayes theorem, followed by the publication in 1812 of *Théorie analytique des probabilities* by the French mathematician Marquis Pierre Simon de Laplace, provided the primary basis of contemporary probability theory. Diverse problems, such as gambling strategies, military strategies, determining mortality rates, and debating the existence of God, were the subjects of early analytical explorations and precursors to the new science of risk assessment.

Among the 20th and 21st century scholars who contributed to probability and risk assessment as advocated in this book are Harold Jeffreys (1957), R.T. Cox (1946), C.E. Shannon (1948), George Pólya (1954), Howard Raiffa (1996), and E.T. Jaynes (2003). Probability theory in the context of Jaynes’ extended logic includes as special cases all the results of the conventional “random variable” theory, and it extends the applications to useful solution of many problems previously considered to be outside the realm of probability theory. The goal of such investigators as Pólya and Jaynes was to formulate a probability theory that “could be used for general problems of scientific inference, almost all of which arises out of incomplete information rather than ‘randomness’.” This interpretation of probability is the nugget that makes it possible to perform meaningful risk assessments of complex systems about which there is little information on their threat environment and vulnerabilities.

The widespread, formal application of risk assessment to critical infrastructure began in earnest in the late 1900s. Applications in the insurance and financial fields were more statistical (actuarial) than probabilistic, more experience-based than inferential, more qualitative than quantitative. Only when societies began depending more on technological systems involving large inventories of hazardous materials did investigators begin to look for more scientifically based ways to assess risks. The particular need was for a method of assessing the likelihood of catastrophic events that could do great harm to public health and the environment.

A.2 Structuring the Scenarios

In many respects the heart of a quantitative risk assessment is the scenarios that arise from trying to answer the risk triplet question, what can go wrong. The scenarios are what tie the physical aspects of the risk assessment to the analytical process. A scenario is defined as a sequence of events, starting with an event known as the initiating event (an event that upsets an otherwise normally operating system), or an initial condition, and then proceeds through a series of events until the system either corrects itself or the scenario of events is terminated at a damaged, degradated, or destroyed state. Thus, structuring the scenarios is a representation of the logic of how a system responds to different types of threats, that is, to different initiating events or initial conditions. Of course, the goal of the analyst is to define the initiating event or initial condition set such that it is complete, or at least complete in the practical sense that all of the important initiators have been identified. The practice of developing scenarios has evolved into a general theory of structuring scenarios.

The theory of structuring scenarios has its primary roots in the field of reliability engineering and analysis developed initially by the United Kingdom and Germany in the 1930s and 40s. The United Kingdom introduced the concept of mean time to failure for aircraft in the 1930s (Green, 1972). They used such information to infer reliability criteria for aircraft and the proposing of maximum permissible failure-rates as a basis for establishing levels of safety for aircraft. For example, in the 1940s performance requirements were being given for aircraft in terms of accident rates that should not exceed, on average, one per 100,000 hours of flying time.

The Germans applied such principles of reliability analysis as the product rule to the development of the V1 missile during World War II to solve serious reliability problems with that weapon system. The product rule has to do with the reliability of components in series and accounting for differences in reliability of individual components. Until the methods of reliability were employed, the design philosophy of the V1 missile was based on the “weak link” theory—a chain cannot be made stronger than its weakest link. Eventually, it was realized that a large number of fairly strong “links” could be more unreliable than a single “weak link” if reliance is being placed on them all. Going from a design philosophy based on the weak link theory to basing it on reliability theory resulted in a vast improvement in the V1 missile reliability.
In the years following World War II the United States adopted the analysis methods of the United Kingdom and Germany and greatly expanded the use of practical tools to improve the reliability of missiles (Kumamoto and Henley, 1996;

Neufeld, 1990). Activities where the methods of reliability analysis and engineering made a major impact were defense systems, commercial aerospace, electronics, chemical plants, power plants, and electrical distribution systems. Reliability analysis became a cornerstone discipline of the evolving field of systems engineering because of its ability to link interacting systems.

Reliability analysis involved many of the steps of risk assessment: (1) defining of the system and its success state, (2) characterization of the hazards involved, and (3) evaluation of the probability of failure of subsystems and components. Reliability analysis and engineering evolved rapidly in such fields as defense and aerospace, making major contributions to the performance of complex systems. Examples are methods for importance ranking of components and subsystems of total systems, the determination of the life of components subject to such phenomena as fatigue, the quantification of the impact of repair, and the development of solid technical foundations for preventive maintenance programs to optimize system performance. Reliability analysis and risk assessment have much in common. The emphasis is what's different. In risk assessment the emphasis is on "what can go wrong" and in reliability analysis the emphasis is on what to do to make the system run the way it is designed. The output of a reliability analysis is *reliability* and what's driving it, while the output of a risk assessment is the *risk* of something bad happening and what contributes to it. Both employ many of the same analytical tools; the perspectives are just different and thus each has a set of analytical tools unique to that perspective. For example, risk assessment gives a great deal of emphasis to quantifying uncertainties as risk assessment applications are very often with respect to rare events of high consequence on which there is very limited data. On the other hand, reliability analysis focuses much more on factors contributing to good performance than on risk assessment.

One of the most important contributions that the reliability sciences have made to systems engineering in general and risk assessment in particular are transparent methods for graphically representing complex systems that can be transformed into analytical models. For example, reliability analysis used block diagrams to describe how components in a large system were connected. From these block diagrams, Watson at Bell Laboratories developed the fault-tree technique, which he applied to the Minuteman Missile launch control system, and which Boeing later adopted and also computerized. These diagrams in combination with the tools of switching algebra and probability theory have provided a powerful tool for displaying and quantifying the "fault paths" of systems, subsystems, components, human actions, procedures, etc.

The advantage of these techniques in the field of safety analysis, as opposed to just reliability, began to be recognized in the 1960s. F.R. Farmer of the United Kingdom proposed a new approach to nuclear power plant safety based on the reliability of consequence-limiting equipment (Farmer, 1964). Holmes and Narver, Inc., a U.S. engineering firm under contract to the then U.S. Atomic Energy Commission performed a series of studies on nuclear reactor safety and reliability. The final report in the series advocated, with examples, the need for much greater use of advanced systems-engineering methods of modeling the reliability of safety systems. The authors made explicit reference to the use of logic tools, such as fault-tree methodology, which has its roots in "switching theory" developed by the telecommunications field (Holmes and Narver, Inc., 1967). At about the same time, a Ph.D. thesis was published that proposed a methodology for probabilistic, integrated systems analysis for analyzing the safety of nuclear power plants (Garrick, 1968).

The breakthrough in quantitative risk assessment (or probabilistic risk assessment as it is generally labeled in the nuclear field) of technological systems came in 1975 with the publication of the *Reactor Safety Study* by the U.S. Atomic Energy Commission under the direction of Professor N.C. Rasmussen of the Massachusetts Institute of Technology (USNRC, 1975). This project, which took 3 years to complete, marked a turning point in the way people think about the safety of complex facilities and systems. The *Reactor Safety Study* provided a basis for a wide range of applications for risk assessment, not only for nuclear power plants and other technological systems (e.g., chemical and petroleum facilities, transportation systems, and defense systems), but also for environmental protection, health care, and food safety.

Besides fine-tuning fault tree analysis for safety applications, the Reactor Safety Study introduced another extremely important graphic tool to facilitate the structuring of scenarios, the event tree. Fault trees and event trees in combination provided a critically important one-two punch in the theory of structuring scenarios. An *event tree* starts with an initiating event and proceeds to identify succeeding events, including branches that eventually terminate into possibly undesirable consequences. An event tree, therefore, is a cause-and-effect representation of logic.

A fault tree starts with the end-state or undesired consequence and attempts to determine all of the contributing system states. Therefore, fault trees are effect-and-cause representations of logic. An event tree is developed by inductive reasoning, while a fault tree is based on deductive reasoning. A key difference in the two representations is that a fault tree is only in "failure space," and an event tree includes both "failure and success space." The choice between the two is a matter of circumstances and preference, and they are often used in combination; the event tree provides the basic scenario space of events and branch points, and the fault tree is used to quantify the "split fractions" at the branch points.

The *Reactor Safety Study* inspired many first-of-a-kind risk assessments in the commercial nuclear power industry that led to major advancements in the application of quantitative risk assessment. One important example was the probabilistic risk assessments of the Zion and Indian Point nuclear power plants sponsored by the owners and operators of the plants. New methods were introduced in those assessments that have become standards of many quantitative risk assessment applications (PLG, et al., 1981, 1982). The methods included the treatment of uncertainty, a framework of risk assessment embedded in the set of triplets definition of risk (Kaplan and Garrick, 1981), common-cause failure analysis, importance ranking of risk contributors, models for calculating source terms, and improved dispersion models for calculating off-site health effects.

These and other studies have evolved into a contemporary theory of structuring scenarios that is part science and part art. Elements of the theory are a set of principles having to do with issues of completeness and the general structure of scenarios. More details are covered in Chapter 2.

A.3 Steps That Have Evolved for Integrated Quantitative Risk Assessment

Although the scope, depth, and applications of quantitative risk assessments vary widely, they all follow the same steps:

1. Define the system being analyzed in terms of what constitutes normal operation to serve as a baseline reference point.
2. Identify and characterize the sources of danger, that is, the hazards (e.g., stored energy, toxic substances, hazardous materials, acts of nature, sabotage, terrorism, equipment failure, combination of each, etc.).
3. Develop “what can go wrong” scenarios to establish levels of damage and consequences while identifying points of vulnerability.
4. Quantify the likelihoods of the different scenarios and their attendant levels of damage based on the totality of relevant evidence available.
5. Assemble the scenarios according to damage levels, and cast the results into the appropriate risk curves and risk priorities.
6. Interpret the results to guide the risk management process.

These steps provide the answers to the three fundamental questions of risk (the “triplet definition”): what can go wrong, how likely it is to go wrong, and what the consequences will be.

Risk assessments are routinely used in many settings, including the electric nuclear power industry, the chemical and petroleum industries, defense industries, the aerospace industry, food sciences, and health sciences. Industries that are increasingly using formal, quantitative methods of safety analysis include marine transportation and offshore systems, pipelines, motor vehicle, and recreational systems. The space program has stepped up its use of quantitative risk assessment since the *Challenger* accident (NASA, 2002). Other less publicized applications include the risk management program used by the U.S. Army for the disposal of chemical agents and munitions (Boyd and St. Pierre, 2001).

The government agencies most involved in using risk assessments are the U.S. Nuclear Regulatory Commission and the U.S. Environmental Protection Agency. Other agencies becoming active users of risk assessment methods are the U.S. Department of Energy, U.S. Department of Agriculture, U.S. Department of Defense, U.S. Food and Drug Administration, the National Aeronautics and Space Administration, and the U.S. Department of Transportation. The most active practitioners in the private sector are the nuclear, chemical, and petroleum industries, although the scopes of application vary widely—the nuclear industry being the most consistent user of complex and sophisticated quantitative methods.

Risk assessment has many buzzwords (e.g., Monte Carlo analyses, influence diagrams, multiple attributes, common-cause failures, realizations, minimum cut sets, sensitivity analyses, fault trees, event trees, etc.), but the basic principles are few. The principles focus on the development of scenarios describing how the system under study is supposed to work and scenarios indicating how the system can be made to fail, catastrophically or otherwise. The likelihood of events in the scenario must be linked to the supporting evidence. Events are propagated to an end state that terminates the scenario (i.e., the consequence). Other principles may be applied to aggregate the various end states into the desired set of consequences.

The results of risk assessments are easy to interpret, including corrective actions having the biggest payoff in terms of risk reduction. Although the literature suggests many different risk assessment methodologies, in fact the differences are primarily in scope, application, boundary conditions, the degree of quantification, figures-of-merit, and quality. Like many other scientifically based methodologies, quantitative risk assessment is founded on relatively few basic principles.

A.4 Application to Nuclear Power: A Success Story

A.4.1Why Risk Assessment

The simple answer to “why risk assessment” for nuclear power plants is that nations and the world have to make decisions about the best energy mix for the future of planet Earth. Risk to people and the environment is a fundamental attribute of societal decision-making. But there is more to it than just decision-making. Early in the development of nuclear power, it became clear that the large inventories of radiation in the nuclear reactors contemplated for generating electricity and the stigma of the dangers of the fission process carried over from the Hiroshima and Nagasaki atomic bombs required a level of safety analysis beyond standard practices. The nuclear power industry was forced to seek new methods of safety analysis of nuclear electric power plants to overcome the “fear anything nuclear” syndrome that prevailed in the minds of some members of the public. New methods were needed to provide answers to the questions, what can go wrong with a nuclear power plant, how likely is it, and what are the consequences. The traditional methods of safety analysis, while somewhat effective in answering questions about what can go wrong and what are the consequences, profoundly failed to adequately answer the question having to do with the likelihood of accidents. The likelihood question held the key for being able to quantify nuclear electric power plant risk. In short, for society to have access to nuclear energy systems that have the potential to provide safe, reliable, and relatively inexpensive electric power, the industry was forced to come up with a more convincing safety case than was possible with past methods of analysis.

The nuclear electric power industry, stimulated by the Reactor Safety Study and the early industry studies on the Zion and Indian Point nuclear power plants, has been the leader in the development and widespread use of quantitative risk assessment (QRA). The U.S. nuclear electric power industry gave birth to the term “probabilistic risk assessment” (PRA). The term “probabilistic safety assessment” (PSA) is sometimes used in the international nuclear community as equivalent. The label that appears to be best received across different industries is quantitative risk assessment. In this discussion, the terms quantitative risk assessment, probabilistic risk assessment, and just plain “risk assessment” are used interchangeably.

Risk assessment has survived and flourished in the U.S. nuclear electric power industry because it became an exceptional tool to make better decisions. QRA was able to satisfy the desire of nuclear plant owners to have a decision tool that quantitatively allowed the evaluation of various options that had multiple input variables. The most important variables to the nuclear plant owners were cost, generation, and risk (public health, worker health, and economic). While QRA started out as a tool to address the public health risk, it facilitated evaluating an entire spectrum of variables. The industry’s recovery from the Three Mile Island Unit 2 accident in 1979 was greatly aided by the use of quantitative risk assessment because of the ability to better focus on the real safety issues. In fact, the industry has enjoyed an impeccable safety record since embracing contemporary methods of quantitative risk assessment. And safety is not the only benefit that has resulted from the widespread use of risk assessment in the nuclear power industry. Risk assessment provides the ability for plant personnel to balance cost, generation, and risk. While there is no U.S. Nuclear Regulatory Commission (USNRC) requirement for an existing nuclear power plant to maintain a risk assessment, the plants do so following general industry guidelines. The USNRC does have a requirement for a prospective licensee to submit a PRA with the application for any proposed new nuclear electric power unit in the U.S.

A.4.2Legacy of Nuclear Safety

Today, nuclear power plant safety analysis employs the most advanced methods available for assessing the health and safety of the public. The only significant impact on public health risk comes from the release of fission products from the reactor core following accidents during power operation. The release of fission products from the reactor core is heavily influenced by operator actions both before and after the initiating event that leads to the accident. The PRA includes potential failures of equipment and operators, both before and after the initiating event. Each PRA in the United States is specific for each nuclear power plant because of the differences between the plants.

Many of the methods used for nuclear power plants have been adopted by such high tech industries as space flight, defense systems, chemical plants, refineries, offshore platforms, and transportation systems. While the probabilistic concepts currently spearhead the level of sophistication of the analyses, there are basic tenets and themes that have guided the safety management of nuclear electric power plants from the beginning. One of the most fundamental of these basic tenets is the concept of multiple barriers.

Multiple barriers are a concept of providing enough barriers between radiation and the environment to provide assurance that the likelihood of simultaneous breach of all barriers is remote. Examples of barriers in a nuclear power plant are high containment capacity fuel with cladding, an isolated reactor coolant system, primary reactor building containment, secondary building containment, and exclusion distance. Effective defense mechanisms developed for nuclear plant safety include improved operator training methods and symptom based operator procedures. Other defense mechanisms include automatic control systems, single failure criteria (no single failure threatens fuel integrity), and recovery capabilities from equipment malfunctions. QRA provides the ability to determine what risk levels are achieved by each barrier and at what cost. The value of each barrier is placed in the context of the overall risk.

A.4.3Historical Development of Nuclear Power Plant Safety

Nuclear plant safety has two major fronts—the physical system itself and the analysis of the physical system. On the physical system front, improvements in safety design included the advent of secondary containment systems (~1953), the inclusion of backup safety systems known as engineered safety features, especially with respect to emergency core cooling systems and electric power (~late 1950s and early 60s), and the introduction of separate and independent safety trains (~1970s). In the 1980s and 1990s, the nuclear power plants initiated programs for scram reduction based on a complete review and analysis of operating transients. As scrams were reduced, public health risk was also reduced because there were fewer departures from normal steady state operation. Also, in the 1980s and 1990s, each nuclear power plant implemented the concept of "symptom based procedures" for accident control and installed improved simulators for operator training.

On the analysis front, many events took place leading to a greatly improved understanding of the safety of nuclear power plants. It was demonstrated that the consequences of accidents had little meaning without a better understanding of their likelihood. It became clear that it was not enough to do worst case and maximum credible accident analysis. Everyday transients followed by multiple failures of equipment and mistakes by operators were more likely to result in reactor core damage than previously defined so-called design basis accidents.

The need for probabilistic analysis was recognized as early as the mid-1950s. However, detailed investigations of the probability of reactor accidents did not begin until about 1965. The first major reactor safety study to highlight the need for PRA of reactor accidents was WASH-740, "Theoretical Possibilities and Consequences of Major Accidents in Large Nuclear Power Plants" (AEC, 1957). Speculative estimates were made in WASH-740 that a major reactor accident could occur with a frequency of about one chance in a million during the life of a reactor. The report went on to observe that the complexity of the problem of establishing such a probability, in the absence of operating experience, made these estimates subjective and open to considerable error and criticism. While not offering many specifics, this study did stir interest in probabilistic approaches and many studies were soon to follow. These included British and Canadian efforts, probabilistic analyses of military reactors, and several studies sponsored by the then U.S. Atomic Energy Commission. The earlier referred to thesis by Garrick (Garrick, 1968) was written about the same time advocating a probabilistic approach to assessing nuclear power plant safety. But as also noted earlier, it was the Reactor Safety Study (USNRC, 1975) that spearheaded the movement towards the application of probabilistic risk assessment.

By the 1980s the question was no longer "why", but how soon could a QRA be developed for every nuclear power plant in the U.S. That goal has essentially been reached. The benefits of QRA for U.S. nuclear power plants have been demonstrated in terms of the reduction in frequency of core damage events (one reactor core lost in approximately the first 450 reactor years of experience versus zero reactor cores lost in over 2000 actual reactor years of experience since the TMI accident) and improved generation with a reduction in the cost of electricity. The most important benefit is nuclear power plants with reduced public health risk. QRA has not only been effective in calibrating the risk of nuclear power, but has provided better knowledge of the worth of safety systems and allowed the allocation of safety engineering resources to the most important contributors. Effective risk management of nuclear electric power plants in the U.S. has become a reality, not just a goal.

A.4.4Nuclear Power Accident Experience

There have only been two accidents worldwide that have resulted in severe core damage of a nuclear power plant designed to generate electricity. The accidents involved the Three Mile Island Unit 2 plant near Harrisburg, Pennsylvania, in the United States and the Chernobyl Nuclear Power Station in the Ukraine of the former Soviet Union. Both accidents permanently damaged the nuclear reactors involved, but only the Chernobyl accident resulted in known fatalities and injuries. The on-site consequences of the Chernobyl accident were very serious, as an estimated 30 people are believed to have died from acute doses of radiation and some 300 people required hospital treatment for radiation and burn injuries. No off-site fatalities or injuries have yet been attributed to the Chernobyl accident, although the latent effects are yet to be quantified.

It is important to put these two very serious accidents in context with the safety experience of the nuclear power industry. There are approximately 440 nuclear power plants in the world. Nuclear energy is just over 5% of the world primary energy production and about 17% of its electrical production. In the United States there are some 103 nuclear power plants operating providing approximately 20% of the nation's electricity. The worldwide experience base is approaching 10,000 in-service reactor-years of which about 3000 reactor-years is U.S. experience. The experience base is likely beyond 10,000 reactor-years if all types of reactors are included such as research, test, weapons, and propulsion reactors. Some 70% of the nuclear power plant experience worldwide involves light water reactors for which only one accident has occurred, Three Mile Island. This safety record is most impressive. The challenge is to keep it that way.

A.4.4.1 Three Mile Island, Unit 2 (TMI-2) Accident. The TMI-2 nuclear power plant, located near Harrisburg, Pennsylvania, went into commercial operation in December 1978. The plant was designed to generate approximately 800 megawatts of electricity and used a pressurized water reactor supplied by the Babcock and Wilcox Company. The accident occurred on March 28, 1979.

Routine mechanical malfunctions with the plant resulted in an automatic shutdown (“feedwater trip”) of the main feedwater pumps, followed by a trip of the steam turbine and the dumping of steam to the condenser. The loss of heat removal from the primary system resulted in a rise of reactor system pressure and the opening of its power-operated relief valve. This action did not provide sufficient immediate pressure relief, and the control rods were automatically driven into the core to stop the fission process.

These events would have been manageable had it not been for some later problems with such systems as emergency feedwater. Perhaps, the turning point of the accident was that the opened pressure relief valve failed to close and the operators did not recognize such. The result was the initiation of the well studied small loss of coolant accident, known as the small LOCA. The stuck open valve together with some other valve closures that had not been corrected from previous maintenance activities, created a shortage of places to put the decay heat loads of the plant. The response of the plant was the initiation of high pressure emergency cooling. Reactor coolant pump high vibration and concern for pump seal failure resulted in the operators eventually shutting down all of the main reactor coolant pumps and relying on natural circulation in the reactor coolant system. It was during the time that the main reactor coolant pumps were off, some 1 to 3 hours, that the severe damage to the core took place. At about 2 hours and 20 minutes into the accident, the backup valve known as a block valve to the stuck-open relief valve was closed. This action terminated the small LOCA effect of the stuck-open relief valve. While the accident was then under some level of control, it was almost 1 month before complete control was established over the reactor fuel temperature when adequate cooling was provided by natural circulation.

The consequences of the accident were minimal in terms of the threat to public health and safety, but the damage to the reactor was too severe to recover the plant. The accident did confirm the effectiveness of the containment system to contain the fission products escaping from the reactor vessel.

A.4.4.2 Chernobyl Nuclear Power Station Accident. The Chernobyl nuclear power plant involved a 1000-megawatt (electrical) boiling water, graphite-moderated, direct cycle reactor of the former Soviet Union. The Chernobyl accident occurred on April 26, 1986, and was initiated during a test of reactor coolant pump operability from the reactor’s own turbine generators. The purpose of the test was to determine how long the reactor coolant pumps could be operated, using electric power from the reactor’s own turbine generator under the condition of turbine coast down and no steam supply from the reactor. However, the experimenters wanted a continuous steam supply so they decided to conduct the experiment with the reactor running—a serious mistake. The test resulted in a coolant flow reduction in the core and extensive boiling. Because of the inherent properties of this particular reactor design (on boiling, the fission chain reaction increases, rather than decreases as in U.S. plants), a nuclear transient occurred that could not be counteracted by any control system. The result was a power excursion that caused the fuel to overheat, melt, and disintegrate. Fuel fragments were ejected into the coolant, causing steam explosions and rupturing fuel channels with such force that the cover of the reactor was blown off.

This accident resulted in approximately 30 immediate fatalities from acute doses of radiation and the treatment of some 300 people for radiation burn injuries. The off-site consequences are still under investigation. Latent effects are expected, but they have not been quantified.

In summary, nuclear power suffered a severe setback from both of the above accidents, although public support for nuclear power was already beginning to decline before these accidents occurred. Some nuclear plants under construction were cancelled and no new U.S. nuclear plants were ordered between 1979 and 2007 [2]. The fact that the TMI-2 accident did not result in any radiation injuries or fatalities and the Chernobyl reactor type is no longer in the mix of viable power reactors has not removed the fear that some segments of the public have of nuclear power. However, the superior performance and safety record in the U.S. since these two accidents has allowed the USNRC to approve power upgrades and license extensions for numerous U.S. nuclear power plants.

[2] *Renewed interest in nuclear power has resulted in the first order in 2007 of a new nuclear power plant in the U.S. with high expectations for several orders in 2008 and 2009.*

A.5 An Example of Nuclear Power Plant Quantitative Risk Assessment Results

It is appropriate to provide an example of risk assessment results deriving directly from the full application of the methodology of Chapter 2. In particular, the focus on this example is on results only as opposed to the details of how they came about. The emphasis of the limited scope case studies of Chapters 3 through 6 is on the process of quantitative risk assessment, not necessarily the results.

The best application examples are the many comprehensive risk assessments prepared for nuclear power plants worldwide since the 1970s. The example discussed here takes the form of some actual results of such risk assessments (Garrick, 1989). We have chosen to withhold naming specific plants as changes have been made in the plants to improve safety and the risk results no longer apply, but this does not take away the value of the analysis. In general the risks have become less because of the availability of more operating experience and improved analyses and systems.

The example is given in terms of (1) the bottom line results, (2) importance ranking of contributors to risk, and (3) the use of risk assessment as a design and risk management tool. A consequence of such results is the answer to the important risk management questions. “What in fact is the risk, including the uncertainties involved?” “What scenarios, operator actions, and system failures are driving the risk?” “How should risk assessment be used during design to guide the design of the plant, especially with respect to accident mitigating systems?” “What corrective actions will have the greatest return for reducing the risk?”

A.5.1 Bottom Line Results

Figure A.1 contains the risk curves from a comprehensive risk assessment performed on a U.S. nuclear power plant during the 1980s.

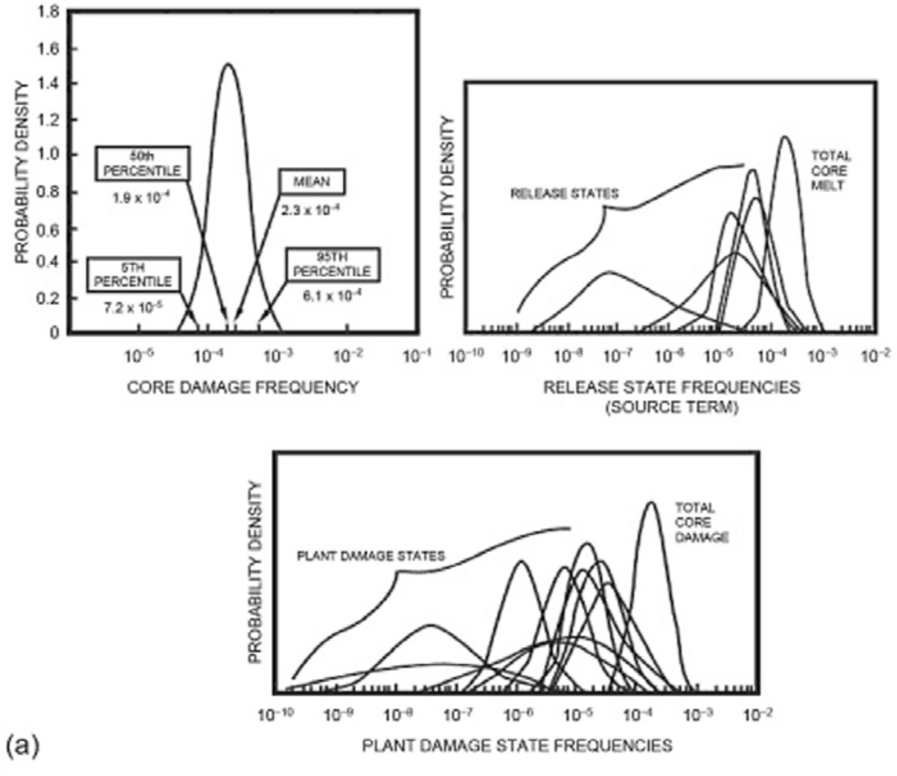

Figure A.1(a) Nuclear power plant risk assessment results.

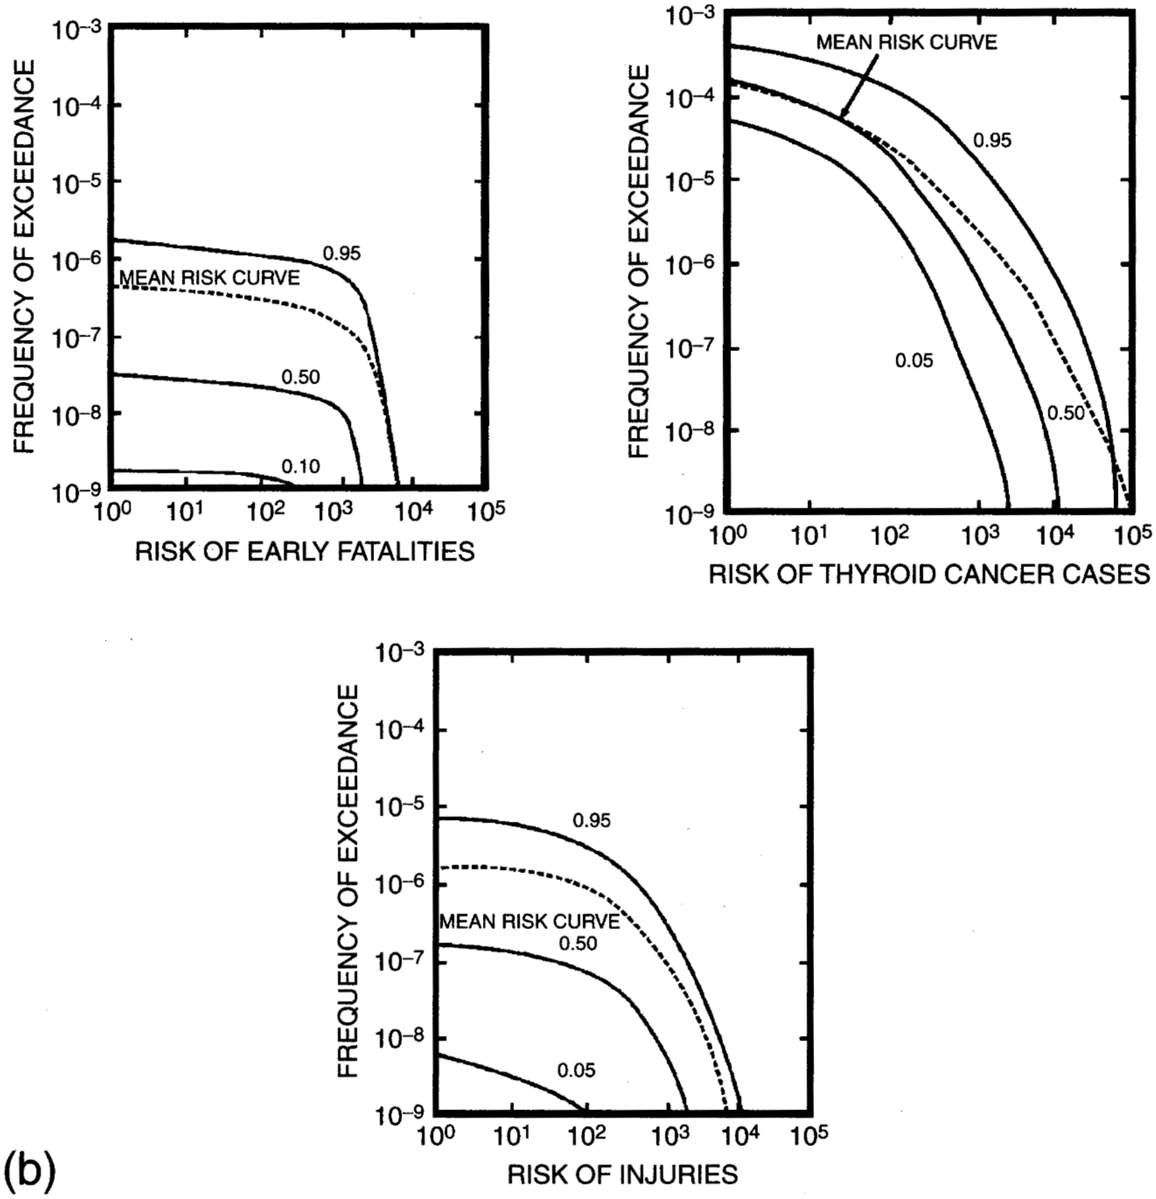

Figure A.1(b) Nuclear power plant risk assessment results.

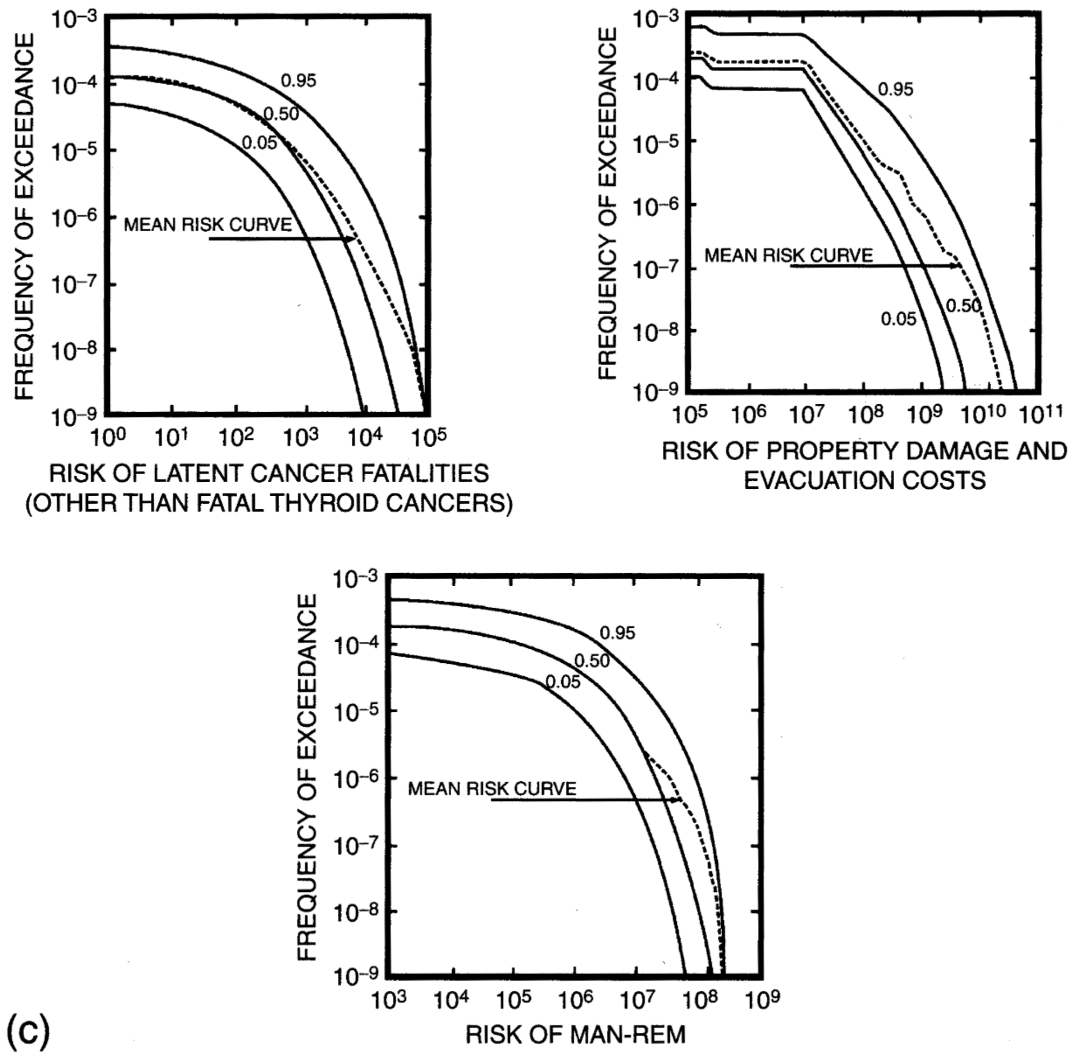

Figure A.1(c) Nuclear power plant risk assessment results.

Figure A.1 represents the integration and assembly of an extensive amount of modeling and analysis covering thousands of pages. The principal elements are initiating events, scenarios, consequences, and likelihoods. The data processing is Bayesian based and involves searches on numerous national databases on human, equipment, and system performance. The results have the forms of Figures 2.9 and 2.10 of Chapter 2. A first impression is that to tell a reasonably complete risk story requires much more than just a number. As a start, it requires a probability curve to communicate the uncertainty in the number. But even a probability curve is not enough to tell the risk story. As can be seen, the risk story told by Figure A.1 involves curves, families of curves, and different representations of families of curves. To make more clearly just what the risk story is, each of the subfigures of Figure A.1 is briefly discussed.

A.5.1.1 Core Damage Frequency. Core damage frequency is currently the most often used measure of risk in contemporary nuclear power plant risk assessments. Simply put, core damage means that the core coolant is not removing heat as fast as it is being generated, and the core begins to degrade through overheating. Core damage includes a range of degrees of damage to the cladding and fuel elements with the most severe being core melt through the reactor vessel. A partially or totally uncovered core as a result of a loss of coolant accident would be an example of a mechanism for creating a core damage event. Figure A.1(a) is a *probability of frequency* (POF) curve for a risk assessment of a specific U.S. nuclear power plant performed in the 1980s timeframe using the methodology of Chapter 2. This curve tells us that the mean frequency of core damage for this particular plant is about once in some 4300 years. It also tells us that there is uncertainty in the core damage frequency. In particular, it tells us that the core damage frequency is uncertain by a factor of approximately 8.5 between the 5th and 95th percentile. As important as it is to know the core damage frequency and its uncertainty, more information is needed before a basis exists to recommend actions that might reduce that frequency. We begin that process by probing into what kind of damage and radionuclide release might take place as a result of a nuclear power plant accident.

A.5.1.2 Plant Damage State Frequencies. Plant damage states describe the different possible states (conditions) of the reactor core at the time of reactor vessel failure. Figure A.1(a) shows in probability of frequency format the various plant damage states that were defined for this particular nuclear power plant. Each plant damage state is characterized by such conditions as timing, pressure, temperature, available coolant, and the status of containment systems. For this risk assessment, eight plant damage states were defined. An example of a plant damage state is an overpressure event with failure of pressure suppression systems. The ninth curve in the figure showing the plant damage state frequencies is the core damage frequency curve. A comprehensive risk assessment of a nuclear power plant consists of three models, a plant model, a containment model, and a site model. The plant damage states are the output of the plant model and the input to the containment model as they represent the threats to the containment system.

A.5.1.3 Release State Frequencies. The release state frequencies, also referred to as release categories, define the source term for the site model and in particular are representations of types, quantities, timings, and elevations of radioisotopes released to the atmosphere. Figure A.1(a) indicates that five release states were defined for this particular assessment. Release states are determined by the state of the containment at the time of core damage, the availability of containment engineered safety features, and the availability of filtration and other mechanisms for radionuclide removal. The release states, or release categories, are the output of the containment model and the input to the site model.

A.5.1.4 Risk Curves for Various Risk Measures. Figure A.1(b) and A.1(c) answer the question of what is the risk in terms of six different risk measures. The risk measures are early fatalities, injuries, cancer cases, latent fatalities, radiation dose, and property damage. These curves known as “frequency of exceedance” curves or “complementary cumulative distribution functions” are becoming increasingly known as “risk curves.” They too are in the “probability of frequency” format, but with the additional feature of consequence as a parameter. Probability is the parameter of the model for displaying uncertainty and is represented by a family of curves, in this case the 5th, 50th, and 95th percentiles. Also shown is the “mean” risk curve. Some of the important observations about the curves are (1) their comprehensiveness in terms of the quantification of a variety of risk measures, (2) the extremely low levels of risk involved when compared with almost any other natural or technological risk, and (3) a clear display of the uncertainties involved.

A.5.2Importance Ranking of Contributors to Risk

While the results such as illustrated in Figure A.1 are extremely important in answering the basic question of what is the risk and what are the uncertainties involved, they only scratch the surface of what is learned from a comprehensive quantitative risk assessment. In the process of evolving to the bottom line results of Figure A.1, information is developed to answer many more questions such as what is contributing to the risk and what actions can be taken to reduce the risk and assure that the risk is reasonably managed. Also, there is the issue of how to use risk assessment during the design process of a nuclear power plant, or any facility, to evolve to a design that is balanced in terms of actions to control the risk. That facet is covered in Section A.5.3.

Figure A.2 is taken from an actual probabilistic risk assessment of a U.S. nuclear power plant. It ranks the scenarios contributing to the different risk measures. The scenario descriptions are greatly abbreviated to simplify the table, but the most important events in the scenarios are identified.

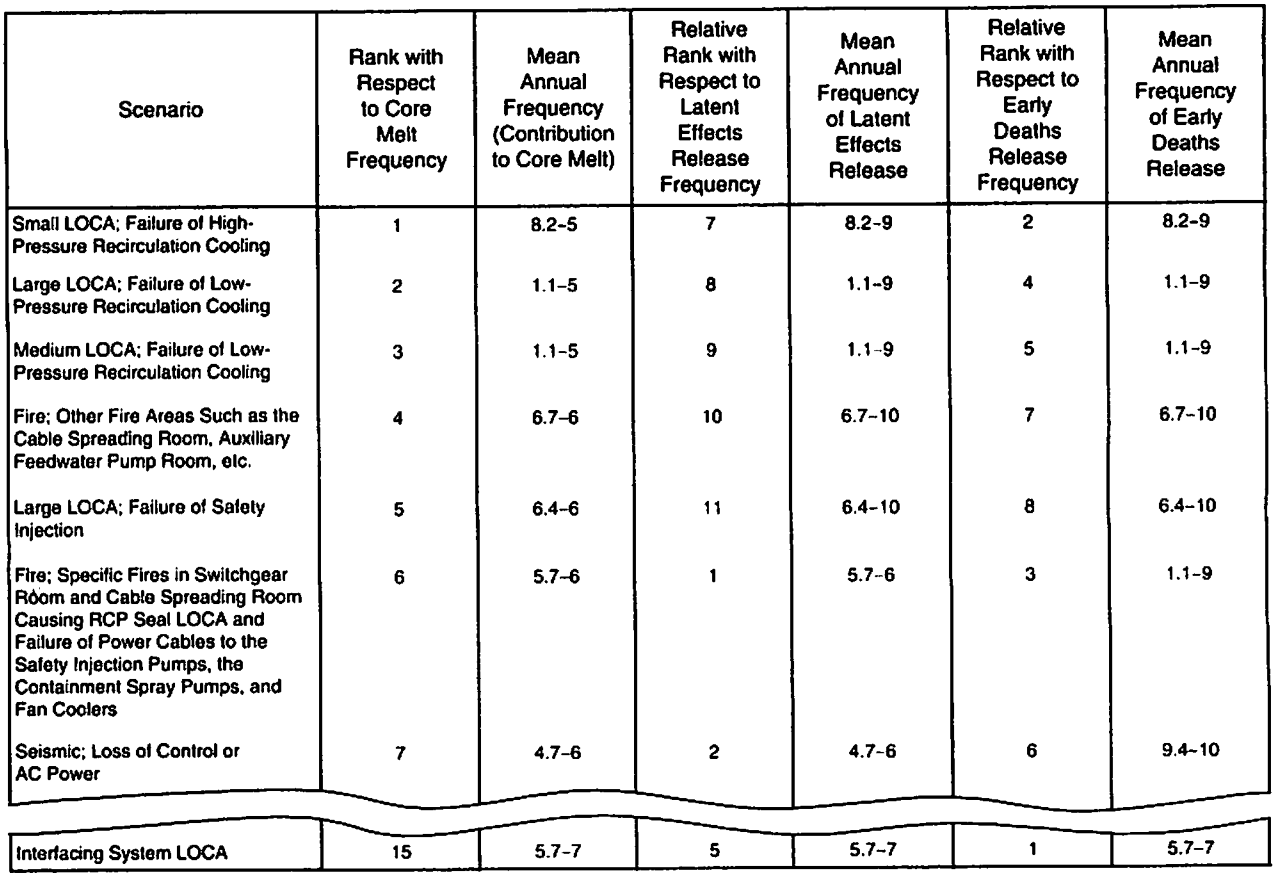

Scenario	Rank with Respect to Core Melt Frequency	Mean Annual Frequency (Contribution to Core Melt)	Relative Rank with Respect to Latent Effects Release Frequency	Mean Annual Frequency of Latent Effects Release	Relative Rank with Respect to Early Deaths Release Frequency	Mean Annual Frequency of Early Deaths Release
Small LOCA; Failure of High-Pressure Recirculation Cooling	1	8.2-5	7	8.2-9	2	8.2-9
Large LOCA; Failure of Low-Pressure Recirculation Cooling	2	1.1-5	8	1.1-9	4	1.1-9
Medium LOCA; Failure of Low-Pressure Recirculation Cooling	3	1.1-5	9	1.1-9	5	1.1-9
Fire; Other Fire Areas Such as the Cable Spreading Room, Auxiliary Feedwater Pump Room, etc.	4	6.7-6	10	6.7-10	7	6.7-10
Large LOCA; Failure of Safety Injection	5	6.4-6	11	6.4-10	8	6.4-10
Fire; Specific Fires in Switchgear Room and Cable Spreading Room Causing RCP Seal LOCA and Failure of Power Cables to the Safety Injection Pumps, the Containment Spray Pumps, and Fan Coolers	6	5.7-6	1	5.7-6	3	1.1-9
Seismic; Loss of Control or AC Power	7	4.7-6	2	4.7-6	6	9.4-10
Interfacing System LOCA	15	5.7-7	5	5.7-7	1	5.7-7

Figure A.2 Comparison of core damage and release frequency contributions.

Besides the fact that the risks are very small, the most important point made by Figure A.2 is that there are different measures of risk, and each is driven by different scenarios. This is an extremely important result of a quantitative risk assessment. This means that when doing risk assessment and ranking the importance of contributors to risk, it is essential to be very clear about what is being used to measure the risk. It is also important to realize that different measures of risk are often necessary to answer the question, what is the risk? Consider the different risk measures of Figure A.2, core damage frequency, frequency of early deaths, and frequency of latent effects. While Scenario 1 is the most important contributor to core damage frequency, it ranks as 15th in importance to the risk measure of frequency of early deaths and 7th in importance to the frequency of latent effects (injuries and deaths). On the other hand, the most important scenario leading to deaths is not the most important contributor leading to core damage and ranks 7th in the list of scenarios important to latent effects.

The above brings up an important issue in making decisions about how to measure risk. That issue is, “which measure is best?” In the case of a nuclear power plant it is clear that if the reactor core is not severely damaged, then there is no risk of radiation exposure to members of the public. On first impression, this would suggest that core damage frequency would be a good measure of risk as it is a surrogate for radiation safety to the public. Furthermore, core damage frequency is much easier to determine with respect to nuclear plant accidents than is the frequency of health effects to the public. That is, there is more confidence (less uncertainty) in the calculation of the likelihood of a core damage event than there is in the calculation of the likelihood of radiation health effects to the public. This is one of the important reasons why core damage frequency has been favored as a measure of nuclear power plant risk. Otherwise, the obvious choice for safety risk would be injuries and fatalities. But there is a need to be cautious about using surrogates to such indicators as health and safety risk as is suggested by Figure A.2. The caution comes from the fact that the risk measure is a frequency of something, not just core damage, injuries, or fatalities. These frequencies are not linearly related. Taking actions to change the frequency of a precursor risk measure has to be examined for how it impacts downstream risk measures. The important fact is the impact could be negative or positive, i.e., it is not always clear whether the impact decreases or increases the value of the downstream risk measure. For example, it would not be good to decrease the core damage frequency and find that the fix to do so results in increasing the frequency of deaths. How could that happen? Suppose we design a plant or make changes to the plant that increases the pressure capacity of the primary system as a way to reduce the core damage frequency. Now we have the situation where we have lowered the core damage frequency for a certain class of core damage scenarios and, most likely, this will result in a lower total core damage frequency. Under these conditions fewer high pressure transients will fail the primary system, but when they do there are now higher pressure transients per transient seen by the containment system. The result is a greater threat to the containment per transient. In other words, depending on the mismatch of the frequencies of the risk measures, we may have increased the failure frequency of the secondary containment and possibly the frequency of injuries and fatalities.

What this all means is that it is not enough to design just for lower core damage frequencies in the case of nuclear power plants. It is important to take a total systems approach and to understand the coupled processes between events that lead to core damage and events that lead to containment failure and off-site consequences.

Other important results from a comprehensive quantitative risk assessment are illustrated in Figures A.3, A.4 and A.5.

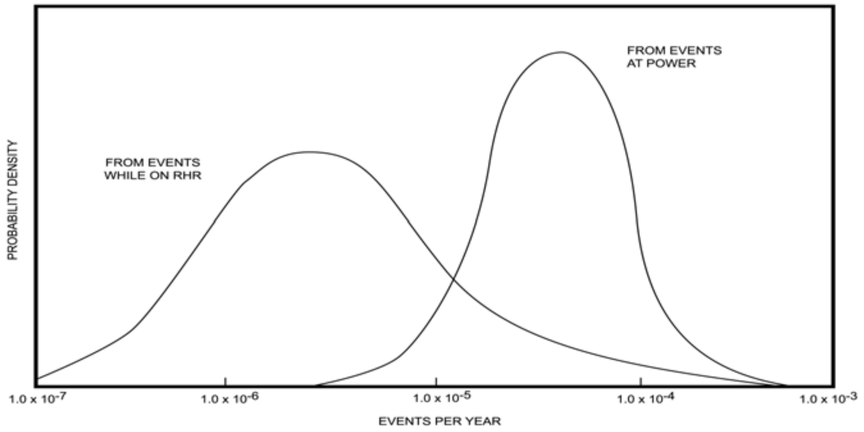

Figure A.3 Probability distribution for core damage.

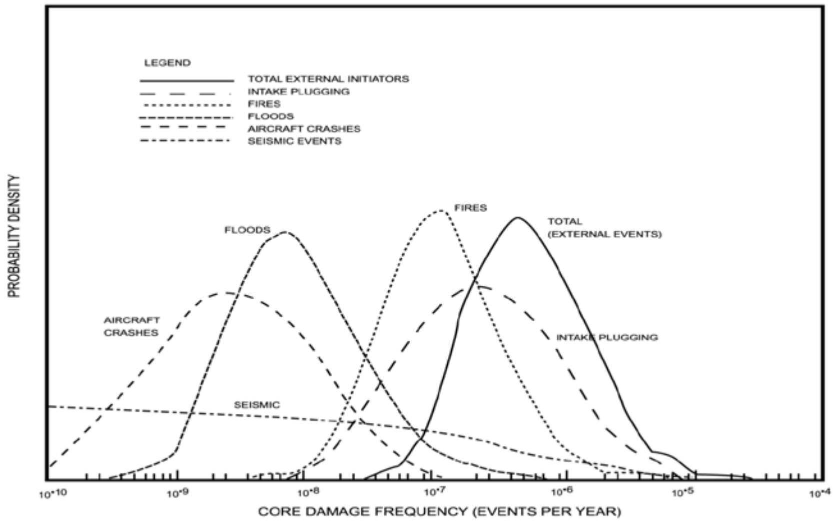

Figure A.4 Contributions of specific external events to CDF.

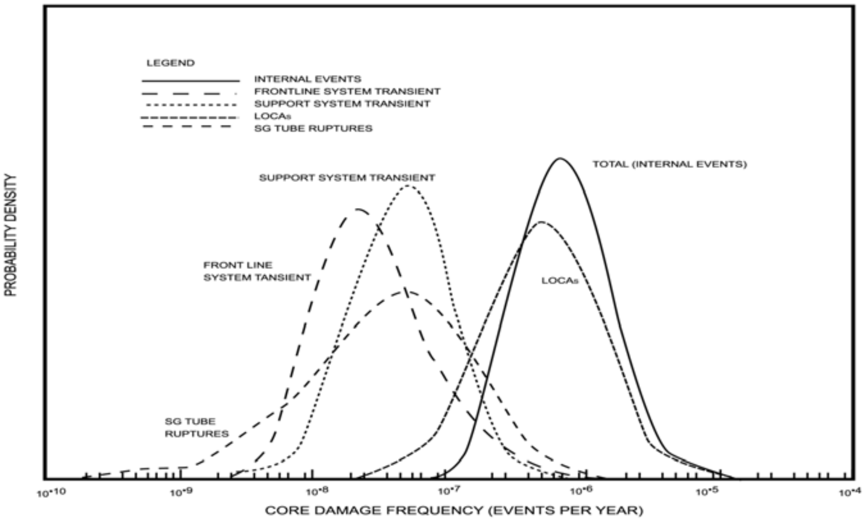

Figure A.5 Contributions of major initiating event classes to CDF from internal events.

Because of radioactive decay of fission products and the heat that they produce, there remains the need to maintain cooling of a reactor core for some time following termination of the chain reaction, i.e., following shutdown of the reactor. Consequently, there remains the risk of fuel damage and the possible release of radiation and radionuclides during times when the reactor is not operating at power. This contribution to the overall risk of exposing workers and the public to radiation must also be a part of a comprehensive risk assessment. Figure A.3 not only presents the core damage frequency of a particular nuclear power plant during off-power conditions, but also provides a direct comparison with the risk of core damage frequency during at-power conditions. More than just comparing the two risks, Figure A.3 quantifies the uncertainties associated with the two contributions to overall risk. As can be seen from Figure A.3, the risk from off-power conditions is very small compared to at-power conditions, but it is certainly not zero. Mean values of the risk of several plants indicate that the contribution to overall core damage frequency of off-power conditions varies approximately from 10 to 30% of the total risk.

In the process of "importance ranking" contributors to risk, it is also essential to quantify any external threats to the plant, i.e., threats that are not due to inherent plant operations but are as a result of such phenomena as external fires, severe weather, aircraft impacts, and earthquakes. Figure A.4 is an example of such an analysis. The most significant external event for this particular plant was the risk of a severe storm creating river debris and other conditions that could lead to plugging the intake lines for secondary cooling. Other visible contributors were external fires, floods, aircraft crashes, and earthquakes.

Figure A.4 makes a very important point about just how much analysis of a contribution should be performed to reduce the uncertainty in the risk measure. It also provides strong evidence for the value of doing uncertainty analysis. As can be seen, the uncertainty range of seismic events covers many orders of magnitude. This amount of uncertainty by itself could be viewed as unacceptable by many analysts, and the push would be on for doing a great deal more work; more work than is necessary for purposes of quantifying the risks. Even with all the uncertainty, Figure A.4 tells us that seismic is not a significant contributor to the overall risk considering the contribution from other sources. If it turned out to be a major contributor, then efforts to reduce the uncertainty could be justified. Summing up the point, if the uncertainty of a contributor varies over many orders of magnitude, say five between the 5th and 95th percentile, and the risk over that interval varies from 10-12 to 10-7, then a risk that is already in the 10-4 range due to other contributors isn't going to be significantly impacted.

Presentations similar to Figure A.4 can be made from many different perspectives. For example, Figure A.5 displays the contribution of different categories of initiating events to the risk. An initiating event is an event that disturbs an otherwise normally operating system and triggers a possible accident scenario. Figure A.5 indicates that for this particular plant, the initiating event category of loss of coolant accidents was the major contributor to core damage frequency. Other important contributors were transients of support systems (emergency power, equipment cooling systems, ventilation systems, etc.), frontline system transients (turbines, feedwater pumps, primary coolant, etc.), and steam generator tube ruptures. Once the scenarios have been ranked in terms of importance to risk, it is usually straightforward to identify the events and equipment items within that scenario most important to risk.

A.5.3Risk Assessment as a Design Tool

Figure A.6 illustrates how risk assessment can be used to achieve maximum benefit from the performance of safety systems during the evolution of a nuclear power plant design. As with Figures A.1 through A.5, Figure A.6 is actual results taken from a nuclear power plant design project.

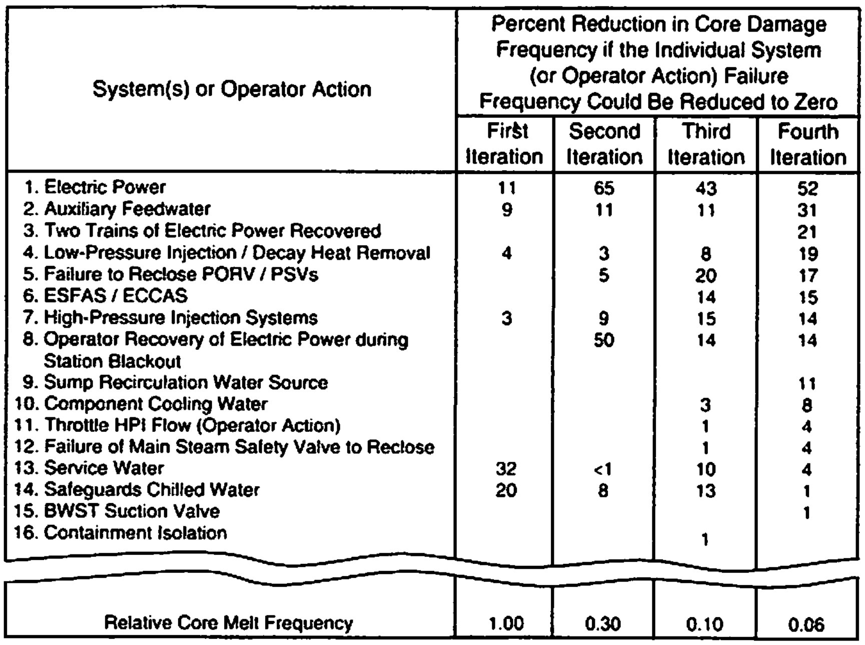

System(s) or Operator Action	Percent Reduction in Core Damage Frequency if the Individual System (or Operator Action) Failure Frequency Could Be Reduced to Zero			
	First Iteration	Second Iteration	Third Iteration	Fourth Iteration
1. Electric Power	11	65	43	52
2. Auxiliary Feedwater	9	11	11	31
3. Two Trains of Electric Power Recovered				21
4. Low-Pressure Injection / Decay Heat Removal	4	3	8	19
5. Failure to Reclose PORV / PSVs		5	20	17
6. ESFAS / ECCAS			14	15
7. High-Pressure Injection Systems	3	9	15	14
8. Operator Recovery of Electric Power during Station Blackout		50	14	14
9. Sump Recirculation Water Source				11
10. Component Cooling Water			3	8
11. Throttle HPI Flow (Operator Action)			1	4
12. Failure of Main Steam Safety Valve to Reclose			1	4
13. Service Water	32	<1	10	4
14. Safeguards Chilled Water	20	8	13	1
15. BWST Suction Valve				1
16. Containment Isolation			1	
Relative Core Melt Frequency	1.00	0.30	0.10	0.06

Figure A.6 Contributors to core damage for four phases of risk management.

The first iteration of the risk assessment determined a core damage frequency and identified the major system and operator activities contributing to the risk. Design flaws were identified with several systems, including the service water system and the safeguards chilled water system—systems critical to cooling safety and support systems during emergency conditions. As a result of the risk assessment, several design changes were made to reduce their contribution to risk. Meanwhile, the core damage frequency was reduced by a factor of over three. The second iteration several months later identified a different set of contributors, but against a much lower risk baseline due to the improvements in the design. Again, there were some standout contributors, most notably the systems for the recovery of the loss of off-site electric power during station blackout.

The third iteration occurring later in the design process began to manifest more balance in the performance of safety systems and less opportunity for any single fix having a major contribution to reducing the risk. In this case it was necessary to consider several design changes to have a significant impact on the risk. However, by the third iteration the mean core damage frequency had been reduced by a factor of 10. By the fourth iteration and the near completion of the design, the calculated core damage frequency had been reduced by almost a factor of 17 and the safety systems were much more balanced in their contribution to protecting the plant. There are some things worth pointing out in the fourth iteration. The 52% contribution from electric power may appear to be a source for further reducing the risk. It is not a good source for corrective action because of its pervasiveness throughout the plant and its association with thousands of subsystems and components. No single or few component, or subsystem, fixes would materially improve the core damage frequency.

An analogy that sometimes clarifies the above process is the mental exercise that comes from thinking of a lake with rocks rising above the surface of the water and viewing the rocks as contributors to risk. What happens when the big rocks are removed from the lake? First, they are no longer there; second, the lake level drops as a result of their removal, and third new rocks appear because of the lower lake level. Repeating this process is very analogous to what is taking place with the risk assessment and design process portrayed in Figure A.6. As the core damage frequency is reduced by removing the important contributors, other contributors begin to appear that become candidates for removal or modification. Eventually you get to a point of diminishing risk benefit.

The point of this example is to illustrate the robust amount of information that comes from a comprehensive risk assessment and the options it provides for cost-effective risk management. The case studies of Chapters 3 through 6 of this book involve much more limited scope assessments and were primarily for the purpose of demonstrating the six- step process of quantitative risk assessment.

A.6 References

AEC (U.S. Atomic Energy Commission). 1957. “Theoretical Possibilities and Consequences of Major Accidents in Large Nuclear Power Plants,” WASH-740, March.

Boyd, G.J., and G. St. Pierre. 2001. “Risk Management Program for the Disposal of Chemical Agents and Munitions,” presented at the Society for Risk Analysis: Special Symposium on Quantitative Risk Assessment, sponsored by the Family Foundations of Chauncey Starr and B. John Garrick; B. John Garrick Foundation for the Advancement of the Risk Sciences, Laguna Beach, CA, May 31-June 2.

Cox, R.T. 1946. “Probability, Frequency, and Reasonable Expectation,” *Am. Jour. Phys.*, 14, 1-13.

Farmer, F.R. 1964. “The Growth of Reactor Safety Criteria in the United Kingdom,” *Proceedings of the Anglo-Spanish Nuclear Power Symposium*, Madrid, Spain.

Garrick, B.J. 1968. “Unified Systems Safety Analysis for Nuclear Power Plants,” Ph.D. Thesis, University of California, Los Angeles.

Garrick, B.J. 1989. “Lessons Learned from 21 Nuclear Plant Probabilistic Risk Assessments”, *Nuclear Technology*, Vol. 84, March.

Green, A.E., and A.J. Bourne. 1972. *Reliability Technology*, John Wiley. London, United Kingdom.

Holmes and Narver, Inc. 1967. “Reliability Analysis of Nuclear Power Plant Protective Systems,” prepared for U.S. Atomic Energy Commission, Washington, D.C., HN-190.

Jaynes, E.T. 2003. *Probability Theory: The Logic of Science*, Cambridge University Press, United Kingdom, Cambridge.

Jeffreys, Sir Harold. 1957. *Scientific Inference*, University Press, Cambridge, United Kingdom, Cambridge, Second Edition.

Kaplan, S., and B. J. Garrick. 1981. “On the Quantitative Definition of Risk”, *Risk Analysis* 1(1): 11-27.

Kumamoto, H., and E.J. Henley. 1996. *Probabilistic Risk Assessment and Management for Engineers and Scientists*,. IEEE Press, New York, 2nd Edition, pp. 1-54.

(NASA) National Aeronautics and Space Administration. 2002. “Probabilistic Risk Assessment Procedures Guide for NASA Managers and Practitioners,” prepared for Office of Safety and Mission Assurance, National Aeronautics and Space Administration, Washington, D.C.

Neufeld, J. 1990. “The Development of Ballistic Missiles in the United States Air Force 1945-1960,” Office of Air Force History, U.S. Air Force, Washington, D.C., pp. 169-215.

PLG (Pickard, Lowe and Garrick, Inc.). 1981. Westinghouse Electric Corporation, and Fauske and Associates, Inc. “Zion Probabilistic Safety Study,” prepared for Commonwealth Edison Company, Chicago, Illinois.

PLG (Pickard, Lowe and Garrick, Inc.). 1982. Westinghouse Electric Corporation, and Fauske and Associates, Inc. “Indian Point Probabilistic Safety Study,” prepared for Consolidated Edison Company of New York, Inc., and the New York Power Authority.

Pólya, G. 1954. *How to Solve It, Mathematics and Plausible Reasoning*, Princeton University Press, Volumes I and II.

Raiffa, H. 1996. *Decision Analysis*, McGraw-Hill Primis Custom Publishing, Columbus, Ohio.

Shannon, C.E. 1948. “A Mathematical Theory of Communication,” *Bell System Technical Journal*, Vol. 27, pp. 379-423 and 623-656, July and October.

USNRC (U.S. Nuclear Regulatory Commission). 1975. “Reactor Safety Study: An Assessment of Accident Risks in U.S. Commercial Nuclear Power Plants,” WASH-1400 (NUREG-75/014).