Chapter 1 Introduction

1.1 The Double-edged Sword

Nuclear reactor safety is essential. The huge power potential inherent in fission is both the reason power reactors are built and the reason the risk is high. The decision to use nuclear power is, ultimately, a public decision. Because of the technical expertise required, the public relies on the engineers and scientists to ensure adequate safety. Few products of technology evoke as much emotion and fear as nuclear technology. For many, it is inherently bad and is to be avoided at all costs. But following such a route would be costly indeed. Much has been written on the benefits of nuclear technology and it need not be repeated here. Suffice it to say that the non-power aspects of nuclear technology includes medical and industrial applications to the tune of $300 billion in the US, power applications were estimated at $57 billion in the US [TUS94] and the Canadian figures are roughly 1/10 of that. Nuclear technology is big business, even if we are currently in a down period. It is big business because there are significant benefits in the use of nuclear technology, benefits that are sufficiently large to pursue in spite of the large potential for severe accidents. In essence, nuclear is no different than all human activities. All activities involve some risk. What is acceptable in terms of risk depends on the benefits. Few things that we do on a day to day basis are as risky as hurtling down a narrow strip of levelled ground in a metal container containing 60 litres of explosive liquid towards someone else in a similar device, using a painted strip as a guide to avoid collision. Yet we do it. Presumably, the benefit is worth the risk.

Once we have decided to employ a technology, the job at hand is to minimize the risk, minimize the cost, and maximize the benefit. These objectives are usually competing and ensures that the job of the designer is "interesting". It is essential to note that tradeoffs are inherent in the nature of the problem. It is not acceptable to require absolute safety at all costs. In fact, it is nonsense to require absolute safety. Nothing is absolutely safe. And we do not have infinite resources. The unrestrained pursuit of additional safety at some point incorrectly and unjustly diverts resources (time, people, money, natural resources ...) away from other important programs (health care, education, transportation, ...).

We need a methodology, then, to quantify risk, safety, benefit, etc., and to permit design, construction and operation to take place on a rational and justifiable basis. This course is an attempt to elucidate that methodology, a methodology employed by the nuclear industry and other industries such as the space and aircraft industries.

1.2 That's incredible

Given a design, the basic methodology can be stated quite succinctly:

Show that the consequences of the event are within acceptable limits

or

Show that the probability of an event (normal or accident) is too incredible to consider.

Acceptable limits are defined with respect to the event frequency, For example, frequent occurrences, like minor faults, should not stress the system or invoke protective systems. Very infrequent events, like a large loss of coolant, are permitted to push the physical systems into plastic deformation but not allow a radioactive release beyond a prescribed limit.

Incredible is defined as sufficiently low, say one in a million. Anything above that frequency typically gives rise to varying degrees of concern as shown in table 1.1.

Table 1.1 Acceptance of annual fatility risk levels. Source [MCC81, table 18-3, pg 370]

Annual fatality risk level, yr^{-1}	Conclusion
10^{-3}	This level is unacceptable to everyone Accidents providing hazard at this level are difficult to find When risk approaches this level, immediate action is taken to reduce the hazard
10^{-4}	People are willing to spend public money to control a hazard (traffic signs/control and fire departments) Safety slogans popularized for accidents in this category show an element of fear, i.e., "the life you save may be your own"
10^{-5}	People still recognize People warn children about these hazards (drowning, firearms, poisoning) People accept inconvenience to avoid, such as avoiding air travel Safety slogans have precautionary ring: "never swim alone," "never point a gun," "never leave medicine within a child's reach"
10^{-6}	Not of great concern to average person People aware of these accidents but feel that they can't happen to them Phrases associated with these hazards have element of resignation: "lightning never strikes twice," " an act of God"

[a] Extracted from H. J. Otway and R. C. Erdmann, *Nucl. Eng. Design* 13, 365 (1970).

So, safety or its negative counterpart, risk, is a function of the frequency of occurrence of an event and the consequence of that event.

1.3 Risk

Safety concerns are ultimately expressed in terms of risk. Risk is customarily defined [MCC81] as:

$$Risk = \sum_i expected\ frequency\ of\ event_i\ X\ expected\ consequence_i \tag{1}$$

which reflects the increase in risk when either the number of events or the magnitude of the events increased. This is by no means a unique definition; for instance, if one wanted to amplify the importance of events with large consequences, risk could be defined as:

$$Risk = \sum_i expected\ frequency\ of\ event_i\ X\ (expected\ consequence_i)^k \tag{2}$$
$$where\ k > 1$$

We seek to minimize risk. We do so by choosing the least risky path to achieve the desired goal. But

lowering risk is usually expensive and, since we have finite resources, we need to balance the cost versus the benefit. This is done by setting quantitative safety targets. The target levels of acceptable risk are set with respect to the alternative ways of achieving the same goals. For instance, acceptable levels of risk for nuclear power plants should be set at levels comparable to the level of risk inherent in coal and oil fired plants. The public is, however, risk adverse to things nuclear. Thus we find that the acceptable level of risk for nuclear power has been set substantially below that of alternative means of large scale power production. This has ensured that nuclear power is safer than the alternatives (and indeed safer than most human activities), but this safety has come at a significant social cost. One can argue that the funds spent on the extra safety should have been spent elsewhere.

Reference [MCC81, Chapter 17] discusses the risk-benefit assessments. Figure 1.1 illustrates that it dealing with risk (ie providing safety) becomes more and more expensive as the risks become smaller - a form of diminishing returns on our efforts to make the world a safer place to live. Conversely, the social cost increases as the risk level increases. We seek to minimize the total cost (assuming that the true cost can be properly quantified). Starting from the right side of figure 1.1, the high social cost of very risky things and the relatively low cost of implementing safer systems leads society to invest wisely in these safer systems. As we progress to consider endeavours of lower and lower risk, the increasing cost of implementation of safer systems begins to outweigh the benefit derived from the safer systems. At some point, we have to say "enough". But how do we know when enough is really enough?

1.4 Historical Development

Quantification of "enough" implies quantifying the consequences and quantifying the frequencies of possible events. In short, we need to analyse the safety aspects of the endeavour in question. There has always been a recognition of the role of probability and consequence in determining the risk of a design even if it was not explicitly stated. But, in the early 1900's and before, because our analysis capability was limited and because failure data was not readily available, risk was reduced by over-design. This works but there is an opportunity cost to this approach. A 10 ton automobile might offer increased safety but at what cost to the environment and to occupants of lesser vehicles? Further, because analysis capability was limited, improvements occurred more as a result of "leaning by mistakes" than as a result of pre-production design and analysis. This is acceptable for products that can be exhaustively tested to failure (like automobiles) but it is not acceptable for the nuclear industry or similar industries where it is neither financially possible nor socially acceptable to test complete systems to failure.

It is only recently that failure rate data has become more available. Consequently, prudent engineering required a more deterministic approach: ensure protection against prescribed events. The probabilistic approach, however, provides a rational framework for the deterministic approach and, thus, it is pedagogically useful to cast our study of safety design in those terms first.

1.5 Probabilistic Safety Analysis

Probabilistic Safety Analysis (PSA) seeks to categorize each event by probability of occurrence and then demonstrate that certain criteria are meet. The nuclear industry uses two general types of acceptance criteria for PSAs: Binning and Averaging.

- Binning techniques are based on limiting the consequences for any event based on frequency. Examples are the ASME code and C-6 discussed in chapters 3 and 4.

- Averaging techniques are based on setting a limit on the frequency of a given outcome, which we will call a "safety goal". For example the expected frequency of the release of XX TBq of radioactivity be less than 10^{-5} events/year or that the core damage frequency be less than 10^{-5} events/year.

Both criteria use the PSA methodology developed in this course. The safety goal methodology requires the summation of the frequency of all events that exceed the stated criteria (set a few orders of magnitude below the desired limit).

Despite different acceptance criteria, these PSAs proceed using the following methodology:

- define the acceptance criteria,
- generate a set of design basis accidents to consider,
- analyze the frequency of the event,
- and finally show that the appropriate frequency based criteria are meet.

Each is discussed in turn in the following.

1.5.1 Safety Criteria

Each event is associated with a criteria against which the event is to be judged. The engineering industry has established this over the years and is epitomized by the ASME and ANSI standards. These standards are concerned with material stress limits. The nuclear industry goes well beyond these standards by considering radioactive releases, as discussed in detail in chapter 3.

1.5.2 Design Basis Events

The task here is to define all the possible initiating events that are deemed necessary to analyze. The range is everything from normal operation to accidents involving major core releases. These form the Design Basis Accidents or DBA, discussed in detail in chapter 4. The worst conceivable accidents are investigated for completeness but their probability is so low (by design) that the are not part of the DBA set.

1.5.3 Probability Risk Assessment

Since events are classified by the frequency of occurrence, the reliability of systems have to be measured or analyzed. Event scenarios, called Event Trees (ET), are developed. Each branch of the ET needs an associated probability if the event and its consequence is to be quantified. Fault trees (FT) are commonly used to determine failure probabilities.

The sequence, then, is to define the accident events to be analyzed (DBA), construct the event trees (ET) supported by the fault trees (FT) probabilities. If an event sequence is "incredible", then no further action is required. This process is illustrated in figure 1.2.

1.5.4 Safety Analysis

For each branch of the ET that is "credible", ie. has a frequency higher than a predefined cutoff, safety analysis must be performed, usually by computation and experimentation, to determine if the consequences are within acceptable limits or not. Safety analyses are very complex and require extensive

knowledge of an event. These analyses are beyond the scope of this course and a black box approach will be used. Chapter 6 discusses safety analysis. If the limits are not exceeded no further action is required. If they are, something has to be done to mitigate the issue. That something is design.

1.6 Safety Design

All designs must be subjected to the above methodology, even passively or inherently safe ones. The "better" the design, the easier it will be to meet the acceptance criteria. But even the most benign design must be **shown** to be benign. As one analyses a given design, weaknesses and areas for improvement show up. We might find that reactors with negative void coefficients of reactivity are not necessarily safer than those with positive coefficients. We might find that most equipment faults of consequence are caused by secondary and supportive systems, not the reactor and reactor coolant proper. We might find that most accidents are cause by human error, not machine error. We might find that **all** designs, even passively safe ones, have failure modes (like loss of reactor power regulation) that are not passively safe. However, we won't find anything unless we look and we can't judge what we find unless we are able to quantify our findings. This course is about how to do just that. Some key CANDU system designs are discussed in chapter 7.

The subject of "safety design" is a combination of safety system design and safety analysis. Design is the process by which a system is engineered to perform its intended function. Ideally, we would like to be able to work backwards from the design criteria to define the actual design, that is, from a performance specification to a component and system specification (geometry, materials and operating parameters). But the calculations are far too complex and convoluted for that. Instead, we use past experience and accepted practices to conceive of an initial design and proceed to analyze that design to see if it meets the performance specifications. Obviously this is an iterative process. In the nuclear industry, practical design exercises rely heavily on previous designs and new designs tend to be evolutionary rather than revolutionary for at least two reasons: cost and performance assurance. It has been estimated that the overall cost of taking a reactor concept from paper to a commissioned prototype power reactor is about $1 billion. This alone biases the design process to lean heavily on past designs. But apart from the cost, overall process and safety performance is a strong function of accumulated operating experience and laboratory testing. New designs tend to be "buggy" at first and the leap of faith required is, more often than not, too big to surmount without some crisis to drive designers into the unknown. "If it isn't broken, don't fix it!" rings true. Design changes have to be carefully managed if quality is to be maintained. A superb new design executed for the first time is usually inferior to a mediocre but well established design. Because quality assurance tends to be expensive, a large part of the total cost of a new design can be directly or indirectly attributed to the assurance of quality. This is not conservatism for the sake of conservatism; rather, it is a progressive and controlled approach to design and can be summed up as simply good engineering practise. That having been said, we must keep a balanced approach and remain open to innovation.

This course is concerned with both design and analysis. The following chapters develop the analysis tools to allow the student to analyze a design, determine weak points and assess alternatives to determine if a system is safe enough. The student should be able to answers questions about the amount of redundancy required for adequate safety. The key tool used to answer these questions is Probabilistic Safety Analysis. The material by Dr. Archie Harms and the course project permit the consideration and exploration of design alternatives using the tools developed herein.

1.7 Actual Practice

Probabilistic Safety Analysis (PSA) as outlined above has proven to be very effective in ferreting out design and operation inadequacies. But it has not been completely successful on two fronts. One, we can only analyze events that we can conceive. What about the unknown? Two, PSAs are sensitive to the choice of branch points in the cut sets of the event and fault trees and are sensitive to the measured equipment failure probability data. Nature is perverse and Murphy was an optimist. As we shall see, the safety criteria used has its roots in a probabilistic approach but for practical purposes, the criteria is deterministic in nature and is firmly founded in the principles of good engineering practice and experience. This has, in more recent years been augmented and complemented by probabilistic analysis. Thus, actual practice has two parallel streams: the deterministic assessment path and the probabilistic path, as illustrated in figure 1.3.

In Deterministic Safety Analysis the acceptance criteria is not based on probability, but on a number of assumed faults. Typically a single/dual mode failure criteria is used. The acceptance criteria is more stringent for the more probable single failure and less stringent for the less probable single failure. Typically they are rooted in probabilistic arguments and are very simple to understand and implement.

For these reasons, the basis of safety analysis remains the same basis as for good design itself. A good design is a safe design; the cost of downtime, worker injury, litigation, and repair more than outweigh the cost of achieving a good design to begin with. Safety does indeed pay. The probabilistic approach, then, fits within the standard engineering design practice. Figure 1.4 is an overview of the design process from a very generic stance. Can you see where the PSA amd the deterministic assessment fit in?

Figure 1.4 is but one way to view the whole process. We'll see other views as well, such as that of the IAEA and the AECB in Chapter 8. We shall see that the views are complementary. All views revolve around the common sense approach that is inherent in good engineering practice: start with a good design, follow established safety and design practices, and provide protection against the risks.

1.8 Learning Outcomes

In each chapter the course objectives (learning outcomes) are set down. Learning by objectives has received some "bad press" since some lecturers have a tendency to be overly specific in their statement of objectives and some students have a tendency to be overly narrow in their learning of only that explicitly stated in the objectives. This is, of course, inappropriate. Herein, the tone of the objective statements is set to be specific enough to serve as a guide to expectations but not so specific as to be questions in disguise. The outcomes are meant to be a guide for the student and teacher alike. The list is by no means exhaustive but it is hoped that it is complete enough to indicate the type and extend of learning to be mastered.

The classifications in the objective statements refer to Bloom's taxonomy [BLO71] for the cognitive domain as given in figure 1.5. These classifications are important in that they indicate the type of understanding that is to occur, ie, whether the student is to just memorize facts or is to achieve some higher level mental ability. The weight of each classification is

a = "must"

b = "should"
c = "could"
indicating the importance of the objective to the understanding of the overall course.

The objectives are keyed to the concepts to be learned and not to specific course content since the content is just one of many ways to elicit understanding in the student.

The overall objectives for the course are as follows:

Objective 1.1	The student should be able to explain the overall theme of the course and relate the roles played by deterministic and probabilistic safety analysis.					
Condition	Closed book written or oral examination.					
Standard	100% on definitions, answer may be given using word descriptions, diagrams or graphs as appropriate.					
Related concept(s)	Overall concept map for the course					
Classification	Knowledge	Comprehension	Application	Analysis	Synthesis	Evaluation
Weight	a	a	a			a

Objective 1.2	The student should be able to explain the role of design basis events, event trees and fault trees in PSA.					
Condition	Closed book written or oral examination.					
Standard	100% on definitions, answer may be given using word descriptions, diagrams or graphs as appropriate.					
Related concept(s)						
Classification	Knowledge	Comprehension	Application	Analysis	Synthesis	Evaluation
Weight	a	a	a	a		

Objective 1.3	The student should be able to calculate event tree and fault tree frequencies and probabilities.					
Condition	Open book examination or Workshop project investigation.					
Standard	75 %					
Related concept(s)						
Classification	Knowledge	Comprehension	Application	Analysis	Synthesis	Evaluation
Weight	a	a	a			

1.9 Course Methodology

Every lecturer is faced with the issues of just when and in what order material should be presented for best effect. The temptation is to present the material from the top down, proceeding from the general, in all its conciseness and beauty, to the specific. This is not a good approach for two reasons. First, people learn from the bottom up, from the specific to the general. Second, skills take time to learn and iteration is necessary. Hence, from a pedagogical point of view, it is preferable to first develop key mathematical and other procedural concepts and techniques in isolation and to then integrate those concepts into a cohesive philosophy. In this manner, the student can focus on acquiring the knowledge base and perfecting the skills of sub areas before attempting to understand the interrelation and integration of the concepts. This also permits the student to start the skills practice early since it does take some time to become fluent. Thus, probability theory and simple reactor sub-systems are addressed first (Chapter 2). By the end of chapter 2, the student should be comfortable with concepts such as failure rates, availability, reliability, test frequencies, dormant and active systems, and probability evaluation for simple systems.

While these skills are being acquired and refined through assignments, exploration of the historical and philosophical basis for nuclear safety might be a welcome relief from computation. Hence the overall approach to safety design is explored and chapter 3 deals with safety criteria used to evaluate the events to be analyzed, which is the subject of chapter 4. By this time, the student should be ready for more "skill" type material; fault trees and event trees are covered in chapter 5. Chapter 6 discusses safety analysis but this is a huge topic onto itself and can only be hinted at within this course. By this time the student should have a reasonable grasp of the overall safety design picture and can meaningfully address practical systems. Thus, chapter 7 reviews and explores the key safety systems of CANDU reactors.

Even though a large fraction of this course is concerned with probabilistic analysis, we should never lose sight of the fact that reactors are safe by design, not by analysis. The analysis is merely to demonstrate the good design. The final chapters of the course provide a wrap up and a look at reactor safety from this fairly lofty and general perspective.

1.10 Exercises

1. Where does figure 1.3 fit into figure 1.4?

2. If you had to take one of the following two risks, which risk would you prefer:
   a. 1 chance in 1000 of losing $1
   b. 1 chance in 1,000,000 of losing $1000?

3. If you had to take one of the following two risks, which risk would you prefer:
   a. 1 chance in 1000 of losing $1000 or
   b. 1 chance in 1,000,000 of losing $1,000,000?

4. If you had to take one of the following two benefits, which benefit would you prefer:
   a. 1 chance in 1000 of receiving $1 or
   b. 1 chance in 1,000,000 of receiving $1,000?

5. If you had to take one of the following two benefits, which benefit would you prefer:
   a. 1 chance in 1000 of receiving $1,000 or
   b. 1 chance in 1,000,000 of receiving $1,000,000?

6. Where do your choices fall on the risk plot of figure 1.6? Are you averse to risk with large consequences?

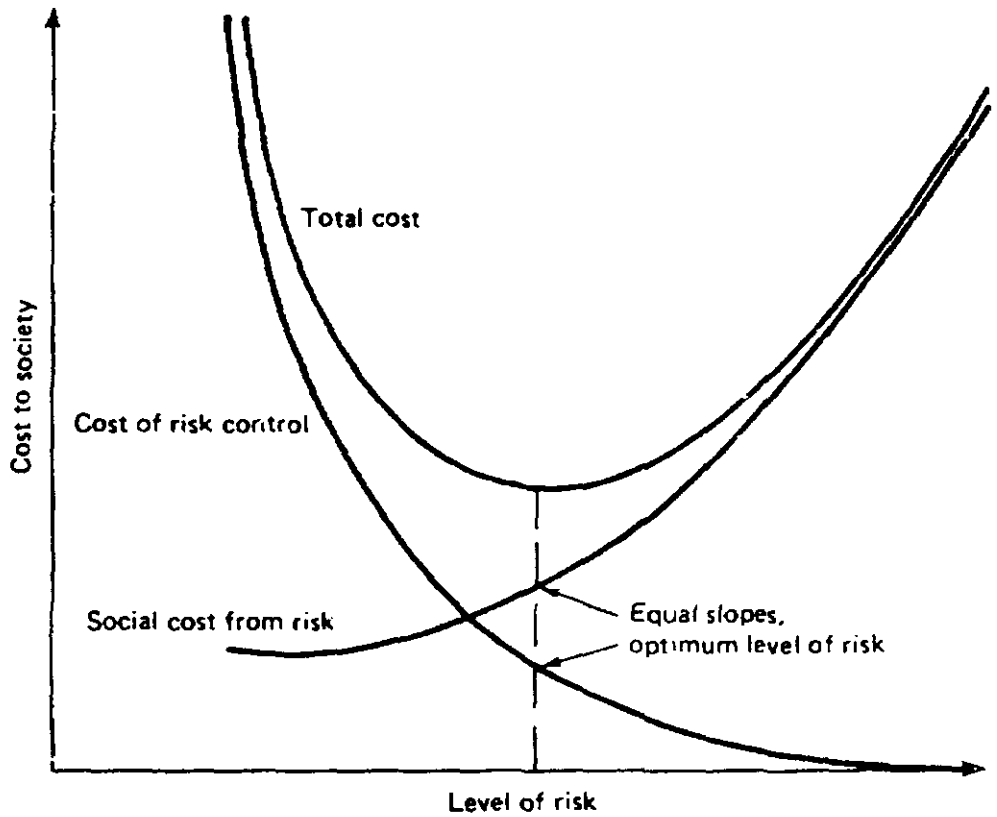

Figure 1.1 Social and control costs versus level of risk [Source: MCC81 figure 17.1]

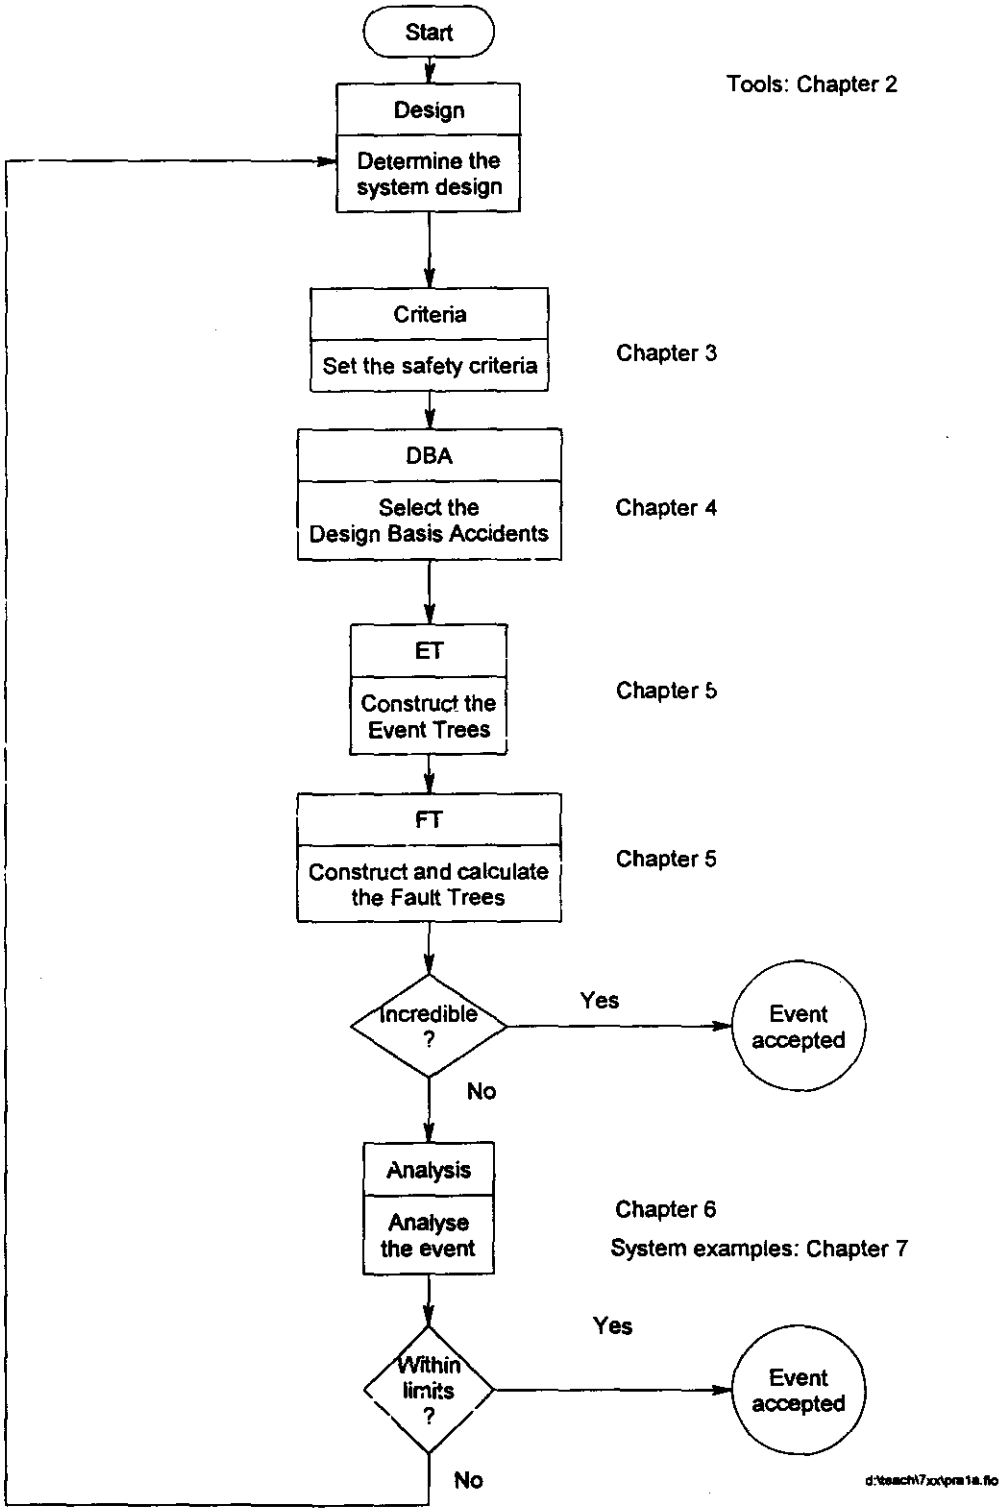

Figure 1.2 Probabilistic Risk Analysis Overview

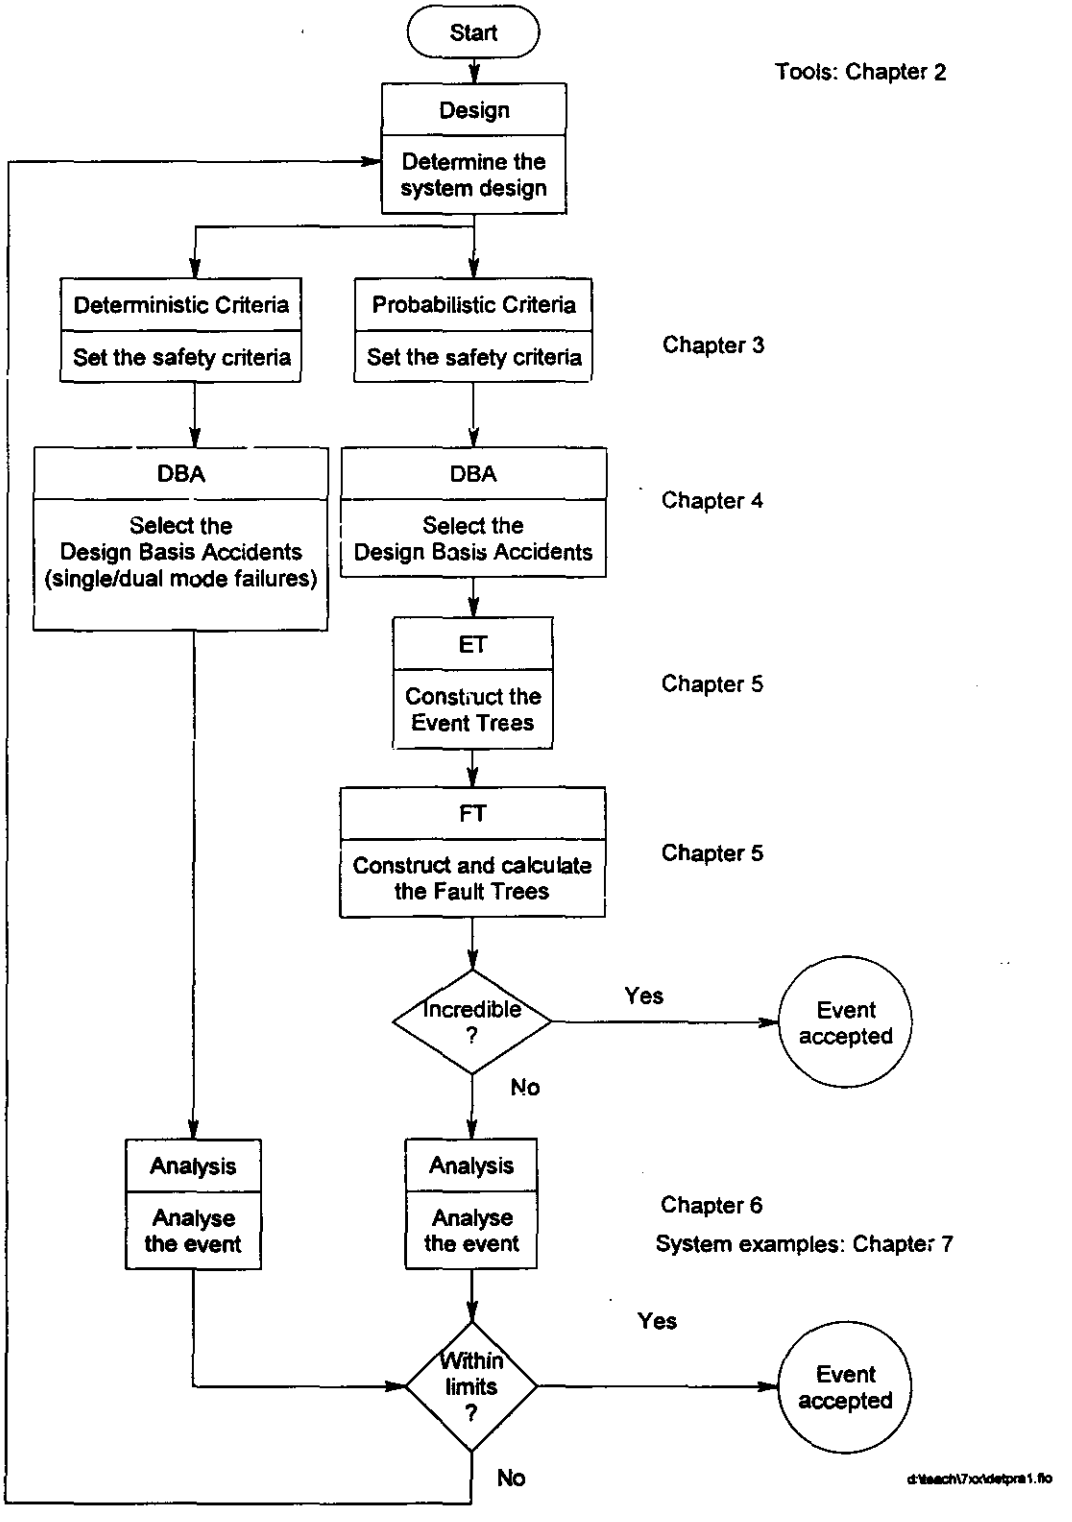

Figure 1.3 Safety Analysis Overview

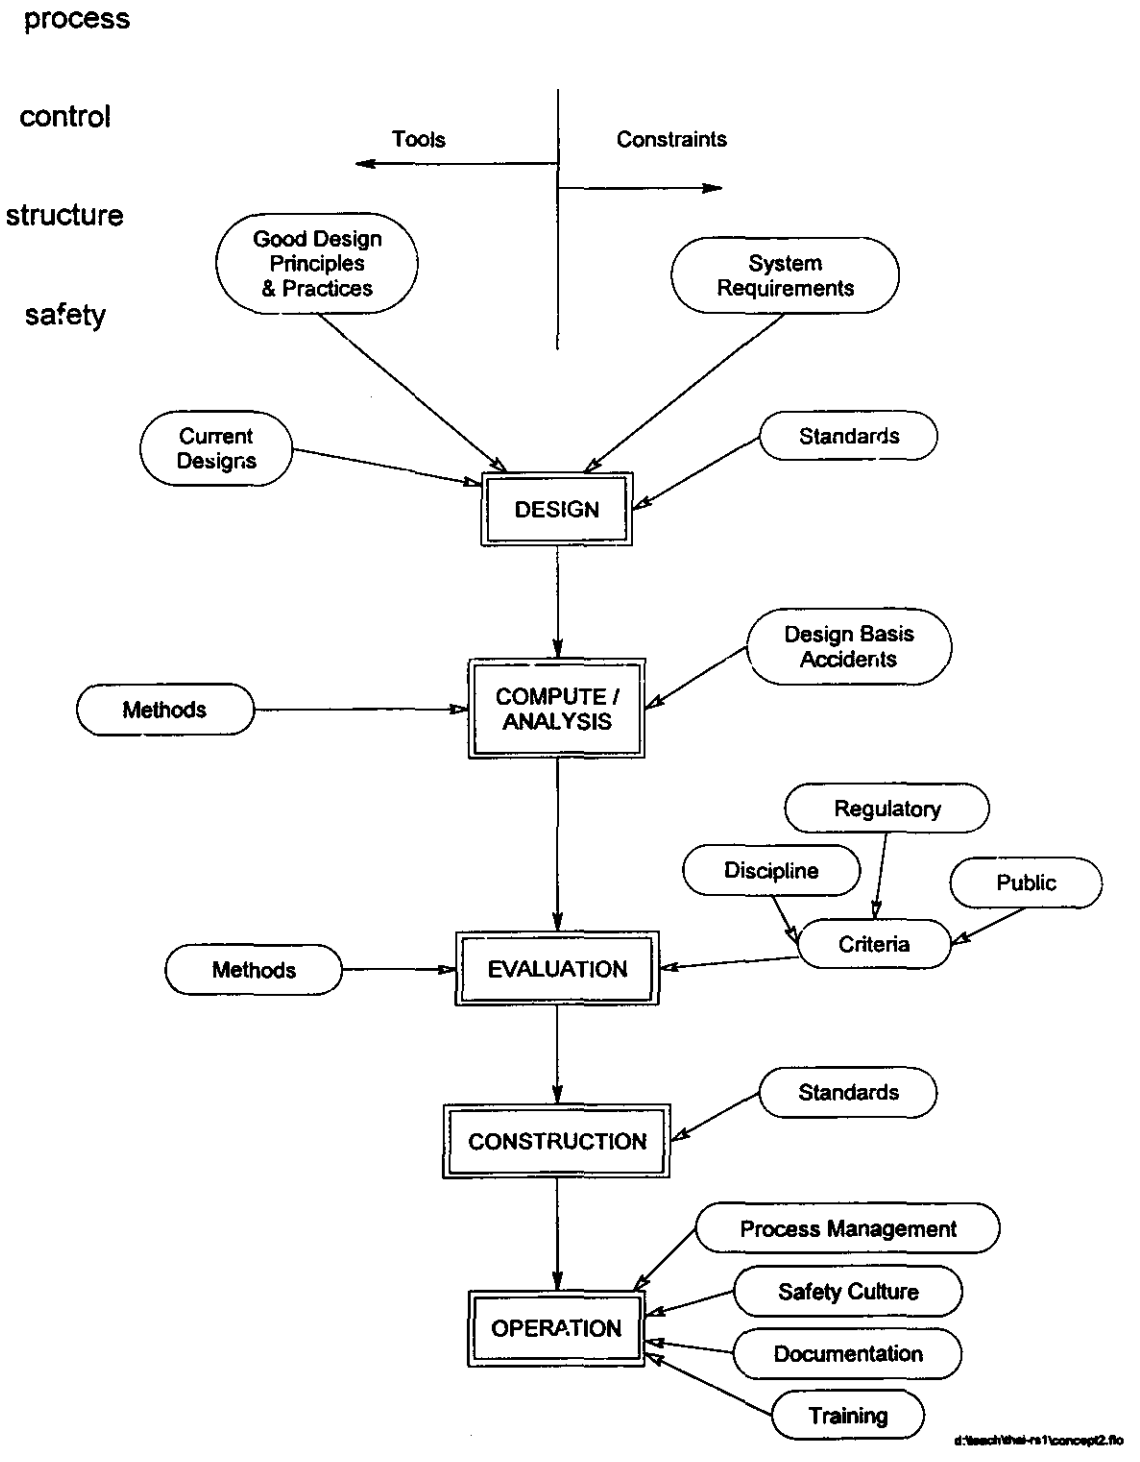

Figure 1.4 Overview of the design process.

Handout Master 12.5
Objectives in the Cognitive Domain

Operationalizing the Taxonomy of Objectives in the Cognitive Domain

Taxonomic Categories and Subcategories	Verbs to Use in Objectives	Examples of Appropriate Content in Objectives
1.00 Knowledge	Define	Vocabulary words
1.1 Knowledge of specifics	Distinguish	Definitions
1.2 Knowledge of ways and means of dealing with specifics	Acquire Identify Recall	Facts Examples Causes
1.3 Knowledge of universals and abstractions	Recognize	Relationships Principles Theories
2.00 Comprehension	Translate	Meanings
2.1 Translation	Give in one's own words	Samples
2.2 Interpretation	Illustrate	Conclusions
2.3 Extrapolation	Change	Consequences
	Restate	Implications
	Explain	Effects
	Demonstrate	Different Views
	Estimate	Definitions
	Conclude	Theories Methods
3.00 Application	Apply	Principles
	Generalize	Laws
	Relate	Conclusions
	Choose	Methods
	Develop	Theories
	Organize	Abstractions
	Use	Generalizations
	Restructure	Procedures
4.00 Analysis	Categorize	Statements
4.1 Analysis of elements	Distinguish	Hypotheses
4.2 Analysis of relationships	Identify	Assumptions
4.3 Analysis of organizational principles	Recognize	Arguments
	Deduce	Themes
	Analyze	Patterns
	Compare	Biases
5.00 Synthesis	Document	Positions
5.1 Production of a unique idea	Write	Products
5.2 Production of a plan	Tell	Designs
5.3 Derivation of a set of abstract relations	Produce	Plans
	Originate	Objectives
	Modify	Solutions
	Plan	Concepts
	Develop	Hypotheses
	Formulate	Discoveries
6.00 Evaluation	Justify	Opinions
6.1 Judgments in terms of internal evidence	Judge	Accuracies
	Argue	Consistencies
6.2 Judgments in terms of external criteria	Assess	Precisions
	Decide	Courses of action
	Appraise	Standards

Adapted from N. S. Metfessel, W. Michael, and D. Kirsner, Instrumentation of Bloom's and Krathwohl's taxonomies for writing educational objectives, Psychology in the Schools, 1969, 6, 227–231.

Figure 1.5 The cognitive domain

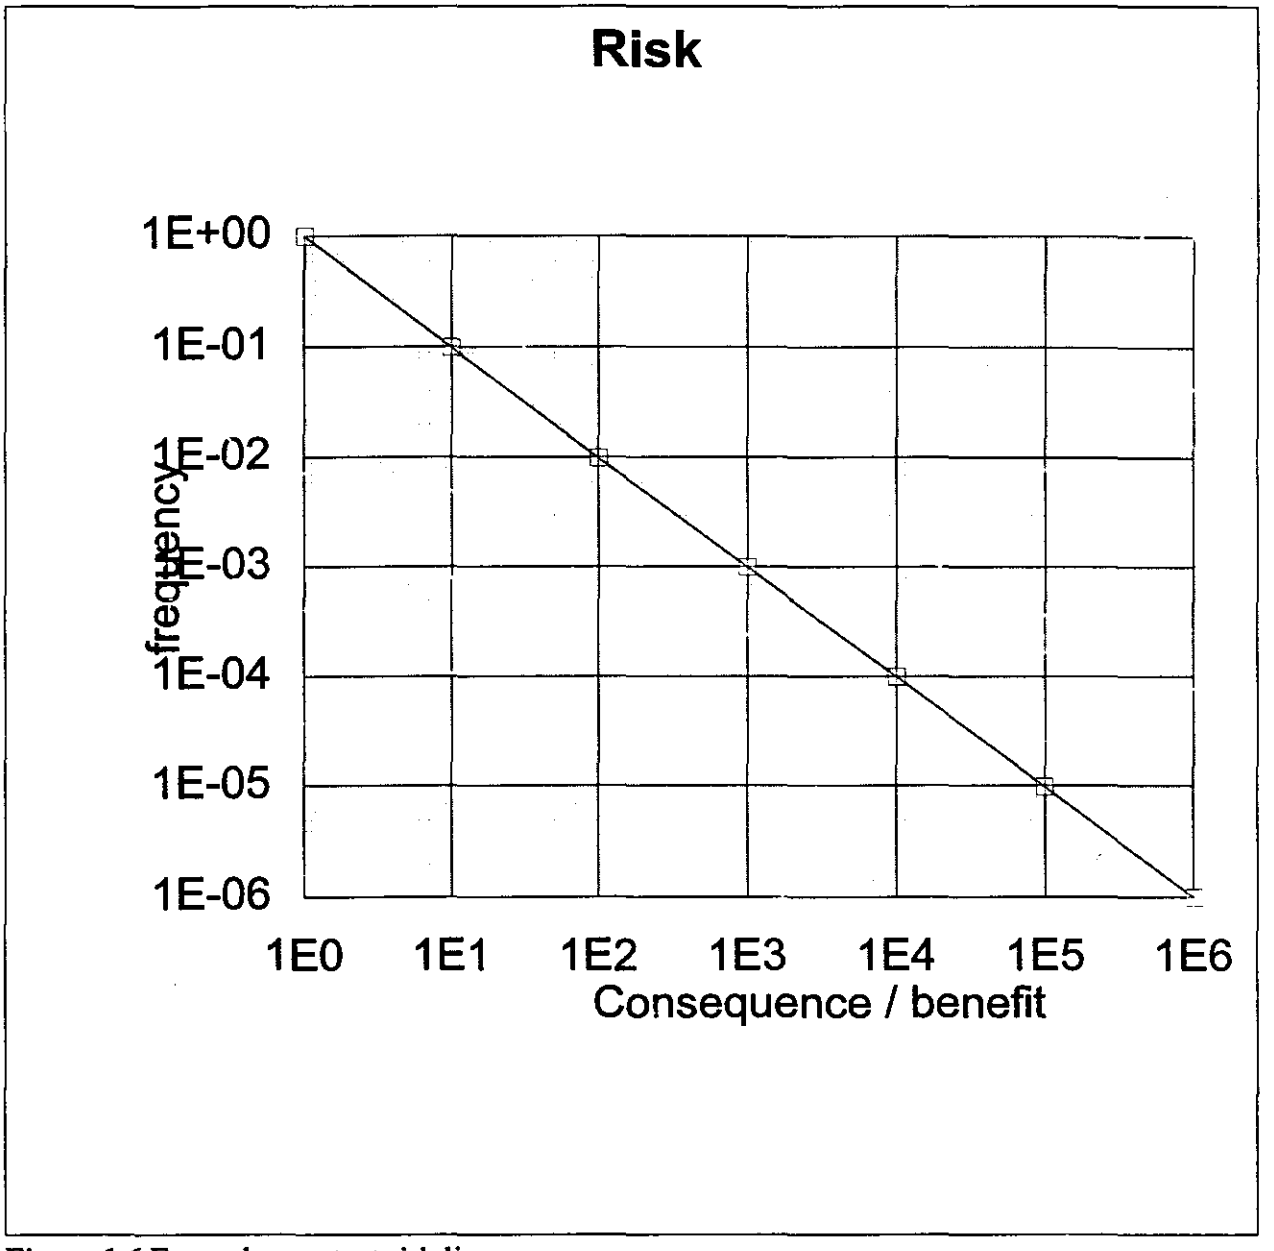

Figure 1.6 Example constant risk line