

# ENVIRONMENTAL REGULATION: DEVELOPMENTS IN HOW TO SET REQUIREMENTS AND VERIFY COMPLIANCE

**Rex V. Brown**

*George Mason University, Fairfax, VA, USA*

**Keywords:** safety regulation, acceptable risk, decision analysis, probabilistic risk assessment, documented and judgmental safety assessment

## Contents

- [1. Introduction](#)
- [2. How Requirements Should Be Set](#)
- [3. How to Verify Compliance with Requirements](#)
- [4. Concluding Comments](#)
- [Acknowledgments](#)
- [Related Chapters](#)
- [Glossary](#)
- [Bibliography](#)
- [Biographical Sketches](#)

## Summary

Current regulation often appears to promote safety management practices that are neither economical nor adequate for environmental protection. However, the underlying paradigms are beginning to shift in an encouraging direction. Driving this shift is the growing acceptance in regulatory agencies, especially nuclear, of personal decision analysis as a guide to rational action. However, the shift is slowed by the conflicting priorities of key players. For example, regulatory agencies may be charged only with protecting the environment, whereas the public expects their decisions to balance protection with economic and other social concerns.

The effectiveness of environmental regulation is geared to a web of uncertain causal connections between *means* and *end*. The end is the advancement of overall social welfare; the means are the concrete steps a regulated facility takes to meet this end (such as equipment design, or a maintenance procedure). Between them is a hierarchy of consequences in several "tiers". For example, design and maintenance affect equipment reliability; which affects accident risk; which affects release of pollutants; which affects public health; which ultimately affects overall social welfare. Each such item can be affected by others at lower tiers. Uncertainty at lower tiers produces uncertainty at higher tiers.

*Regulatory requirements* can be set anywhere in this hierarchy. The facility can be told what to do in great detail, i.e. means. Or it can be required to simply "serve the public interest", i.e. the end. Or it can be required to assure that accident risk is acceptable (say, less than one in ten thousand probability), i.e. in-between. The closer the requirement is to means (and further from ends), the easier it is to check compliance; but the less confident one can be that it, in fact, serves the public interest.

It has been sound common practice to set requirements at more than one tier ("defense-in-depth"). However, requirements have usually been set more stringently on means than on the higher tiers, which leads facilities to be preoccupied with satisfying onerous prescriptions, and to neglect risks that these do not address. The reverse practice, where requirements are progressively loosened at lower tiers, is more promising it, since it assures that the facility's dominant priority is the ultimate goal of safety. Within reasonable limits, the facility can then use its discretion to decide how most conveniently to achieve that goal. Intermediate requirements are set comparably.

Requirements can only be enforced if *compliance* with them can be checked, and this demands realistic assessment of performance of the item in question. The prevailing practice of what has been called probabilistic risk assessment (PRA) suffers from certain limitations, at least as a decision aid, that newer so-called comprehensive safety assessment (CSA) might avoid.

- CSA formulate uncertainty as a regulators personal probability; whereas PRA produces frequencies, difficult for regulators to interpret and use for unique events.
- CSA draws on all available knowledge, however soft, including the regulators experience and judgments, and on multiple approaches to an assessment, including PRA; whereas PRA draws only on well-documented data (such as experiments or panels of experts) in a single model.
- CSA addresses all recognized sources of risk (including intangibles, such as safety culture); whereas PRA addresses only those risks that can be well documented (such as equipment failure).
- On the other hand, the overtly subjective element in CSA often makes it more vulnerable than PRA to accusations of bias and manipulation in controversial cases.

Subjectivity in comprehensive safety assessment should decline over time, as more impersonal methods for documenting elusive risks mature. This would reduce, and perhaps eventually close, the gap between CSA and PRA. However, for now, decider-specific subjective judgment is unavoidable in responsible regulatory decisions.

## 1. Introduction



### 1.1. Problem Background

Even in advanced societies, the regulation of risk or its converse, safety--is often seriously flawed and may lead to environmental disaster, as at Chernobyl, or immense economic waste, as in the US nuclear waste program. Hazards are misdiagnosed, inadequately protected against, and such protection as there is, is unnecessarily costly. For example, the nuclear industry argues that the US Nuclear Regulatory Commission (NRC) imposes needlessly costly safety measures that threaten the survival of the nations nuclear energy option. At the same time, public interest groups argue that NRC fails to require prudent and cost-beneficial measures. There is merit in both complaints.

The main causes are inappropriate public policy and research to support it. The first focuses on controlling the means of safety management rather than the ends. The second relies only

on documented evidence, often in the form of what is commonly called Probabilistic Risk Assessment (PRA), thereby neglecting much important knowledge.

This chapter seeks to develop guidance for environmental regulators responsible for authorizing hazardous activities, much of the guidance at variance with conventional wisdom. The argument is illustrated primarily in the context of US nuclear reactor regulation, which is already at the technological leading edge of safety assessment, but it is applicable, in principle, to other safety regulation.

### **1.1.1. The Social Purpose of Safety Regulation**

Left to themselves, people can harm the rest of society. Thus, individuals drive dangerously, incinerators pollute, and power companies operate unsafe reactors. In response, society empowers regulators to limit such activities. For example, NRC sets safety requirements for reactor operations and closes them down if it finds that requirements are not met.

Clarity about social objectives should clearly precede determining which regulatory practice will achieve them. Ideally, private rights and public interests will be balanced, but this rarely happens. For example, the regulator charged with environmental protection may act as a single-minded advocate of safety, in effect acting both as prosecutor and judge.

The ultimate purpose of environmental regulation is surely not to minimize risk in any absolute sense, but to serve public interest. Regulatory actions have to take into account conflicting societal interests, including the cost and availability of energy and business viability. For example, if society values a statistical life at US\$5 million industry dollars, the regulatory process should result in industry spending up to US\$5 million to save an expected lifewhether or not regulation explicitly requires it to do so.

### **1.1.2. Deficiencies in Common Practice**

Critics contend that, in practice, environmental regulation is often costly, inconsistent and fails to protect the environment adequatelyall at the same time. Consider recent US examples:

- NRC declared a reactor to be dangerously unsafe and required costly improvements, although a major PRA indicated it to be among the safest in the United States.
- Local Alaskan regulators denied oil company BP permission to build a causeway into the Arctic Ocean, on the grounds of harm to the fish. A Republican administration overruled them on the grounds of public interest. A Democratic Congress protested improper "political interference" by the administration. BP abandoned the causeway, claiming that the permitting process was too unpredictable.
- Environmental Protection Agency (EPA) regulation required a nuclear repository to isolate radioactive waste for 10 000 years, with 90% probability. There was broad scientific consensus that, based on current evidence, a Nevada site met that test. Yet, the Department of Energy spent many years and over a billion dollars to make sure, while waste piled up dangerously.

Critics on all sides agree that current practice needs improvingbut not on how.

### 1.1.3. Policy Criteria

A sound regulatory policy has several desiderata.

*Sound regulator judgment:* It must go without saying that good regulatory practice should abide by the best judgment of the regulator; let us call him or her R. This requires that R:

- make use of any knowledge R already uses, augmented by any new knowledge;
- draw on value judgments that R adopts; and
- combine knowledge and values logically.

In an ideal world, all parties to a decision would share the same knowledge and values and come to the same conclusion. Failing that ideal, we will try to help all parties to make up their own minds responsibly on the basis of whatever they know and value, by developing tools to aid Rs judgment.

*Institutional fit:* Good regulatory practice needs to take institutional realities into account. In particular, decisions and the grounds for them must be legally and politically acceptable, which suggests that the rationale should be transparent and reviewable by others. Institutional fit may interfere with sound judgment by excluding important *unsubstantiated* knowledge and lead to what has been called "organizational foolishness," unacceptable to a reasonable individual.

Other institutionally desiderata are:

- to be able to take account of the cognitive needs and capacities of all parties;
- to adapt to existing practices, rather than make radical and disruptive change; and
- to be predictable by a potential "regulatee."

### 1.1.4. Distinct Regulatory Tasks

There are three distinct regulatory responsibilities to be aided:

- define objectives
- specify what is required of regulatees
- verify compliance.

Although these tasks are distinct, they interact. For example, deciding what requirements to specify depends on how feasible it is to verify compliance with them.

Different organizations and people within them may perform these three roles. The legislature may determine ends, senior agency officials may specify requirements, and agency staff may verify compliance. Several parties in these categories may be involved in making the final enforcement decision, as in the earlier oil construction example. Their positions will inevitably differ, whether aided or not.

## 1.2. Methodological Perspectives

Two significant decision-aiding paradigms are available (possibly complementing each other).

### **1.2.1. "Impersonal" Decision Analysis**

A paradigm that has dominated safety management practice seeks findings that are impersonal, in the sense that they do not depend on the subjective judgment of any individuals. It relies only on documented knowledge and is limited in scope to issues on which some documented knowledge is available.

For example, the costbenefit analysis widespread since the 1960s has focused on costs and benefits that can be measured with little controversy. Similarly, the dominant safety assessment practice since the mid 1970s has been PRA, which considers only documented data thus following classical statistical traditions.

Reliance on "objective" or impersonal findings is well suited to public scrutiny and judicial proceeding. On the other hand it does not do justice to the experience and perceptiveness of competent people active in the field. It may also show consistent biases compared with sound professional judgment. For example, costbenefit analysis has typically undercounted benefits of environmental protection, which tend to be more intangible and less objectively measurable than costs. This favors business interests, by appearing to show that promising regulations are not cost-beneficial. Conversely, PRA may overstate safety, by omitting undocumented sources of risk.

### **1.2.2. "Personal" Decision Analysis**

An alternative and in some ways complementary paradigm for evaluating regulatory (and other) options is that of statistical decision theory and its practical implementation personal decision analysis (PDA), which is basically the formalization of professional judgment.

Any reasonable informal argument can be accommodated and enhanced logically with a quantitative PDA model. The general methodology is well established. It is based on the quantification of personal uncertainty and value in the form of subjective probability and utility and derives their logical implications for action.

## **1.3. Structure of Remaining Chapter**

Section 2 addresses how requirements should be set, considering both their form (e.g., to require some acceptable risk) and content (e.g., where to set the acceptable risk). It draws on the causal linkages between regulatee activities (means) and social ends, and suggests how to model them. Section 3 addresses how to verify compliance with such requirements, through safety assessment. Section 4 offers some concluding comments, including discussion of political issues and technical trends.

## **2. How To Set Requirements**



### **2.1. Basic Issues**

### 2.1.1. Theory versus Practice

Evaluating regulatory options has two parts: assessing the impact of alternative safety management practices and assessing the impact of regulation on those practices. *Requiring* a safety management practice does not assure that it will be *adopted* as Three Mile Island (TMI) and Chernobyl testify.

The social meta-problem is to devise a regulatory regime that, *as implemented*, leads to "regulatee" action that serves the public interest. Specifying requirements that, *if met*, would serve public interest may not achieve it, if regulated behavior does not conform. Regulatees who manage to persuade the regulator that they comply with a requirement may not actually do so. Conversely, they may "play it safe" by exceeding requirements.

### 2.1.2. Relevant versus Verifiable Requirements

Effective regulation requires both what is important and what is economically enforceable. Regulatory requirements may focus on controlling readily verified features of a regulated activity (like hardware reliability), but neglect features that better protect the public (like a good corporate culture). Beware of "looking for your keys under the street lamp, rather than in the shadows where you lost them." What is needed is a torchlight to help search the shadows of poorly illuminated risksanalytic tools that draw thoroughly on intangible experience and judgment.

### 2.1.3. Quantitative, Qualitative, and Ambiguous Requirements

Regulatory requirements may be qualitative, at differing degrees of specificity (e.g., low: do not degrade fish habitat, or high: locate a repository in a specific geologic medium). Alternatively, requirements may be quantitative (e.g., allow at most one oil spill in ten years; or assure that groundwater travel from a nuclear repository will not exceed 1000 years).

Qualitative and quantitative performance requirements (especially probabilistic) have complementary appeals. Qualitative requirements are fairly simple for regulator and regulatee to use, but may be subject to controversial interpretation. Quantitative requirements involve more effort, delay, and expertise, but are more verifiable. Both have a role to play. Qualitative requirements can set the stage for operational interpretation in quantitative terms. This chapters attention is primarily on numerical requirements, but analogous reasoning applies to qualitative requirements.

Current law and regulation often specify requirements that are not only qualitative, but also ambiguous for example, they may require that risks be "as low as reasonably achievable" (ALARA). This is imprecise enough to give the regulator discretion to adapt to improved knowledge and changing circumstances (like relaxing requirements if energy becomes scarce), but that discretion is also open to inconsistency, arbitrariness, and abuse.

R may also take into account requirements additional to those specified in regulation i.e. to serve unstated objectives. Some of these objectives may be perfectly legitimate, like addressing deficiencies in regulation (e.g., undervaluing short-term consequences in nuclear waste disposal). Other objectives may be less defensible (e.g., avoiding the bureaucratic



certainly be dangerous to rely it on alone, and that is virtually never done. Here and elsewhere additional requirements are specified.

### 2.2.3. Means: Specific Measures

A more common alternative is to go to the other extreme of prescriptive requirements, corresponding to the bottom of the Figure 1. R specifies action that regulatees are required to take. For example, NRC requires light water reactor vessels to conform to a certain engineering design. Nuclear waste disposal regulation requires that an acceptable repository meet certain specifications.

*Problems:* Means requirements alone often have consequences no better than public interest alone, in both costs and benefits to society. Although a means requirement can be verified, its results are more elusive, and important health and safety considerations may be omitted. Furthermore, R may be technically less qualified than the regulatee to judge what is safe practice. According to experienced industry safety engineers, NRC orders to shut down reactors may increase risk, rather than decrease it..

There are many alternative ways of achieving any given level of safety. Decreasing one risk may compensate for increasing another. R is usually neither competent nor motivated to pick the safety measure that best meets social needs *other than* safety (e.g. economic). It may, for example, be cheaper to enhance safety by improving security procedures than by installing an additional power generator (though NRC may require the latter because it is easier for inspectors to verify).

A common argument for not setting requirements at high tiers of performance is that "they are too uncertain." However, there is no technical reason why the issue of uncertainty should stand in the way of high-tier performance requirements, provided these are expressed probabilistically. Prescribing specific action simply moves the uncertainty from whether there is compliance to what compliance will achieve. Both will improve with better data and analysis.

### 2.2.4. Trading Off Sub-Ends

A given public interest requirement is implied by requirements one tier below in the hierarchy (i.e., on public health, industry economics, and other social interests), along with value tradeoffs among these sub-ends. Published safety goals for reactor regulation, for example, require spending up to US\$1000 to avoid one man-rem emission, and value a death at US\$1 million. The latter tradeoff has also been used at EPA to guide measures in hazardous laboratory experiments.

*Problems:* It is not clear who should make the value tradeoff judgments. Although it might seem reasonable for an elected legislature to represent public taste, it would be politically dangerous for a legislator to commit himself, say, to a specific value of a life.

Even if a costbenefit rule that incorporated value tradeoffs were adopted, it would be difficult to assess the actual impacts a proposed activity would have on cost and benefit sub-ends. For example, controversial technical judgment would be needed to verify an industry

claim that so many lives will be saved. A regulatee usually has more resources to argue for its special interests than the regulator has to second-guess their argument. (Remember how the defense dominated the prosecution in legal firepower in the O.J. Simpson murder trial?) Regulatory rulings can be challenged endlessly, regardless of cost, if noncompliance cannot be clearly demonstrated.

### **2.2.5. Requirement on One Sub-End: Acceptable Risk**

It is common practice to set requirements on just the environmental protection sub-end(s) of public interest, say, in the form of acceptable public health risk. For example, official regulatory guidance specifies that the probability of a prompt fatality in the vicinity of a nuclear plant should not exceed one in a million. The regulation of healthcare facilities requires that a nursing home be scored on a number of safety dimensions such as fire protection and human access where the maximum score for each reflects its relative importance. The sum of these scores must exceed some threshold, which implies value tradeoffs.

*The acceptability of acceptable risk.* As noted, setting a requirement on safety alone disregards at least explicitly tradeoffs with other social interests. Some decision analysts advise against "acceptable risk" as a basis for regulation, because it ignores criteria other than safety (like prosperity). The proper degree of safety thus varies with context. Nevertheless, a fixed acceptable risk can be an appropriate guide for safety choices that are homogeneous in those other respects (provided its context dependence is recognized). Nuclear waste siting may be such a case.

Society is not a unitary actor. Multiple players, including legislature and business, are involved in setting and enforcing compliance with regulations. Acceptable risk has bureaucratic appeal, in that it is compact, easily understood, and lends itself administratively to division of labor between complementary safety regulators (like NRC) and agencies (like the US Office of Management and Budget).

### **2.2.6. Requirements at Intermediate Means/Ends Tiers**

Between the two extremes of requiring the public interest (or its immediate determinants) and prescribing means, there is the middle ground of specifying acceptable risk at an intermediate tier. Acceptable risk could be "one in 10 000 probability of core melt", or "10% chance that radioactive releases from a waste repository will reach the environment in 10 000 years". At a lower tier, acceptable risk could be some probability of component failure. To the extent that the requirements at any one tier are comprehensive (which they may not be), they imply, along with tradeoff judgments, some requirement at the next tier up. (At the top two tiers, requirements on health, economics, etc. imply some requirement on public interest). The arguments for acceptable risk at one tier over another are intermediate between those for ends and means. The lower the requirement, the more verifiable it is, but the less relevant.

### **2.2.7. Motivations Affecting Choice of Requirement Tier**

The location of a requirement in the means/ends hierarchy has important motivational

implications. The lower the tier, the less the regulatees are motivated to act responsibly. The regulatees attention is focused on what they can be readily held accountable for, such as engineering features. As a result they usually control these well. The more the regulatees are pre-occupied with meeting narrow requirements, the less attention they can afford (or feel obliged) to devote to less easily enforced requirements (such as organizational effectiveness). Serious reactor incidents, like TMI and Chernobyl, are more often due to poor organization than poor engineering.

More generally, performance that is not required is liable to be neglected if it interferes with the regulatees other interests, like meeting production goals. Moreover, regulatory staff will not be motivated to exercise responsible discretion in pressing for safety measures that are not specifically required, nor to relax those measures that *are* required, even if that is socially desirable.

A prime example is requiring acceptable accident risk rather than acceptable public safety, which is one tier up. Utilities typically expect NRC to evaluate their safety primarily by the probability of an accident (e.g. core melt), rather than, say, minimizing population dosage if an accident occurs even if that protects public health more cost-effectively. They are thus motivated to concentrate their scarce safety resources on reducing accidents rather than minimizing their consequences. Since accidents can be evaluated more readily, than, say, population evacuation in the event of an accident, it is better controlled. To some extent R can compensate informally for inadequate requirements. For example R can exercise managerial discretion in interpreting ambiguous requirements, thereby mitigating dysfunctional effects. O R can informally promote a practice that is not required, such as a safety culture.

### **2.3. Hybrid Requirements**

R may set requirements at several tiers, instead of just one, in the meansends hierarchy, as a kind of "defense-in-depth", on the grounds that incremental assessment effort at any one tier has decreasing effect (i.e. "decreasing marginal utility"). It is more cost-effective to spread effort among several tiers than to devote the same effort to one tier. This is not uncommon practice. The NRC has a health goal as well as three separate lower-tier goals for accident, containment, and release impact (e.g., less than one in 10 000 accident probability). NRC also specifies three-tier plant security: a general security performance objective, a physical system that protects against a design basis threat, and some specific prescriptions.

#### **2.3.1. "Progressive Tightening" Mix**

Under much current regulatory practice, requirements are set more strictly as they get more prescriptive (i.e., lower in the meansends hierarchy). This effectively causes means that prescription dominates safety regulation. This "ratcheting" is a major source of industry frustration. Lower-tier requirements tend to be set so high (whether by regulation or at the initiative of inspection staff) that meeting them assures that the regulatee will *exceed* higher-tier requirements.

For example, the lower-tier NRC accident, containment and release requirements imply a probability of individual death well below the higher-tier fatality requirement. The

probability of fatality is approximately the probability of an accident, times probability of release given accident, times probability of a death given release. The three latter acceptable probabilities are commonly set at one in 10 000, 10, and 1000, respectively, which implies an acceptable probability of a fatality at one in 100 million. In fact, regulation sets acceptable probability of a fatality at one in a million; i.e. acceptable risk at the lower tier is 100 times more stringent than at the higher tier.

There can be a plausible rationale for controlling lower order variables more stringently than is needed to meet higher-tier requirements. For example, avoiding an accident has value beyond public safety, in terms of industry economic interests and NRC political interest. (Although the Three Mile Island core damage accident in 1969 had negligible health consequences, it caused severe public relations problems for NRC.)

On the other hand, it might plausibly be argued that accident risk need not be regulated at all, since hardheaded business considerations will force regulatees to attend adequately to safety. R could then concentrate on containment, siting, and other safety requirements that influence the *consequences* of an accident. Since these protect society's interests rather than the utility's interests, there is greater need for regulation to promote them. However, society may not trust industry even to act in its own best interest without regulatory constraint.

### **2.3.2. "Progressive Loosening" Mix**

A more promising version of hybrid requirement is to make requirements progressively *less* stringent, as they go down the figure 1 hierarchy from ends to means. Weak requirements are set at low tiers, but regulatees have the flexibility to decide which to exceed in order to meet higher-tier requirements. Unlike progressive tightening, this practice appropriately focuses safety management attention on ends rather than means.

For example, accident risk, release containment and release lethality may be set at modest levels that do not jointly meet a tougher public health requirement, one tier up. Therefore, at least one of them must be raised, but the regulatee can choose which. At a still lower tier, any combination of core cooling design and cooling procedure that assures acceptable overall core cooling system performance (see Figure 1) would be acceptable (provided their weaker individual requirements are met).

*Implication:* Progressive loosening has better prospects of just meeting top-level requirements

## **2.4. Formal Validation**

The above informal regulatory argument can be quantified and validated in a PDA model of the means-ends risk hierarchy. This permits requirements to be specified more precisely.

### **2.4.1. Quantifying the Model**

Each term in the Figure 1 hierarchy of risks can be defined as a variable expressed as a function of variables at the tier below (e.g., a product in the third tier). Any variable, whether input directly or inferred from others, is normally uncertain, except at the bottom (and

increasingly uncertain further up the hierarchy). Variables are represented as personal probability distributions, reflecting all available knowledge and judgment, however soft.

*Inferring tier probabilities from tiers below.* Probability for a variable is, in principle, inferable from the (joint) probability distribution of the variables at the tier below, and, by extension, to all probabilities in the sub tree below them. For example, as noted, the chance of a fatality depends jointly on the chances of reactor accident, containment failure, and deaths per release. The inference is only exact if the linking functions are identities. The functions can be made identities by including an error term. For example, in Figure 1, if there are determinants of public health (at tier 2) other than accident, release and lethality (at tier 3), it would be covered by an error term, "public health other than deaths." Another error term could cover modeling error (i.e., where an inaccurate function links determinants).

#### **2.4.2. Compensation Among Variables**

Overall safety depends on the interplay of sources of risk that can compensate for each other through the tradeoff functions at each tier. In principle, therefore, a single requirement could be strong enough to make other requirements at the same tier unnecessary. For example, any public health requirement (tier 2) could be met (at tier 3) by setting accident safety requirements strict enough, without any release requirement. Similarly, requirements only need to be set at one tier to assure that all higher-tier requirements will be met provided all requirements at the lower tier are specified. However, some redundancy of requirements is needed to allow for imperfect compliance (see next section).

#### **2.4.3. Specifying versus Achieving Compliance with Requirements**

The meansends model traces the consequences of safety management (means) as practiced, and of the performance of intermediate risk variables that the means lead to. Setting regulatory requirements on these risk determinants does not assure that the requirements will be met. Even if compliance with equipment reliability and human performance requirements (tier 5) *would* assure acceptable accident risk (tier 4), each may not perform as required.

In order to evaluate any set of regulatory *requirements*, their *actual* impact on risk determinants must be assessed. Thus, a formal evaluation of regulatory requirements would call for modeling a layer of causal linkages between requirements and behavior. Such analytic complexity is not explicitly addressed in figure 1, and may be best handled informally.

### **3. How to Verify Compliance with Requirements**



Once regulatory requirements have been specified, it remains for R to determine whether the hazardous activity complies with them. At present, R normally, and quite properly, uses "best judgment" in determining and justifying that compliance. In other words, everything R knows or can find out tangible and intangible is logically taken into account in R's judgment. We now address how R might make that judgment.

#### **3.1. The Meaning of Compliance**

### 3.1.1. Interpreting Safety Assessment

The meaning of complying with a "means" requirements is straightforward: R merely has to be satisfied that the regulatee has done what was required. Requirements set at any higher tier of the meansends hierarchy are normally probabilistic. R must confirm that realistic probability distributions at that tier are acceptable (i.e., meet requirements). What constitutes a "realistic" probabilistic safety assessment (PSA) has two alternative interpretations, according to intended use. (Both obey the formal rules of statistical theory).

### 3.1.2. Desiderata of Safety Assessment for Public Scrutiny

If decisions have to withstand public or adversarial scrutiny, assessment needs certain properties:

- The assessment rationale is transparent enough to be reviewed by others and revised to reflect alternative positions
- Conclusions are supported as far as possible by "impersonal" data, not liable to manipulation by interested parties

The latter data would include documented knowledge, derived from a replicable research process, such as well-documented expert elicitation, or be uncontroversial, in the sense that any reasonable person would come to the same conclusion based on available evidence.

### 3.1.3. Desiderata of Safety Assessment for Decision Purposes

If the object is to help a particular R decide responsibly whether an activity is in compliance, Rs safety assessments needs the following properties:

- They represent Rs personal judgment, based on all data however soft available (or potentially available)
- The realism of the judgment is enhanced by statistical and decision theory logic.
- In any given circumstance, both purposes may operate.

## 3.2. Safety Assessment Paradigms

The major distinction between appropriate approaches to the needs of public scrutiny and professional decision is between impersonal and personal decision analysis, respectively, mentioned earlier. They are typified, respectively, by documented and judgmental safety assessment (which itself can incorporate the former into a comprehensive safety assessment).

### 3.2.1. Documented Safety Assessment (DSA)

*Documented safety assessment (DSA)*

DSA basically limits consideration to documented evidence (such as experimental data on component failure and structured expert panels). It also addresses only those sources of risk on which documented evidence is yet available. For example, it has often omitted risks

initiated externally to the reactor, such as from earthquakes. Gaps in existing documentation are filled, if necessary, by assumption (e.g., that undocumented and analogous documented risks are the same). Thus DSA disregards judgmental knowledge, if the judgment is based solely on Rs undocumented experience, even if it includes Rs considered judgment on how well assumptions fit reality. DSA seeks scientifically validated statements designed to withstand peer review and satisfy institutional need for authenticated assessments. They do not necessarily meet the desiderata for decision purposes. .

The 1975 Rasmussen Reactor Safety Study (WASH-1400) initiated what has been called *probabilistic risk assessment* (PRA), which has dominated the safety assessment of engineered systems, such as nuclear reactors, ever since. Indeed, it is mandated by and illustrated in official US regulatory guidance documents. PRA is a special case of DSA, with some distinctive features. It is basically an extension of a well-established reliability engineering practice, and its exponents have been predominantly scientists and engineers. This practice deals with events more naturally viewed as repeatable (such as component failure) and as having a researchable frequency, than, say, a unique reactor accident (see, [Appendix 1](#)). PRA characterizes *all* uncertainty as frequency.

Exponents of PRA have been extending its scope by the use of unconventional evidence for nonrepetitive sources of safety, such as reviewably elicited expert opinion. The reasonable ultimate goal of PRA, as of DSA more generally, is to be so comprehensive as to render unnecessary any additional *undocumented* input. Since PRA appears to be the only version of DSA in current use, in this chapter PRA and DSA will be used interchangeably.

### 3.2.2. Judgmental Safety Assessment (JSA)

However, until that happens, a responsible R will make use of all the knowledge available, however intangible, to answer unavoidable questions of the form: "What can I say now about whether this reactor will have an accident within a year?"

Judgmental safety assessment (JSA) involves modeling the judgment however formed of a practicing regulator (or other safety manager). This judgment is in the form of a personal probability of an undesirable event (such as an accident at a particular facility). It is based on the evidence, typically undocumented, that R could (and would) draw on in the course of habitual managerial practice. They would include informal observation and intuition based on experience, as well as any considerations at the back of Rs mind.

Rs undocumented judgment may be validated by no more than professional credentials, as with expert testimony in court. However, JSA is in principle explicable, through well-known personal decision theory, which can also make Rs judgment sounder. In principle, any informal statements of judgment can be modeled in formal, personal probability terms. For example:

- "I used to think no unsophisticated intruder could break into a reactor, until it almost happened," could be represented by a "Bayesian updating" model
- "Radioactive release through human intrusion depends on whether an economic use for nuclear waste is found," can be represented by a "conditioned assessment" probability model.

Although published official regulation is typically expressed in impersonal probability terms, some regulatory requirements are already consistent with personal probability. US Nuclear waste disposal regulation, for example, calls for an acceptable "likelihood" of radioactive release. In fact, impersonal probability can be interpreted as personal probability that is so conclusive that all reasonable people would adopt it. Though practically unachievable, it is sometimes a useful construct, as discussed in Appendix.

Jurist Stephen Breyer has suggested to the author that an assessment whose rationale is explicit can be accepted as not "arbitrary and capricious" in legal proceedings. The old adage "Garbage-in-garbage-out" may be quite truedoubtful inputs produce doubtful outputs. However, processing that garbage properly (i.e., logically inferring output from input) can make it more usable by ensuring that no garbage is added in-between.

### **3.2.3. Comprehensive Safety Assessment (CSA)**

Although not necessarily limited to JSA in forming a judgment, R may, in practice, be reduced to doing so unless some sound means of incorporating DSA can be found. In any case, a responsible R will surely want the assessment to be as firmly and defensibly based as possible. R will therefore want to enhance unaided judgment (JSA) by taking account of any DSA or other data, in what we will call comprehensive safety assessment (CSA). Its methodology, as with JSA, is grounded in established decision theory. However, it can also take DSA into account intuitively, provided R fully understands and can allow for any deficiencies (e.g., unrealistic assumptions).

CSA may serve to structure Rs unaided thinking, make best use of DSA and other new evidence, and guide the search for new data by directing the search in the most productive directions. Unlike pure DSA, it is an enhancement of (rather than an alternative to) Rs unaided judgment. R should not automatically substitute CSA for his own direct best judgment JSA but only consider it as an alternative perspective on safety. CSA technique is not yet mature enough to guarantee improvement over structured or unstructured JSA.

## **3.3. DSA for Regulatory Decisions**

### **3.3.1. DSA as an Approximation to CSA**

JSA adds nothing significant to DSA, for example where the only source of risk is a well-researched engineering process, such as the failure of a standard component. DSA can then serve as CSA and the distinction is practically unimportant. However, even then, R must confirm explicitly that *in his/her judgment* there are no grounds to take issue with the DSA. DSA gets closer and closer to CSA, for a ever-wider range of risky situations, as DSA matures (i.e., as documented knowledge on less readily researchable sources of risk, such as organizational failure, is developed). On the other hand, in many cases (e.g., nuclear waste disposal) there are so many elusive sources of risk that probably little of use that can *ever* be said will come from hard data.

### **3.3.2. Realism of Default Assumptions**

Some logically essential elements of an assessment on which no documented data are

available have to be handled by assumption in DSA. To avoid the *appearance* of judgment, these are typically chosen to be tractable (e.g., a "null" assumption), but not necessarily realistic. Some common assumptions are "conservative" (i.e. pessimistic), for example that no recovery from an accident will occur. However, in practice the dominant effect of assumptions (typically implicit) is over-optimism, for example:

- any risk that is not documented is zero (e.g., risk from external events was ignored in a PRA put forward as a "model" in an NRC regulatory guide)
- failures are independent, neglecting common causes.

Some other assumptions are neutral (but not necessarily realistic), for example:

- this plant is similar to other plants that have been studied
- experimental data is representative of the target process (e.g., failures in the field equate to failures in the laboratory).

### 3.3.3. Avoiding "Uncertainty"

A common defense of DSA is that the JSA alternative, incorporating personal judgment into a risk model, "introduces uncertainty." This is a misconception: Rs uncertainty about any given risk is irreducible, due to limited knowledge, regardless of how it is taken into account, quantitatively or otherwise. Modeling only those parts of the problem that are relatively "certain" does not make the remaining uncertainty go away; it simply conceals it from view. In principle, DSA is no less subject to uncertainty than JSA or CSA. The claim that empirical documentation assures objectivity is a delusion. Indeed, DSAs for the same case have been known to differ by two or more orders of magnitude (due possibly to different arbitrary assumptions).

### 3.3.4. Regulator Confidence

"Above all, do no harm." (Hippocrates). NRC regulators known to this author are generally reluctant to use DSA as a realistic sole test of compliance. (In a variant of an old anecdote, a renowned DSA exponent is asked if he uses DSA on a problem of his own. He answers: "Certainly not! This is an important problem".) Official guidance documents, such as NUREG 1050 and INSAG-3, have not advocated that PRA (DSA) be used generally to verify regulatory compliance.

A recent US\$4 million PRA showed that a certain reactor was one of the safest in the United States, but an NRC regulator disagreed and put it on a "watch list" of dangerous plants. The R observed here, for example, that the PRA omitted intangible sources of risk R knew to be important. He also noted that an alarming number of near misses ("accident precursors") were ignored in the PRA. G. Apostolakis has reported that PRAs (DSAs) do not always distinguish between plants with good and poor organization.

### 3.3.5. Technician Confidence

Many in the community of PRA (DSA) practitioners are well aware of its deficiencies and, in particular, agree that DSA will not do, as it stands, to evaluate overall plant safety

realistically. S. Hirschberg, for example, has advocated rejecting PSAs that produce "unrealistic" findings, in particular accident probabilities below one in a million, on the grounds that too many sources of risk must have been omitted. The common response, however, has been to urge only that the gaps be closed within the DSA paradigm, not that judgment be used to fill them in the mean-time.

### 3.3.6. Motivation

A danger of overlooking the incompleteness of DSA, whether temporary or permanent, is that total risk is understated (which may be a motive for interested parties to keep undocumented risks excluded). More subtly, it will encourage licensees to focus attention on those risk sources that are modeled (e.g., hardware rather than people problems). If utilities believe the regulator will hold them to stricter account on those measures, it may distort the allocation of safety resources. This is particularly dangerous if it focuses utility attention on reducing accident risk rather than its public consequence (which may account for reactors sometimes being located in heavily populated areas).

### 3.3.7. Decision Aiding Potential

*Technical evolution:* In recent years there has been a determined and promising effort in the PRA community to bring more and more sources of risks and methods of assessing them within the scope of DSA. This has already included use of assessment of human reliability and the impact of organizational factors, the incorporation of accident precursors into safety assessment, and reliance on plural evaluation ("multiple lines of evidence").

The contentious aspects of DSA may be mitigated through a set and reviewable procedure for selecting and eliciting expert opinion (e.g., with carefully chosen panels). Though not exactly objective, the reviewability of such a procedure reduces its subjective nature and therefore its susceptibility to manipulation by interested parties. It thus may be more politically acceptable with controversial regulatory decisions. But however impressive progress along these lines may be, it will be a long time before this DSA core can substitute for fully comprehensive safety assessment (CSA). Any attempt to do so is conceptually and practically misguided. As Granger Morgan has noted, "good policy calls for bad science." By "bad science" he presumably means findings that have not (yet) met the rigorous tests they would need, to be accepted as "scientific."

In the long run, DSA and CSA approaches may converge, as enough empirical data is developed for virtually all sources of risk to be addressed in DSA, and for most personal judgment to be swamped in CSA by the DSA component. But, as M. Keynes said, in the long run we are all dead. The distinction between DSA and CSA should always be kept clear, lest we be tempted to interpret DSA as CSA before the bridge between them is completed.

### 3.3.8. Decision Applications

R may feel quite comfortable with PRA for many purposes *other than* realistic safety assessment. DSA has been successfully used in many such applications, for example to prioritize hardware and other risks on which there is substantial empirical data that swamps

softer data. It is useful in evaluating the relative safety of generic reactor designs, which are largely unaffected by regulatee-specific risks (say from deficient management practice), or by other sources of risk that are likely to affect contending designs equally (say external incidents). T.E. Murley notes: "One should view a PRA (DSA) estimate of the core damage frequency for a plant more as figure of merit for that plants design, rather than a prediction of the actual likelihood of a core damage accident at that plant." S. Hirschberg also favors PSA (DSA) for plant-specific prioritization of contributors to risk.

*Other uses of DSA:* When reporting findings and justifying action to the public, institutional pressures may lead R to conceal subjective elements, misrepresent them, or even act at variance with them. Here, we are only concerned with what R personally thinks and favors, or can be led to accept (with such aid as we can offer). If DSA happens to confirm Rs best judgment, derived by other means, so much the better; it can be used as a persuasion device.

### 3.4. CSA as integration of DSA and JSA

Faced with a choice between accepting DSA, as is, or informal judgment (JSA), R will typically choose the latter, but can, in principle, do better than either.

Although DSA provides a promising core for realistic safety assessment, it is intrinsically incomplete, without informed judgment, at least as far as helping R come to a conclusion. As that core expands and increasingly dominates Rs actionable judgment, the better that DSA core can stand alone and withstand legal or other outside challenge. However, DSA is not yet ready as a realistic decision-aiding tool for a broad range of problems, and perhaps never will be.

#### 3.4.1. Graphic Illustration

Figure 2 shows (with illustration from reactor safety) what kinds of hard and soft knowledge are to be combined, but it does not address how.



Figure 2. Elements of comprehensive safety assessment

The target assessment of CSA, at the top of the figure, is Rs final probability distribution on the "hazard of interest" (number of deaths caused by an accident in a given year). Hard (i.e., documented) elements appear in the left column, undocumented in the right column. The three top tiers of the figure refer to safety assessments at tiers 25 of the meansends hierarchy in Figure 1. The bottom tier represents primary data, documented or undocumented.

Tier two distinguishes those major sources of risk on which there is some documented knowledge (e.g., reactor accident) from those where there is none (e.g., population exposure

in the event of an accident). All such variables (including deaths given exposure), taken together, determine the target (deaths).

Tier three shows documented and undocumented causes of accident risk (e.g., component failure and management error, respectively unless there happens to be documented knowledge on management error). Note that documented risks may have undocumented, as well as documented causes.

As in Figure 1, probabilistic assessments on variables at each tier, along with a functions linking them (e.g., multiplicative), imply corresponding assessments in the tier above, taking into account any uncertainty about the linking function.

Assessment of each of the tier three variables is based on various kinds of evidence, shown in the bottom tier. Arrows confirm that, by definition, assessments of undocumented causes are based only on judgment, itself based on soft data or other evidence (e.g., NRC inspectors observation of slack maintenance practices bears on Rs judgment of serious management failings). On the other hand, assessment of documented causes may draw on both hard and soft knowledge. For example, R may judgmentally adjust more or less plausible assumptions bearing on documented causes (e.g., component robustness experiments and the assumption that the component manufacturing process is stationary may be modified by the judgment that the process has deteriorated since the experiments were made).

How these items of evidence are combined will vary. Within decision analysis methodology, for example, alternative approaches to assessing the same variable would be pooled or reconciled. Appendix 2 shows how an integration of DSA and JSA of this kind might be quantified.

### **3.4.2. Reactor Backfit Case**

A backfit evaluation was needed for a pressing NRC decision on whether to close down a certain reactor or to require that a costly safety feature be installed. The author and colleagues conducted a CSA for R, in this case a senior NRC official.

It built on an earlier large-scale DSA, whose action implications were not clear to R. It had produced an "expected core damage frequency" of  $5E-5$  (i.e., 5:100 000), which comfortably met a published safety goal of E-4. However, it addressed only internally initiated events (like a pipe break) and assumed (among other things) that plant management followed all required safety procedures. A 90% "uncertainty range" of  $2E-5$  to E-4 was reported (still meeting the safety goal), but it reflected uncertainty only about certain parts of the DSA model (i.e. parameter values).

R had judged this exercise not to be very useful, as it stood. He was reluctant to accept its optimistic output and was also not clear how to "splice" in other knowledge and opinion that he had access to. For example, the DSA made no use of the fact that there had been a number of near misses (accident precursors) that R was aware of. In fact, R was considering closing down the reactor, on the basis largely on inspection reports and on his or others personal judgment.

The CSA performed extended the DSA model to attempt to address all potential sources of risk, including operator error and externally initiated events, and also allowed for the possibility of modeling error. The assessments were elicited from NRC staff familiar with the plant and from research contractors who had worked on the original DSA. "Frequency" in the DSA was reinterpreted as an ideal impersonal probability (see above). Rs own judgment was introduced by means of an interactive computer program, which allowed him to second-guess and override specific inputs.

Based on this exercise, R assessed the probability of core damage to be  $3E-5$ , with a margin of uncertainty of  $3E-6$  to  $2E-4$ . Since the upper end of the distribution failed the E-4 safety goal R decided to close down the plant pending either a safety "backfit" or additional evidence that would bring the uncertainty range entirely within the safety goal. This conclusion was consistent with Rs intuition, though it was more detailed. (Contrasting DSA approaches to comparable backfitting decisions have been reported, e.g. by the International Atomic Energy Agency).

### 3.4.3. Cost Implications

A common, if understandable, misconception is that, because CSA extends the scope of DSA, it must be more costly. Paradoxically, the opposite is usually the case. CSA can be aggregated at whatever level available resources and timeliness permits. Default inputs are judgmental, based on whatever the assessor knows, including any already available DSAs. When a pressing regulatory decision has to be made, the supporting CSA will typically be done with a few man-weeks of effort, whereas DSAs commonly take at least several months.

The CSA for the above reactor case cost under US\$50 000. It built on a DSA that cost US\$4 million. Had that DSA not been available, Rs assessment would have been limited to a JSA that directly models Rs knowledge (perhaps augmented by the judgment of subordinate managers and staff). A "quick-and-dirty" CSA, which might still be better than nothing, could consist of R assessing the proportion of total risk omitted in a DSA and adjusting it accordingly. (The informal reasoning to be modeled might be something like "DSAs for plants of this type typically understate total risk by 50200%, but this DSA seems more complete than most, so lets say 30150% for performance evaluation purposes." This is a great margin of error, but not atypical.)

### 3.4.4. Methodological Implications

The central methodological issue is what knowledge base to use (i.e., how much of the possibly relevant knowledge available should be drawn upon in making the assessment?) The position taken here is that it should include everything that a sensible R would normally take into account, however ill documented. By all means supplement or override this knowledge with harder data when and if you can. But when an assessment has to be made (say to inform a current policy decision), do not ignore any knowledge available at the time.

## 4. Concluding Comments



### 4.1. Development needed

Much remains to be done to make the ideas presented here fully operational as tools for specifying or verifying requirements, primarily for methodological research on the integration of DSA and JSA into CSA, in parallel with the independent development of DSA, JSA and CSA. For example, the informal guidance on specifying a desirable mix of requirements at different tiers presented in Section 2 can be made rigorous and reviewable by implementing and quantifying a model along the lines of Figure 1.

#### **4.1.1. DSA Methods**

There is much to be said for continuing to develop DSA technically along lines that its exponents are well aware of and already working on. The PRA (DSA) enterprise has a highly developed technology, with an experienced community of specialists proficient in reliability engineering. It would be disruptive and unnecessary to re-orient it in directions that call for quite different expertise (e.g., subjective probability elicitation, typically found among decision analysts). In time, no doubt, the exponents of JSA and DSA will gain enough mastery of each others skills that CSA will become a unified field, much as statistics is coming to encompass classical frequentists and personal probabilists, after several decades of separate development. For now, both DSA and CSA can be used to support managerial safety evaluation: DSA presents a relatively researchable piece of the evaluation; CSA attempts to model the whole evaluation.

#### **4.1.2. "Assessment Uncertainty"**

There are other issues, not addressed here, that are worth developing. Regulations can reasonably require and sometimes do not only that a current safety assessment should be acceptable, but that there be some assurance that new evidence will not undo that acceptability. This higher-order compliance may be specified informally, as "reasonable assurance," as in nuclear waste disposal regulation (US Code of Federal Regulations 10 CFR, part 191), or it can be specified formally, in terms of a second order probability distribution. A comparison of assessment uncertainty between corresponding elements of DSA and JSA would also be used to weight them for pooling purposes in CSA.

Specifying and verifying such "assessment uncertainty" requirements calls for ambitious analysis and challenging personal judgment. It involves assessing how firm individual elements are, or conversely, how liable to revision in the face of new evidence. The assessment uncertainty methodology to do this is still primitive.

#### **4.2. Political Issues**

In Utopia, a wise, responsible, and all-powerful ruler could prescribe in detail the action citizens should take to assure the best balance of society's interests, in any field of endeavor. For example, a judge could decide on a defendant's guilt or innocence on his own, after weighing all the evidence. By the same token, a regulator could specify what safety-related action would produce the greatest public good (e.g., by considering the causal linkages in Figure 1). Since we do not live in Utopia and do not trust our betters with such power, society has opted for a largely adversarial system. Representatives for competing interests analogous to counsels for defense and prosecution battle it out. If we are lucky, the resolution approximates what the Utopian dictator would have determined.

In safety regulation, we have a flawed example of this adversarial practice. Regulators act to some extent as advocates for health and safety; industry fights for its economic interests. However, regulators often have controlling power (NRC can close down nuclear plants if it wants), which puts them in the awkward position of being both prosecutor and judge. Alternatively, the government in power may tip the balance, one way or the other. For example, the Republican Quayle Commission on Competitiveness in the United States in the 1980s served as both defendant and judge, by paying predominant attention to the interests of industry.

To some extent, it may be that power is now allocated among competing advocates to assure a constructive tension that produces a reasonable balance for society. Alternatively, some new super-agency might be established as an arbiter. It would acknowledge the regulatory agency's health and safety orientation legitimate, but would have discretion to trade off the competing public interests and overrule that agency, if necessary.

### 4.3. Paradigm Shifts

This chapter has suggested a regulatory philosophy intended to help a regulator to serve the public interest properly. Although it has been oriented toward helping an individual regulator make responsible decisions, official directives limit his discretion. They include formal regulations and published guidelines, and less formal guidance programs. Clearly, the practical implementation of this philosophy will be contingent on official encouragement indeed direction to practicing regulators. In turn this should stimulate change within the safety management community, including the training of regulatory professionals and support technicians.

Adopting the perspective presented in this chapter would appear to have the following action implications, each representing a major change in common practice:

- *For legislators.* Regulatory objectives would be defined in terms of social interest and comprehensive knowledge.
- *For environmental agency heads.* Requirements would be set probabilistically on a hierarchy of means and ends, with compliance verified in terms of responsible staff judgment. Compliance tests would be more stringent for ends than means.
- *For regulatory staff.* Compliance determinations would be based on staff judgment, substantiated with as much documentation as possible, and a logically articulated rationale.
- *For safety managers.* They would maximize probability of compliance by basing safety management decisions on comprehensive analysis and knowledge, documented as thoroughly as possible.

### Acknowledgements



The work on which this chapter is based was supported, in part, by the Nuclear Regulatory Agency, under contract NRC-04-90-369. The author thanks Dr. T.E. Murley, previously Head of the NRC Office of Nuclear Reactor Regulation, for many stimulating discussions. The opinions expressed, here, however, are those of the author alone.

[Appendix 1: Objective Risk?](#)

[Appendix 2: Quantified Illustrations of Comprehensive Safety Assessment](#)

**Related Chapters** 

*Related Links will be activated soon!*

**Glossary** 

**Compliance:** Determination that a regulatory requirement is satisfied.

**Costbenefit analysis:** Weighing the documented pros and cons of a choice.

**CSA:** Comprehensive safety assessment; probabilistic and based on all hard and soft available knowledge.

**DSA:** Documented safety assessment; probabilistic and based solely on validated evidence.

**Ideal probability:** Personal probability based on unlimited, but realistic knowledge.

**Impersonal (assessment, etc.):** Not specific to any individual.

**JSA:** Judgmental or personal safety assessment.

**Meansends hierarchy:** A causal scheme of risk determinants, with public interest(s) at the top and specific regulatee action at the bottom.

**NRC:** US Nuclear Regulatory Commission.

**PDA:** Personal decision analysis, attributed to an individual.

**Personal probability:** Attributed to an individuals judgment, with decision theory interpretation.

**Plural evaluation:** Making the same evaluation different ways.

**PRA:** Probabilistic risk assessment DSA as traditionally practiced by the US safety assessment community.

**PSA:** Probabilistic safety assessmentsafety assessment expressed in probabilistic terms.

**R:** Regulator, the decider whose judgment is to be modeled.

**Regulatee:** The regulated facility or activity.

**Requirement:** A constraint on regulatee behavior, embodied in regulation.

**Risk:** The prospect of an undesirable incident.

**Tier:** Level in a meansend hierarchy.

**Bibliography** 

Ahearne JF. (1999). Responsibilities of a probabilistic safety analyst. *Journal of Risk Research* 2(4), October. [Top-level perspective of an influential member of the US Nuclear Regulatory Commission.]

Apostolakis G. (1992). The concept of probability safety assessments of technological safety. *Reliability Engineering and System Safety* 38, 326. [The position of a leading academic advocate of documented safety assessment.]

Breyer S. (1994). *Breaking the Vicious Cycle: Towards Effective Risk Regulation*. Cambridge, MA: Harvard UP. [Policy paper by a leading jurist, later Supreme Court justice.]

Brown R.V. How decision makers are misled when quantitative methods are misused. *Interfaces*, in press. [Cautionary tales by decision analysis practitioner.]

Brown R.V. and Ulvila J.W. (1988). Does a reactor need a safety backfit? Case study on communicating

- decision and safety analysis information to managers. *Risk Analysis* **8**(2), 271282. [Example of an application of the approach suggested in this chapter.]
- Carnegie Commission on Science, Technology and Government. *Risk and Environment: Improving Regulatory Decision Making*. (1993). New York: Carnegie Commission. 150 pages. [Authoritative evaluation of regulatory practice, by leading figures in law, government, business and academia].
- Covello V.T. and Mumpower J. (1985). Risk analysis and safety management: an historical perspective. *Risk Analysis* **5**(2).
- Fischhoff B. (1994). Acceptable safety: a conceptual proposal. *Risk: Health, Safety and the Environment*, winter, 328. [Thoughts of a leading behavioral scholar of safety analysis.]
- Hirschberg S. (1992). Prospects for probabilistic safety assessment. *Nuclear Safety* **33**(3), JulySeptember, 365377. [An authoritative European perspective.]
- Kaplan S. and Garrick B.J. (1981). On the quantitative definition of safety. *Risk Analysis* **1**, 1127. [A traditional DSA view.]
- March J.G. and Shapira Z. (1982). Behavioral decision theory and organizational decision theory. *Decision-making: An Interdisciplinary Inquiry* (ed. G.R. Ungson and D.N. Braunstein), Boston: Kent Publishing Co. pp.92-115. [A seminal reference on the sources, prevalence, and consequences of "organizational foolishness."]
- McGarity T.O. (1987). Regulatory analysis and regulatory reform. 65. *Texas Law Review*. pp 1243-8. [A seminal work that discusses the use of costbenefit analysis in legal decisions.]
- Murley T.E. (1999). *The Role of the Nuclear Regulator in Promoting and Evaluating Safety Culture*. Nuclear Energy Agency, Paris: Office of Economic Cooperative Development, June. [Reflections on the need to take account of a critical source of undocumented safety by an ex-head of reactor regulation at NRC.]
- National Research Council. (1983). *Risk Assessment in the Federal Government; Managing the Process*. Washington, D.C.: National Academy Press.
- Pratt J.W., Raiffa H., and Schlaifer R. (1995). *Introduction to Statistical Decision Theory*. Cambridge, MA: MIT Press. [Basic reference on the theory and algorithms of applied decision analysis.]
- Quade, E.S. *Analysis for Public Decisions*. (1982). New York: North-Holland. 380 pp. [Broader range of analytic methods than addressed here]
- Raiffa, H. (1968). *Decision Analysis: Introductory Lectures on Choices Under Uncertainty*. Reading, MA: Addison-Wesley. 96 pp. [A classic, nontechnical text on the logic of decision analysis by a founder of the field.]
- Rasmussen N.C. *Reactor Safety Study: An Assessment of Accident Risks in US Commercial Nuclear Power Plants*. (1975). Ref: NUREG-75/014. WASH-1400. Washington DC: US Nuclear Regulatory Commission [The pioneering study that launched PRA as a dominant policy tool.]
- Ruckelshaus W.D. (1983). Science. *Science, Risk and Public Policy* **221**(4615), 10261028. [An insightful perspective of an ex-administrator of the US Environmental Protection Agency.]
- US Nuclear Regulatory Commission. (1986). *Safety Goals for the Operation of Nuclear Power Plants: Policy Statement and Republication*. Federal Register vol. 51 pp 30028-33 (162), Washington DC: Government Printing Office. pp. [Official quantitative guidelines for setting nuclear regulatory requirements.]
- US Nuclear Regulatory Commission. (1996). *Probabilistic Risk Assessment Policy Statement* Washington DC: Government Printing Office. Federal Register Vol.60 FR p. 42622. [Official guidance on the use of PRA in nuclear regulation.]
- Wu J.S, Apostolakis G., and Okrent D. (1991) On the inclusion of organization and management influences in PSA of nuclear plants, in *The Analysis, Communication and Perception of Risk*. Garrick BJ and Gekler WC (eds.) New York: Plenum Press, pp 429-439. [An innovative attempt to incorporate a major, but hitherto undocumented, source of safety into documented safety assessment.]

## Biographical Sketches



**Rex Brown** has spent 40 years alternating policy consulting with research. As a consultant he advised senior environmental regulators on decision aiding methods, and this experience has centrally informed this chapter. Other clients include executives in government and business on foreign and domestic public policy, international business, safety management, regulation, capital investment, defense, education, and law. As an academic, he has held university appointments in public policy, management, statistics, psychology, economics, and systems engineering. His background is interdisciplinary, with a focus on modeling judgment, prescriptive decision theory, and the human context of applying it.

**Education:** Harvard DBA (Statistical Decision Theory), Incorporated Statistician (UK), Cambridge BA and MA (Economics and Anthropology).

**Academic appointments:** George Mason University (Distinguished Senior Fellow, School of Public Policy; Research Professor of Systems Engineering), University of Michigan (Business), and Harvard Business School (Managerial Economics and Marketing). Visiting faculty: London School of Economics (Psychology), Carnegie-Mellon (Social and Decision Sciences), University College London (Statistics), and Cambridge U. (Management). Faculty Associate: Dartmouth College (Arctic Studies).

**Research:** Books: *Tools of Rational Choice* (in press), *Decision Analysis for the Manager*, *Research and the Credibility of Estimates*, *Teaching Decision Skills to Adolescents*, *Marketing Research and Information Systems*. Some 80 archival papers in *The Statistician*, *Journal of Risk and Uncertainty*, *Acta Psychologica*, *Theory and Decision*, *Annals of Operations Research*, *Journal of the Royal Statistical Society*. Topics: decision models and aids, environmental management, survey methods, program evaluation, risk analysis, market research, organizational behavior, teaching aids, inference, estimation, AI, use of science.

**Application:** Chairman of Decision Science Consortium, Inc. (consulting and research). Employee of Decisions and Designs, Inc. and METRA UK (industrial market research). Executive clients have been in nuclear regulation, environmental management, resource allocation, capital investment, business policy, industrial marketing, sales forecasting, foreign and public policy, battle and intelligence management, legislation.

**Credits and dissemination:** British Institute of Statisticians Annual Award, 1967. Social Science Research Council (UK) Senior Fellowship, 1976. Twenty-five NSF, NIH, ONR, etc. research grants/contracts. Editorial Board *Decision, Risk and Policy*, *Journal of Decision Analysis*, and *Journal of Behavioral Decision Making*; peer review panels. Expert witness in judicial and legislative proceedings. Listed among most cited authors in judgment and decision research (White and Griffiths, 1981). Articles in *Harvard Business Review*, *Washington Post*, and other media. Guest on National Public Radio and local TV. Subject of articles in *Chance* and *Illustrated London News*.

## EOLSS - REQUIREMENTS AND COMPLIANCE

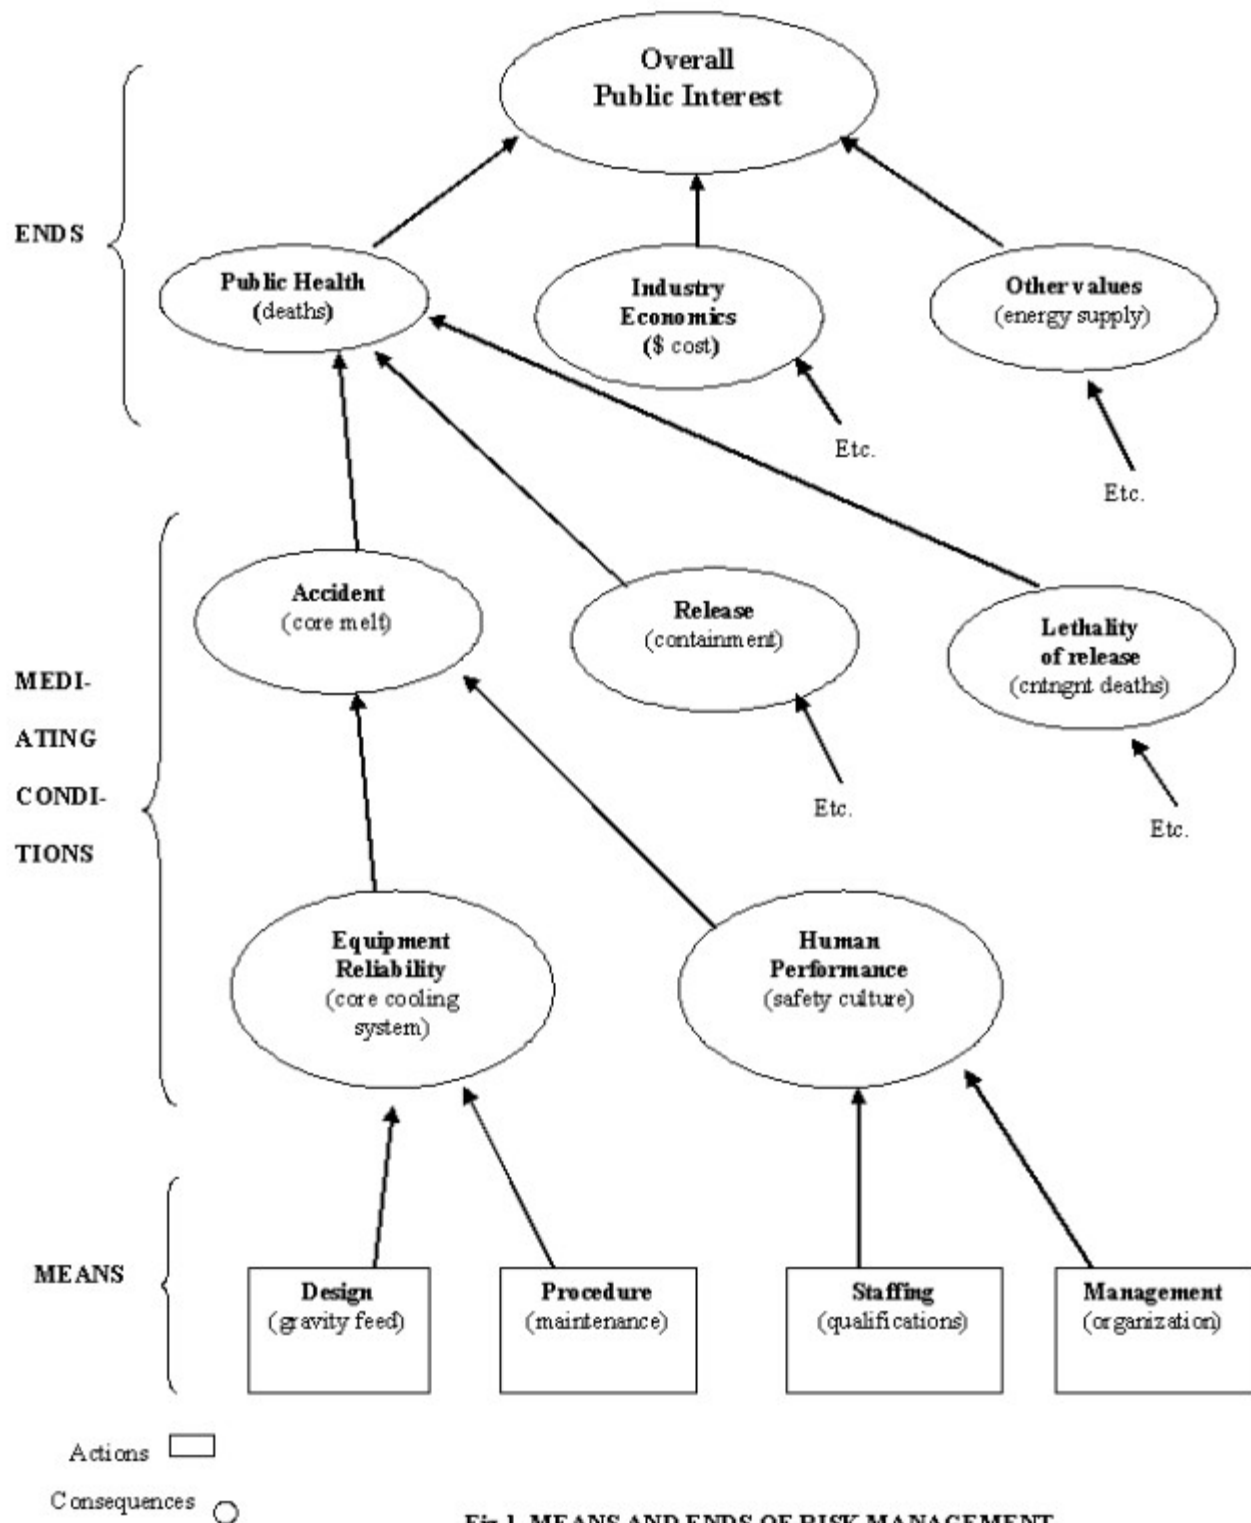
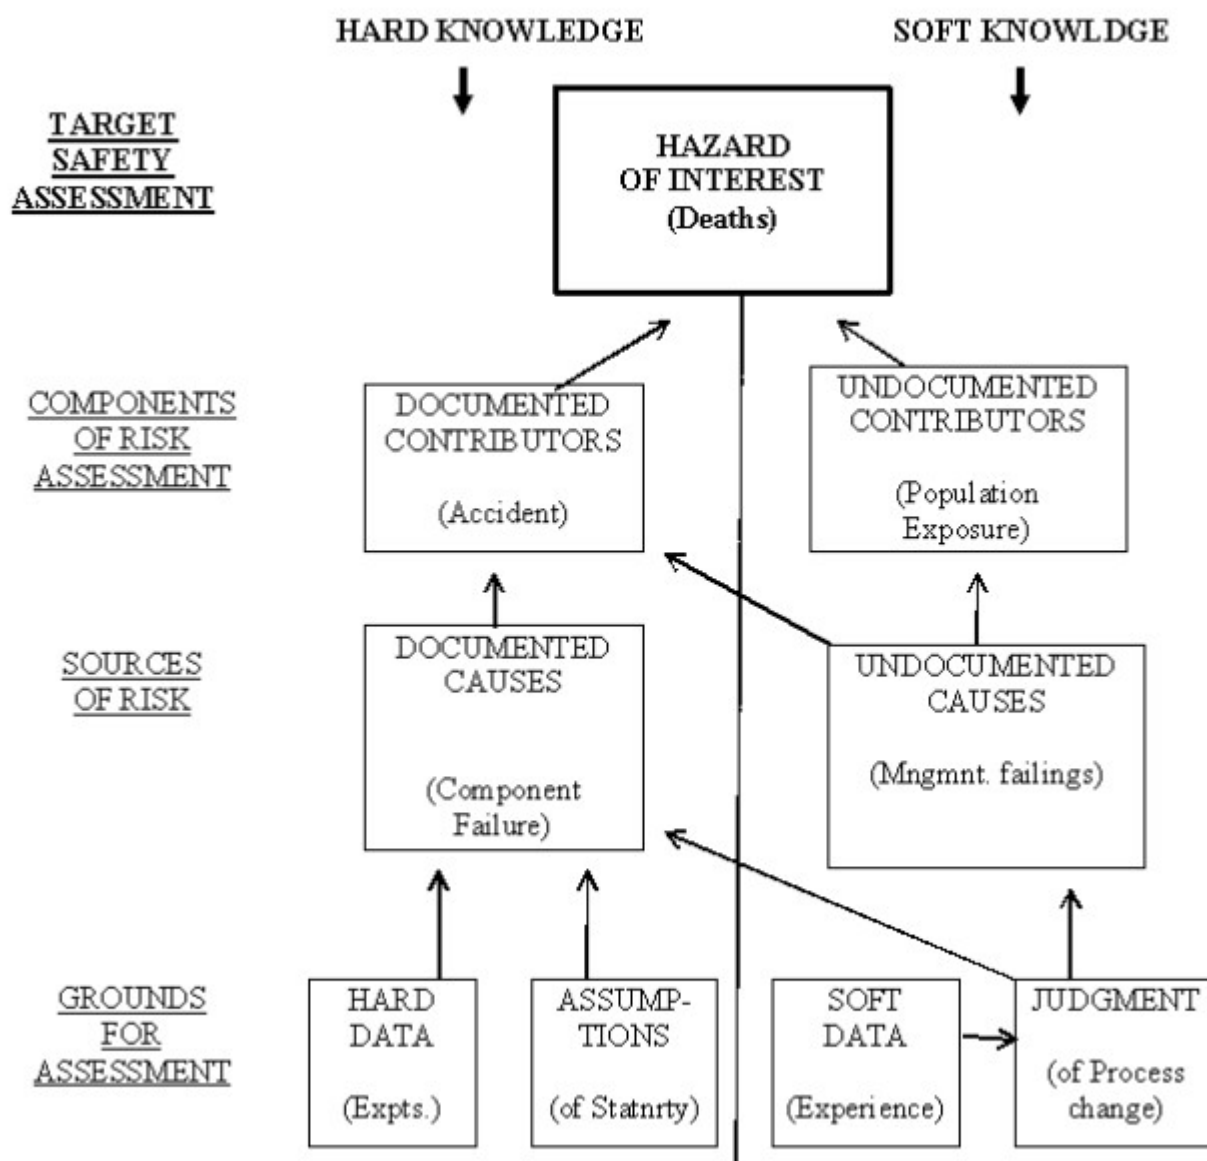


Fig 1. MEANS AND ENDS OF RISK MANAGEMENT

(Reactor safety example)

## EOLSS - REQUIREMENTS AND COMPLIANCE



**Figure 2. ELEMENTS OF COMPREHENSIVE PSA**  
(Reactor safety example)