

SHOULD ENVIRONMENTAL REGULATIONS REQUIRE SPECIFIC ACTION, ACCEPTABLE RISK, OR COST-EFFECTIVENESS?¹

Rex V. Brown
Institute of Public Policy
George Mason University

Abstract

Environmental regulations are often in a form that results in needless cost *and* inadequate social protection. The narrowest form of regulation requires some specific action; the broadest forms require that the activity be cost-beneficial to society. Requiring that risk at different levels be acceptable lies between those extremes. Broad requirements are more socially relevant, but narrow requirements are more verifiable. Hybrid regulations are common, but the narrowest requirements are normally made the most stringent. We propose an alternative regulatory approach, which would better balance the health, economic and other interests of society. It is based on the reverse principle, i.e. that the broadest requirements are made the most stringent. Requirements would be set and verified according to an unconventional extension of Probabilistic Risk Assessment. The approach is validated and implemented by means of a hierarchical personal probability model of the means and ends of social policy. The argument is illustrated from nuclear and oil regulation.

Key words: regulatory requirements, acceptable risk, decision analysis

¹ The ideas in this paper were first presented in a consulting report to the Director of the Office of Nuclear Reactor Regulation, Dr. T.E. Murley (Brown, 1994).

Introduction

Background

Many argue that US regulation is costly, inconsistent and fails to protect the environment adequately – all at the same time. There are certainly many controversial or anomalous cases. Here are some examples:

NRC declared a reactor to be dangerously unsafe and demanded costly improvements, although according to a costly Probabilistic Risk Assessment it was among the safest in the US.

Local Alaskan regulators denied an oil company permission to build a causeway into the Arctic Ocean, on the grounds it would harm the fish. Bush administration overruled them on Public Interest grounds. Democratic Congress balked, claiming political interference. BP abandoned the causeway: a permit was too unpredictable.

EPA requires a nuclear dumpsite to isolate radioactive waste for 10,000 years, with 90% probability. There is consensus that a Nevada site meets that test comfortably, at least on current evidence. Yet DOE is still spending many years and a billion dollars to make sure (but may not be motivated to find contrary evidence). Meanwhile the waste is dangerously exposed (relatively).

Problem to be addressed

Function of regulation. Left to themselves, people may do harm to the rest of society. Cars are driven dangerously, advertisers mislead, incinerators pollute and power companies operate unsafe reactors. So regulators are given power to put limits on regulatees' activity, supposedly to ensure that private rights and public interests are acceptably balanced. For example, NRC puts safety requirements on reactors and closes them down if requirements are not met.

The regulatee may be required to take (or avoid) specific action, or to meet acceptable risk or performance standards, or to trade off costs and benefits to serve the public interest. Regulatory requirements may be qualitative (e.g. "safeguard public health!" or "do not degrade fish habitat!"). Or they may be quantitative (e.g. assure at most 1 in 10,000 chance of core melt; at most one oil spill in ten years.)ⁱ

The purpose of environmental regulation should not be to minimize risk in any absolute sense. Conflicting societal interests, like the cost and availability of energy, and business viability are also to be served.

Deficiencies in current regulation. To be effective, regulation should require what is both desirable and enforceable. Much regulatory practice falls short in one or both respects. As a result, economic resources are wasted without necessarily protecting the environment

adequately. Nuclear power industry advocates argue that NRC imposes needlessly costly safety measures that threaten the survival of the nation's nuclear energy option. On the other hand, public interest critics argue that NRC fails to require prudent and cost-beneficial measures. There may be merit in both complaints.

Relevant vs. verifiable requirements. Regulatory failings often have to do with conflict between what performance is important and what can be verified. Regulatory requirements have typically focused on controlling readily verified features of a regulated facility (like frequent reports and other simple but burdensome activity). They neglect other features that may more efficiently protect safety (like a responsible organizational culture), but are more difficult to check up on. For example, nuclear plant evaluations (PRAs) have typically assessed only the risks that are readily measured (like hardware reliability), and even then have relied only on hard evidence (like experiments or documented history)--thereby overlooking important soft data. This is "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 risks—i.e. analytic tools that draw thoroughly on experience and judgment.

Issues to be addressed. To specify and enforce regulatory requirements more effectively, certain issues need to be addressed. What constitutes successful regulation? What requirements should be set? How should they be defined and verified? What guidance, constraints and authority should be given to regulators? In particular, what should regulators require of facilities regulated and how can they ensure that these requirements are met? These issues interact: what is required is influenced by what can be enforced.

Illustrative case. Nuclear safety regulation provides most of the illustrative material for the theses in this paper, though the arguments are generalizable to other environmental regulation.

Choosing among types of requirement

What is successful regulation? The object of regulation is not to minimize harm. The regulator, it is true, is typically charged with attending to a particular public interest, notably environmental protection, and may be viewed as (and behave as) a single-minded advocate. However, regulatory actions have to be constrained to take economics and other affected interests into account. Otherwise the regulator is both prosecutor and judge (and he does, in fact, often take on both roles).

The object is surely to maximize public interest. When we come to evaluate how well a regulatory system is working we want it to meet some cost beneficial test. Suppose society equates a statistical life to five million industry dollars. Then, when regulation has worked its magic, we want power companies to be spending up to \$5 million to save a life – whether or not they are specifically required to.

What is actually required may take several forms.

Requiring that the public interest be explicitly optimized. It would be tempting to require regulatees to simply serve the public interest, by observing specified cost-benefit value

trade-offs, and then to check that they have done so satisfactorily. Federal navigable waterways statutes are sometimes of this form (section 404 of CFR xxx).

However, as a practical matter, this is extremely difficult to implement effectively, and dangerous to rely on, on its own. Who would make the value judgments that trade off, say, safety against economic welfare? It is a brave (or foolhardy) congressman who would dare urge publicly that a human life is worth \$5 million, or any other specific sum (Porter, 1991).

Even if a cost-benefit rule were adopted, it would be difficult to enforce, i.e. to establish reliably and comprehensively what the impacts on cost and benefit of a proposed activity would be. (For example, too much controversial judgment is needed to verify a regulatee's claim that so many lives are saved). A regulated corporation has more resources to produce a plausible argument favoring its special interests than the regulator has to second guess the argument on society's behalf. (Remember the imbalance of legal firepower in *The People vs. OJ Simpson?*). Regulatory rulings can be endlessly challenged, regardless of cost, since compliance cannot be unambiguously checked.

Micromanaging regulatee actions. So why not go to the other extreme of prescriptiveness? Have the regulator determine what action it considers cost beneficial and then require companies to do it. For example, NRC requires reactor vessels to meet certain engineering design specifications (10CFR?)

The practical consequences of this form of requirement are often even worse, in terms of both societal cost and benefit. The requirements can be verified, to be sure; but more elusive and perhaps more important health and safety considerations may not be addressed. The regulatee's attention is focused on what he can be readily held accountable for, such as engineering specifications. These may be well controlled, but at the expense of less easily enforced requirements, such as organizational effectiveness. However, serious reactor incidents are more often due to poor organization than poor engineering, as witnessed by TMI and Chernobyl.

Furthermore, the regulator may be less qualified than the regulatee to judge technically what is safe practice. According to experienced industry safety engineers, an NRC order to shut down a reactor, may increase rather than decrease the risk of an accident. More significantly, there are many ways of achieving any given level of safety (or other aspects of plant performance). Decreasing one source of risk may adequately compensate for increasing another. The regulator is neither competent nor motivated to pick the safety measure that best meets the regulatee's needs other than safety, like cost. It may 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).

Intermediate requirement: requiring acceptable risk. Between the two extremes of requiring the public good and detailed prescription, there is a middle ground: to specify

acceptable risk. I.e. some low probability of an undesirable event (such as a component failure or a radioactive release) is mandated. Ideally, the threshold probability would be set to prompt cost-beneficial regulatee action. However, there is no guarantee that it will do so—even if the level of risk could be verified (which is still problematic)

Alternative levels of risk to be controlled. The risks whose acceptability is required can be specified at different levels, which trade off relevance and importance against verifiability. High levels of safety performance (e.g. low accident probability) are what society really cares about, but they are difficult to measure. Lower levels of performance, (e.g. system reliability and core cooling effectiveness) are more verifiable, but less relevant.

For example, a low-level requirement that "component failure frequency is to be less than 5%" is easier to verify than the higher level requirement "probability of radioactive release is to be less than 10%". To meet higher level requirements, the regulatee has to persuade the regulator that the appropriate safety goal has been met (a negative), but he can use his own discretion and resources to figure how that is best or most cheaply done (a positive). [*Ratcheting*]

Considerations in choosing among requirements.

The alternative forms of requirement lie on a continuum of ends versus means, with public interest at the top, specific action at the bottom, and acceptable risk, at various levels, in between. Our task is to pick one—or more—requirements to regulate on, in the hierarchy that implies (see figure 1).

[FIGURE 1 ABOUT HERE: MEANS-ENDS HIERARCHY FOR REACTOR SAFETY]

Motivational effects. Controlling risk, even informally, at the top level provides a safeguard against anti-social industry priorities. A regulatee is motivated to focus its safety resources on whatever it believes it will be held accountable for. This encourages him to neglect other risk precautions -- except to the extent that his economic interests dictate.

A prime example is accident risk, as measured by core damage "frequency", the most common indicator of plant risk reported to NRC. Utilities typically expect their safety to be evaluated primarily by this core damage risk. Then they are motivated to concentrate their scarce safety resources on reducing that, at the expense of, say, containing emissions in the event of an accident or minimizing population dosage given an emission. However, the latter requirements may protect public health more cost-effectively. Furthermore, if a risk assessment mainly reflects hardware risk, utilities will be motivated to save effort on "soft" but critical factors, like organizational culture.

Regulation should surely focus on assuring those safety safeguards that the regulatee is not particularly motivated to provide, like a remote plant location. In practice, most regulatory attention goes to controlling other risks, which can be evaluated more readily, like reactor accident. However since that danger has the worst economic consequences for the company, it may be the one the regulator need worry about least.

Uncertainty and probability. A common argument for not setting requirements at high levels 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-level performance requirements, provided

they are expressed probabilistically. Prescribing specific action simply moves the uncertainty from whether there is compliance to uncertainty about what compliance will achieve. Both will improve with better data and analysis.

Multiple requirements. A common regulatory response to these choices among levels of requirement is “all of the above”, or as it is sometimes put “defense in depth”. NRC makes requirements at several levels for reactor safety. For example, it requires all of the following: specific maintenance procedures; less than 1/10,000 probability of core melt; and spending \$1000 to avoid one man-rem from being emitted. The trouble is that there is no defensible rationale for how stringently these requirements are to be set, which leads, for example, to well-founded industry accusations of “ratcheting”, i.e. erring on the side of maximum stringency.

Current "tightening" mix of levels. Under current regulatory practice, requirements seem to become stricter as they get more prescriptive, which is a major source of industry frustration. Lower level requirements tend to be set so high (whether by regulation or by inspection staff pressures) that, if a licensee meets them, he is bound to *exceed* higher level requirements.

For example the lower-level accident, containment and release consequence goals in USNRC (1986) imply a probability of individual death well below the published higher-level health goal. That probability 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 1/10,000, 1/10 and 1/1000, respectively, which implies an acceptable probability of a death is 1/100M. In fact, regulation sets that at 1/1M. I.e. the acceptability limits at the lower level is 100 times more stringent than at the higher.

I will argue below that a mixed strategy is the way to go, but with stringency *decreasing* with lower performance levels--a critical difference from common practice.

Conceptual framework

A rationale for the design of a set of regulatory requirements can be based the rational analysis of choice and uncertainty embodied in “decision analysis”. Any relevant knowledge and judgment available to the decider is quantitatively modeled and their logical action implications are inferred (Raiffa, 1968, Pratt et. al 1995).

A means-ends model of public interest and risk

The above qualitative discussion, for example, can be cast in the following format. Given that we wish the regulatee to take action that serves the public interest, what are the causal linkages between action and public interest? And where inside such a hierarchical network of means and ends should requirements be set?

Simple example: reactor safety. Figure 1 shows such a hierarchy (schematically and incompletely) for reactor safety. It is a relatively simple case, where the regulated hazard is a single possible event, in this case an accident. More complex cases, which involve multiple hazards, such as various environmental impacts can be handled comparably.

Factors determining regulatory consequence. The ultimate end, the overall public interest, sits alone at the top. Directly controllable measures or actions that may affect public interest – the means - are at the bottom. In between are the various stages through which the controllable means exercise their influence on the public interest ends.

Examples, going down the chain, would be:

- public interest
- potential harm to be regulated, such as to health (e.g. fatality).
- probability of accident (e.g. core melt)
- an untoward event (such as equipment failure) and the human response to that event;
- specific actions (like instituting a maintenance procedure)

Performance at any one of these levels is completely determined by performance at the level below, provided all determinants at that level are addressed. For example, the probability of a fatality depends jointly on the probabilities of reactor accident, containment failure, and deaths per release.

Quantitative model

Deterministic form. Although this causal hierarchy can be left in a qualitative form, it can be expressed precisely and quantitatively. Each “factor” can be defined as a variable, which can be expressed as an exact function of variables at the level below. To make it an exact identity, an error term (“other components”) may need to be included.

Probabilistic model. In any given situation to be evaluated, the values of variables, at any level, are not normally known with certainty. Any realistic statement about them must therefore be intrinsically probabilistic, increasingly so as one moves up from means to ends.

All variable probabilities are inferable from the (joint) probability distribution of the variables at the level below, and, by extension, to all probabilities in the subtree below².

Personal interpretation of model output. The model needed is one that assesses realistically the socially relevant consequences of the regulated activity. E.g. It derives the risk and social value (if any) of a given reactor, operated in a certain way. Its purpose is to enhance and make explicit what a public policy maker knows or can learn which bears on his taking a responsible regulatory position.

This means that the model as implemented should reflect all available knowledge and judgment, no matter how “soft”. Any probabilities figuring in the model should be realistic measures of uncertainty about what is being assessed. They are personal probabilities as commonly used in decision analysis. This modeling perspective is not generally adopted within the risk analysis community.

² By the well established statistical theory of the distribution of functions of random variables.

Verifying compliance with requirements

Requirements must be enforceable to be useful. The above philosophy would indicate that verification should use *all* available knowledge and professional judgment. However, to demonstrate compliance, for legalistic regulatory reasons, the required probabilities, etc., would need also to be supported by authoritative data, or at least a legitimized assessment procedure (Brown, 1990).

“Probabilistic Risk Assessment” (PRA)

Traditional PRA addresses only sources of risk on which there are hard data (however soft their relation to reality), and it disregards other knowledge even about those sources. Currently PRA is limited to modeling risk as a function of what can be "objectively" documented, e.g. with observed frequencies (USNRC-1990). It is adapted from reliability engineering practice, which typically deals with replicable hardware systems, where this reasonably captures what is relevant. However, important sources of reactor risk do not fit this mold well, notably human deficiencies. Perhaps PRA should be read as “Partial Reliability Analysis”, reserving PSA for Probabilistic Safety Analysis?

Comprehensive Uncertainty Evaluation (CUE)

A more comprehensive form of PSA than PRA would help a regulator assess *his* probability of a unique hazard. It addresses all sources of risk, and integrates the assessor's best judgment with PRA and any other information, however tentative or controversial. The resulting probability should be the most realistic an assessor can make. (It may not be a frequency, as in traditional reliability models). Its irreducible subjectivities may still make it unpersuasive to others, but they must nevertheless still cope with the same uncertainties. I like to call it CUE—for Comprehensive Uncertainty Evaluation—to underscore the point that an uncertainty evaluation can be based on any diagnostic cue, however weak (Brown, 1998).

Although comprehensive in scope, a CUE assessment need not be costly. It can be aggregated at whatever level resources and timeliness require. Default inputs are judgmental, based on whatever the assessor knows, including available PRAs. One such safety backfit evaluation for NRC, cost under \$50,000 (Brown and Ulvila 1988). It modeled directly the knowledge of a team charged with deciding if a reactor is acceptably safe, in just a few days.

Acceptable risk requirements

Within the means-ends hierarchical model, one can, in principle, pick any probabilistic variables at any level, and specify “performance requirements”; i.e. put acceptability constraints on their distributions. An acceptable risk could be a requirement that a hazardous event have less than a specified probability of occurrence, e.g. 1/10,000 of core melt. Or the requirement could be that some quantity not exceed some limit with specified probability, e.g. 10% chance that radioactive releases from a nuclear dump reach the environment in 10,000 years. (The lowest “means” level of requirement is not probabilistic because the regulatee controls it -- it is prescribed action).

For any completely specified means-ends model setting acceptability requirements on all variables at effectively implies an acceptability requirement at all higher levels.

Implications for regulatory requirements

Compensation among requirements.

As the above modeling framework indicates, overall safety depends on the interplay of many sources of risk that can compensate for each other. A major advantage of performance requirements at high levels is that they allow tradeoffs among requirements at a lower level

The tradeoff could be formalized by a combination rule; i.e. some function of individual performance requirements must meet a certain threshold. This has some of the plusses of a higher performance requirement, without giving up verifiability. A precedent (which does *not*, however, include probabilistic requirements) is provided by fire protection regulations for Nursing Homes. Different safety features of the home are awarded points and a license is granted if total points exceed some threshold.

On the face of it, compensation means a single requirement could be made strong enough that other requirements at the same level are unnecessary. For example, a low enough probability of core damage could make emission containment unnecessary. By a similar argument, requirements only need to be set at one level to assure implicitly that all higher level requirements would be met -- provided *all* bases are covered.

Mixing requirements

It may make sense, however, to set requirements at several levels in the means-ends hierarchy, as a kind of vertical defense-in-depth (DID). The idea is not novel. The NRC safety goals (USNRC, 1986) provide a notable precedent for mixed-level DID, in having a health goal and also three separate goals for accident, containment and release impact. Another, in the Code of Federal Regulations, is nuclear plant security. 10 CFR 73 specifies three levels: a general security performance objective; a physical system which protects against a design basis threat; and some specific requirements (prescriptions) The rationale is that incremental assessment effort at any one level has decreasing effect so that it is more cost effective to spread effort among several; and pooling several assessments (i.e. at different levels) reduces uncertainty.

The case for current practice. As noted above, current NRC practice makes requirements more demanding at lower than at higher levels. A plausible rationale for keeping, say, core damage risk more stringently controlled than needed to meet higher level requirements might be that industry's economic interests and NRC political interest dictate a tougher core damage standard than that required by safety alone. A core damage accident (like TMI) is undesirable for public relations reasons that go beyond public safety (the "hullabaloo" factor).

On the other hand, it might be argued that industry should be left unregulated to make its own decisions as far as core damage risk is concerned, with only minimal checks by NRC, on the grounds that hard-headed business considerations will force them to assure adequate safety. Regulatory control could then concentrate on containment, siting and other safety measures 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.

Proposed "progressive loosening" mix of levels. I suggest that a quite different type of mixed-level requirement may be appropriate; that is requirements get progressively *less* stringent, as they go down the means-ends hierarchy, from health objectives to prescriptions. This would make regulatees meet certain *minimum* requirements at each level, but leave them the flexibility to decide which of these requirements to *exceed* in order to meet higher level requirements.

For example, provided that that accident risk, release containment and release lethality all meet moderate requirements, the performance of any one of them can be raised to meet the tougher public health requirement. At a lower level (see figure 1), any combination of core cooling design and procedure which assures acceptable core cooling system performance would be acceptable, provided that both design and procedure met minimal standards.

It is true that compliance with higher order performance requirements will still be hard to verify, and that is a significant disadvantage. However, relying solely on low level prescriptions seems to have unacceptable consequences. It runs the risk of critical gaps in safety protection, through the omission of other low-level requirements that are difficult to verify (e.g. organizational culture). *Some* regulation is needed, and the “loosening mix” form may be the least objectionable.

Conclusions

In Utopia, a wise, responsible and all-powerful ruler could prescribe in detail the action citizens should take to assure the best combination of consequences for society. A judge could decide on a defendant’s guilt or innocence on his own after weighing all the evidence. By the same token, a risk manager could specify what risk-related action would produce the greatest public good (for example, by considering the causal linkages between the bottom and top of figure 1).

Since we do not live in Utopia and do not trust our betters with such power, we have opted for a largely adversarial system. Representatives for competing interests, typified by counsels for defense and prosecution, battle it out. If we are lucky, the resolution approximates what the Utopian dictator would have determined.

In risk management, we have a flawed example of this institutional practice. Regulators act to some extent as advocates for health and safety; industry fights for its economic interests. However, they 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. Conversely, some argue that the Quayle Commission on Competitiveness in the 80s 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 and a few honest brokers, that the resulting constructive tension produces a reasonable balance. The government in power may tip the balance one way or the other, as we have recently seen, but elections prevent the will of the people from being distorted too much.

In this paper, I have suggested a regulatory philosophy that tries to serve the public interest in the light of this political and institutional reality. It has focused mainly on how to specify regulatory requirements to achieve a given level of public health and safety cost-effectively.

It has not directly addressed how to set that level in a way that takes balanced account of competing economic and other public interests (the top two rows of figure 1). I believe the arbiter at that level should be higher than the regulatory agency (OMB? Some new

super-agency?), so as not to compromise the agency's legitimate health and safety orientation. Moreover, in keeping with the philosophy of this paper, I suggest that that (or another) higher authority have some discretion to trade off the competing public interests and be able to overrule the regulatory agency on appeal. The existing court system would not be the appropriate arbiter here, since the driving considerations would be political not judicial.

References

Brown, R.V. Assessment uncertainty technology for making and defending risky decisions. *Journal of Behavioral Decision Making*, 3, 213-228, 1990.

Brown, R.V. Toward a prescriptive science and technology of decision aiding. *Annals of Operations Research, Volume on Choice Under Uncertainty*, 1989, 19, 467-483.

Brown, R.V. *Thoughts on Regulatory Philosophy*. Report to Director of Nuclear Reactor Regulation, Nuclear Regulatory Commission. June 1994.

Brown RV. *Combining judgment and data to assess risk*. Working paper. GMU. 1998.

Brown, R.V., and Ulvila, J.W. Does a reactor need a safety backfit? Case study on communicating decision and risk analysis information to managers. *Risk Analysis*, 1988, 8(2), 271-282.

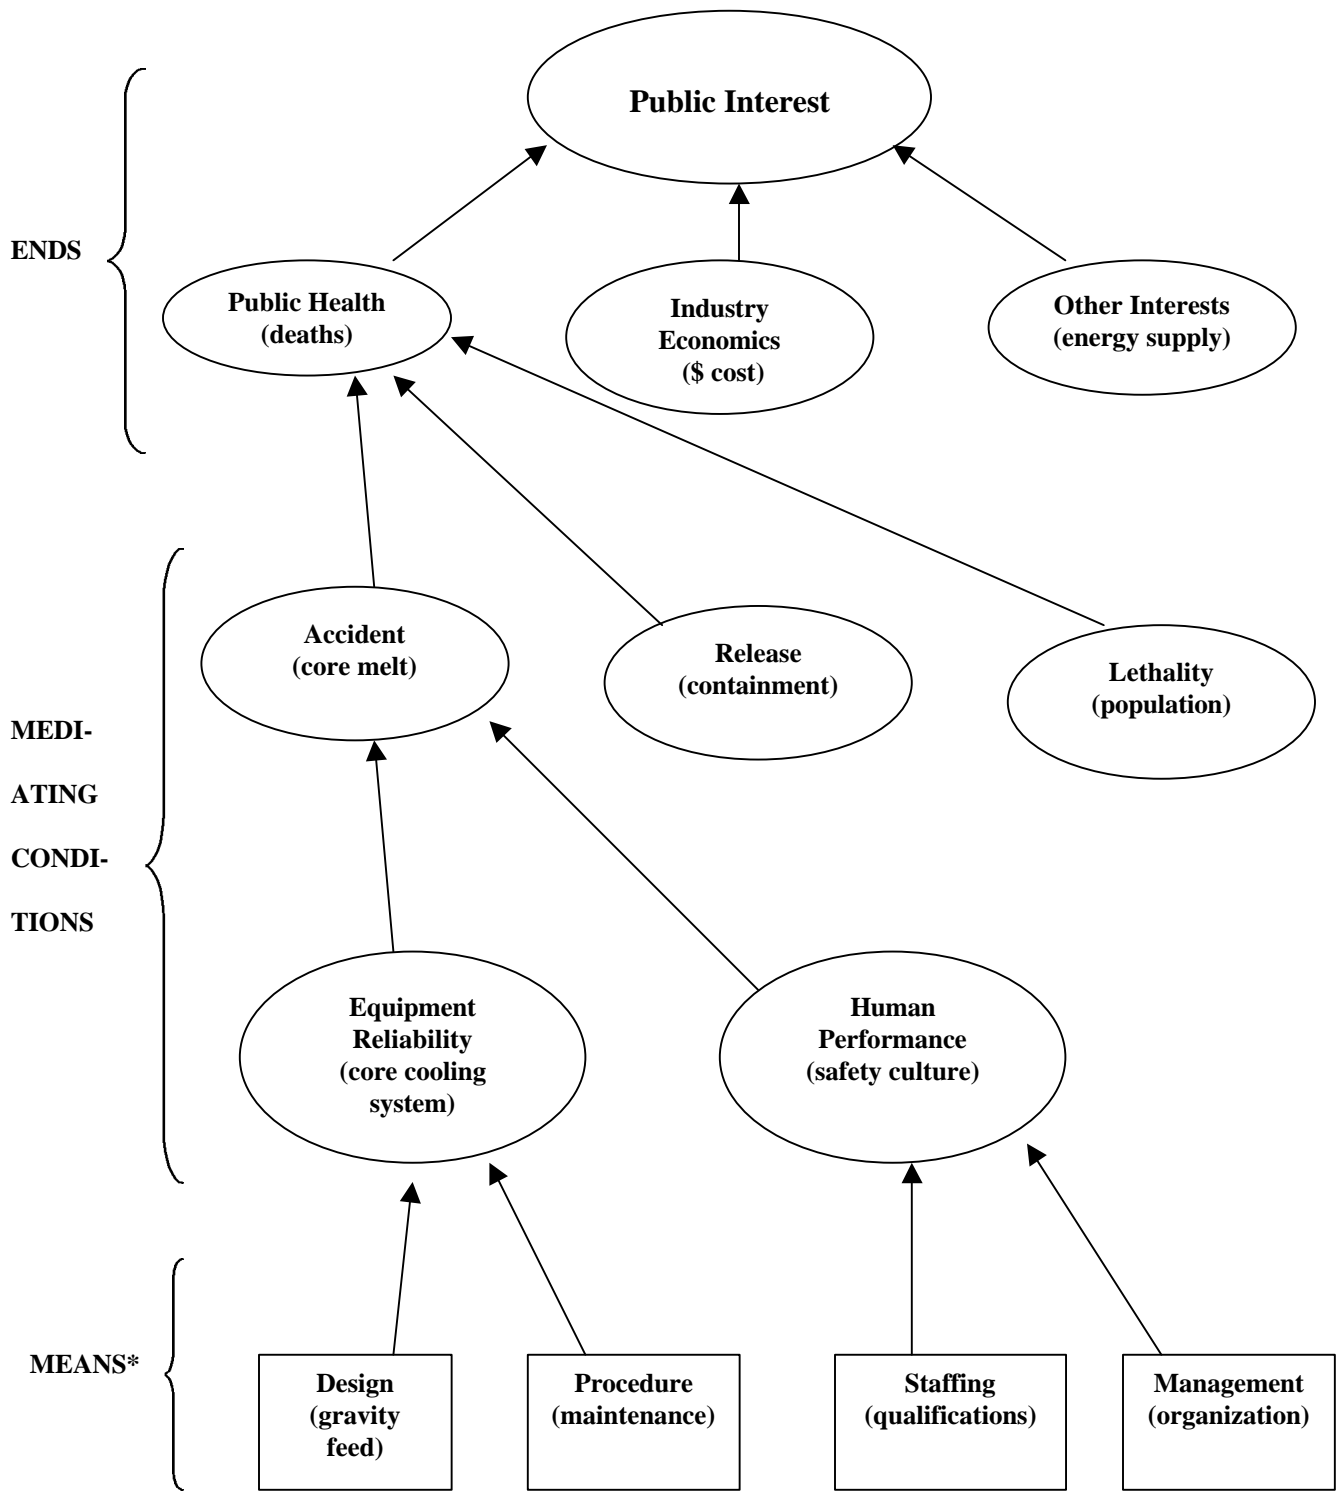
Porter, Henry. Profile: A decision analyst looks at his work. *Chance*. 1991.

Pratt, J.W., Raiffa, H., and Schlaifer, R. *Introduction to statistical decision theory*. MIT Press. 1995.

Raiffa, H. *Decision Analysis: Introductory Lectures on Choices under Uncertainty*. Reading, MA: Addison-Wesley, 1968.

USNRC, Safety Goals for the operation of Nuclear Power Plants: Policy Statement and republication. FR 51 (162) 30028-33, Aug 1986

USNRC, Severe accident risks: an assessment of five US nuclear power plants. NUREG-1150. Washington, DC, 1990.



* Verifiable actions

FIGURE 1
HIERARCHY OF MEANS AND ENDS OF REACTOR SAFETY
 (Partial Schematic)

ⁱ **End note.**

Current law and regulation which embody "acceptable risk" are often qualitative and ambiguous, e.g. requiring risks "as low as reasonably achievable" (ALARA). This gives the regulator discretion to adapt to improved knowledge and to changing circumstances (like easing up if energy becomes scarcer). On the other hand, that discretion is open to inconsistency, arbitrariness and even abuse. Qualitative and quantitative performance requirements (especially probabilistic) have very different, but complementary, appeals. Qualitative requirements (e.g. "not to constitute an unreasonable risk to public health and safety") are fairly simple for licensee and NRC to use, but are subject to controversial staff interpretation. Quantitative/probabilistic requirements involve much more effort, delay and expertise, but are more verifiable. Both have a role to play. Qualitative requirements can set the stage for an operational interpretation in quantitative terms.