Life Safety Risk Criteria in Seismic Decisions

Keith A. Porter, M.ASCE[1]

This paper addresses acceptable life-safety risk within the context of risk-mitigation decision-making. The gathering of risk information and risk-mitigation alternatives raises two questions: is a current level of risk acceptable, and if not, how does one determine if a particular mitigation measure is reasonable and prudent? A procedure is proposed for decision-makers to select among competing risk-management alternatives and to defend the selection. The procedure applies four tests to each alternative: necessity (is the existing risk unacceptably high); efficacy (is the alternative capable of mitigating risk to acceptable levels); efficiency (the relationship between cost and benefit of the risk-mitigation alternative); and how well the alternative provides for social equity. The necessity test relies on the decision-maker's subjective judgment of acceptable risk. Though no universally accepted measure of acceptable risk exists, guidance can be provided by legal and other precedent regarding the peril in question, and by considering levels of risk from related perils deemed acceptable by other authorities and decision-makers. The efficacy test examines whether a proposed mitigation alternative results in an acceptable level of risk, so its criteria are identical to the necessity test. The efficiency test can depend on subjective judgments of what is a reasonable cost for an incremental reduction in life-safety risk, or a reasonable cost per statistical life saved. Guidance on reasonable levels of efficiency can be found in legal and behavioral precedent, and by theoretical means using decision theory. The equity test relies on an examination of quantitative risk to people in the various social and economic groups affected by the mitigation decision. The necessity, efficacy, efficiency, and equity (NE^3) tests are illustrated with a case study that considers the decision to add seismic bracing to automatic sprinklers in highrise buildings. In the study, it is determined that the measure is justifiable in terms of the NE^3 tests. It is noted that risk-management decisions are made within a broader context. Consequently, the NE^3 tests may not be adequate to resolve fully the risk-management decision, but they can form a quantitative basis for informed discussion.

INTRODUCTION

ACCEPTABLE RISK AND RISK-MANAGEMENT DECISIONS

Substantial sums are spent every year to assess, transfer, and mitigate earthquake risk. The annual market for loss-estimation software and services is on the order of hundreds of millions of

---

[1] G.W. Housner Postdoctoral Fellow, Dept. of Applied Mechanics and Civil Engineering, California Institute of Technology, Pasadena CA 91125-4400

From Taylor, C.E., and E. VanMarke, 2001. *Acceptable Risk to Lifeline Systems from Natural Hazard Threats*, Technical Council for Lifeline Earthquake Engineering, American Society of Civil Engineers

dollars. California property owners pay on the order of $1 billion annually for earthquake insurance (Roth, 1997). Annual expenses on earthquake risk reduction are probably around the same order of magnitude. Every time somebody spends money to evaluate, transfer, or mitigate risk, the decision to do so amounts to an assessment that his or her current risk may be unacceptably high.

These decisions tend to be motivated by estimates of worst-case scenarios and a general feeling on the part of the decision-maker that the worst case represents an unacceptable level: the probable maximum financial loss is too high, or the risk of fatalities that could occur in a large, rare earthquake is too great. Often these judgments are made on an intuitive level, with minimal attention to or understanding of the probabilistic nature of risk.

Worst cases are only part of overall risk. Lesser, more frequent losses can also be considered when making risk-management decisions. How can the full range of risk be accounted for in a systematic and defensible way? The present paper proposes a method to examine, test and defend a risk-mitigation decision to improve life-safety. (Economic-risk mitigation decisions are not addressed here. The interested reader is referred to Earthquake Engineering Research Institute, 2000; Roth, 1994; or Porter, 2000.)

DEFINITION OF LOSS AND RISK

Before addressing the question of risk-management decision-making, it is worthwhile to define the terms peril, loss and risk. A peril is a hazard or activity that can cause a loss. Loss is the measure of some undesirable outcome such as financial cost, number of deaths, or number of injuries. Single perils can cause various levels of loss with varying probability. For example, catastrophic releases of radiation from a nuclear power plant are far less likely than small accidents that result in no fatalities. Consequently, risk can be measured in terms of both the severity of loss and the frequency (or probability) of that loss. Risk can be expressed as a continuum (e.g., the probability of exceeding any level of loss over a range of interest), or as a discrete relationship (i.e., only one or more particular levels of loss are considered, and their associated probabilities quantified). Figure 1 illustrates these two methods of quantifying risk.

Sometimes risk is quantified with only one or two parameters, rather than as a continuum or a probability distribution with many possible discrete values of loss. One common one-parameter measure is the mean annual loss, such as the average number of fatalities per year from some ongoing peril. Another way to quantify risk is with a single point on the risk curve such as the probable maximum loss, often defined as an amount of loss associated with a particular small probability of exceedance in given long period of time. While such point estimates of risk have their uses, the present paper primarily treats acceptable risk in terms of compound measures of severity and frequency within the context of a risk-management decision.

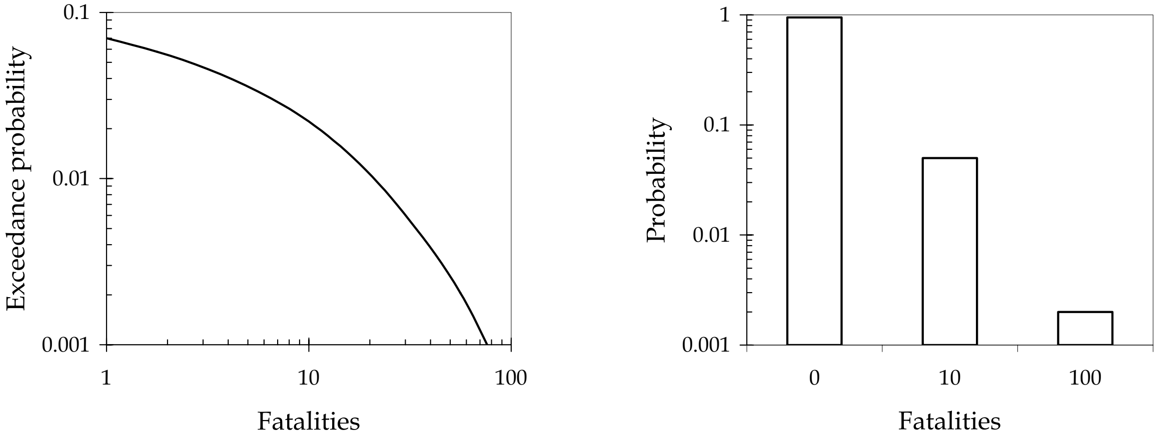

Figure 1. Two alternative depictions of risk. The continuous curve on the left shows the probability of exceeding a given number of fatalities. The discrete relationship on the right shows the probability of three particular levels of loss.

FOUR TESTS OF A LIFE-SAFETY RISK-MITIGATION ALTERNATIVE

FRAMEWORK: A RISK-MANAGEMENT DECISION SITUATION

In a simple decision situation, there is a single decision-maker who can select among two or more alternatives: the do-nothing alternative, or one of several competing risk-mitigation alternatives. Suppose that the decision-maker begins with the following information set: a measure of current risk (i.e., under the as-is situation), the cost of the mitigation efforts, and a measure of the post-mitigation risk under each alternative. With this information in hand, the alternatives can be explored with a series of four tests, which are proposed here as the basis for a quantitative and defensible selection of the best alternative. The four tests are of necessity, efficacy, efficiency, and equity (or NE^3 as a convenient acronym). These tests are detailed below, but in summary they can be understood as follows.

The necessity test examines the question of whether any mitigation at all is called for: that is, does as-is risk exceed acceptable levels? This requires that the decision-maker has formed a judgment of what level of risk is acceptable. If the decision-maker does not already have a predefined level of acceptable risk, then precedent and analogy can provide guidance on what level of risk other authorities consider to be acceptable. If as-is risk exceeds the decision-maker's chosen level of acceptable risk, then some mitigation is necessary and the do-nothing alternative can be ruled out.

If risk mitigation is shown to be necessary, each mitigation alternative is tested for its efficacy, that is, whether it reduces risk to an acceptable level. If it does not do so, it can be excluded from further consideration. If all measures failure the efficacy test, then new options must be developed and explored.

If a measure passes the efficacy test, it can be evaluated for its efficiency, that is, whether it reduces risk at reasonable cost. Finally, in many situations social equity must be considered, that is, whether the risk-mitigation measure equitably protects the various imperiled groups.

These four tests will not always completely justify a risk-mitigation measure in the eyes of all interested parties. Externalities can impinge on a risk-management decision. These can include other perils, alternative uses of the funds, the larger business or political situation, etc. However, the NE^3 tests provide structure and information to the discussion of larger policy decisions.

TESTS OF A LIFE-SAFETY RISK-MITIGATION MEASURE: NECESSITY

As noted above, the first test of a life-safety risk-management decision is of necessity, that is, whether as-is risk exceeds acceptable levels. This test requires a preexisting standard of acceptable risk. If the decision-maker has not already established an acceptable-risk level, he or she must do so, either from internal considerations or by considering what other authorities have previously decided about acceptable risk. Unfortunately, no universal or societal measures of acceptable life-safety risk currently exist. Even the idea that society can in some universal way consciously accept risks may not be valid. Such a societal measure of acceptable risk requires that consistent group preferences exist, a requirement that Arrow’s (1963) Impossibility Theorem denies. Acceptable risk is therefore a personal, subjective judgment that can vary from decision-maker to decision-maker.

Regardless of the theoretical justice of that assertion, in many practical decision situations the decision-maker will be called upon to justify to others his or her selected threshold of acceptable risk, whether those others are stockholders, employees, colleagues, regulators, juries, constituents, the press, or others. In such cases, strong precedent of risk deemed acceptable by other authorities can provide valuable guidance and powerful defense.

Testing necessity against precedent

For the case of life-safety risk associated with earthquakes, historic model building codes provide the strongest indication of the risk various engineering authorities have accepted. By adopting these codes, cities and other jurisdictions have made model codes the de facto legal standard of acceptable life-safety seismic risk for buildings.

Central among the variety of building codes in the United States in 2000 are the International Building Code (International Code Council, 2000), the NEHRP Provisions (Building Seismic Safety Council, 1997a), and the SEAOC Blue Book (Structural Engineers Association of California, 1999). All three codes agree that it is not possible to assure complete safety, and explicitly accept a certain level of risk, although they are silent on quantifying that risk.

Earlier editions of the NEHRP Provisions (Building Seismic Safety Council, 1992, 1994) do shed light on quantitative fatality risk. They provide for a level of risk equivalent to approximately $1.0*10^{-6}$ probability of death per person per year from round-the-clock occupancy in an engineered building in NEHRP's highest-seismicity region. For normal office occupancy, say 45 hours per week, the risk is $0.28*10^{-6}$ per year (for detail on these figures, see Porter et al., 1998). With some exceptions, these levels of risk refer only to buildings built to the standards of the Provisions, not in general to older buildings built to different standards. Nonetheless, the round-the-clock risk figure of 10^{-6} probability of death per year is a useful benchmark of acceptable life-safety risk. A level of life-safety risk substantially higher could reasonably imply the necessity of seismic risk mitigation.

To put this risk in perspective, the risk of death from a single airplane flight during the 1990s was $0.125*10^{-6}$ (Federal Aviation Administration, 1997). The risk of death from motor vehicle accidents in 1995 was $159*10^{-6}$ (Rosenberg et al., 1996). Thus, the death risk from working in an engineered building in a high-seismic region for one year is approximately equal the death risk from a roundtrip airplane trip, and 1/500th the risk of death from automobile accident during the same period.

Testing necessity using acceptable risk from analogous perils

The much greater acceptable risk associated with automobiles raises the question of what kinds of hazards or activities one can use to infer acceptable risk for another peril. That is, what constitutes a reasonably analogous peril? If a peril is very similar to the threat under consideration, it is reasonable to consider that the two perils will agree in terms of acceptable risk. The closer the resemblance of the perils, the stronger the case for comparable acceptable risk. What are the bases for comparison?

Starr (1969, 1972) found a strong relationship between publicly acceptable risk and whether the risk resulted from voluntary exposure to the associated peril. He argued that his analysis showed that public tolerance for voluntary risks allowed for 1,000 times the risk associated with involuntary risks that offered the same dollar benefit per person. A good analogy should therefore have a similar degree of voluntariness and personal benefit.

Sometimes it is difficult to say whether a person is exposed to a particular peril voluntarily or not. Burke (1981) sites the examples of cigarette smoking and automobile accidents, which involve both voluntarily and involuntarily exposed persons. The involuntary nature of earthquake risk is similarly equivocal. Clearly one can largely avoid earthquakes by moving to a region of low seismicity. To resolve the dilemma of what is voluntary, the effort required to avoid a peril must be considered.

Slovic et al. (1980, 1981) examined a variety of hazard characteristics in addition to voluntariness. They listed a large number hazards and activities and performed a number of interviews to ascertain public perceptions of each hazard’s riskiness. In addition to judging riskiness, the interviewees were asked to describe each hazard in terms of 18 characteristics such as: number of people exposed, the degree to which the interviewee was personally imperiled, threat to future generations, the threat of global catastrophe, etc.

The authors then performed factor analysis on the riskiness data, and found that the 18 characteristics could be represented well by three general factors. That is to say, much of the perceived riskiness of the hazards could be predicted by knowledge of the three factors underlying the 18 characteristics. (Factor analysis examines statistical correlations among many observed attributes in an attempt to discern a smaller number of attributes that are not directly observed. For a discussion of factor analysis, see for example Rummel, 1970.)

Factor 1, which the authors dubbed *dread risk*, is associated with lack of control, high catastrophic potential, reactions of dread, and the belief that the risks are not easily reducible. Factor 2, called *unknown risk*, is associated with risks that are unknown, unobservable, new, and delayed in their manifestation. Factor 3 is related to the number of people exposed and the interviewee’s personal experience. (The authors do not name factor 3). Factor 1 was found to be the most important of the three.

The implication for earthquake risk is that, for comparison purposes, if a level of acceptable risk has been established for a different peril that has a similar degree of dread, unknownness, and public and personal nature, then the earthquake peril probably has about the same level of acceptable risk. In practice, this means an analogous risk should have an involuntary nature, similarly large potential for fatalities, be delayed in time by years or decades, and be societal in exposure.

Acceptable risk measures from analogous perils represent good proxies for acceptable earthquake risk, particularly if the frequency and severity relationship for single events can be compared. Good examples include major industrial accidents, epidemic diseases, and toxic waste spills. Conversely, perils that are voluntary, produce small numbers of fatalities in single incidents, and occur frequently, such as automobile accidents, hunting, and non-epidemic diseases, are poor analogies to earthquake risk for purposes of evaluating acceptability.

Helm (1996) examined a variety of industrial and other technological perils that can produce large numbers of fatalities. Such risks match the general characteristics of earthquakes: they are involuntary, potentially catastrophic, infrequent, and societal in exposure. Helm assessed the tolerability of these perils as a function of frequency and severity (number of deaths), and found an inverse linear relationship between the severity of loss and tolerability. For example, 100 fatalities with an annual probability of 10^{-5} are as tolerable as 1,000 fatalities with an annual probability of 10^{-6}. He found that there are four general regions of the frequency-versus-severity space that characterize the tolerability of risk. (He uses the expression *tolerable* risk rather than acceptable risk because literally speaking no fatalities are acceptable.) The four regions, as illustrated in Figure 2, are as follows:

1. Intolerable. High frequency and severe consequences exceed local acceptability of deaths from industrial and other accidents. In this region, “risk cannot be justified except in extraordinary circumstances.”
2. Possibly unjustifiable. Risk is "tolerable only if risk reduction is impractical or if its cost is grossly disproportionate to the improvement gained." This is the upper portion of the region Helm denotes ALARP (as low as reasonably practicable), meaning the risk is tolerable as long as all reasonably practical steps are taken to reduce the risk further.
3. Lower ALARP. Risk is non-negligible, but is "tolerable if cost reduction would exceed the improvement gained."

4. Negligible. Below the negligibility line, frequency and severity are low enough for the risk to be considered broadly acceptable.

Helm's frequency-severity-tolerability relationship is particularly useful for helping to assess the necessity of risk mitigation, for several reasons. First, it allows one to characterize risk acceptability in both its dimensions of frequency and severity. Second, it acknowledges that vast gray areas of acceptability exist. Third, it recognizes that costs and benefits of a peril are relevant when the risk is moderate, but become irrelevant as the risk increases. Finally, it acknowledges that for moderate to high-risk perils, there is a distinction between reasonable and unreasonable cost for risk reduction.

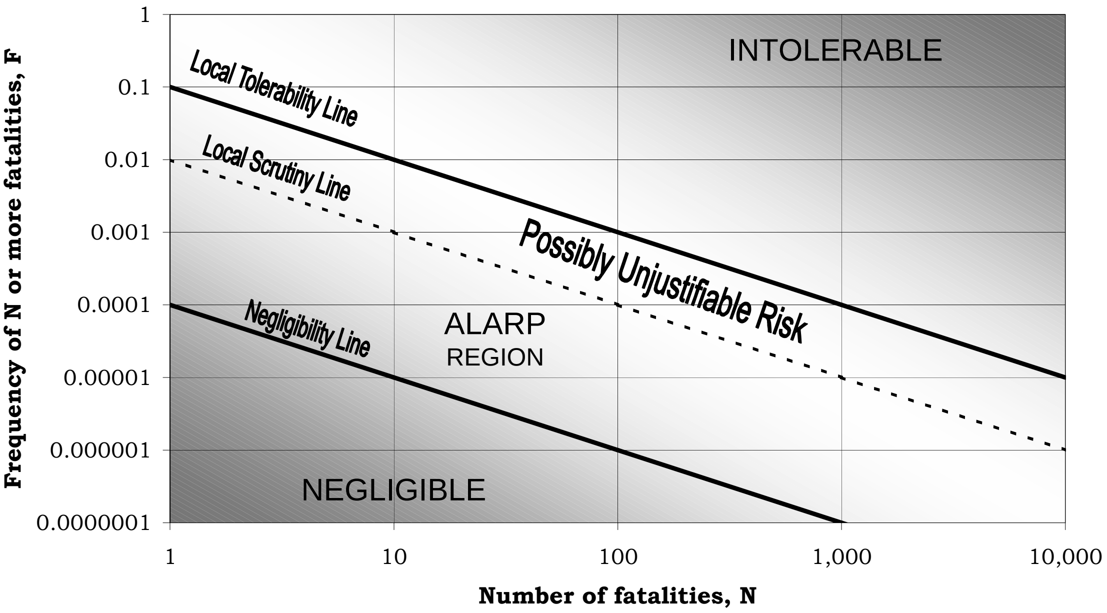

Figure 2. Tolerable risk as a function of severity (after Helm, 1996).

Figure 2 raises a question of the scope of the risk in question. Should a decision-maker assess risk in terms of the isolated system over which he or she has control (e.g., one building), or within some broader context, such as considering all similar buildings that could be affected by the same earthquakes? For one building, the number of possible fatalities in an earthquake ranges from zero to approximately the maximum occupancy of the building. This is the loss over which the building owner has some control. So perhaps considering Figure 2, one should plot the seismic risk for one building as a curve relating the number of deaths in the building and the frequency of that number of deaths.

Regionally however, the seismic risk could involve the thousands of similar structures damaged by the same earthquakes that affect the decision-maker's building. From the regional viewpoint then, the seismic risk could be orders of magnitude larger than that of one building. If these two risk relationships were plotted on Figure 2, the curve from the building owner's perspective could be far below the curve from a societal viewpoint. The former curve could lie within the ALARP or even negligible region, while the latter implied an intolerable risk. Clearly, viewpoint matters.

The question of scope is resolved by observing that Helm’s chart provides no information about tolerability of risk from an individual’s point of view. The tolerability regions in Figure 2 are only defined from the local society’s viewpoint. Hence, if one uses Helm’s chart as a guide to the tolerability of an earthquake-related risk, one must quantify societal risk from the peril in question. For example, if the decision-maker is considering the life-safety risk from damage to an unreinforced masonry building, then to get any guidance from Figure 2, the decision-maker must consider the regional risk from similar unreinforced masonry buildings.

TEST 2: EFFICACY

If a peril exceeds acceptable risk levels, one must determine whether the alternatives under consideration adequately reduce the risk. If the risk-mitigation measure does not actually mitigate the risk to an acceptable level (or to a tolerable level, in Helm’s terminology), then the decision and its associated expense will be difficult to defend.

A classic example is the 1912 Titanic disaster, in which 1522 people died when the luxury liner, en route from Southampton England to New York, struck an iceberg in the North Atlantic and sank. The vessel had been provided with enough lifeboats for only one-quarter the ship’s complement of passengers and crew. It is impossible to justify such a life-safety risk-mitigation measure that clearly and predictably does not achieve its safety objective, at least on any basis other than that of meeting regulatory requirements. (In this example, questions of equity also arise in view of how seats on the lifeboats were actually accorded by class, age, and gender.)

Half-measures are not the only potential cause of failure to pass the efficacy test. Other perils may be causally related to or probabilistically dominant over the one under consideration. If a fatality can occur in a particular situation as a result of two or more causes, and the risk-mitigation measure in question addresses a significantly less-likely cause, then the measure is unlikely to be effective. For example, it may make little sense to improve the seismic reliability of hospital equipment if it is far more likely that the building will collapse than that the equipment will fail. In such a situation, seismic rehabilitation of the equipment without strengthening the supporting structure will probably not save lives.

If as-is risk is intolerably high and an alternative explicitly intended to mitigate that risk does not reduce it to tolerable levels, then the decision-maker should search for other alternatives either to supplement or replace the alternative under consideration. Otherwise, the alternatives that pass the efficacy test can be examined for their efficiency.

TEST 3: EFFICIENCY

Consider once again Helm’s (1996) frequency-fatalities graphic (Figure 2). It acknowledges a wide gray area of acceptable risk: if a risk lies in the upper ALARP region (possibly unjustifiable risk), the decision-maker is called upon to mitigate the risk unless it is impractical to do so or if the cost is grossly disproportionate to the improvement gained.

Since the improvement to be gained is measured in fatalities avoided, this raises the uncomfortable issue of what cost is “grossly disproportionate” to the number of statistical lives saved. Valuing human life is a notoriously difficult topic from a legal and ethical standpoint, but often some means of addressing the subject is necessary, such as in the punitive phase of wrongful-death lawsuits, in the promulgation of costly safety regulations, and so on.

Similarly, the advent of reliability-based building codes forced code authors to consider the economic value of life safety[2]. This is an important and possibly contentious point. Building codes explicitly allow buildings to suffer structural damage at the maximum earthquake design levels of ground shaking. It is certainly possible to design a building to remain elastic, that is, suffer no structural damage, in a large earthquake. The code could conceivably require design to this level, and doing so would greatly reduce the probability that life-threatening structural damage would occur, but the resulting designs are considered by code writers to be too uneconomical.

The American Society of Civil Engineers (1995) states in its commentary on its design standard ASCE 7 (Minimum Design Loads for Buildings and Other Structures), “The design limit state for resistance to an earthquake is unlike that for any other load within the scope of ASCE 7…. The reason is the large demand exerted by the earthquake and the high cost of providing enough strength to maintain linear elastic response in ordinary buildings.” The NEHRP Provisions (Building Seismic Safety Council, 1994 and 1997b) likewise explicitly acknowledge this tradeoff. Thus, current design levels are based on a conscious balance of cost and safety.

The question remains, what is a reasonable cost per statistical life saved? There is no general agreement on how to value a statistical life. Probably there cannot ever be, because the value of life is a subjective matter. Courts and various authors have made attempts, however. Hoffer et al. (1998) in a report commissioned by the Federal Aviation Administration (FAA) recommends that the FAA should be willing to make investments or take regulatory action to reduce fatalities when the costs of those investments or regulations is as high as $2.7 million per statistical death avoided. Fischhoff et al. (1981) report figures varying between $200,000 and $4 million, which in 2000 dollars equate with $400,000 to $8 million. Needleman (1982) examines several methods to value a statistical human life, including: lifetime earning potential (the human capital approach); questionnaires to ascertain people’s willingness to pay to reduce their own risk; observed willingness to pay to reduce risk; and observed willingness to take on additional risk for extra pay. He finds that the last method produced the most reliable estimates of people’s valuation of small changes in their own risk, and that the upper bound of that value is equivalent to 20 times their annual salary per statistical fatality. Considering current U.S. salaries, this figure tends to agree with those quoted by Fischhoff (1981) on an order-of-magnitude basis.

Howard (1980) presents a decision-theory approach to valuing small changes in the probability of fatality. He considers small increments of risk and the amount of money

---

[2] To place a value on life safety is *not* the same as placing a value on a *particular* human life. The question addressed by the efficiency of a life-safety measure is not one of deciding how much the life of your grandmother is worth, but rather what cost is reasonable to reduce the risk so that one fewer (unknown) person will die over a particular period of time or in a given number of trials.

individuals should be willing to pay to reduce their own risk by small amounts, or the amount of money individuals should demand to undertake small quantities of additional risk. Howard's approach contrasts with Needleman's (1982) study in that it is normative, i.e., Howard examines what people *should* do if they want to act consistently with their own stated preferences and available information. Needleman's study by contrast is descriptive in that it explores what people actually do or say about their approach to risk.

In his work on microrisk in medical decision analysis, Howard (1989) proposes two units to measure small quantities of fatality risk: the micromort (a unit denoted by μmt), a 10^{-6} probability of death; and microhazard (denoted by μhz), a 10^{-6} probability of death per year from a continuing hazard. He shows that, for single-incident risks of death up to about 0.1%, the value one places on a unit risk of one's own death should be fairly constant. He illustrates the point with charts like Figure 3, which shows the payment a hypothetical person should require to accept risk of death *p*, as well as the payment an individual should be willing to make to avoid risk of death *p*.

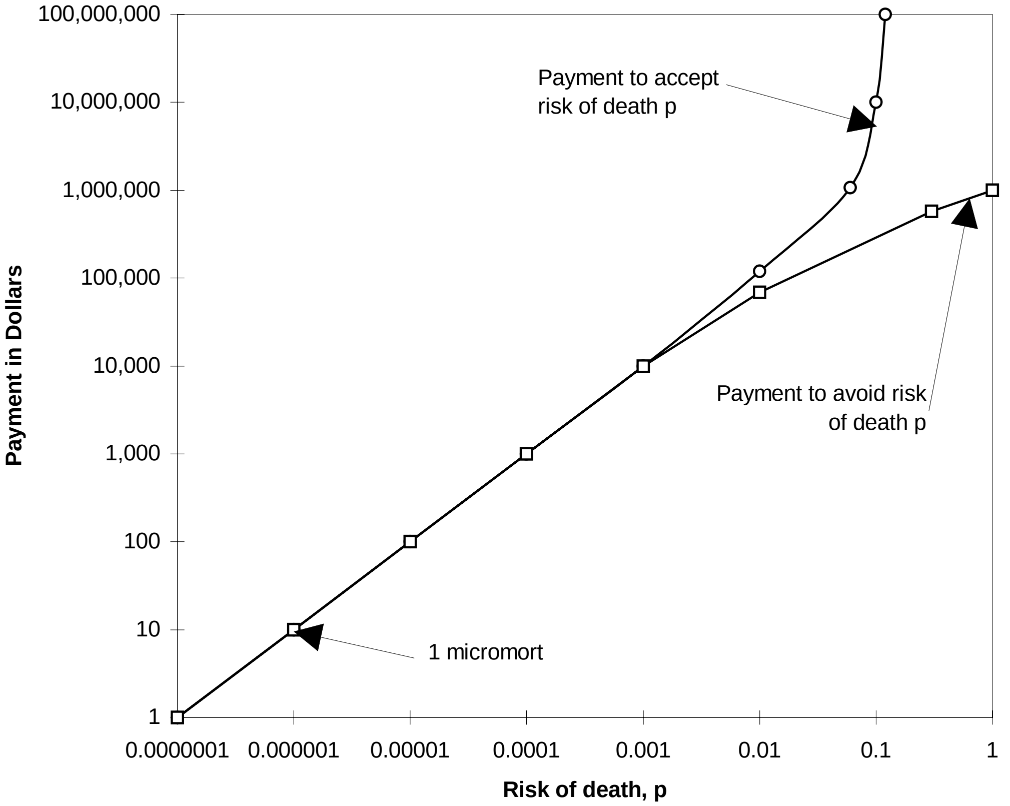

Figure 3. Payment to accept or avoid death risk p (after Howard, 1989).

Before proceeding, it would be worthwhile to explain the two curves of Figure 3 in more detail. The upper curve refers to payment an individual should require to accept addition risk above what he or she currently faces, for example, hazard pay to accept an unusually dangerous job, or savings in airfare to fly on a slightly less-safe airline. The lower curve refers to money an individual should be willing to pay to eliminate risk that is already integrated into his or her life. An example would be paying for optional airbags in a new car.

The curves in Figure 3 diverge above 1,000 to 10,000 μmt. Above this level, an individual loses the ability to pay more to avoid additional risk, and no amount of money paid to a person can induce him or her to accept greater death risk. However, as long as death risk is below about 10,000 μmt, as is the case with building seismic safety, the curves coincide and are linear. Their slope is measured in dollars per μmt.

Howard points out that an individual's wealth state, risk attitude, time preference, and lifestyle as it relates to life expectancy drive his or her dollar value per μmt, but that on average, value can be estimated as a function of age, sex, and annual consumption in dollars. He calculates that these considerations lead to personal exchange rates per μmt and per μhz as shown in Figure 4 and Figure 5, respectively. In the figures, annual consumption is defined as after-tax income above survival level. (More precisely, annual consumption is the minimum amount of money that one would accept annually, after taxes and bare survival, for the rest of one's life, in exchange for the uncertain and varying income he or she will otherwise receive.)

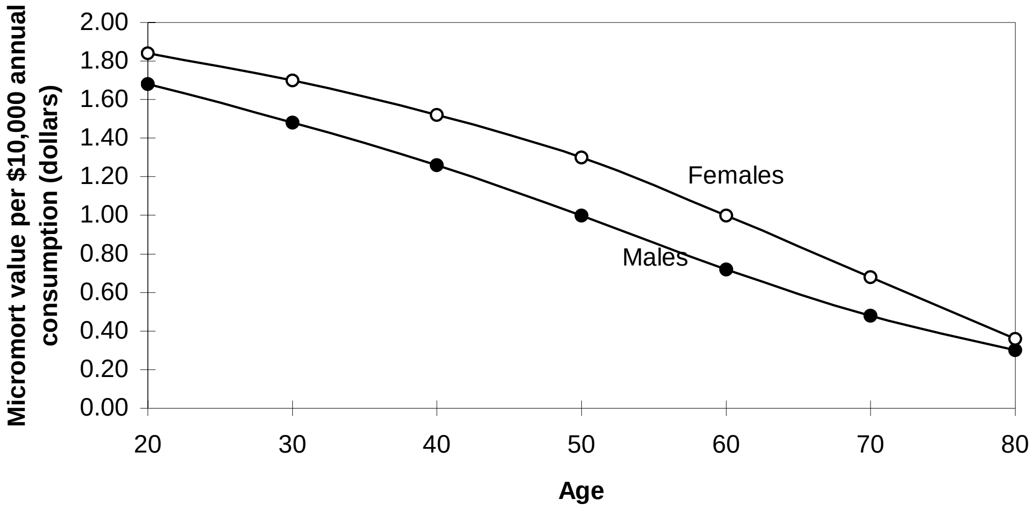

Figure 4. Micromort dollar value (after Howard, 1989).

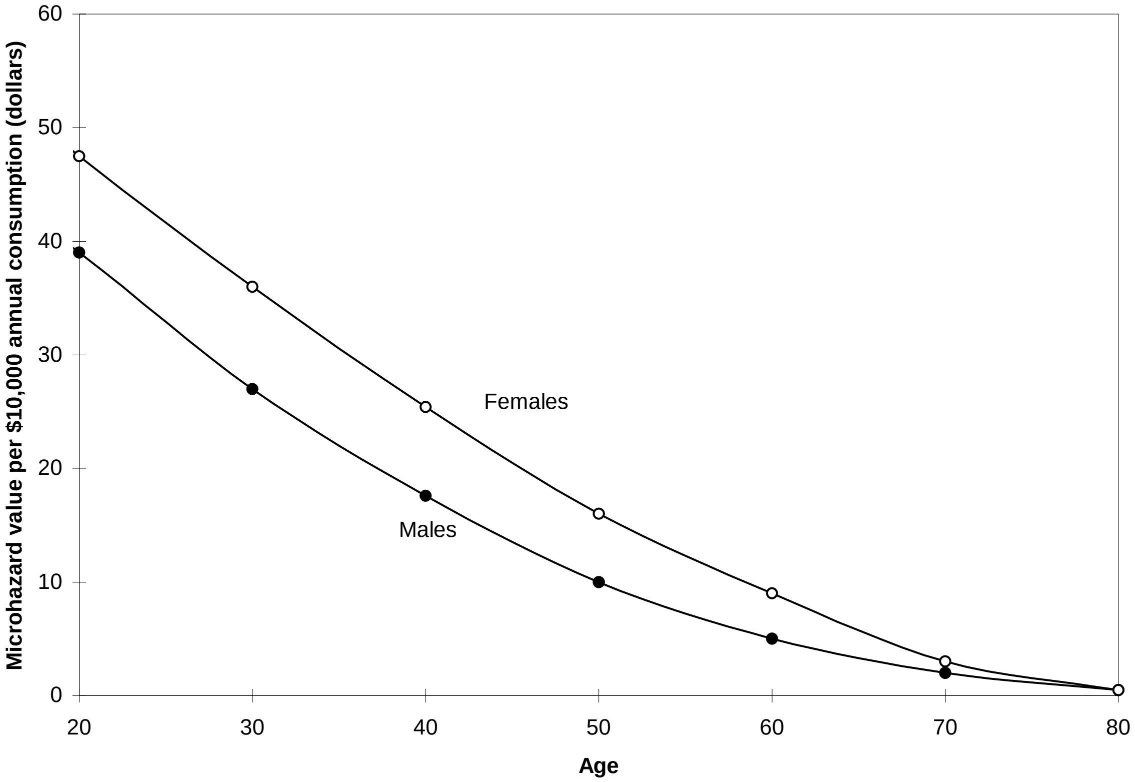

Figure 5. Microhazard dollar value (after Howard, 1989).

Consider the implications of Howard's system for a risk-mitigation measure that reduces annual death risk to an individual by some quantity *x* μhz. Referring to Figure 5, one could estimate that a 50-year-old woman with $45,000 constant annual consumption should be willing to pay a one-time sum of approximately $72*x* ($16.00/μhz/$10,000 * $45,000 * *x* μhz) to purchase the safety provided by the risk-mitigation measure. Looked at another way, she should be willing to pay approximately $6.00**x* per year to be protected from the risk for that one year (from Figure 4, $1.30/μmt/$10,000 * $45,000 * *x* μhz).

A reasonable cost per statistical life saved can be inferred from the Howard methodology: $*y* per μmt avoided is equivalent to $*y* million per statistical life saved. A value of $0.40 to $1.80 per μmt per $10,000 annual consumption, in a nation where typical annual consumption is on the order of $5,000 to $100,000, implies that individuals should value their own lives at the micromort or microhazard level on the order of $200,000 to $20,000,000.

These figures tend to agree with those used in promulgating U.S. federal regulations associated with improving life safety, when the value of death risk reduction is explicitly considered. As part of its 1990 debate over the Clean Air Act, the United States Congress created a commission (Commission on Risk Assessment and Risk Management, 1996) charged, among other tasks, with reviewing federal procedures to assess health risk. The commission reviewed a variety of federal programs that addressed cost per statistical life saved. It found:

> For agencies that explicitly value death risk reductions, the implied value of a statistical life ranges from $1 million to $10 million. For agencies that do not explicitly value death risk reductions, but instead base decisions on an "acceptable" cost per life-saved, the implicit value of a statistical life can be far higher. One study of EPA regulatory decisions that affected cancer risks found regulations promulgated that cost over $50 million per life saved. The Office of Management and Budget study of such behavior, involving a broader range of causes of death, found even higher costs per life saved, as did a recent Congressional Budget Office study of drinking-water standards.

Thus, there is general agreement among a wide variety of studies over the past 20 years on reasonable cost per statistical life saved for risk reduction at the microrisk level. In the United States in 2000, it appears that a figure far in excess of $10 or $20 million per statistical life saved can be considered to be grossly disproportionate to the improvement gained. Likewise, a cost per statistical life saved less than $1 to $10 million could defensibly be argued to be efficient based on a variety of sources: government precedent, interviews of ordinary people, and on the reasonable application of decision theory. At other times and in other countries, different thresholds may apply, but the same principles can be used to determine the appropriate threshold of efficiency.

TEST 4: EQUITY

Whenever the decision-maker is responsible for managing risk imposed on others, the question of social equity enters into consideration. A risk-mitigation alternative that benefits one group of people exposed to the peril in excess or at the expense of another group can be readily attacked based on notions of social equity, particularly if the latter group is politically or economically disadvantaged. For example, in the United States, special effort is often expected to protect children, the elderly, people with disabilities, and the poor. The defense of a proposed risk-mitigation alternative is greatly strengthened by explicit consideration of equity to these disadvantaged groups. The equity test can be thought of in three steps:

1. Identify distinct classes or groups of people exposed to the peril. In any particular case, appropriate groupings should be evident, but some guidance can be provided here. Building codes identify groups of people for whom safety is specially considered through the mechanism of importance factors. These importance factors are used to calculate design loads (which relate to safety), and are arranged by building occupancy. The International Building Code (International Code Council, 2000) for example provides special protection to children in schools and day-care facilities; to occupants of health care facilities, jails and detention facilities; and to emergency responders (police, fire, and rescue).
2. Evaluate the risk to each group under the measure being considered.
3. Determine whether the measure reduces the risk by appropriately proportional amounts. For cases where the measure affects classes who do not expect special protection, appropriately proportional means equally. For groups that can expect special protection, appropriately proportional can mean either equally or somewhat more than for others.

The last point leaves the question of special protection unresolved. This reflects the fact that some authorities (such as building codes) consider special protection for certain classes of people and some authorities do not. The importance factors of the International Building Code (International Code Council, 2000) reduce by approximately half the probability of life-threatening damage in schools, hospitals, and jails, and by approximately 2/3 the probability of such damage in emergency-response facilities (Building Seismic Safety Council, 1992, pg. 47.) In contrast, the Hoffer et al. (1998) study for the Federal Aviation Administration does not differentiate between categories of exposed people in its recommendation of appropriate costs per statistical death avoided. (Howard [1980, 1989] deals with how much an individual should be willing to pay to improve his or her personal safety, and does not address equity between classes of people.)

Thus, the notion of special protection for certain classes of people is not universal. On the other hand, equal protection for everybody else appears to be universally recognized: none of the building codes, government studies, and other sources examined here suggest that there are certain groups of people who deserve less protection than the general public does. The decision-maker is therefore left with imperfect guidance on the equity test, to which he or she must add personal and perhaps legal judgment.

CONTEXT, INFORMED DISCUSSION, AND MAINTENANCE

Each of the NE^3 tests involves values that vary between societies and over time. Hence the precise measures of acceptable risk and reasonable efficiency discussed above must be reviewed for applicability when considered outside the U.S. or at a future time. However, the general principles proposed for determining a threshold value of acceptable risk and reasonable cost per statistical life saved can be generally applied by considering available precedent, analogy to related risks, and appropriate application of decision theory.

This brings up the issue of context. A risk-mitigation decision most likely takes place within a larger business or social context. Competing risk-mitigation alternatives often exist, in which case the most readily defensible one will probably be that which best passes the NE^3 tests. Within a corporation, that defense can often be made in a small committee by considering the company's profitability, availability of funds to undertake the risk mitigation, the company's size and public profile (and hence the degree to which it is subject and sensitive to outside scrutiny), and any corporate policy regarding the protection of human life.

Within a public agency, the decision to act must be made in a much larger and less cohesive forum. Here, the additional context includes public perception of risk, the availability of funds, the public's willingness to spend, competing programs that improve health or safety, public distrust of decision-makers and of technical experts, and other preexisting and potentially adversarial relationships between interest groups. In situations of public risk-management, therefore, the NE^3 tests represent only a starting point for selection of a particular risk-mitigation alternative. Nonetheless, they provide a valuable basis for informed discussion among the interested parties, in which the larger contextual issues can be addressed.

Finally, context can include a physical or business environment that can erode the efficacy of the risk-mitigation measure, so ongoing maintenance may be required. Proper maintenance depends on the nature of the mitigation measure. Safety procedures require ongoing, effective training and testing. Engineering systems can require physical maintenance and regular testing. Emergency plans and contractual arrangements require periodic review and practice. For examples illustrating the failure to maintain risk-management measures, the interested reader is referred to Martin et al. (1989) and Petroski (1981).

ILLUSTRATION: LIFE-SAFETY EQUIPMENT

THE PROBLEM OF FRAGILE AUTOMATIC SPRINKLER SYSTEMS

The procedures discussed above are illustrated in a recent study using the example of automatic sprinklers in highrise buildings (Porter et al., 1998). Sprinklers are known to suffer frequent damage in earthquakes, particularly in cases where long runs of sprinkler pipe lack adequate seismic bracing. The problem is compounded by the fact that earthquakes can result in numerous fire ignitions, which can overwhelm fire departments. Furthermore, physical damage to the building can block building egress and trap building occupants inside, particularly in highrise buildings where egress through windows is impossible. Thus, the sprinklers can be damaged just at a time when a fire is much more likely to occur, professional firefighters are much less likely to be able to respond, and occupants are less likely to be able to escape the building on their own.

The fragility of the sprinklers can be reduced by providing 4-way diagonal bracing and compression struts to unbraced sprinkler lines, with one such brace typically installed for every 12 feet of pipe. In addition, the mechanical equipment, electrical equipment and emergency-generator equipment necessary to run emergency fire pumps in the event of electrical failure must be properly anchored and braced, and any unbraced suspended ceilings and light fixtures similarly braced against sway. Undertaking these measures does not guarantee that automatic sprinklers will not fail in an earthquake; sprinklers are still fragile components that can be damaged by interaction with above-ceiling components, sidesway of ceilings, differential building motion at expansion joints, etc., but the sprinkler reliability is greatly improved.

The question facing building owners then is whether to perform these potentially costly seismic strengthening measures. A similar question faces the local public offices, namely whether to require that building owners perform these measures. The decision of whether to brace these lines therefore represents a good case study for the NE^3 tests.

TESTING THE NECESSITY AND EFFICACY OF SEISMIC STRENGTHENING

Modeling as-is and post-mitigation risk

If a large earthquake were to strike a major metropolitan area, unbraced sprinkler systems in many highrise buildings could fail, and fires in some of those buildings could kill trapped building occupants. These fatalities would probably be viewed by the public as the result of a single common engineering failure, namely that of unbraced sprinkler systems. From this viewpoint, the basis for assessing as-is risk for an individual building should be taken to be the risk associated with all sprinkler systems in highrise buildings in a single, large geographic area.

This risk can be estimated by considering the faults present in the geographic area, the rates at which earthquakes of various sizes occur, the proximity of exposed buildings to the earthquake faults, the seismic vulnerability of the buildings and their sprinkler systems, and the occupancy and potential for fire ignition within those buildings.

A variety of loss-estimation software packages exist that can perform the intensive calculations required to estimate the overall life-safety risk. In the subject study, the seismic risk-estimation software EQEHAZARD was used (EQE International, 1993). (Since this study was conducted, public-domain software has been developed that can likewise perform such an analysis. See National Institute of Building Sciences et al., 1997.) The loss-estimation software contains a model of regional earthquake fault activity, including locations of all faults and fault zones and their associated seismicity. The software models earthquakes on all local faults in 0.1 magnitude increments, from a threshold magnitude of 5.0 to the maximum magnitude of which the fault is believed capable. The rupture location is simulated incrementally from one end of the fault to the other. The seismic shaking intensity at each building is then simulated using mathematical functions relating seismic shaking intensity, magnitude, fault distance, and soil conditions.

The probability of sprinkler failure and post-earthquake fire ignition, given an estimate of seismic shaking intensity, was modeled using fault-tree analysis. In a fault tree, an undesirable event of interest (the *top event* in the tree) is described as the consequence of subsidiary events through the logical conjunctions *and*, *or*, and *not*. The top event occurs if the conditions specified by the subsidiary events are met. Each subsidiary event is similarly related to even more-basic events, until all the basic events are ones whose probabilities can be estimated with some confidence. Then the probabilities of the top event can be estimated based on the basic-event probabilities and all the intervening logic.

For example, consider the simple fault tree shown in the legend of Figure 6. Event 6 is the top event. It occurs if event 4 and event 5 both occur. The probabilities for both basic events are known and are independent – that is, the occurrence of event 4 has no effect on the probability of event 5 occurring. The probability of event 6 is then the product of the two basic-event probabilities. The interested reader is referred to Frankel (1988) or Henley et al. (1981) for further detail.

The fault tree used to model the fragility of sprinkler systems in highrise buildings is shown in Figure 6. Ovals at the bottom of the tree indicate basic events whose probability can be modeled using available data. The basic-event probabilities are expressed as functions of ground shaking intensity, and are considered to be conditionally independent of each other, given the shaking intensity. Rectangles are upper events whose probability is determined from the probability of lower events. P-codes, e.g., P_{21}, are used to identify the event's probability distribution. The basic-event probabilities were developed from statistics in a database of historic equipment failure, as described by Swan et al. (1998).

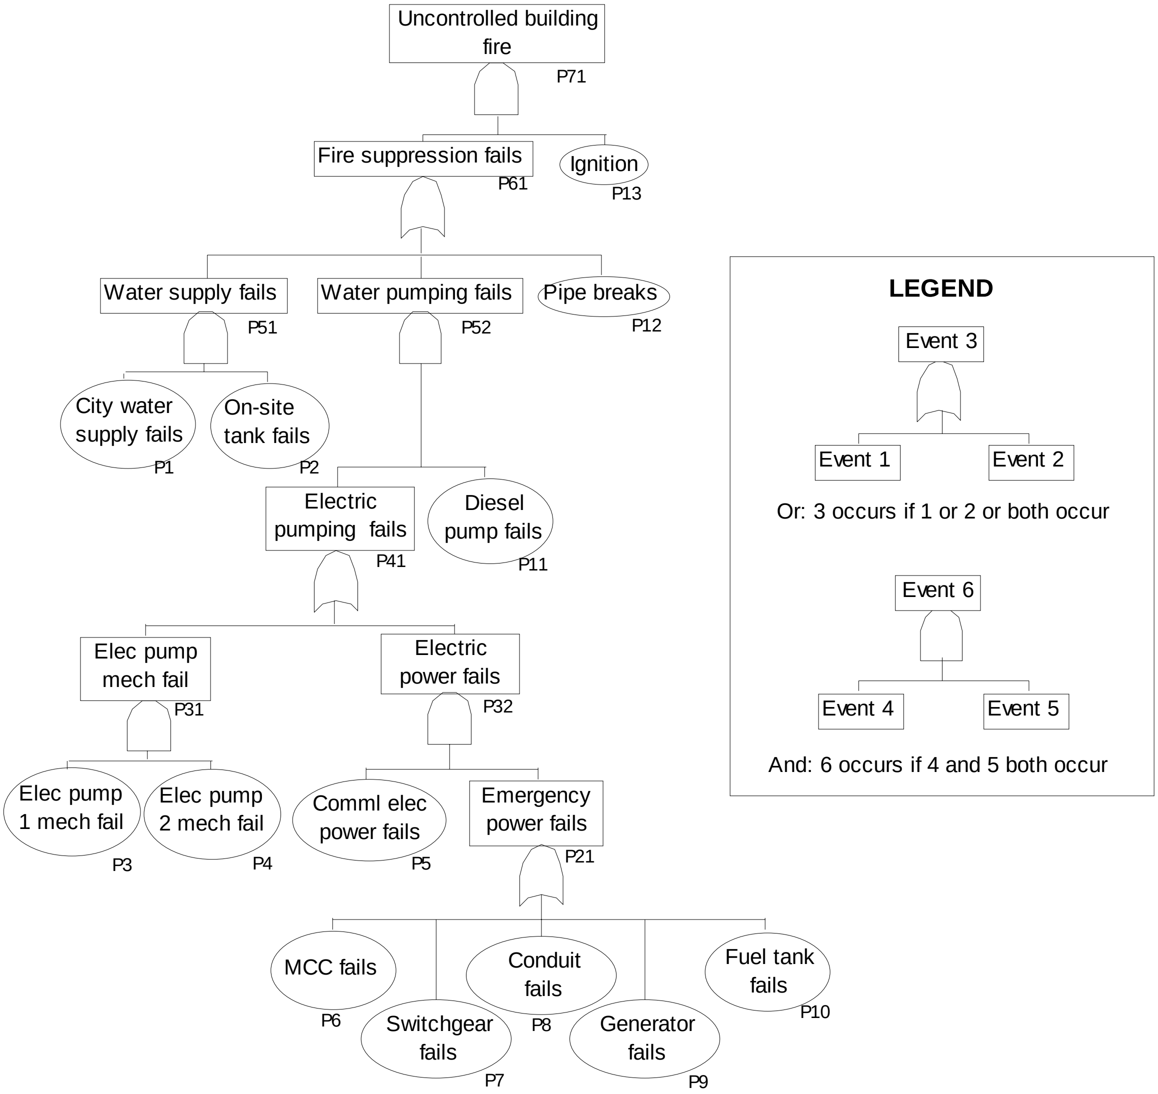

Figure 6. Fault tree used to estimate the probability of uncontrolled building fire.

The result of the fault tree analysis is the probability that an uncontrolled building fire occurs in a particular building, as a function of ground-shaking intensity. The probability that a particular occupant is killed is estimated as the probability that an uncontrolled building fire occurs, times the probability that the ignition occurs at a floor level lower than the occupant. In this analysis, the location of the ignition was modeled as equally probable at every floor level. (In many cases building occupants above the fire would escape via exit stairways. No data are readily available to quantify the likelihood that occupants could and would escape in time, hence the conservative assumption that debris, damage, or other causes trap everyone above the ignition floor and prevent them from exiting the building quickly.)

To model the exposure for the entire San Francisco Bay area, an inventory of highrise buildings was created using tax-assessor information for all the counties adjoining the bay, plus

Santa Cruz County. San Mateo County was not included because of inadequate tax-assessor data. The building inventory included area, height, and street address for each of 714 buildings, together with the building's use. The latitude and longitude of each building was determined from its street address, and the soil conditions at each site determined from soil maps. The number of occupants in each building at any particular time was estimated based on the building's area and use. For example, offices in the United States are occupied primarily in the daytime, with approximately one person per 300 square feet of building area.

Using this seismicity information, seismic attenuation relationships, knowledge of the buildings and occupants at risk, and the results of the fault-tree analysis described above, the loss-estimation software calculated the relationship between the number of fatalities and the annual probability of exceeding that number. This process was performed twice: once assuming that all buildings in the inventory had unbraced automatic sprinklers, and then again assuming that all the buildings had braced sprinklers. (Details of the exposure data, construction costs, sprinkler-system fragility calculations, earthquake faults, etc., can be found in Porter et al., 1998.)

Estimated fatality risk

Figure 7 shows the results of the analysis overlain on Helm's (1996) risk-versus tolerability diagram. Two pairs of risk relationships are shown: one pair for daytime occupancies (offices, retail stores, etc.), and one pair for nighttime occupancies (apartment buildings, condominiums, hotels, etc.). The upper line in each pair shows results assuming unbraced sprinkler lines, and the lower lines show the results assuming braced sprinkler lines.

The figure shows that the decision to brace automatic sprinkler lines passes both the necessity and efficacy tests. Unbraced sprinkler lines in both daytime and nighttime occupancies exceed the risk associated with local tolerability. Bracing those lines is shown to reduce risk to the upper-ALARP region, that is, the region where risk is tolerable only if risk reduction is impractical or if its cost is grossly disproportionate to the improvement gained. (Practical strengthening schemes cannot reduce the risk to the lower-ALARP region, because of the inherent fragility of these systems.)

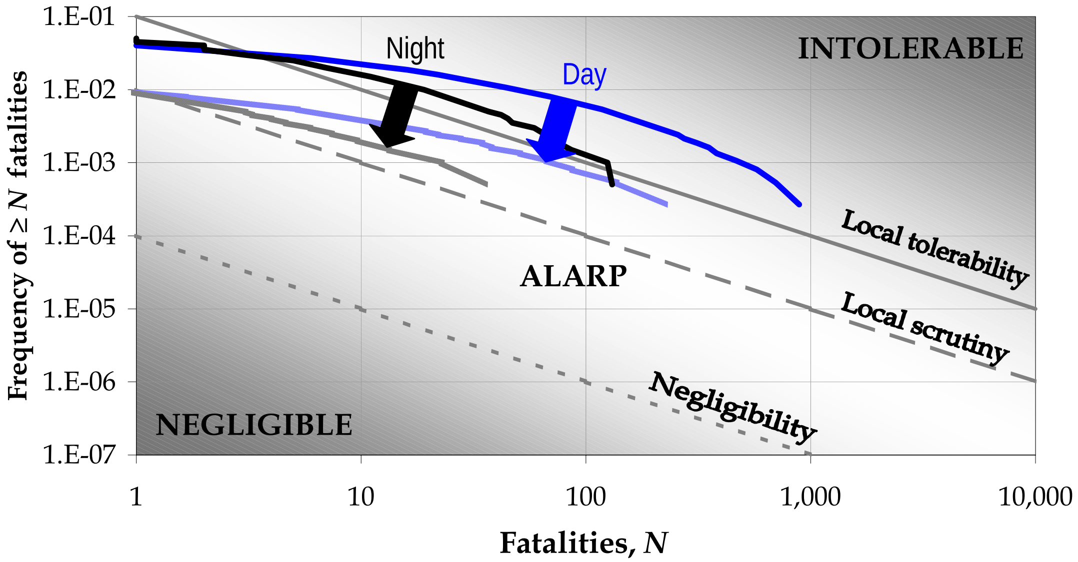

Figure 7. Demonstration of the necessity and efficacy of bracing automatic sprinklers.

TESTING THE EFFICIENCY OF THE SEISMIC STRENGTHENING MEASURE

The efficiency of the rehabilitation measure (the third of the NE^3 tests) is estimated by calculating the cost per statistical life saved. Note that under Helm's (1996) framework, the efficiency test is called for when a risk is already in the ALARP region, not when it is in the intolerable region. As calculated by the model, the unbraced sprinkler lines lie in the intolerable region, in which "risk cannot be justified except in extraordinary conditions." As such, the efficiency test is not strictly called for. However, the efficiency test can strengthen the defense of a decision to improve safety, and so it is examined here.

The cost to anchor or brace the relevant components was estimated to be $2.80 per square foot in 1996, based on data provided by the Federal Emergency Management Agency (1994) and RS Means Co. (1994). Retrofit details took into account recommendations of the National Fire Protection Association (1994). The total cost for the strengthening is approximately $449 million for all 714 buildings. Savings associated with reduced physical damage were also taken into account. Physical damage repair costs were based on the fraction of the building affected by fire and the building replacement cost per square foot.

The results of the efficiency test are shown in Table 1. The table shows results for each of three types of occupancies: buildings primarily occupied in the daytime (offices, etc.), nighttime (dwellings), and 24-hour occupancies (hospitals). The cost to strengthen the sprinkler systems, denoted as rehabilitation or rehab, is shown in the second column. The third column shows the same cost on an annualized basis, using an interest rate of 7.25% minus 2.18% annual inflation, amortized over the design life of the building (50 years). Repair of physical damage before and after strengthening is shown in columns 4 and 5. The net amortized strengthening cost (strengthening cost minus damage reduction) is shown in column 6. The mean annual number of fatalities is shown for unbraced and strengthened conditions in columns 7 and 8. The net cost per life saved is thus calculated as the net amortized cost of strengthening, divided by annualized reduction in fatalities.

As the table shows, the seismic strengthening of automatic sprinklers has an estimated mean net cost of $0.4 million to $7.0 million per statistical life saved, or equivalently, $0.40 to $7.00 per micromort. These figures are within the range of acceptable cost of life-safety improvement for the 1990s in the U.S., and therefore pass the efficiency test.

Table 1. Annualized costs and fatalities associated with failure of automatic sprinklers in highrise buildings. (All dollar figures in millions)

A: Occupancy	B: Total rehab cost	C: Amortized rehab cost	D: Annual damage, unbraced	E: Annual damage, w/rehab	F=C–D+E: Net amort rehab cost	G: Ann deaths, unbraced	H: Ann deaths, w/ rehab	I=F/(G-H): Cost / life saved *
24-hour	$5.0	$0.28	$ 0.26	$ 0.03	$0.05	0.13	0.01	$0.39
Day	$302.4	$16.74	$ 14.09	$ 1.77	$4.42	2.50	0.29	$1.99
Night	$136.6	$7.56	$ 3.23	$ 0.38	$4.71	0.75	0.08	$7.09

* Also equals $/μmt, (i.e., net amortized retrofit cost / reduction in annual death risk)

TESTING THE EQUITY OF THE SEISMIC STRENGTHENING MEASURE

It appears that all occupancy types benefit efficiently from seismic strengthening of sprinklers, so tenants of one occupancy type should be treated the same as any other type. Hence, equity between tenants of various occupancies is maintained. It must be acknowledged that landlords pay for the retrofit, and that damage reduction only partly offsets the capital costs, whereas tenants reap the benefit of increased safety. However, the situation is the same for structural aspects of seismic design requirements. Building codes accept a tradeoff between the interests of landlords and tenants in the service of public safety.

One possible problem with equity is that landlords are likely to pass on the strengthening costs to tenants in the form of increased rent, which could be a serious burden for low-income residential tenants. Many highrise buildings in the San Francisco Bay area serve as residential hotels. For tenants whose income is near subsistence levels, marginal rent increases can have severe and immediate life-safety consequences.

One of the scenarios for a decision to strengthen inadequately braced sprinklers is a city council considering such a requirement for public-safety reasons. In such a situation, it might be difficult to ensure equity between low-income tenants and others without some public assistance to offset the differential burden on the former group.

CONCLUSIONS

The foregoing case study suggests that the decision to strengthen unbraced automatic sprinkler systems is justified based on tests of necessity, efficacy, efficiency, and equity. It was found that unbraced sprinkler systems in highrise buildings pose a life-safety risk in excess of that considered tolerable for risks from industrial and other technological perils; this assessment considers the frequency and number of fatalities associated uncontrolled building fires in all highrise buildings in the San Francisco Bay area. It was also found that seismically strengthening the sprinkler systems costs less than a reasonable maximum according to a variety of standards, both considering cost per statistical life saved and cost per unit reduction in fatality risk. Finally, the measure reasonably assures social equity, with a caveat for cases where the strengthening cost would be passed on to low-income tenants.

More generally, the methodology illustrated here can be used to make and defend other risk-management decisions that affect life safety. The methodology applies acceptable risk as part of the overall process of making decisions regarding the reduction of risk from some particular peril. The process calls for four tests to justify the selection of any particular risk-management alternative. To recap, these tests, whose first letters form the acronym NE^3, are:

1. Necessity. This test requires a demonstration that the peril exceeds acceptable-risk levels. Since acceptable risk is a subjective measure, no consensus exists –perhaps none can exist – on how it may be definitively determined. However, for decision-makers who have not already established their own standards of acceptable risk, precedent and analogy can provide guidance on levels of risk considered by others to be generally acceptable, for the same or a closely related peril. To be applicable, acceptable risk from a related peril must be similar to the one under consideration. Key factors of comparison include lack of control, number of people exposed, high catastrophic potential, delayed risk, and the personal experience with the peril held by the person making the judgment.
2. Efficacy. This test is applied to each proposed risk-mitigation alternative to determine the degree to which the alternative reduces risk below acceptable levels. The efficacy test is used to eliminate from consideration half-measures and measures that do not address the probabilistically dominant hazard.
3. Efficiency. While people are generally willing to spend money to reduce risk to themselves or others, there are limits of how much is a reasonable expense per unit reduction in fatality risk. Significant precedent for some types of risk exist, as expressed by court judgments, government actions, decision theory, and other sources.
4. Equity. A viable risk-management alternative must preserve equity between groups affected by it, and not excessively advantage or disadvantage one group compared with another. Public policy decisions typically carry the burden of extra protection for children and the elderly, the poor, the disabled, and certain other groups. Business decisions can be subject to similar considerations. Explicit consideration of equity issues can greatly strengthen the defense of a particular risk-management alternative.

The NE^3 tests are primarily screening tools. If a peril poses an unacceptable risk under the necessity test, then the remaining tests act negatively, i.e., to exclude proposed alternatives from further consideration. They do not add new options. The exploration of possible alternatives to manage a technological risk requires imagination on the part of the analyst or decision-maker and a willingness to consider a variety of risk-management strategies.

Strategies can include technical modifications (changes in the strength, ductility, flexibility, redundancy, or other interdependency of engineered components or systems); training of operators or occupants in emergency procedures; the establishment of emergency-response agreements; and perhaps other measures. The analyst or decision-maker should be willing to explore a variety of these means, in order to produce a set of alternatives sufficiently broad that at one or more will pass the NE^3 tests. The preferable alternative is most likely the one that best passes the efficacy, efficiency, and equity tests.

Finally, the NE^3 tests may not be adequate to resolve fully the risk-management problem, because risk-management decisions typically take place within a broader context of limited resources, competing political or economic interests, risk from other perils, and a variety of other personal, business, legal, or public-policy considerations. However, careful consideration of all the options available, followed by analysis of each using the NE^3 tests, can provide a strong quantitative basis for subsequent informed discussion among decision-makers and other interested parties on the proper selection and maintenance of a risk-management alternative.

ACKNOWLEDGMENTS

This research was funded in part by the Multidisciplinary Center for Earthquake Engineering Research (formerly the National Center for Earthquake Engineering Research), whose support is gratefully acknowledged. Several engineers of EQE International contributed to the research, most notably Craig Taylor, Charles Scawthorn, Neil Blais, Robert Sheppard, Gayle Johnson, and Sam Swan.

REFERENCES CITED

American Society of Civil Engineers (ASCE), 1995, *Minimum Design Loads for Buildings and Other Structures (ANSI/ASCE 7-95)*, New York: American Society of Civil Engineers, 213 pp.

Arrow, K.J., 1963, *Social Choice and Individual Values, 2nd ed.*, New York: John Wiley and Sons

Building Seismic Safety Council, 1992, *FEMA 223: NEHRP Recommended Provisions for Seismic Regulations for New Buildings and Other Structures, Part 2 - Commentary*, Washington, DC: Building Seismic Safety Council, 237 pp.

Building Seismic Safety Council, 1994, *FEMA 303: NEHRP Recommended Provisions for Seismic Regulations for New Buildings and Other Structures, Part 2 - Commentary*, Washington, DC: Building Seismic Safety Council, 335 pp.

Building Seismic Safety Council, 1997a, *FEMA 302: NEHRP Recommended Provisions for Seismic Regulations for New Buildings and Other Structures, Part 1 - Provisions*, Washington, DC: Building Seismic Safety Council, 337 pp.

Building Seismic Safety Council, 1997b, *FEMA 303: NEHRP Recommended Provisions for Seismic Regulations for New Buildings and Other Structures, Part 2 - Commentary*, Washington, DC: Building Seismic Safety Council, 362 pp.

Burke, F.E., 1981, “Decisions About Public Dangers – A Model Structure for a Valuation Process” *Technological Risk,* N.C. Lind, ed., Waterloo Ontario: University of Waterloo Press, 119-151

Commission on Risk Assessment and Risk Management, 1996, *Risk Assessment and Risk Management in Regulatory Decision-Making; a Draft Report for Public Review and Comment,* Washington, DC: US Congress

Earthquake Engineering Research Institute, 2000, *Financial Management of Earthquake Risk, EERI Endowment Fund White Paper,* Oakland, CA: Earthquake Engineering Research Institute, 114 pp.

EQE International, 1993, *EQEHAZARD-P Portfolio Analysis Software,* Oakland CA: EQE International

Federal Aviation Administration, 1997, *The Aviation Safety System,* http://www.faa.gov/publicinfo.htm

Federal Emergency Management Agency, 1994, *FEMA-74, Reducing the Risks of Nonstructural Earthquake Damage, 3rd ed.,* Washington D.C.

Fischhoff, B., S. Lichtenstein, P. Slovic, S.L. Derby, and R.L. Keeney, 1981, *Acceptable Risk,* New York: Cambridge University Press, 185 pp.

Frankel, E.G., 1988, *Systems Reliability and Risk Analysis, 2nd ed.,* Boston MA: Kluwer Academic Publishers, 429 pp.

Helm, P., 1996, "Integrated Risk Management for Natural and Technological Disasters," *Tephra, June 1996, vol. 15, no. 1,* New Zealand Ministry of Civil Defense, 4-13

Henley, E.J., and H. Kumamoto, 1981, *Reliability Engineering and Risk Assessment,* Englewood Cliffs NJ: Prentice-Hall Inc., 568 pp.

Hoffer, S., F. Berardino, J. Smith, and S. Rubin, 1998, *Economic Values for Evaluation of Federal Aviation Administration Investment and Regulatory Decisions, Report FAA-APO-98-8,* Washington, D.C.: U.S. Department of Transportation, Federal Aviation Administration, 114 pp.

Howard, R.A., 1980, "On Making Life and Death Decisions," *Societal Risk Assessment,* R.C. Shwing and W.A. Albers, eds., New York: Plenum Press, 89-106

Howard, R.A., 1989, "Microrisks for Medical Decision Analysis," in *International Journal of Technology Assessment in Health Care*, 1989, 5, pp. 357-370

International Code Council, 2000, *International Building Code,* Falls Church, VA, 756 pp.

Martin, M.W., and R. Schinzinger, 1989, *Ethics in Engineering, 2nd ed.,* San Francisco: McGraw-Hill, Inc., 404 pp.

National Fire Protection Association (NFPA), 1994, *Automatic Sprinkler System Handbook, 6th Edition,* Quincy MA: NFPA

National Institute of Building Sciences (NIBS) and Federal Emergency Management Agency (FEMA), 1997, *HAZUS Earthquake Loss Estimation Methodology: Technical Manual, Volumes I, II, and III, NIBS Document Number 5201,* Washington, DC: Federal Emergency Management Agency

Needleman, L., 1982, "Methods of Valuing Life" *Technological Risk,* N.C. Lind, ed., Waterloo Ontario: University of Waterloo Press, 89-99

Petroski, H., 1981, *To Engineer is Human,* New York: St. Martin's Press, 247 pp.

Porter, K.A., 2000, *Assembly-Based Vulnerability of Buildings and Its Uses in Seismic Performance Evaluation and Risk-Management Decision-Making, a Doctoral Dissertation,* Stanford CA: Stanford University, 147 pp.

Porter, K.A., C. Scawthorn, C. Taylor, N. Blais, 1998, *Appropriate Seismic Reliability for Critical Equipment Systems: Recommendations Based on Regional Analysis of Financial and Life Loss, MCEER-98-0016,* Buffalo, NY: Multidisciplinary Center for Earthquake Engineering Research, State University of New York, Buffalo, 104 pp.

Rosenberg, H.M., S.J. Ventura, J.D. Maurer, et al. 1996, *Births and deaths: United States, 1995. Monthly Vital Statistics Report, vol. 45 no. 3 supp. 2,* Hyattsville MD: National Center for Health Statistics

Roth, R.J., 1997, *Earthquake Basics: Insurance*, Oakland, CA: Earthquake Engineering Research Institute, 16 pp.

RS Means Co., 1994, *Means Square Foot Costs, 15th ed.,* Kingston, MA: RS Means Company, Inc., 442 pp.

Rummel, R.J., 1970, *Applied Factor Analysis,* Evanston: Northwestern University Press, 617 pp.

Slovic, P., B. Fischhoff, and S. Lichtenstein, 1980, “Facts and Fears: Understanding Perceived Risk,” *Societal Risk Assessment,* R.C. Shwing and W.A. Albers, eds., New York: Plenum Press, 181-214

Slovic, P., B. Fischhoff, and S. Lichtenstein, 1981, “Perceived Risk: Psychological Factors and Social Implications,” *The Assessment and Perception of Risk,* London: The Royal Society, 17-34

Starr, C., 1969, "Social Benefit versus Technological Risk," in *Science,* vol. 165, pp. 1232-1238

Starr, C., 1972, "Benefit-Cost Studies in Sociotechnical Systems," in *Perspective on Benefit-Risk Decision-Making,* Washington, DC: Committee on Public Engineering Policy, National Academy of Engineering

Structural Engineers Association of California (SEAOC), 1999, *Recommended Lateral Force Requirements and Commentary, Seventh Edition,* Sacramento: SEAOC, 440 pp.

Swan, S.W., and R. Kassawara, 1998, “The Use of Earthquake Experience Data for Estimates of the Seismic Fragility of Standard Industrial Equipment,” *ATC-29-1, Proc., Seminar on Seismic Design, Retrofit, and Performance of Nonstructural Components,* Redwood City CA: Applied Technology Council, 313-322