Approaches to Quantitative Risk Assessment with Applications to Physical Protection of Nuclear Material

Gebhard Geiger
Technische Universität München
Fakultät für Wirtschaftswissenschaften
Munich, Germany

and

Anselm Schaefer
Institute for Safety and Reliability GmbH
Garching, Germany

Abstract

Violations of physical protection combined with threats of misuse of nuclear material, including terrorist attack, pose increasing challenges to global security. In view of this situation, we exploit recent advance in theoretical and applied risk and decision analysis to attain methodological and procedural improvements in security risk management, especially quantitative risk assessment and the demarcation of acceptable risk. More precisely, we employ a recently developed model of optimal risky choice to compare and assess the cumulative probability distribution functions attached to safety and security risks. Related problems such as the standardisation of risk acceptance criteria frequently used in physical protection can also be approached on this basis. With regard to nuclear and radiological threats, the paper suggests possible applications of the improved methods to the safety and security management of nuclear material, cost efficiency of risk management practices, and the harmonisation of international safety and security standards of physical protection. An example selected from the security risks of spent nuclear fuel transport will be presented in some more detail to demonstrate the practical force of the approach.

1. Introduction

During the past decade, the need for improving physical protection (PP) of nuclear material has been felt increasingly in science and international security. Concerns about the security of nuclear material have been raised in view of a wide distribution of orphan radiation sources (Ortiz *et al*. 1999) and a dramatic increase in the illicit trafficking of nuclear materials worldwide, doubling the 1996 annual rate of reported incidents in less than half a decade (Nielsson 2001), with threats of international nuclear terrorism rapidly developing (IAEA 2001; Committee on Science and Technology for Countering Terrorism 2002, Chap. 2).

We consider the physical protection (PP) of nuclear and radioactive material as a security risk assessment and management task. While "nuclear safety" means the prevention of nuclear accidents or mitigation of accident consequences, the term "security" refers to measures to prevent the loss, theft or unauthorised transfer or use of radiation sources or radioactive materials. We specifically address the problem that security risks are difficult to approach within the framework of quantitative risk analysis since potential violations of PP are hard to predict and assess in probabilistic terms.

Exploiting recent advance in theoretical and applied risk and decision analysis, we begin with the outline of a suitable model of quantitative risk assessment. Our approach offers methodological and procedural advantages for the demarcation of acceptable risk, and the management of nuclear safety hazards and security threats. More specifically, we employ a recently developed model of optimal risky choice to compare and assess the cumulative probability distribution functions attached to nuclear safety and security risks. We then proceed to argue that related problems such as the standardisation of risk acceptance criteria frequently used in PP can also be approached on this basis. With regard to nuclear and radiological threats, possible applications of the improved methods will be discussed. They include the safety and security management of nuclear material, the cost efficiency of risk management practices, and the harmonisation of international safety and security standards of PP. An example selected from the security risks of spent nuclear fuel (SNP) transport will be presented in some more detail to demonstrate the practical force of the approach.

Our analysis intersects with a variety of approaches to risk analysis and its applications to nuclear safety and security that have been developed in the literature. They include probabilistic risk analyses and empirical studies of individual and societal risk perception and acceptance patterns, risk-benefit analyses, approaches to nuclear risk management, and applications to the risks of nuclear facilities and SNF transport (McCormick 1981; Royal Society Study Group 1992; Jorissen and Stallen 1998; Fullwood 2000; Committee on High-Level Radioactive Waste 2001). However, for the sake of definiteness, we do not discuss our approach within the broader contexts of the available theories of individual and societal risk bearing in detail here. As for a review of these issues and the relevant literature, we refer to Geiger (2001, 2002a, c). More narrowly circumscribed problems of this kind associated with the assessment and management of nuclear and radiological risk will be mentioned below within the particular contexts into which they belong.

2. Probabilistic Risk Analysis and Non-Expected Utility Theory

2.1 Utility theory as a risk assessment framework

The risk of a random event E is often defined, in quantitative terms, as the probabilistic expectation of damage or loss from E. However, in experimental contexts, this definition makes sense only if E occurs repetitively in a series of trials so that the relative frequency of E is likely to approach the probability of E ("law of large numbers"). In applied risk analyses, a broader definition of risk is required, adequately characterising probabilistic events that may occur in a non-repetitive fashion. In fact, many security risks in science and society arise within the contexts of one-shot decision tasks. An example involving a potential terrorist attack on a single SNF shipment has been described by Múnera *et al.* (1997).

Utility models of decision making under risk provide suitable approaches to this sort of problem, with many applications in economics and the engineering sciences, including transport risks of hazardous material (Chankong and Haimes 1983, Chap. 3; French 1988; Evans and Verlander 1997). Utility theory conceptualises risky choice as the acceptance or rejection of lotteries offered to a person. The possible outcomes x of any such lottery are values of a numerical random variable, with gains $x \geq 0$, losses $x \leq 0$, and probability $p(x)$ or, equivalently, cumulative probability distribution $F(x) = \Sigma_{y \leq x}\, p(y)$, and analogously for continuous distributions with probability densities $f(x)$. In particular, we ambiguously use the symbol "p" to denote probability functions and the risks, or risky courses of action, they represent. The gains (losses) involved can be amounts of money, fatalities prevented (incurred), radiation doses averted (received), or multiple relations between them (Keeney and Raiffa 1976; Chankong and Haimes 1983). The decision maker is further supposed to assess the likely gains and losses x in terms of the utility u, which will depend exclusively on x only in idealised cases, however. In applied risk analyses, assessments of the outcome x in utility terms rather tend to vary with p and the probability ε that p gets resolved within a given period of time which is characteristic of the decision problem in point. Accordingly, the concept of utility must be defined in terms of a parameter family of suitable probability-dependent utility functions $u_\varepsilon(p, x)$. A simple and empirically realistic account of utility developed by one of us (Geiger 2001; 2002a, b) uses the following parameters to specify $u_\varepsilon(p, x)$:

Aspiration level x_0: The outcome is evaluated as a gain ($x \geq x_0$) or loss ($x_0 \geq x$) with reference to some neutral point x_0 (Kahneman and Tversky 1979) which may be positive, negative or zero. For instance, in radiological applications such an aspiration level x_0 may be the maximum admissible effective radiation dose or dose rate per person specified by a nuclear regulatory agency. For computational purposes it is often convenient to transform the outcome axis $x \rightarrow x - x_0$ so that

$$x_0 = 0 \tag{1}$$

Reference risk s, or *status quo* (Pratt 1988): Depending on the particular application, the *status quo* is the actor's present state of wealth or health involving some degree of uncertainty (economics, health care, etc.), or extant risk of injury or fatality (natural hazards, hazardous technologies) with probability distribution $s(x)$. For example, in radiological applications s may be chosen as the individual life-time dose risk or dose rate of exposure to natural radiation.

Persistence ε of the status quo s: If a risk p is assessed against a given reference risk s, it may be important for the decision maker to know which of the two is likely to be resolved first. A suitable estimate is provided by the probability ε that p gets resolved prior to the resolution of s ($0 \leq \varepsilon \leq 1$). Conversely, $1-\varepsilon$ is the probability of persistence of p in the presence of the *status quo* risk s.

The overall utility U_ε of a risk p is the average of the utilities of the likely gains and losses $x_1, ..., x_n$,

$$U_\varepsilon(p) \;=\; \Sigma_{i\leq n}\, p(x_i) u_\varepsilon(p, x_i), \qquad 0 \leq \varepsilon \leq 1,\; n \geq 1 \tag{2a}$$

$$U_\varepsilon(x) \;=\; u_\varepsilon(p, x), \qquad p(x) = 1,\; p(x') = 0 \text{ for } x \neq x'. \tag{2b}$$

In the special case $u_\varepsilon(p, x_i) = u_\varepsilon(x_i)$ in which the utility function does not explicitly depend on p, $U_\varepsilon(p)$ reduces to the familiar expected utility (EU). Otherwise, the expression (2a) is non-linear in the probabilities (generalised expected utility, or non-EU; see Fishburn 1988; Quiggin 1993).

We first infer a few of its general properties before we specify $u_\varepsilon(p, x)$ for particular parameter values below. Figure 1 illustrates the concept of "decreasing marginal utility", meaning that $u_\varepsilon(p, x)$ is strictly increasing in x and concave. $U_\varepsilon(p)$ is the average of a two-component probability distribution p defined for x_1, x_2, with the mean $\mu(p) = x_1p(x_1) + x_2p(x_2)$ and the so-called *certainty equivalent* $c_\varepsilon(p)$ lying on the x-axis between x_1 and x_2. As indicated in the figure, the certainty equivalent is implicitly defined by

$$U_\varepsilon(p) \;=\; u_\varepsilon(p, c_\varepsilon(p)) \;=\; U_\varepsilon(c_\varepsilon(p)) \qquad (3)$$

meaning that you are indifferent between receiving either the lottery p with possible monetary prizes x_1 and x_2, or the sure amount c_ε of money. Similarly, one has $c_\varepsilon(p) < \mu(p)$ in Figure 1, that is, you prefer receiving a sure amount of money c', $c_\varepsilon(p) < c' < \mu(p)$, to the lottery p although c' is smaller than the expected value $\mu(p)$ of that lottery ("risk aversion"). Equality $c_\varepsilon(p) = \mu(p)$ obtains in case of vanishing risk aversion, or risk neutrality, where the utility curve is a straight line. The certainty equivalent is thus a measure of the decision maker's attitude toward risk. The important point here is that in the utility approach to risk assessment a risk p is uniquely mapped into its certainty equivalent so that random (risky) and non-random (riskless) outcome variables are consistently assessed in a common quantitative framework. In particular, one has $U_\varepsilon(p) > U_\varepsilon(q)$ exactly if p is preferred to q (i. e., $c_\varepsilon(p) > c_\varepsilon(q)$), and $U_\varepsilon(p) = U_\varepsilon(q)$ exactly if p and q are indifferent (i. e., $c_\varepsilon(p) = c_\varepsilon(q)$).

2.2 An approach to non-EU

Scaling the outcome axis according to (2) and considering (3b), one has (Geiger 2002a, b),

$$u_\varepsilon(p, 0) = 0, \qquad 0 \le \varepsilon \le 1 \qquad (4a)$$

$$U_\varepsilon(s) = U_\varepsilon(0) = u_\varepsilon(s, c_\varepsilon(s)) = c_\varepsilon(s) = 0, \qquad 0 \le \varepsilon \le 1 \qquad (4b)$$

$$u_0(p, x) = u_0(s, x) = \text{constant in } p \qquad (4c)$$

$$u_\varepsilon(s, x) = u_0(s, x), \qquad 0 \le \varepsilon \le 1, x \in \_. \qquad (4d)$$

Equation (4b) means two things. First, the utility of receiving nothing with certainty is $U_\varepsilon(0) = 0$. Second, doing nothing and, thus, remaining in the *status quo*, amounts to receiving nothing with certainty so that $U_\varepsilon(s) = U_\varepsilon(0) = 0$. The certainty equivalent of s accordingly vanishes for all ε. The utility indifference of s and 0 (or x_0) is consistent with the notion that decision makers maximise the utility of changes in wealth rather than that of wealth levels and that such changes are assessed relative to the aspiration level (Kahneman and Tversky 1979). Correspondingly, a risks p is referred to as *neutral* if p and s are indifferent.

Since $\varepsilon = 0$ obtains in the application considered below, we only need to determine the x-dependence of $u_\varepsilon(p, x)$ in the special case (4c) here. The calculation of $u_0(p, x) = u_0(s, x)$ draws upon the following result which holds in a broad class of generalised expected utility models for all ε, $0 \le \varepsilon \le 1$ (Geiger 2002a, b),

$$U_\varepsilon(p) = U_\varepsilon(q) \quad \Leftrightarrow \quad u_\varepsilon(p, x) = u_\varepsilon(q, x), \qquad x \in \_. \qquad (5)$$

The equivalence (5) means that indifferent probability distributions have identical utilities as functions of x. This result has some fortunate practical consequences. Once $u_\varepsilon(p, x)$ has been determined for some p, the utility functions $u_\varepsilon(q, x)$ of all distributions q indifferent to that p are known as well. One can exploit this situation by restricting the analysis to gambles p with only two possible outcomes x_1, x_2 to determine $u_\varepsilon(q, x)$ for all multi-component probability distributions q indifferent to p. To see this, let $x_1', \ldots, x_n'$ be the possible outcomes of q and observe that although p and q have finite numbers of outcomes, the functions $u_\varepsilon(p, x)$ and $u_\varepsilon(q, x')$ are respectively defined for all real x and x'. If now the decision maker is indifferent between p and q, he only needs to know $u_\varepsilon(p, x)$ to calculate $U_\varepsilon(q)$,

$$\Sigma_{i \le m} q(x_i')u_\varepsilon(p, x_i') = \Sigma_{i \le m} q(x_i')u_\varepsilon(q, x_i') = U_\varepsilon(q) \qquad (6)$$

To determine $u_0(p, x)$ according to (4c), assume first that s is a two-outcome lottery involving a loss $x_1^s < 0$ and gain $x_2^s > 0$ respectively obtained with probabilities $s_1 > 0$ and $s_2 = 1 - s_1 > 0$. Define the (negative) standard score of the neutral point x_0 in the normalisation (1),

$$z(s) = \frac{\mu(s) - x_0}{\sigma(s)} = \frac{\mu(s)}{\sigma(s)} = \frac{\mu(s)}{\Delta^s \sqrt{s_1 s_2}}, \qquad \Delta^s = x_2^s - x_1^s > 0 \qquad (7)$$

where σ is the standard deviation. If, however, the *status quo* is a multi-component risk, it is always possible to use (2a), (5) and (6) to construct, by numerical approximation, a two-component probability distribution s' so that s and s' are indifferent (Geiger 2002a, App. B). Hence, (7) describes the general case of a *status quo* risk. Now put $z(s) = z_0$ (the lower index "0" in the parameters z_0 and x_0 denotes universal constants and has nothing to do with the particular $\varepsilon = 0$). One then has for arbitrary p and $x \ge 0$ (Geiger 2002a),

$$\begin{aligned} u_0(p, \pm x) &= u_0(s, \pm x) \\ &= -\,u_0(s, \_1)\,\frac{2x + z_0^2(1+x) - \sqrt{z_0^4(1+x)^2 + 4z_0^2 x}}{2 + z_0^2(1+x) + \sqrt{z_0^4(1+x)^2 + 4z_0^2 x}}, \end{aligned} \qquad (8a)$$

with $u_0(p, 0) = 0$, and the negative branch of the utility function

$$-\frac{u_0(p, -x)}{u_0(p, x)} = \frac{\sqrt{1+z_0^2} + z_0}{\sqrt{1+z_0^2} - z_0} = A_0. \qquad (8b)$$

Figure 2 shows $u_0(p, x)$ normalised to $u_0(p, -1) = -1$ for $z_0 = 0.33$. The gross pattern of the utility curve is S-shaped, that is, concave (risk averse) for gains and convex (risk prone) for losses. The exception is a small neighbourhood of the origin (encircled) in which the converse is the case, that is, convexity (risk proneness) for gains and concavity (risk aversion) for losses. The curve is steeper for losses than for gains, corresponding to $z_0 > 0$ and

$A_0 > 1$ in (8b). Figure 2 is in conspicuous qualitative agreement with broad evidence from experimental decision analysis (Luce 2000; Starmer 2000).

For arbitrary two-outcome gambles p, one defines the parameter $z(p)$ in a fashion analogous to (7)

$$z(p) = \frac{\mu(p)}{\sigma(p)} = \frac{\mu(p)}{\Delta\sqrt{p_1 p_2}}, \qquad \Delta = x_2 - x_1 > 0 \tag{9}$$

Given s and z_0 as above, one straightforwardly verifies that

$$\mu(p) - z_0\sigma(p) = 0 \;\Rightarrow\; U(p) = 0 \tag{10}$$

For each pair x_1, x_2 of possible outcomes with $x_1 < 0 < x_2$, the particular p^0 satisfying (10) is obtained by solving $\mu(p^0) = z_0\sigma(p^0)$ for p^0,

$$p_1^{\,0} = \frac{2x_2 + z_0^2\Delta - z_0\sqrt{z_0^2\Delta^2 - 4\,x_1x_2}}{2\Delta(1 + z_0^2)} \tag{12}$$

Considering Equation (6) and the result that arbitrary risks are characterised by the properties of their two-component counterparts to which they are indifferent, we can now give a simple characterisation of risk acceptance on the basis of (10) and (11). Since (10) implies $U(p) = 0$ and, hence, a vanishing certainty equivalent for p if $\mu - z_0\sigma = 0$, one has

$z_0 > 0$ or

$z_0 < 0$ depending on whether the decision maker is risk averse or risk seeking, and $z_0 = 0$ for risk neutrality. The agent's neutral point of risk tolerance is thus implicitly defined by (10), why neutral risks may be called *marginally acceptable*. We refer to the particular value $z(p^0) = z_0$ as the agent's *limit of acceptable risk*, or, briefly, *critical risk acceptance*. Similarly, p is called *acceptable* (*unacceptable*) exactly if $z(p) \geq z_0$ ($z(p) < z_0$). In this sense the various acceptability properties of risks are evaluated with reference to the decision maker's aspiration level and *status quo.* Moreover, the decision maker's attitude toward risk is not a constant, but generally varies with p even for fixed parameters x_0 and z_0. This variability is further increased if $\varepsilon \neq 0$. It corresponds to the observed coexistence of different risk attitudes in decision makers (aversion, proneness, neutrality towards risk) even for one and the same set of possible outcomes (Geiger 2002a).

Although the preceding conceptualisations have been introduced within the context of a probabilistic model of individual utility-oriented choice, they are in surprising qualitative and quantitative agreement with familiar social risk acceptance attitudes (Geiger 2001, 2002c) such as arise in voluntary and involuntary exposure to, and perceptions of the social costs of, collective and catastrophic risks (Starr *et al*. 1976; Okrent 1981; McCormick 1981, Chaps. 17, 18; Stallen *et al*. 1996).

3. Nuclear and Radiological Security Risk Analysis

3.1 Comparative and quantitative risk assessment

The non-EU model outlined thus far involves three governing parameters x_0, z_0 and ε which in applied risk analyses can often be measured or at least estimated with some confidence (see Sec. 4). They confer considerable flexibility on the model, which is not even fully exploited below when we restrict the analysis to $\varepsilon = 0$. In fact, once x_0, z_0 and ε

have been fixed, the utility function $u_\varepsilon(p, x)$ and, hence, the certainty equivalent $c_\varepsilon(p)$ can be determined for arbitrary risks p in each particular application. Risks can thus be quantitatively and consistently assessed and compared not only for one given set of parameters, but also in different decision contexts, with different x_0, z_0 and ε. For instance, dose risks of radiation from a given amount of one and the same radioactive material may be evaluated in different countries with reference to different tolerability limits (aspiration levels), or on the basis of different population densities making an impact on the *status quo* risk. Such country-specific differences will lead to different certainty equivalents. The corresponding differences in dose risk assessment of one and the same radioactive source are made explicit and precise in this way. Accounting for such differences may be instructive and useful in harmonising international radiological safety and security standards.

3.2 Standardisation of risk acceptance limits

Measures to mitigate the consequences of rare, severe catastrophic events have often been chosen to keep the risk of damage below cumulative life-time exposures to relevant comparable risks. In the area of radiation protection, for instance, international standards have been developed which assure that measures with major impact on public life are only taken if the doses averted are comparable to, or higher than, the average cumulative life-time exposure to natural radiation (ICRP 1991).

Such cumulative life-time risk concepts can also serve as a basis for judgements on the tolerability of risks related to the vulnerability of systems of physical protection by very unlikely but extreme actions such as terrorist attacks. Moreover, their use within a methodological framework of decision and utility theory provides new possibilities for defining tolerability thresholds, for developing protection standards, and for optimising solutions in physical protection.

Below we apply this new approach to an example of risks from severe accidents with, and attacks against, shipments of highly radioactive material. Our reasoning is based on a dose limit approach that can be straightforwardly extended to include life-time natural radiation doses. It is further based on the use of risk profiles from state of the art probabilistic risk analysis, and the evaluation of such information in terms of non-EU. The example indicates that the approach is viable in principle and that it may be used to translate radiation protection requirements into judgements about the tolerability of risks and the need for standardisation and improvement.

3.3 Cost efficiency of risk management measures

A similar argument applies to measurements of cost efficiency of risk management practices (McCormick 1981, Chap. 17; Royal Society Study Group 1992, Sec. 6.3). As an example, consider the problem of minimising the risk of transporting SNF by suitably rerouting highway shipments (Glickman and Sontag 1995). The operating costs per number of assemblies shipped are generally known. Further assume that the accident probabilities associated with each alternative route can be estimated (Sprung *et al*. 2000, Chap. 8). Then the difference between the certainty equivalents of any two transport risks gives the amount of risk reduction that can be achieved by changing from a less expensive, more risky route to a safer, or more secure, though more expensive one. In more general terms, the trade-off in risk reduction per dollar invested can be assessed by calculating and comparing the certainty equivalents of risks that are managed at different costs at different safety (security) levels.

It should be emphasised that although in this example SNF transport risk is assessed in utility terms, the relevant quantity measuring risk and risk reductions is the certainty equivalent, contrary to the conventional approach in terms of averages of radiological incident consequences (e. g., McCormick 1981, p. 359). It is commensurate to dose (rem, Sv) rather than utility, which after all is a dimensionless and purely theoretical (i. e., non-observable) quantity. It is one of the practical advantages of the present non-EU approach to risk analysis that it turns risk directly and consistently into a measurable quantity that can be priced, that is, whose monetary trade-off can be uniquely specified.

3.4 Mean value vs. certainty equivalent

The present non-EU theory admits an approach to risk assessment more subtle than the one provided by the probabilistic expectation of loss. To see this, let x_0 be a given maximum admissible radiation dose per person per unit time, and let μ_0 be the mean effective dose received by individuals per unit time of exposure to the normal *status quo* radiation. In our example below, μ_0 is the average dose risk per person per shipment of the incident-free truck transport, along a given travelling route, of a given number of assemblies of SNF. One generally has

$\mu_0 < x_0$. However, to make the example compatible with the conceptualisations of Section 2, one rescales the radiation dose x received,

$$x \;\rightarrow\; -x + x_0 \qquad (13)$$

so that $x_0 = 0$, $x < 0$ for doses “larger” than admissible (i. e., detriment is negative), and $\mu_0 > 0$. Accordingly, $z_0 > 0$ for the critical risk acceptance. The situation is shown in Figure 3. The risk p with mean $\mu(p) > 0$ and $c_0(p) > 0$ is acceptable, that is, $z(p) > z_0 > 0$, whereas p' has positive mean value but is marginally acceptable. There is also a parameter regime $z_0 > z(p'') > 0$ within which risks with mean $\mu(p'') > 0$ are nonetheless unacceptable. The average outcome is above the aspiration level, but the utility $U_0(p'')$ and the certainty equivalent $c_0(p'')$ are negative in those cases. The various situations depicted in Figure 3 demonstrate that it may indeed be misleading to compare simply mean doses with the acceptance limit and incident-free case to assess the tolerability of radiological risks.

3.5 Security risk management

In contrast to the safety risks of technical systems, security risks are notoriously hard to specify in probabilistic terms since they involve intentional human action. Nevertheless, one can take a “What-if” approach to the assessment and management of security risks to which the non-EU model may apply. The approach is based on the distinction between probabilities for scenarios, or security incidents, and probabilities for their likely consequences (cf. Kaplan and Garrick 1981; Múnera *et al*. 1997). One arbitrarily sets the probability of any such incident equal to 1, and concentrates on the probabilistic assessment of its consequences. Our utility model can then be employed to assess the potential loss or damage that may arise provided the incident occurs. This kind of restricted security risk analysis is a useful risk management approach, especially one to assess the cost efficiency of risk mitigation and damage prevention measures (Valentin 1999). Here we use it to draw, in a systematic fashion, conclusions from nuclear safety risk analysis to nuclear security risks.

4. Application: SNF Transport Security Risks

4.1 Modelling security risks

SNF transport has repeatedly been considered to be a potential target of terrorism (Múnera *et al*. 1997; Chapin *et al*. 2002). Yet the security risks associated with shipments of nuclear and radioactive material are hard to evaluate quantitatively since in applications the probability of an attack cannot usually be specified in any meaningful way, not even in instances in which terrorist attacks on highway or railway traffic are frequent (Múnera *et al*. 1997). Since probabilistic assessments of spent fuel security risks would nevertheless be highly desirable from a risk management perspective, we choose an approach which is somewhat more restricted in scope than the attempt to assign probabilities to attacks. We concentrate on the assessment of the likely consequences conditional on the occurrence of a security incident. We start from the available quantitative analyses of the safety accident risks of SNF transport. We further assume that the probabilities of the radiological consequences of a terrorist attack on a truck or train shipment of nuclear material are, to some extent at least, similar to those of potential transport accidents so that, in first approximation, the former can be modelled by the latter. The assumption is based on the fact that the impact on the population of the bombing of, or use of anti-tank weapons against, a transport cask is subject to the same random constraints as those of safety accidents such as weather conditions, route parameters and geographical variations in population number. It is this impact to which the present approach applies. On the other hand, the radiation dose and amount of nuclear material released from a "successful" bombing of a spent fuel shipment may be larger than the ones released even from severe transport accidents (Lyman 1999). Experts have repeatedly disputed this hypothesis (e. g., Chapin *et al*. 2002). But even if the critics are right, it should be subject to probabilistic analysis for the same reason and in the same way as the SNF accident risks.

Now the impact of an increased amount and modified inventory of the radioactive material released from an attack can be adequately treated by modifying the source term magnitude and fractions of failed rods and released radionuclide inventory in the computer calculations of the dose risk probabilities (cf. Sprung *et al*., 2000, Sec. 2.5). Detailed specifications of the source terms are beyond the scope of our consequence analysis, however. But we can determine the certainty equivalent as a function of the probability that an accident will be severe enough to cause a spent fuel cask to fail and release radioactivity to the atmosphere. Sufficiently large values of that probability then provide a quantitative estimate of the radiation exposure of the public, and its probabilistic distribution, arising from a successful attack. In this way, security risk management measures suitable to decrease this probability can also be assessed with regard to the dose risk reduction they allow.

4.2 The governing parameters

To provide an example of applied PP security risk analysis, we use the data, parameters and results of Sprung *et al.* (2000) on spent fuel truck and rail transport in the US, and related US NRC documents. These results give a detailed probabilistic account of incident-free and accident dose risks, or complementary cumulative distribution functions (CCDFs), to which our approach can be directly applied. Recall that the CCDF is defined as $F_c(x) = \Sigma_{x<y} p(y) = 1 - F(x)$, where x is the individual dose received. The calculations of the CCDFs have been based on data and computer models specifying numerous diverse constraints and parameters governing release and exposure such as wind speed and direction, population density, US highway and railway accident statistics, package inventories, and cask structural impact responses, to name a few.

We evaluate, in terms of its (dis-)utility and certainty equivalent, the risk to the population exposed to the plume of radioactive material released in a hypothetical terrorist attack on a generic Type B SNF cask. Of the four Type B casks considered by Sprung *et al.* (2000), we mainly apply our model to the generic steel-lead-steel cask. For the sake of definiteness, we calculate the governing parameters x_0, z_0 and ε required by our approach for 1 shipment of 1 assembly of pressure water reator (PWR) SNF by truck transport on US interstate highways involving short stops only (refuelling, recreation of crew, meals, but no sleep; see Sprung *et al*. 2000, Chap. 8). Larger packages and numbers of assemblies shipped per truck or railway waggon are assumed to increase the risk of the emitted radiation dose roughly in proportion to the size of the shipment. Risk is measured in units of dose (person rem). It is proportional to the transport route length and decreases inversely with the travelling speed. To calibrate our model in terms of x_0, z_0 and ε, we use the representative route data provided by Sprung *et al*. (2000, Chap. 3) for the incident-free radiation risk, but turn to one of the illustrative real route cases for our security incident analysis (Sprung *et al*. 2000, Sec. 8.10).

Since population doses from SNF truck transport emitted along routes of arbitrary length are proportional to the overall route length, the population dose from a reference route is required. The example of the 420 km-route studied by Mills and Neuhauser (1999) is typical of SNF truck routes in the US. The authors have subdivided the total route into 29 segments of variable length L_i (km), each passing through an urban, suburban or rural area with average near-route number N_i of persons exposed to radiation with dose rate r_0 during time $\_t_i$, where $\_t_i$ is the duration of exposure of the population in the neighbourhood of route segment i while a truck is travelling along segment i, $1 \leq i \leq 29$ (for further details of the population statistics and distribution along transport routes see also Mills and Neuhauser 1998, 2000). To obtain the maximum admissible dose x_0 from a single shipment, we choose the average dose rate on the basis of the US NRC maximum admissible individual dose rate

$$r_0 = 0.1 \text{ rem person}^{-1} \text{ a}^{-1}$$

so that

$$x_0 = r_0 \Sigma_{i \leq 29} \_t_i N_i$$

The average travelling speed of an SNF transport truck has been estimated by Sprung *et al*. (2000, Tab. 3.3, p. 3-7) at 55 mph = 88 km h^{-1}. Hence, $\_t_i$ =
L_i /88 km h^{-1} and

$$x_0 = (r_0/88 \text{ km h}^{-1}) \Sigma_{i\leq 29} L_i N_i$$

$$= 0.01 \text{ person rem} \tag{14}$$

where $\Sigma_{i\leq 29} L_i N_i$ = 76500 km according to Mills and Neuhauser (1999, Tab. II). Altogether, the value of 0.01 person rem for x_0 corresponds to the (hypothetical) overall population dose released by a single SNF truck shipment of 1 assembly emitting at the rate r_0 while the truck covers a route length of 420 km at an average speed of 88 km h^{-1}.

In a similar fashion, we choose the statistical parameters of the *status quo* risk as the incident-free total transport dose risk μ_0 with standard deviation σ_0 for the 420 km-route case (Mills and Neuhauser 1999, Tab. III)

$$\mu_0 = 0.008 \text{ person rem} \tag{15a}$$

$$\sigma_0 = 0.006 \text{ person rem} \tag{15b}$$

The use of (15) within our analysis raises the following two problems. First, the population numbers applied to calculate the numerical values (15) are based on detailed geographical population data and, hence, are different from the near-route values N_i entering Equation (14) (see Mills and Neuhauser 1999 for discussion of their Tables II and III). Fortunately, the difference is not significant, however. Mills and Neuhauser applied a Chi-Square test to the dose risk distributions each based on one of the alternative population statistics, with the result indicating that the two distributions are roughly the same. We exploit this insensitivity of the overall population dose risk to the underlying population statistic, approximating the real incident-free dose risk s by a two-component distribution $p_1^0 = p_2^0 = 0.5$ with the possible outcomes $\mu_0 \pm \sigma_0$. Then, by construction, p^0 has the mean μ_0 and standard deviation σ_0. The second problem refers to the utility indifference of s and p^0. As has been mentioned above, the construction of an indifferent two-component probability distribution from an arbitrary discrete *status quo s* requires a numerical iteration procedure describe elsewhere (Geiger 2002a, App. B). Application of this procedure to the data of Mills and Neuhauser (1999, Tab. III) shows that our calculation of p^0 does not improve significantly if it is carried beyond the first approximation step.

The critical risk acceptance z_0 now is

$$z_0 = (-\mu_0 + x_0)/\sigma_0 = 0.33 \tag{16}$$

Equation (15) means that although (15a) is positive, its contribution to z_0 is negative since doses received are considered to be detriments and, hence, negative. On the other hand, since (15a) is still below the maximum admissible dose (14), $-\mu_0 + x_0$ and z_0 are altogether positive. Observe that in contrast to μ_0, σ_0 and x_0, the parameter z_0 is invariant to changing route length.

In many applications, the degree ε to which an attack would seem more likely than the incident-free case, can be assumed to be very small so that

$$\varepsilon \;\_\; 0. \tag{17}$$

We once more emphasise that the chance ε that a risk gets resolved prior to the resolution of other risks that have also been committed to can be very important for the assessment of one risk in the presence of others. As for a detailed account of this ε-dependence in terms of two-stage lotteries, see Geiger (2002a). More generally, the problem is treated under the rubric of *status quo*- or background-dependent decision making under risk in the literature (e. g., Pratt 1988).

The utility function $u_0(p^0, x)$ corresponding to Equations (14) to (17) is shown in Figure 2.

4.3 Security incident risk

We first treat the truck transport of a generic steel-lead-steel PWR SNF cask along the illustrative real route of 4800 km with “no sleep” considered by Sprung *et al*. (2000, Tab. 8.7, p. 8-29, and Subsec. 8.10.1). Except for the route length, all parameters and input data to the calculation of the dose risk CCDF are as in the incident-free case. Figure 4 shows various CCDFs from a set of Monte Carlo samples of dose risk including the mean, 5th, 50th (median), and 95th percentile curve of the set (after Sprung *et al*. 2000, p. 8-30). We choose the CCDF of mean values for application. The maximum admissible dose rate x_0 is increased from (14) roughly in proportion to the route length by

4800 km/420 km = 11.4,

$$x_0^* \ = \ 0.114 \text{ person rem}$$

The dose risk distribution of Figure 4 is proportional to the truck accident probability which has been estimated at

$$p_{acc} \ = \ 1.8\_10^{-3}$$

on the average per trip of 3000 miles, or 4800 km (US NRC 2000, p. 18). Dividing the CCDF values by p_{acc}, one gets the dose risk conditional on the occurrence of an accident or, in our interpretation, a violation of PP.

The CCDF $F_c(x)$ shown in Figure 4 has the mean

$$\mu \ = \ 9.53\_10^{-7} \text{ person rem}$$

(Sprung *et al*. 2000, Tab. 8.8, p. 8-36). It is discontinuous at $x = 0$ since by definition $F_c(0) = 1$ while from Figure 4 $\lim_{x\to 0}F_c(x) = F_c^* \_ \ 10^{-7} << 1$. We first calculate $c_0(F_c)$, which is simply the certainty equivalent of the dose risk from a truck accident, with no security incident being involved. To this purpose, we must give F_c an indifferent two-component representation p similar to p^0 entering (15) so that $U_0(F_c) = U_0(p)$ with the governing parameter $z(p)$. As for the details of the construction of such a p from a continuous CCDF, see Geiger (2002b). We partition $F_c(x)$ into two components p_1, p_2 respectively summing up the probabilities of all negative and all positive x-values in the normalisation (13),

$$p_2 = F_c(0) - F_c(x_0{*}) = 1 - F_c(x_0{*}) \_ 1 - 8\_10^{-8} \_ 1 \qquad (18a)$$

$$p_1 = 1 - p_2 \_ 8\_10^{-8} \qquad (18b)$$

with the outcomes $x_1' < 0$, $x_2' > 0$ so that $\mu' = p_1x_1' + p_2x_2'$, where dashed symbols denote quantities in the normalisation (13). To determine x_1' and x_2', observe that the area under the mean CCDF in Figure 4 is μ (Sprung *et al.* 2000, p. 8-14). We put $\mu' = -\mu + x_0{*} = \mu_+'$ $+ \mu_-'$ and $\mu_+' \_ -\mu_+ + x_0{*} \_$

$-x_0{*}F_c{*} + x_0{*}$, with $\mu_+ \_ x_0{*}F_c{*}$ being the area under the F_c-curve between 0 and $x_0{*}$, that is, the "positive" contribution to μ for doses x smaller than the maximum admissible value $x_0{*}$. We get

$$x_2' = \frac{\mu_+'}{p_2} = \frac{-\mu_+ + x_0{*}}{p_2} = \frac{-x_0{*}F_c{*} + x_0{*}}{p_2} \_ x_0{*} \qquad (19a)$$

$$x_1' = \frac{\mu_-'}{p_1} = \frac{\mu' - \mu_+'}{p_1} = \frac{-\mu + \mu_+}{p_1} \_ -1 + x_0{*} \qquad (19b)$$

and, analogously to (10),

$$z(p) = \frac{-\mu + x_0{*}}{(x_2' - x_1')\sqrt{p_1p_2}} \_ x_0{*}p_1^{-1/2} \_ 360 >> \quad z_0 \qquad (20)$$

Equation (20) confirms the conclusion arrived at by Sprung *et al.* (2000, p. 8-18) on the basis of comparison of expected values, namely that, for any truck shipment, incident-free dose risks greatly exceed accident dose risks. It follows that the latter are acceptable in the technical sense of the risk acceptance terminology introduced in Subsection 2.2. Accordingly, from (18a) and $U_0(p) \_ u_0(p, x_2') \_ u_0(p, c_0(p))$,

$$c_0(F_c) = c_0(p) = -x_0{*}F_c{*} + x_0{*} > 0. \qquad (21)$$

Although because of (16) z_0 is risk averse, and $z(p)$ is even more so according to (20), $c_0(p)$ is slightly larger than the risk neutral value $\mu' = -\mu + x_0^*$, which means risk proneness. However, this apparent inconsistency vanishes considering the convexity of the utility curve in the neighbourhood of the origin (x_2' small, p_2 _ 1) that is indicated in Figure 2.

We proceed to assess the dose risk for the case that an accident, or, in our interpretation, a violation of PP has occurred. Assuming that $F_c(x)$ is roughly proportional to the average accident probability, the relevant dose risk distribution function is $G_c(x)$, with

$$G_c(x) \;=\; F_c(x)/p_{acc}, \quad x > 0 \qquad (22)$$

$$G_c(0) \;=\; F_c(0) \;=\; 1$$

$G_c^* = \lim_{x\to 0} G_c(x) = F_c^*/p_{acc}$ _ 5_10^{-5}

To calculate $c_0(G_c)$, we proceed as in Equations (18) to (20), everywhere replacing F_c by G_c, and μ by μ/p_{acc}. Observing that p_1 _ 4_10^{-5} and p_2 _ 1, and neglecting small terms, we find

$$z(p) \;=\; 16 \;>>\; z_0$$

$$c_0(G_c) \;=\; c_0(p) \;=\; -x_0^* G_c^* + x_0^* \;>\; 0.$$

Although $z(p)$ has decreased considerably from (20), it is still larger than the critical value (16), with the certainty equivalent being positive. This means that, provided a violation of PP has occurred, the resulting dose risk is still tolerable when assessed on the basis of the incident-free case.

The situation changes if we make the dependence of the CCDFs on the severity of an attack explicit. Let p_{no} be the probability that the shipment occurs without an incident severe enough to cause a release of radioactivity to the atmosphere. Then $F_c^* = p_{acc}(1 -$

p_{no}) (Sprung *et al*. 2000, p. 8-64). Given F_c^* _ 10^{-7}, the probability of no release of radioactivity is

$$p_{no} = 1 - 5\_10^{-5} \qquad (23)$$

However, in the event of a terrorist attack, p_{no} will generally decrease from its accident value (23), and $F_c(x)$ will accordingly increase. We incorporate this effect by introducing the variable η, letting η vary between the accident value (23) $\eta = 1 - p_{no} = 5\_10^{-5}$ (weak attack) and $\eta = 1$ (heavy attack). This procedure implies two things. First, we admit a parallel vertical shift of the CCDF depending on the severity of an attack. Second, to provide a simple PP incident model, we neglect modifications of the source terms that may change the shape of a CCDF in the security incident case. We have estimated the implications of this neglect both qualitatively and numerically in an order-of- magnitude fashion, as will be discussed in some more detail in the concluding section. Altogether, we consider the distribution function

$$H_c(\eta, x) = \frac{\eta}{1-p_{no}} G_c(x), \qquad x > 0 \qquad (24)$$

$$H_c(\eta, 0) = G_c(0) = F_c(0) = 1, \qquad 5\_10^{-5} \leq \eta \leq 1,$$

$$H_c^*(\eta) = \lim_{x\to 0} H_c(\eta, x) = \frac{\eta}{1-p_{no}} \frac{F_c^*}{p_{acc}} = \eta$$

It gives the dose risk conditional on the occurrence of an attack, with the incident severity η, that is, the chance η of release of radioactivity to the atmosphere. This implies $H_c(1, x) = 10^7 F_c(x)$ and $H_c^*(1) = H_c(1, 0) = 1$. Alternatively, if $\eta = 1 - p_{no}$, this leads back to the case $H_c(1-p_{no}, x)$ _ $G_c(x)$ treated above. Analogously to (18), we get

$$p_2 = 1 - H_c(\eta, x_0^*) \ \_ \ 1 - \eta$$

$$p_1 \ \_ \ \eta$$

$$z(p) = \frac{-\eta + x_0^*}{\sqrt{\eta(1-\eta)}}$$

The parameter $z(p)$ falls below the critical value (16) for $\eta > 0.02$, but remains positive for $\eta < 0.1$, and so does the mean

$$\mu' = -\frac{\mu\eta}{p_{acc}(1-p_{no})} + x_0^* = -\eta + x_0^*$$

(Fig. 5). For $0.02 < \eta < 0.1$ (shaded area), the certainty equivalent is negative and, hence, the risk H_c unacceptable although the expected dose $\mu\eta p_{acc}^{-1}(1-p_{no})^{-1}$ in case of a successful attack is still smaller than the maximum admissible dose $x_0^* = 0.1$ person rem. The certainty equivalent is shown as a function of μ' in Figure 6.

The results shown in Figures 5 and 6 mean that attack risks with low probability $\eta < 2\%$ of severe consequences are acceptable, and attack risks with moderate or high probability $\eta \geq 10\%$ of severe consequences are unacceptable, independently of whether they are assessed in terms of their expected doses or certainty equivalents. The two assessment modes yield contradictory results in the intermediate probability range of the order of a few percent (shaded areas) where risk acceptance attitudes may be particularly controversial in social perceptions of nuclear security threats. The present approach may then prove helpful in clarifying (e. g., public) disputes about the tolerability of the security risks in point. It may also prove useful for calculating the costs that need to be incurred to decrease η below the threshold to non-acceptance.

5. Discussion and Extensions

The previous section illustrates various points we have made with regard to quantitative risk assessment as applied to PP of nuclear material. Above all, conventional estimates of the tolerability of nuclear and radiological risk can be improved beyond what purely statistical analysis can achieve, namely, the comparison of expected dose and admissible dose limit. This conclusion reflects the common knowledge from risk and decision analysis that the expected value of a random variable may be misleading as a risk indicator (neglect of risk aversion, inappropriateness to one-shot decision problems, underrating the importance of tail probabilities, etc.). In particular, between the high and low risk regimes there is an intermediate domain in which mean doses below the admissible limit may nevertheless be non-acceptable. This result is clearly an outcome of the present utility approach which assesses a given risk in terms of the CCDF as a whole rather than on the basis of one single statistical parameter. Intuitively, risks identical by expected value, but different by distribution and, especially, by tail probabilities, are not generally indifferent. Our approach to risk assessment makes the utility differences explicit in a quantitative, realistic, consistent and, after all, computationally simple fashion.

In a sense, security risks are different from safety risks. By this we mean the possibility that one and the same set of likely consequences may be acceptable or not, depending on whether the consequences are elicited by an accident or by a security incident. Within the present conceptual framework, the difference arises from the need to evaluate security risk consequences given the occurrence of an incident, whereas no such need exists for safety risks, at least as long as accident probabilities can be estimated and security incident probabilities cannot. Conditioning CCDFs on the occurrences of hazardous events not only increases the expected value and (dis-)utility of loss, it may also transform risk attitudes qualitatively. To illustrate this conclusion, consider once more the complementary cumulative distribution F_c of Figure 4 and its successive conditionings

(22) and (24). Let $0.02 < \eta < 0.1$. Then F_c , G_c and H_c are all below the admissible dose limit x_0*, with F_c and G_c being acceptable, but not H_c. It is the "What-if" perspective of the present approach to security risk assessment that makes the difference.

The latter result offers an explanation for the common observation that risks with unknown or unspecified probabilities tend to meet with unusually low degrees of tolerance in the public. The effect is often attributed to the risks of rare catastrophic events in general, but may also be specific of security risks such as nuclear terrorism. In situations that are uncertain in the sense of unknown outcomes and unspecified outcome probabilities, people cannot but assess risks from a "What if" perspective, thereby systematically overestimating even low-dose risks that would otherwise be accepted. The significance of this conclusion for risk management, public policy making and international standardisation of nuclear security practices would seem obvious.

The relationship between the incident severity η and certainty equivalent $c_0(H_c)$ depicted in Figure 6 provides a simple framework for cost efficiency estimates of measures to reduce SNF transport security risks. Measures such as rerouting or augmented escorting of SNF shipments, or improving the cask structural impact response behaviour will generally reduce η to an extent that can be quantified with some accuracy, and similarly so for the costs involved. The c_0-curve in Figure 6 then gives the corresponding amount of risk reduction.

A more detailed cost efficiency calculation would include the numerous and complex modifications of the CCDFs that can be achieved by technical and organisational measures to protect SNF shipments. One would then have to compare the certainty equivalents of CCDFs of different shape and axis intercepts. To estimate the effects to be expected from more detailed analyses of this kind, we applied our model to some of the many CCDFs Sprung *et al*. (2000) calculated for different types of spent fuel casks, transport modes (truck, rail), routes (length, travelling time, regional climate and weather, near-route population numbers, etc.) as well as size and nuclear inventory of shipments (number of SNF assemblies, PWR and BWR SNF, etc.). What we found from a preliminary overview was that the certainty equivalent largely varies with influence factors such as route length and shipment size to which the dose risk is roughly proportional. Thus, when we estimated the certainty equivalents of the risks of different transport modes (truck, rail) for equal route lengths on a per assembly and order of magnitude basis, no significant departures from the certainty equivalent ratios of F_c, G_c and H_c described above could be observed. We therefore conclude that the present treatment of SNF transport security incidents, though computationally simple, covers basic quantitative features of this type of risk of violation of PP of nuclear material.

References

Chankong, V., Haimes, Y. Y., 1983. Multiobjective Decision Making: Theory and Methodology. North Holland, New York-Amsterdam, 1983.

Chapin, D. M., Cohen, K. P., Davis, W. K., Kintner, E. E., Koch, L. J., Landis, J. W., Levenson, M., Mandil, I. H., Pate, Z. D., Rockwell T., Schriesheim, A., Simpson, J. W., Squire, A., Starr, C., Stone, H. E., Taylor, J. J., Todreas, N. E., Wolfe, B., Zebroski, E. L., 2002. Nuclear Power Plants and their fuel as terrorist targets. Science 297, pp. 1997-1999.

Committee on High-Level Radioactive Waste through Geological Isolation, 2001. Disposition of High-Level Waste and Spent Nuclear Fuel: The Continuing Societal and Technical Challenges. National Academic Press, Washington, D.C.

Committee on Science and Technology for Countering Terrorism, 2002. Making the Nation Safer: The Role of Science and Technology in Countering Terrorism. National Academic Press, Washington, D.C.

Evans, E. W., Verlander, N. Q., 1997. What is wrong with criterion FN-lines for judging the tolerability of risk? Risk Analysis 17, pp. 157-168.

Fishburn, P. C., 1988. Non-Linear Preference and Utility Theory. John Hopkins UP, Baltimore.

French, S., 1988. Decision Theory. Ellis Horwood, Chichester.

Fullwood, R. R., 2000. Probabilistic Safety Assessment in the Chemical and Nuclear Industries. Butterworth-Heinemann, Boston.

Geiger, G., 2001. A dynamic account of rational decision making under uncertainty: the case of risk assessment in hazardous technological systems. In: Matthies, M., Malchow, H., Kriz, J. (Eds.). Integrative Approaches to Natural and Social Dynamics. Springer, Berlin, pp. 305-318.

Geiger, G., 2002a. On the statistical foundations of non-linear utility theory: the case of status quo-dependent preferences. European Journal of Operational Research 136, pp. 449-465.

Geiger, G., 2002b. Pragmatic constraints on rationality: an axiomatic approach to non-expected utility theory. Manuscript, submitted for publication.

Geiger, G., 2002c. Risk acceptance from non-linear utility theory. Manuscript, submitted for publication.

Glickman, T. S., Sontag, M. A., 1995. The tradeoffs associated with rerouting highway shipments of hazardous materials to minimize risk. Risk Analysis 15, pp. 61-67.

IAEA, 2001. Nuclear Security & Safeguards. IAEA Bulletin 43, No. 4.

ICRP (International Commission on Radiological Protection), 1991. Principles for Intervention for Protection of the Public in a Radiological Emergency. Annals of the ICRP 22, No. 4.

Jorissen, R. E., Stallen, P. J. M. (Eds.), 1998. Quantified Societal Risk and Policy Making. Kluwer, Dordrecht.

Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. Econometrica 47, pp. 763-291.

Kaplan, S., Garrick, B. J., 1981. On the quantitative definition of risk. Risk Analysis 1, pp. 11-27.

Luce, R. D., 2000. Utility of Gains and Losses: Measurement-Theoretical and Experimental Approaches. Lawrence Erlbaum Publishers, London.

Lyman, E. S., 1999. A critique of physical protection standards for transport of irradiated materials“. Paper presented at the 40th Annual Meeting of the Institute of Nuclear Materials Management, Phoenix, AZ.

McCormick, N. J., 1981. Reliability and Risk Analysis: Methods and Nuclear Power Applications. Academic Press, New York.

Mills, G. S., Neuhauser, K. S., 1998. Urban risks of truck transport of radioactive material. Risk Analysis 18, pp. 781-785.

Mills, G. S., Neuhauser, K. S., 1999. Statistical evaluation of population data for calculation of radioactive material transport accident risk. Risk Analysis 19, pp. 613-619.

Mills, G. S., Neuhauser, K. S., 2000. Geographical resolution issues in RAM transport risk analysis. RAMTRANS 11, pp. 295-299

Múnera, H. A., Canal, M. B., Muños, M., 1997. Risk associated with transportation of spent nuclear fuel under demanding security constraints: the columbian experience. Risk Analysis 17, pp. 381-389.

Nilsson, A., 2001. Security of material: the changing context of the IAEA's programme. IAEA Bulletin 43, No. 3, pp. 12-15.

Okrent, D., 1981. Industrial risks. Proceedings of the Royal Society of London A376, pp. 133-148.

Ortiz, P., Friederich, V., Wheatley, J., Oresegun, M., 1999. Lost and found dangers: orphan radiation sources raise global concern. IAEA Bulletin 41, No. 3, pp. 18-21.

Pratt, J., 1988. Aversion to one risk in the presence of others. Journal of Risk and Uncertainty 1, pp. 395-413.

Quiggin, J., 1993. Generalized Expected Utility Theory. Kluwer, Boston.

Royal Society Study Group, 1992. Risk: Analysis, Perception and Management. The Royal Society, London.

Sprung, J. L., Ammerman, J. D., Breivik, N. L., Dukart, R. J., Kanipe, F. L., Koski, J. A., Mills, J. S., Neuhauser, K. S., Radloff, H. D., Weiner, R. F., Yoshimura, H. R., 2000. Reexamination of Spent Fuel Shipment Risk Estimates, Vols. 1,2. Sandia National Laboratories, Albuquerque, NM.

Stallen, P. J. M., Geerts, R., Vrijling, H. K., 1996. Three conceptions of quantified societal risk. Risk Analysis 16, pp. 635-644.

Starmer, C., 2000. Developments in non-expected utility theory: the hunt for a descriptive theory of choice under risk. Journal of Economic Literature 38, pp. 332–382.

Starr, C., Rudman, R., Whipple, C., 1976. Philosophical basis for risk analysis. Annual Review of Energy 1, pp. 629-662.

US NRC, 2000. Discussion Draft: An Updated View of Spent Fuel Transportation Risk. US NRC, Washington, D. C.

Valentin, J., 1999. What if? ICRP guidance on potential radiation exposure. IAEA Bulletin 41, No. 3, pp. 45-48.

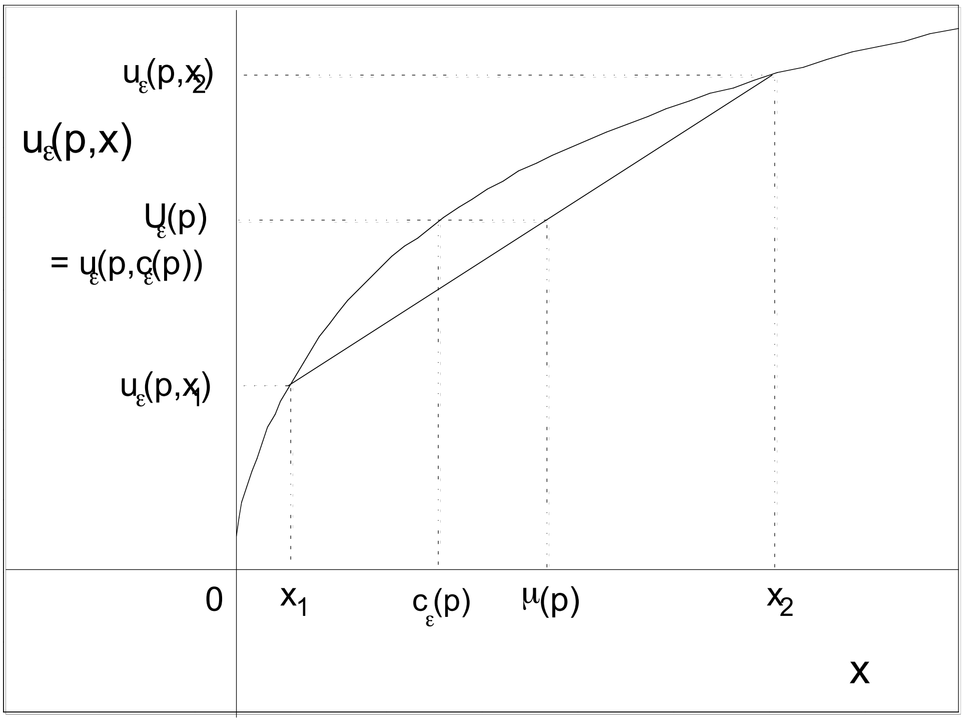

Fig. 1

Fig. 1. Utility function $u\_(p, x)$ and expected utility $U\_(p)$ of a two-point probability distribution p defined for x_1, x_2, with the mean $\mu(p)$ and certainty equivalent $c\_(p)$.

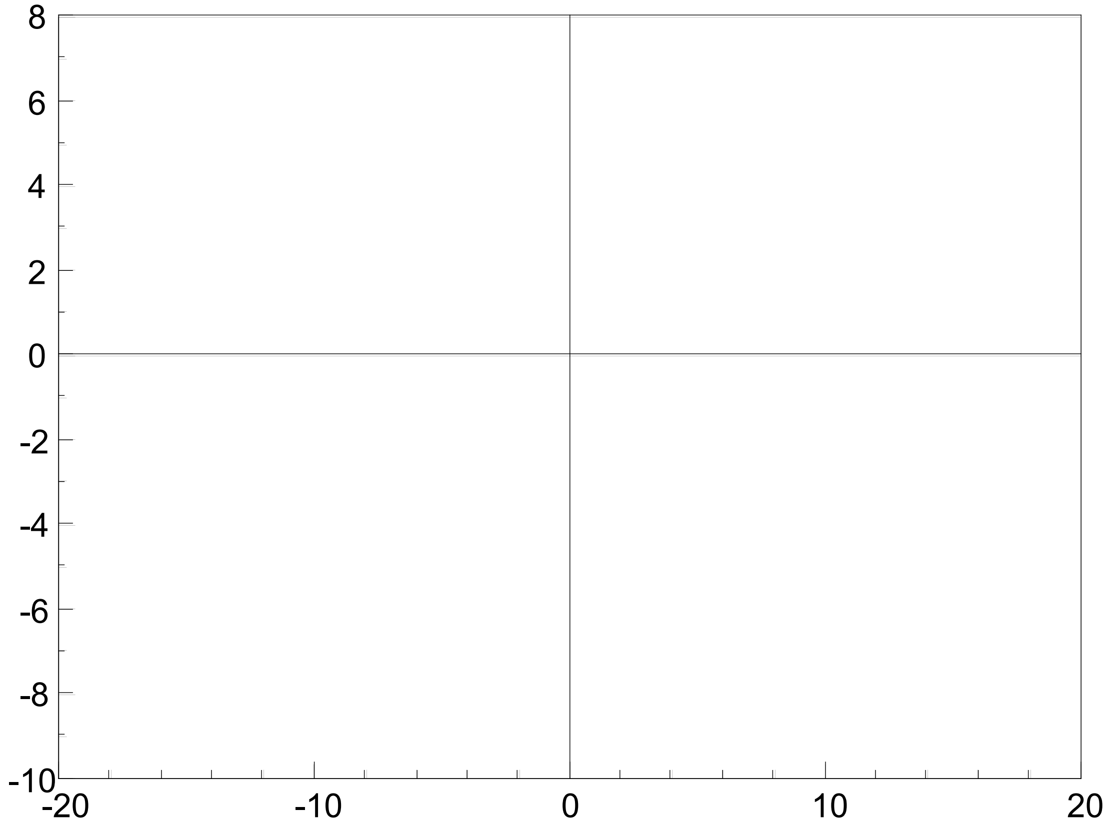

Fig. 2. Utility function $u_0(p^0, x)$ of marginally acceptable risk p^0 with parameter z_0 = 0.33.

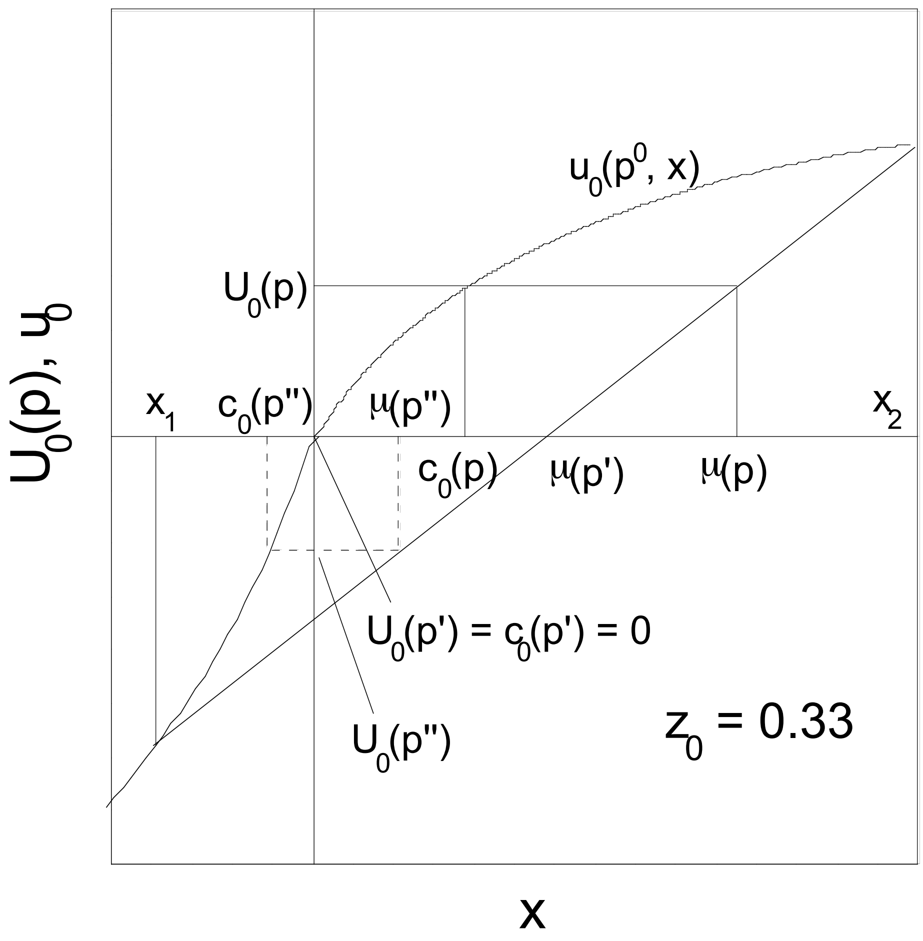

Fig. 3

Fig. 3. Risks with positive mean μ. Positive (solid lines) and negative (dashed lines) certainty equivalent.

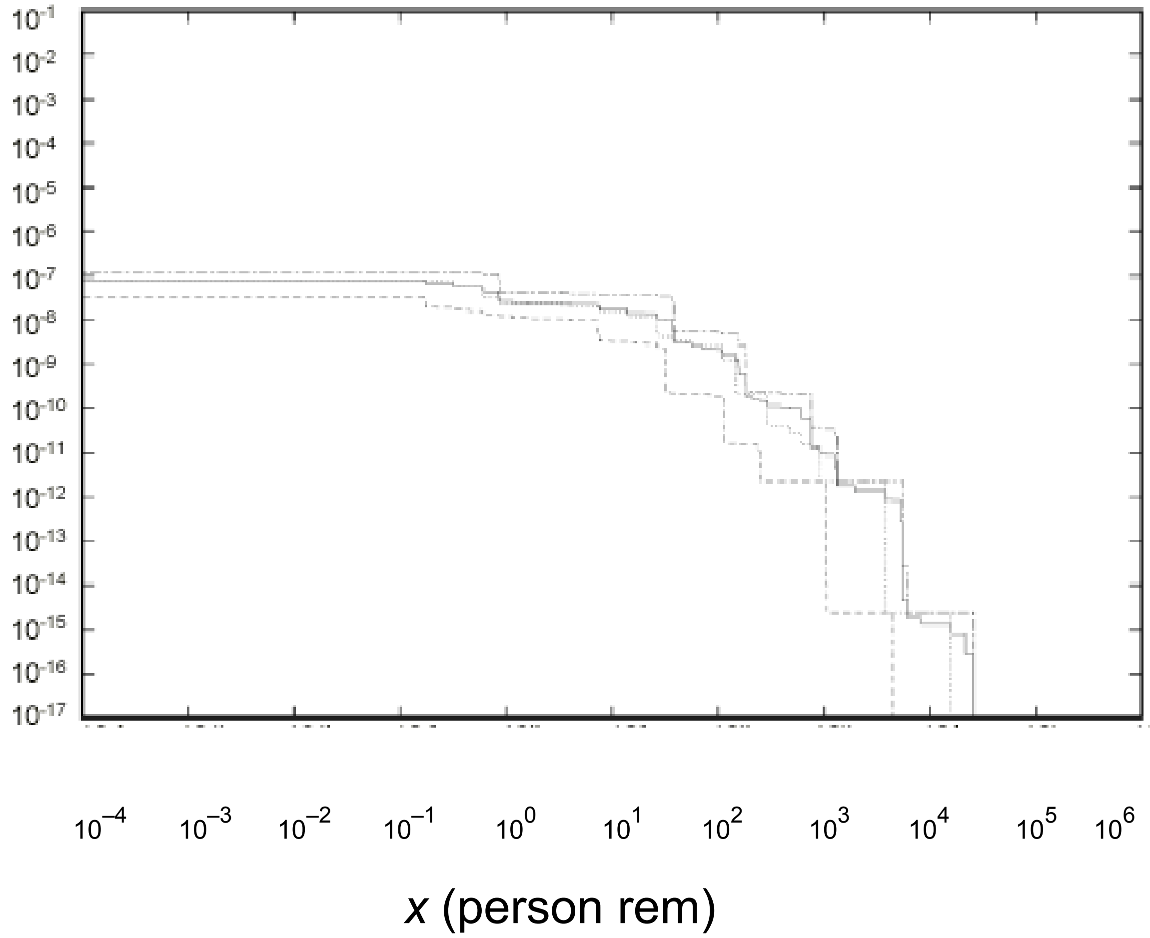

Fig. 4

Fig. 4. CCDFs from a set of Monte Carlo samples of dose risk, with 5th, 50th (median), mean and 95th percentile curve of the set (after Sprung *et al*. 2000, p. 8-30).

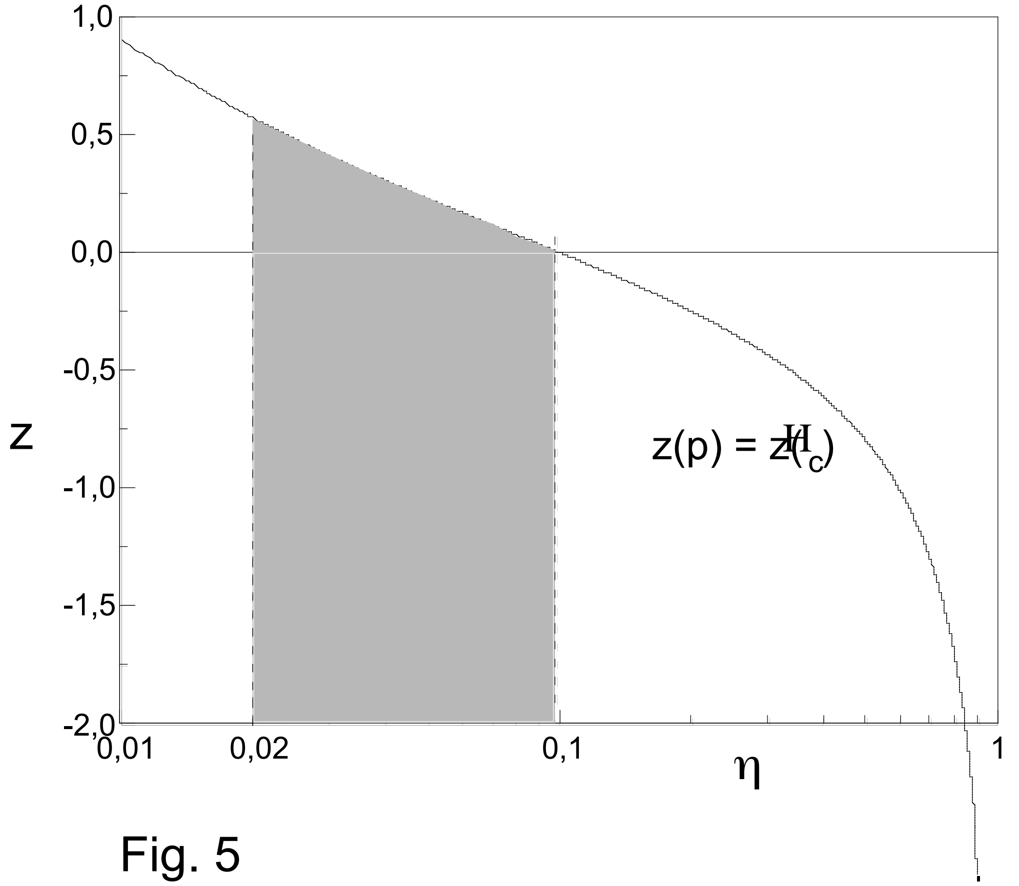

Fig. 5

Fig. 5. Parameter $z(p)$ associated with the risk H_c as a function of incident severity η. For $0.02 < \eta < 0.1$ (shaded area), the certainty equivalent is negative and, hence, the risk H_c unacceptable although the expected dose
$\mu \eta p_{acc}^{-1}(1-p_{no})^{-1}$ is smaller than the maximum admissible dose
x_0*_ 0.1 person rem.

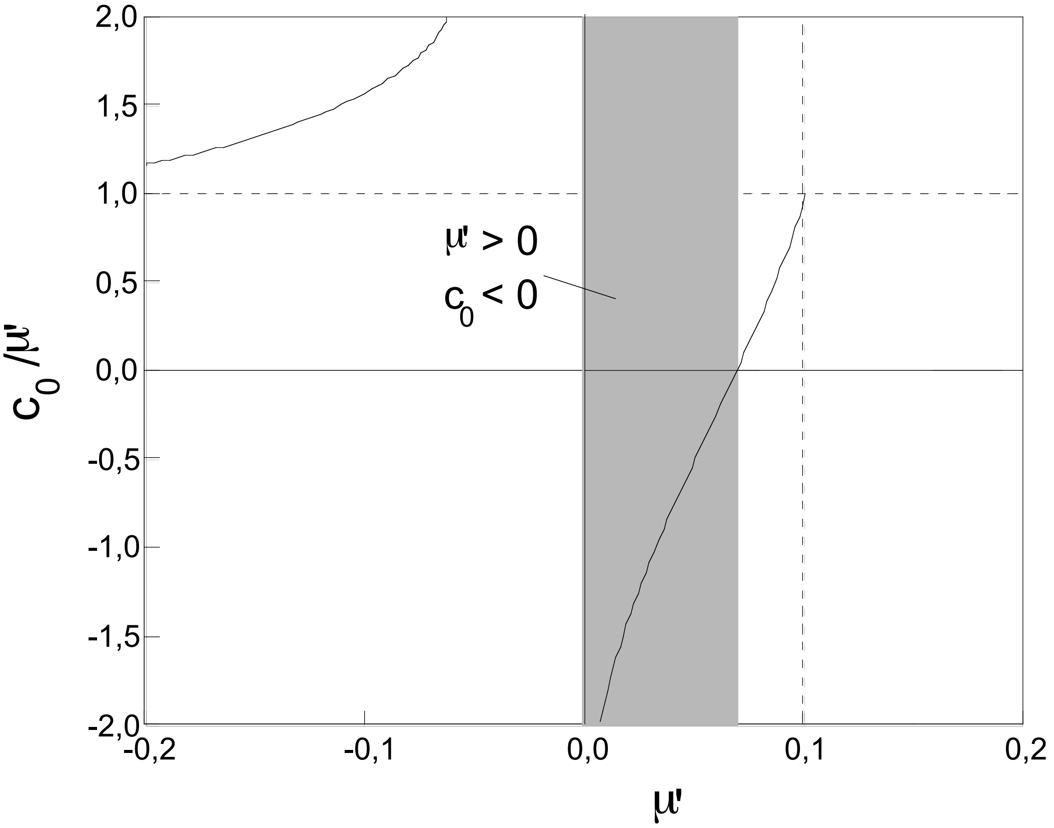

Fig. 6

Fig. 6. Ratio of certainty equivalent c_0 to expected value μ' as a function of μ' in the normalisation (16).