Safety, Reliability and Risk Analysis: Theory, Methods and Applications – Martorell et al. (eds)

The optimization of system safety: Rationality, insurance, and optimal protection

R.B. Jongejan & J.K. Vrijling
Delft University of Technology, Delft, The Netherlands

ABSTRACT: The optimization of system safety involves a tradeoff between risk reduction and investment cost. The expected utility model provides a rational basis for risk decision making within the context of a single decision maker. Yet the costs of risk reduction and the consequences of failure are generally distributed over a number of individuals. How to define optimality in terms of (different) individual preferences? This paper applies insights from welfare and financial economics to the optimization of system safety. Failures of technological systems are typically uncorrelated with market returns, so that technological risks could—at least theoretically—be financed against expected loss. It will be shown that mechanisms for risk transfer such as insurance sometimes allow us relax assumptions that would otherwise have to be made concerning the exact shape of individual utility functions and the interpersonal comparability of utility. Moreover, they allow us to lower risk premiums as insurance premiums are typically lower than certainty equivalents. The practical significance of these results is illustrated by a case study: the optimal design of a flood defense system.

Keywords: system safety, optimization, expected utility, insurance.

1 INTRODUCTION

Quantitative risk analysis is frequently used as a basis for balancing investments against risk reduction, a practice that is often linked to the expected utility framework. Von Neumann and Morgenstern (1944) and Savage (1974) construct a utility function that describes the preference ordering of a rational individual, and show that the individual, faced with uncertainty, ranks actions on the basis of the expected utility of their consequences. The idea that people maximize the expected utility of wealth rather than expected wealth itself was first introduced by Bernoulli (1738). Amongst other, the Joint Committee on Structural Safety has embraced the expected utility framework for the comparison of decision alternatives (Faber et al., 2007), it underlies the Life Quality Index methodology that can be used to optimize measures that increase life expectancy (Rackwitz, 2004), and it serves as the basis for the optimization of flood defenses in the Netherlands (Van Dantzig, 1956; Vrijling et al., 1998).

Unfortunately, the expected utility model is not as rational as is often suggested when it is applied in a context in which a single decision affects multiple individuals. It will be shown however that the complexities of collective choice can sometimes be avoided when mechanisms for risk transfer operate efficiently. Moreover, it will be shown that mechanisms for risk transfer, such as insurance, can have a profound impact on the optimization of system safety.

The text is organized as follows. First, the expected utility framework is introduced to discuss the interplay between insurance and prevention for the simple one person case. Insurance theory typically deals with this subject out of concern for moral hazard (e.g. Arrow, 1971; Kunreuther, 1996), but the underlying logic applies equally to the optimal design of technological systems such as flood defenses. The expected utility framework is then extended to a setting in which a single decision affects multiple individuals. A case study, the optimal design of a flood defense system, will finally illustrate the practical implications of our results.

2 THE ONE PERSON CASE

2.1 *The rational utility maximizer*

Let us first consider a single, rational decision maker that lives in a world of perfect certainty. Denote the admissible set of prospects by a. The individual's utility function U would rank these prospects such that:

$$a_i \geq a_j \text{ if and only if } U(a_i) = U(a_j) \qquad (1)$$

But how to rank prospects when outcomes are uncertain? Under the expected utility theorem, a rational decision maker ranks prospects on the basis of their respective *expected* utilities (Von Neumann & Morgenstern, 1944). Let a now be a set of risky prospects. The decision maker's utility function would now be such that:

$$a_i \geq a_j \text{ if and only if } E\left[U(a_i)\right] \geq E\left[U(a_j)\right] \tag{2}$$

It should be noted that this utility function is not unique: we could multiply it by positive constant, or change its base point, without changing the ranking of alternatives. This has important implications when we are to consider the case of decisions that affect multiple individuals. But before considering decisions in a multi-actor setting, let us first consider the relations between insurance and optimal prevention for the relatively simple one person case.

2.2 *Optimal prevention without insurance*

Consider an individual that is confronted with a probability of losing q dollars due to the failure of a technological system. He or she can reduce this probability from unity to p_m by investing m dollars in risk reduction, where the probability p_m of losing q dollars is a strictly decreasing function in m. Denote initial wealth by w. The optimal investment m^* can now be determined by maximizing the expected utility $E[U(x)]$ of total income x, which is a function of investment m.

$$E\left[U(x)\right] = p_m U(w - m - q) + (1 - p_m)U(w - m) \tag{3}$$

Differentiating expected utility with respect to m then yields the optimal investment m^*. This optimal investment depends, amongst other, on the shape of the utility function. A risk averse individual that considers a risk worse than expected loss should for instance invest more than a risk neutral person. Conversely, the certainty equivalent of a risk seeking individual would be lower than expected loss, and such an individual should invest less.

2.3 *Optimal insurance*

Let us now consider a rational utility maximizer that contemplates an insurance purchase. The individual is faced with a fixed probability p_0 of suffering a loss q. Denote the amount of damage that can be reclaimed under the insurance contract by c, the annual cost per unit coverage by z, and initial wealth by w. The optimal amount of coverage c^* can now be determined by maximizing the expected utility $E[U(x)]$ of total income x, which is a function of coverage c:

$$\begin{aligned} E\left[U\right] &= p_0 U(w - q + c - zc) \\ &\quad + (1 - p_0)U(w - zc) \end{aligned} \tag{4}$$

When the individual's utility function is concave (risk averse behavior), the optimal amount of insurance coverage can be found by differentiating $E[U(x)]$ with respect to c. The optimum c^* is bound between zero and maximum loss (after Kunreuther & Pauly, 2006). When $c^* = 0$, the rational utility maximizer would not purchase insurance coverage. And when $c^* > q$, he or she would purchase full coverage ($c^* = q$) as it is impossible for claims to exceed total loss q. Note that the outcome $c^* > q$ can only be obtained when insurance is available at a subsidized rate as it implies that coverage is offered at a price below expected loss.

2.4 *Optimal prevention with insurance*

The loss probability was treated as a constant in the previous sub-paragraph. Let us now assume that it depends on the investment in risk reduction (m), as in subparagraph 2.2. Note that the cost per unit coverage (z_m) will depend on p_m, and hence on m. It will after all be less costly to purchase insurance coverage when the probability of suffering a loss is 1/10000 rather than 1/10. The cost per unit coverage will thus be a strictly decreasing function in m of the investment in risk mitigation ($\mathrm{d}z_m/\mathrm{d}m < 0$). To determine the optimal amount of insurance coverage (c^*) and the optimal investment in risk reduction (m^*), we again have to maximize the expected utility of total income $E[U(x)]$, which is now a function of m and c:

$$\begin{aligned} E\left[U\right] &= p_m U(w - q + c - z_m c - m) \\ &\quad + (1 - p_m)U(w - z_m c - m) \end{aligned} \tag{5}$$

The outcome of the optimization depends on the cost of insurance coverage. But what insurance premium could be expected? The failure of technological systems is generally uncorrelated with market returns. When mechanisms for risk transfer would operate efficiently, technological risks would cost expected loss. With full insurance costing $p_m q$, equation (5) could be rewritten to:

$$E\left[U\right] = U(w - p_m q - m) \tag{6}$$

We would now be able to determine an optimal investment in risk mitigation without having to worry about the exact shape of the decision maker's utility function (as long as it is not convex). In fact, the decision problem could be formulated as if our decision maker was risk neutral (note also the irrelevance of initial wealth w):

$$\text{Max}\,[E\,[U]] = \text{Max}\,[-m - p_m q] \tag{7}$$

Besides the advantage of no longer having to specify a utility function, another important result is that we would overinvest in risk reduction if were to ignore the opportunities for risk transfer. Capital markets allow us to diversify exposures and thereby reduce the cost of risk bearing (provided markets operate efficiently and people are risk averse). An investment in system safety thus cannot be evaluated without concern for the way potential losses are borne.

3 THE BENEVOLESCENT DESPOT: THE MULTI-ACTOR CASE

3.1 *The optimization of total output without insurance*

Many decisions concerning the safety of engineered systems affect a public rather than a single individual. Infrastructures such as tunnels, flood defenses and high speed rail links are generally paid for through general taxation, requiring all taxpayers to contribute. And failures of technological systems often affect multiple individuals as well. How to decide on an optimal investment in risk reduction if this decision is to be based on individual preferences?

The single decision maker framework could be extended to this multi-actor setting by considering what a benevolescent decision maker would decide given the public's appreciation of every decision alternative. But "the public" is merely a collection of individuals. When every individual would prefer prospect a_1 over a_2, we could say without much difficulty that the public prefers a_1. But what if the public is divided and ten percent of the population preferred a_2? The approach that is commonly taken to overcome this dilemma is to consider the aggregate or population-averaged costs and gains of different prospects. This practice is based on the Kaldor-Hicks potential compensation criterion that allows us to consider total output and distribution apart (Pearce & Nash, 1982). But despite its widespread use, Arrow cautions us that "there is no meaning to total output independent of distribution" (Arrow, 1963: 40). Let us evaluate this statement in some greater detail and see what its implications are for the optimization of system safety.

As discussed in paragraph two, a person's utility function is defined up to a positive affine transformation under the expected utility theorem. Without a single scale and the possibility to meaningfully compare the intensities of people's likes and dislikes, it would be impossible to construct a ranking of alternatives on the basis of total public pleasure. But how to meaningfully add people's pleasures and pains? The aggregation of individual preferences presupposes cardinal utility and the interpersonal comparison of utility, two assumptions that are not uncontested. If, however, these assumptions are accepted, a benevolescent despot who would be solely interested in the aggregate utility of the population would rank all alternatives such that:

$$a_i \geq a_j \text{ if } \sum_{k=1}^{n} E\,[U_k(a_i)] \geq \sum_{k=1}^{n} E\left[U_k(a_j)\right] \tag{8}$$

Where n denotes the number of individuals that would be affected by a decision. Unfortunately, additional assumptions are unavoidable if we are to evaluate total output using a single utility function that is based on individual preferences. To see this, let us consider two individuals that are exposed to a hazard that could cause them to lose q_1 and q_2 dollars respectively ($Q = q_1 + q_2$). They can reduce the probability of a loss from unity to p_M by investing m_1 and m_2 dollars in risk reduction ($M = m_1 + m_2$), with $\mathrm{d}p_M/\mathrm{d}M < 0$. Denote their initial wealth by w_1 and w_2 respectively ($W = w_1 + w_2$). The individual investments that maximize the aggregate expected utility of income follow from the maximization of:

$$E\,[U_1(x_1) + U_2(x_2)] = E\,[U_1(x_1)] + E\,[U_2(x_2)] \tag{9}$$

Or

$$\begin{aligned} E\,[U_1(x_1) + U_2(x_2)] = {} & p_M U_1(w_1 - m_1 - q_1) \\ & + (1 - p_M)U_1(w_1 - m_1) \\ & + p_M U_2(w_2 - m_2 - q_2) \\ & + (1 - p_M)U(w_2 - m_2) \end{aligned} \tag{10}$$

$$\begin{aligned} E\,[U_1(x_1) + U_2(x_2)] = {} & p_M (U_1(w_1 - m_1 - q_1) \\ & + U_2(w_2 - m_2 - q_2)) \\ & + (1 - p_M)(U_1(w_1 - m_1) \\ & + U(w_2 - m_2)) \end{aligned} \tag{11}$$

We could construct a single utility function for aggregate consequences (U_{1+2}) when the individual utility functions are linear and the same for both individuals (or when $m_1 = m_2$ and $q_1 = q_2$). The above expression could then be rewritten to:

$$\begin{aligned} E\,[U_{1+2}(x_1 + x_2)] = {} & p_M U_{1+2}(W - M - Q) \\ & + (1 - p_M)U_{1+2}(W - M) \end{aligned} \tag{12}$$

Or (note the absence of wealth effects):

$$\text{Max}\,[p_M(-M-Q)+(1-p_M)(-M)] \tag{13}$$

Though convenient, the assumption that all individuals have linear and equal utility functions is rather unrealistic. After all, it implies risk neutral behavior, and the assumption that the marginal utility of income is constant across individuals. But it seems unlikely that an extra dollar in the hands of a millionaire brings the same amount of pleasure as an extra dollar in the hands of a poor man.

Optimizing the safety of technological systems is clearly not without considerable difficulty. The expected utility framework can be extended to a multi-actor setting, but only when additional (disputable) assumptions are made. But in defense of the optimization of total output, one might argue that its ramifications remain limited when the "unfair" distributive consequences of a large number of decisions cancel out. This, however, need not be the case.

3.2 *Insurance, investment costs and optimal prevention*

The presence of efficient mechanisms for risk transfer could strongly reduce the theoretical and normative dilemmas that surround the use of the expected utility framework in a multi-actor setting. When insurance markets operate efficiently, every risk averse individual would purchase full insurance coverage. The certainty equivalent of every risk averse (or risk neutral) individual could then be replaced by expected loss, irrespective of people's exact risk preferences. But to allow for further simplification, we also have to consider the distribution of investment cost. These will often be concentrated in the hands of a few, but when all individuals would be charged a contribution proportional to their expected loss, the complexities of the optimization of public safety could be strongly reduced. This could for instance be the case when exposures and the taxes to pay for risk reduction would both be proportionate to wealth. And for publicly traded firms, the proportionality between expenditures and exposures is rather straightforward: corporate investments in risk reduction and the costs of unexpected events both affect shareholders in proportion their holdings (when default risk is negligible).

When perfect insurance is available, i.e. when full coverage is available against the actuarially fair premium (expected loss), and when expenditures are distributed in proportion to individual exposures, the optimization of system safety could be reduced to an expected value optimization. To see this, let us again consider two persons that are faced with an event with probability p that could cause them to lose q_1 and q_2 respectively ($Q = q_1 + q_2$). They can reduce this probability by investing m_1 and m_2 in risk reduction ($M = m_1 + m_2$), with $\mathrm{d}p_M/\mathrm{d}M < 0$. When the individuals are risk averse, they would both purchase full insurance coverage, costing $p_M q_1$ and $p_M q_1$ respectively. When each individual's contribution to system safety is set in proportion to his or her exposure (insurance premium), the interests of both would be perfectly aligned. None of them could be made better off by making the other worse off when resources are allocated according to:

$$\text{Min}\,[M + p_M Q] \tag{14}$$

With

$$m_1 = M\frac{q_1}{q_1+q_2} \quad m_2 = M\frac{q_2}{q_1+q_2} \tag{15}$$

When mechanisms for risk transfer operate efficiently, and when costs are borne in proportion to individual exposures, the practical and normative dilemmas associated with the aggregation of individual preferences could be completely overcome. We could then consider gains and losses under the condition of risk neutrality.

It seems unlikely that any of the two conditions that allow us to strongly simplify the optimization of system safety will ever be completely met. While investment cost and exposures might for instance be distributed in a *fairly* proportional way, it seems unlikely that they will ever be distributed in an *exactly* proportional way. And "perfect insurance" presupposes that all consequences of failure can be compensated for, that insurance is offered against the actuarially fair premium (expected loss), and that full compensation is so frictionless that people do not even notice that they have suffered a loss. The optimization of risks to the public thereby typically involves a trade-off between practicality and scientific rigor.

4 CASE STUDY: FLOOD PROTECTION IN THE NETHERLANDS

4.1 *The economic decision problem*

The optimal investment in a flood defense system could be determined by minimizing the sum of the discounted investments in flood defense and the discounted expected value of future losses (after Van Dantzig, 1956). A simplified one-period model without economic growth or sea level rise is presented here for illustrative purposes (it is assumed that the NPV of the costs of delayed programs exceed the NPV of the costs of direct investment):

$$\mathrm{Min}\,[NPV] = \mathrm{Min}\left[M + \int_0^T pQe^{-rt}dt\right] \qquad (16)$$

Where NPV = net present value of total cost (euro); M = total investment in flood defense (euro); p = probability of flood (per year); Q = total damage in case of flood (euro); T = planning horizon (year); r = discount rate (per year).

When overtopping is the only failure mode of a flood defense, and when exceedance frequencies of water levels can be described by an exponential distribution, the failure probability p can be defined as:

$$p = e^{-\frac{h-a}{b}} \qquad (17)$$

Where p = probability of failure; h = dike height (m); a, b = constants (m). When it is assumed that the cost of dike heightening consist of a fixed and a variable part, equation (17) can be expanded to:

$$\mathrm{Min}\,[NPV] = \mathrm{Min}\left[M_0 + M'h + \int_0^T e^{-\frac{h-a}{b}} Qe^{-rt}dt\right] \qquad (18)$$

Where M_0 = fixed cost of dike heightening (euro); M' = variable cost of dike heightening (euro); h = dike height (m); a, b = constants (m).

The optimal standard of protection (h^*) can now be found by differentiating NPV with respect to h. Figure 1 shows how the net present value of total cost changes when flood defenses are raised (fictitious parameter values). The optimal height of the flood defenses would in this case be 5.5 m.

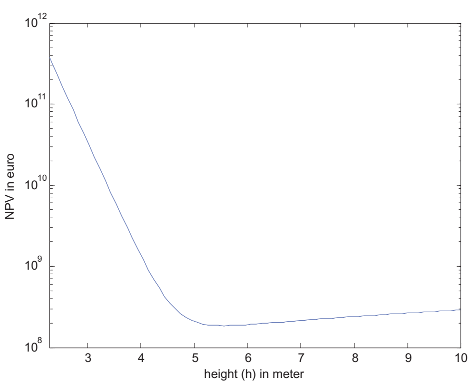

Figure 1. Optimal investment in flood protection under risk neutrality; $M_0 = 40 \cdot 10^6$ euro; $M' = 25 \cdot 10^6$ euro/meter; $a = 2.3$ m; $b = 0.3$ m; r = 0.015/yr; $Q = 10 \cdot 10^9$ euro.

Equation (18) presupposes risk neutrality. But why should the benevolescent despot only consider expected loss when citizens are risk averse? Interestingly, equation (18) shows remarkable resemblance to equation (14) which was based on the availability of perfect insurance and contributions to system safety in proportion to individual exposures. Although the latter condition seems broadly reasonable, the first deserves a closer look.

4.2 *Relaxing the assumption of perfect insurance*

Flood risks in the Netherlands are highly concentrated: a single flood could give rise to billions of euros of damages. Such low-probability, high-consequence events are notoriously difficult to insure and private insurance against large-scale floods is currently unavailable in the Netherlands. The assumption of perfect insurance might therefore seem highly unrealistic. But the assumption of perfect insurance does not presuppose the provision of insurance by private enterprise.

When a government is trusted to provide swift and complete compensation for losses, the assumption of a negligible risk premium (or: efficient insurance) could be fully justifiable (human tragedy apart). After all, the Dutch government has an AAA-rating so that it is able to borrow at the risk-free rate and "sell" coverage at a price that corresponds to expected loss (see also Cummins et al., 1999 for a more elaborate treatment). To see this, consider an unexpected loss C. If the government were to finance this loss by issuing a perpetual at the risk free rate r (a perpetual is a bond without a maturity), it would pay rC annually to the bondholders. The present value of the interest payments would be $rC/r = C$. The loss would have been spread over an infinite time horizon without a surcharge for doing so. When a government is able to spread losses over time without extra cost, the NPV of its irregular loss pattern is equivalent to the NPV of a constant stream of annual losses with magnitude expected loss. The government would thus be able to "sell" coverage at expected loss. Indeed, the Delta Committee assumed that funds would be set aside to cover future losses and it treated the economic decision problem as "an insurance problem" (Van Dantzig, 1956: 280).

So far, it has been assumed that all losses could be compensated for. But as shown recently by the flooding of New Orleans, the human tragedy of a large-scale flood extends well beyond the loss of property and income. Floods can disrupt communities and families, and leave considerable trauma behind. It would be unreasonable if not immoral to argue that decision makers and engineers should ignore these aspects and only consider the economic impacts of the failure of technological systems. The approach that was followed by the Delta Committee was to multiply potential economic damage with a constant to account for intangible losses. Distress and loss of life were thus assumed to be proportionate to economic loss. This practical solution allowed decision makers to evaluate the influence of intangible losses on their optimizations. It also showed what minimum investment in flood protection would be appropriate, as rational individuals would *at least* consider their economic losses. The decision problem was thus reformulated to:

$$\text{Min}\left[M_0 + M'h + K\int_0^T e^{-\frac{h-a}{b}} Qe^{-rt}\,dt\right] \tag{19}$$

Where K = multiplication factor to account for intangible losses.

The optimal height of a flood defense, and hence the optimal investment in flood protection, goes up when the multiple of expected loss increases. As shown in Figure 2, the optimal height of the flood defense is about 6 meters when the certainty equivalent equals five times expected loss.

The above formulation would also be appropriate when full insurance coverage would be available, but only at a price that exceeds expected loss. Reinsurance premiums for low-probability, large-scale losses are typically a multiple of expected loss (Froot, 2001). This multiple of expected loss should not be confused with a pricing rule for (re)insurance contracts. It is merely a descriptive metric that illustrates how costly insurance coverage can be relative to the actuarially fair premium. It offers a relatively simple way to account for costly (re)insurance coverage in the optimization of technological systems: we could replace the certainty equivalent by an insurance premium that equals a multiple of expected loss (see also Kok et al., 2002).

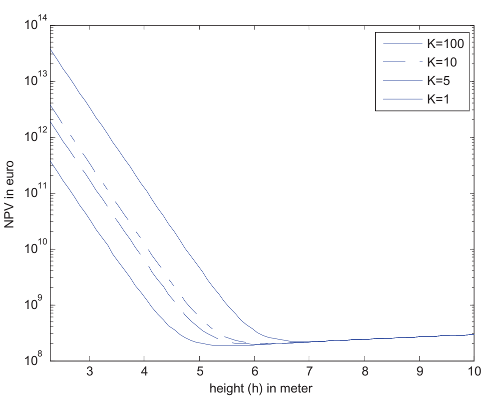

Figure 2. Optimal investment in flood protection under K times expected loss; $M_0 = 40 \cdot 10^6$ euro; $M' = 25 \cdot 10^6$ euro/meter; $a = 2.3$ m; $b = 0.3$ m; r = 0.015/yr; $Q = 10 \cdot 10^9$ euro.

The formulation of the decision problem would have to be adjusted considerably when the availability of flood insurance (or government compensation) would be unrealistic. We would then have to consider each individual's utility function, and the distribution of initial wealth, investment costs, and risks to calculate certainty equivalents and maximize total public pleasure. After all, the use of a single utility function for aggregate gains and losses might now yield grossly invalid results, even when everyone's risk preferences would be exactly similar. Consider for instance the case in which an investment to reduce expected loss would concentrate this risk in the hands of a few. When every individual would be risk averse, the latter situation would be associated with a greater total certainty equivalent than the first. But a risk-averse utility function for *aggregate* gains and losses would suggest the opposite.

To reduce the complexities of the decision problem to manageable proportions, let us assume that wealth effects are absent, that exposures and investment costs are uniformly distributed over the population, and that everyone has the exact same utility function. These assumptions allow us to consider a *single* utility function for *aggregate gains and losses* (no wealth effects). The decision problem can now be reformulated to:

$$\text{Max}\,[E\,[U]] \tag{20}$$

with

$$E\,[U] = -M_0 - M'h - \int_0^T C_{p,Q} e^{-rt}\,dt \tag{21}$$

Where $C_{p,Q}$ = certainty equivalent for a loss Q with probability p (euro).

By assuming the absence of wealth effects, we have assumed that individuals display constant absolute risk aversion (CARA). Despite the practicality of this assumption, it places constraints on individual risk preferences that hardly seem realistic. After all, it seems unlikely that people's pains would depend on the absolute severity of losses, rather than the severity of losses relative to initial wealth. Relaxing the assumption that wealth effects are absent thus seems reasonable. Iso-elastic utility functions describe behavior that corresponds to constant relative risk aversion (CRRA) (Chavas, 2004):

$$U(x) = x^{(1-\varphi)} \quad \text{for } \varphi < 1 \tag{22}$$

Where φ denotes the Arrow-Pratt coefficient of relative risk aversion ($\varphi = -xU''(x)/U'(x)$). This coefficient can be interpreted as an elasticity: it expresses the percentage decrease in marginal utility due to a percentage increase in wealth (Chavas, 2004). Unfortunately, the use of a CRRA utility function requires us to drastically reformulate the objective function. Let us assume that wealth is uniformly distributed over the exposed population. Denote annual per capita income by w. Denote every individual's contribution to flood defense by m (with $m_0 = M_0/n$ and $m' = M'/n$). The maximization of aggregate expected utility now boils down to:

$$\text{Max}\left[np\left(w - m_0 - m'h - \int_0^T qe^{-rt}dt\right)^{1-\varphi} + n\left(1-p\right)\left(w - m_0 - m'h\right)^{1-\varphi}\right] \tag{23}$$

Solving the above formulation of the economic decision problem would require considerable computational effort. It seems questionable whether this effort would be worthwhile given the disputable assumptions that again have to be made (most notably the assumptions of cardinal utility and the interpersonal comparability of utility). The Delta Committee's relatively simple model can assist decision makers in balancing their views on the cost of risk bearing against the cost of flood protection. If cost-benefit analysis is to be used to support decision makers, such an insightful model seems preferable over a more refined yet equally disputable model.

5 CONCLUSIONS

The expected utility framework provides a rational basis for the evaluation of risks and the optimization of system safety within the context of a single decision maker. Yet most decisions about system safety concern risks to a public. When decisions are to be based on the preferences of all potentially affected individuals, the expected utility framework has to be supplemented by several disputable assumptions. When analysts for instance consider population-wide or population-averaged effects, they implicitly make assumptions concerning the comparability of individual preferences. While the consistent optimization of total output might make everyone better off in the long run, this need not be the case.

When risks are unsystematic and markets operate efficiently (or: when full insurance coverage is available against expected loss), and when individual contributions to system safety are set in accordance to individual exposures, it would be appropriate to simply balance total investment cost against reductions in total expected loss. In other cases, analysts have to make assumptions concerning amongst other the interpersonal comparability of utility, and the shape of individual utility functions when attempting to optimize system safety.

Efficient mechanisms for risk transfer reduce optimal investments in system safety. This has been illustrated for the design of flood defenses in the low-lying Netherlands: the availability of an efficient flood insurance program would reduce optimal investments in flood defense to a minimum. Current cost-benefit studies for the Dutch primary flood defenses are risk neutral, which implies that they assume that full coverage is available against the actuarially fair premium. Since that assumption hardly seems reasonable at present, the results of such risk neutral studies could be considered unduly optimistic.

REFERENCES

Arrow, J.K. 1971. Insurance, risk and resource allocation. In K.J. Arrow (ed.) *Essays in the Theory of Risk-Bearing*: 134–143. Amsterdam and London: North-Holland Publishing Company.

Arrow, J.K. 1963. *Social Choice and Individual Values*. New Haven and London: Yale University Press.

Bernoulli, D. 1738. Exposition of a new theory on the measurement of risk (original publication: 'Specimen Theoriae Novae de Mensura Sortis, Commentarii Academicae Imperialis Petropolitanas, Tomus V'). *Econometrica* 1954, 22: 22–36.

Chavas, J.P. 2004. *Risk Analysis in Theory and Practice*. San Diego: Elsevier Academic Press.

Cummins, J.D., Lewis, C.M. & Phillips, R.D. 1999. Pricing Excess-of-Loss Reinsurance Contracts against Catastrophic Loss. In K.A. Froot (ed.) *The Financing of Catastrophe Risk*: 93–141. Chicago and London: The University of Chicago Press.

Faber, M.H., Maes, M.A., Baker, J.W., Vrouwenvelder, T. & Takada, T. 2007. Principles of risk assessment of engineered systems. In *10th International Conference on Application of Statistic and Probability in Civil Engineering (ICASP10)*. Tokyo, Japan.

Froot, K.A. 2001. The market for catastrophic risk: a clinical examination. *Journal of Financial Economics* 60: 529–571.

Kok, M., Vrijling, J.K., Van Gelder, P.H. A.J.M. & Vogelsang, M.P. 2002. Risk of flooding and insurance in The Netherlands. In Wu (ed.) *Flood Defence*, New York: Science Press Ltd.

Kunreuther, H. 1996. Mitigating disaster losses through insurance. *Journal of Risk and Uncertainty* 12: 171–187.

Kunreuther, H. & Pauly, M. 2006. Rules rather than discretion: Lessons from Hurricane Katrina. *Journal of Risk and Uncertainty* 33: 101–116.

Pearce, D.W. & Nash, C.A. 1982. *The Social Appraisal of Projects: A Text in cost-Benefit Analysis*. London and Basingstoke: The Macmillan Press Ltd.

Rackwitz, R. 2004. Optimal and acceptable technical facilities involving risks. *Risk Analysis* 24: 675–695.

Savage, L.J. 1974. *The Foundations of Statistics*. New York: Dover Publications Inc.

Van Dantzig, D. 1956. Economic decision problems for flood prevention. *Econometrica* 24: 276–287.

Von Neumann, J. & Morgenstern, O. 1944. *Theory of Games and Economic Behavior*. Princeton: Princeton University Press.

Vrijling, J.K., Hengel, W.V. & Houben, R.J. 1998. Acceptable risk as a basis for design. *Reliability Engineering and System Safety* 59: 141–150.