31 December 2009

Review of the J-value literature – Final Report

The Health and Safety Executive

NERA
Economic Consulting

Project Team

Michael Spackman

NERA Economic Consulting
15 Stratford Place
London W1C 1BE
United Kingdom
Tel: +44 20 7659 8500
Fax: +44 20 7659 8501
www.nera.com

Contents

Executive Summary

The J-value (or "Judgement-value) terminology was first widely disseminated in two papers published in 2006 in Transactions of the Institution of Chemical Engineers. These are useful and impressive papers, illustrating how the valuation of fatality risks varies across different regulatory regimes, and varies even more in potential and actual application.

The data collected by the authors could have been used to present, more simply, the differences between the "values of a prevented fatality" (VPF) implied by those regulatory conventions and applications. However the chosen comparative measure was the cost/benefit ratio (described as the J-value) of expenditure divided by a VPF that is re-calculated implicitly for each application, using a "Life-Quality Index" (LQI) method developed in the 1990s in Canada.

The LQI method derives a "value of life" as the present value of discounted lifetime welfare. Welfare (or life-quality) is assumed to be proportional to income, to a power q, multiplied by leisure time. q is derived from the national average ratio of working time to leisure time, which is assumed to be at a social optimum. This derivation of a "value of life" is ingenious, and it is useful to have that literature drawn to the attention of UK safety regulation analysts. However the method is too simplistic to be a competitor to the methods now established in the UK and elsewhere for the valuation of fatality risks.

These current methods are based on observation or testing of people's attitudes to risk. There are strong arguments for valuing fatality risks in terms of risk preferences rather than "discounted lifetime welfare". But if the latter approach were nonetheless adopted it would be based in practice on the well established 'Value of a Life-Year', originating in health economics.

While the wide variations in fatality risk valuations, in regulatory conventions and their application, are familiar to analysts working in general safety regulation policy, the J-value formulation may have usefully helped to inform a wider audience. However for policy analysis the J-value ratio and its presentation offer no clear benefits over other measures. The form of presentation in these and other papers lacks in particular the transparency and flexibility needed by users for practical application in government. Nor does it contribute to the currently more challenging problems of valuing illness and non-fatal injuries.

The literature applying a J-value ratio to the *total* costs of a major accident focuses on the concept of an investing institution having a 'constant-elasticity' utility function, relating utility to assets. But utility functions generally apply, and then only in some simple cases, to individuals or households. For institutions there is no good reason to challenge the normal assumption of "risk-neutrality", except where risk is "systematic". Private sector institutions will (or should) be increasingly averse to financial risks that threaten their viability, but this would not be not well portrayed by a constant elasticity utility function. And neither of these issues, of systematic risk or corporate viability, is relevant to government project risks of the kind considered in the J-value literature.

Given the assumption of a constant-elasticity utility function, the J-value papers correctly derive the "risk multipliers" that such functions imply. These ascribe somewhat higher than "risk-neutral" values to large downside impacts. However the papers' further extension into other concepts is flawed by an assumption that "absolute" utility values, as distinct from differences, have meaning, which the authors correctly recognise elsewhere as mistaken.

The J-value literature may appear to offer an analytical basis for larger numbers for the VPF and for the valuation of large accident risks, but it does not really offer a way forward. There may however be other avenues worth considering, in developing a structured approach to the recurring problem of appraising impacts that can justify the expenditure of much more on some safety measures than would be justified by the expected reduction in human costs alone. These could include closer examination of all the public expenditure and other costs associated with some specific hazards and the development of earlier work on "societal concerns".

1. Introduction

This Report is in response to a request by the HSE to review the" J-value approach" to comparing the costs and benefits of safety regulation, focusing on its potential contribution to regulatory decision making.

The J-value concept has been developed by Professor Philip Thomas and colleagues at the Risk Management, Reliability and Maintenance Centre of the City University School of Engineering and Mathematical Sciences. It builds initially on Canadian work, originating in the early 1990s and funded on occasion by the National Science and Engineering Council of Canada, based on a Life Quality Index defined as a function of income and life expectancy.

The City University team have developed the Life Quality Index to establish a J-value (or "Judgement-value") framework (the J-value being a ratio of cost to benefit), which they have applied to a range of UK health and safety contexts. They have also developed a separate analysis of non-human costs and risk aversion. They have obtained funding for this work from the EPSRC and the ESRC and also from the Ministry of Defence.

Chapter 2 below summarises the development of the J-value approach to fatality risks. Chapter 3 provides a commentary on this, including a summary of the conventional Whitehall approach to this issue. Chapters 4 and 5 similarly summarise and comment on the application of the J-value approach to total costs and, especially, its approach to risk aversion. Chapter 6 draws conclusions.

References cited in the main text are listed after Chapter 6. References specific to the Appendices are listed at the end of, or in footnotes to the relevant Appendix.

This review has been undertaken within NERA mainly by Michael Spackman. Mr Spackman was previously a physicist/engineer in the nuclear industry and then the Civil Service, and remains a Chartered Physicist. Subsequently, as an economist, he headed for ten years the Treasury Public Expenditure Economics Group, with responsibilities including development of the Treasury's guidance on appraisal. He was at that time and has been subsequently, as a consultant, closely involved in the development and application of appraisal methodology for health and safety and environmental regulation, mainly within UK government.

2. J-values applied to fatality risks

The J-value as applied to fatality risks is described in Thomas, Stupples and Alghaffar (2006.1 and 2006.2). It is applied to some examples in Thomas and Jones (2006)[*], and further refined in Thomas, Kearns and Jones (2009)[*].

Thomas el (2006.1) explains in its Summary that "the new J-value technique is developed from a life-quality index that is a function of life expectancy, average income and work-life balance" and that the paper uses the method "to assess the degree of consensus on health and safety expenditure amongst regulators across different sectors of the economy." The Summary to Thomas et al (2006.2) records that "the measure of agreement in the regulator's theoretical recommendations found [in the 2006.1 paper] by using the J-value technique is contrasted with the very large disparities found on practical health and safety schemes". It proposes the J-value method "as a common yardstick for assessing health and safety spend for the use of decision makers in all sectors", suggesting that "its adoption could lead to better targeting of health and safety spend in all areas of the economy."

2.1. The Life-Quality Index

The key building block of the J-value approach to fatality risks is a Life-Quality Index (LQI) derived by a method originated, as noted above, by Canadian authors and attributed in Thomas et al (2006.1) specifically to Pandey and Nathwani (2003). The LQI is also described in the Canadian and the J-value literature as "utility".[1]

The LQI is denoted by Q and is derived as a function of expected lifetime leisure time (T years) and earnings or consumption, £G. UK national average figures are taken as the baseline, G being national GDP per capita. The function is specified initially as $Q = aG^{b}T^{g}$, where α, β and γ are constants. With some algebraic manipulation a final version is derived as $Q = G^{q}X_{d}$, where q is derived as described immediately below and X_d is average discounted life expectancy, derived as explained in section 2.2.

The exponent (or elasticity) q is derived by replacing leisure time, T, in the earlier equation by $(1-w)X$, where w is the average fraction of time spent working and X (years) is life expectancy. G is then replaced, temporarily, by kw, where k is pay per unit time (£/y). It is then assumed that "individuals in the nation will have adjusted their 'work-life balance' [*viz* the ratio $w/(1-w)$] so as to optimise their quality of life." Differentiating Q with respect to w shows that Q is maximised when $w = q/(q+1) = w_0$, from which it follows that $q = w_0/(1-w_0)$.

Thomas et al (2006.1) record that Pandey and Nathwani suggest typical values for industrialised countries for working lifetime and length of working week that imply a value for w_0 of 0.144, which they recommend "should be reduced by 10% to account for health care", giving w_0= ~0.125=1/8. This implies a value for q "of q =1/7= ~0.143, the figure used

---

[*] The two asterisked papers were kindly supplied by MOD, the others by the HSE.

[1] "The general idea is that a person's enjoyment of life, or utility in an economic sense arises from a continuous stream of resources available for consumption over the entire life. Therefore, income required to support consumption and the time to enjoy are two determinants of the life quality." (Pandey and Nathwani, 2003, p68).

in this [Thomas et al (2006)] study. Sensitivity analyses were also carried out over $0.12 \leq q \leq 0.2$, and the results were found to be insensitive to changes within this range."

Thomas, Kearns and Jones (2009) develop a subsequent extension of the Canadian work (Pandey, Nathwani and Lind, 2006), in which GDP is equated with a constant returns to scale Cobb-Douglas function, so that GDP = constant x $(\text{capital})^{(1-q)}$ x $(\text{labour})^{q}$. From this it is shown to follow that $q = (1/q)w_o/(1-w_0)$. Thomas, Kearns and Jones derive values for the UK of $w_0 = 0.091$ and $q = 0.546$, giving a value for q of 0.183.

2.2. Discounted life expectancy

"Discounted life expectancy", X_d, is so defined that, when applied as a multiplier to the value of G^q for the current year it yields the present value of that and all future years' life quality, over an appropriate lifetime. It is derived in several steps, initially in terms of an individual and then generalised to a population.

The first step is to establish an exponential discount rate, r_d to apply to life-quality expected in future years. This rate is described as "the real rate of time preference which we shall also term the discount rate". It corresponds broadly with what in mainstream economics literature is often described as "pure time preference" for marginal utility. The values for r_d used in Thomas et al (2006.1) are 0% and 2.5%.

The second step is to apply this discount rate to the life-quality that someone will generate at a future time. It is assumed that income will grow at a (real terms, exponential) rate of r_g, so implying that the "utility" of this income will grow per year at the lower rate qr_g. Thus the present value of this "utility" enjoyed in a future year (contingent on survival into that year) is obtained by discounting a hypothetical flow in that year of the current level of income at the discount rate $r = r_d - qr_g$. r is described as the "net discount rate".

The final two steps introduce life expectancy for an individual and then generalise to the national population, with Thomas et al drawing on UK actuarial tables. X_d is then derived in terms of national average life expectancy and average income – the latter assumption being "to avoid different treatments for high-earning and low-earning groups with regard to safety spend".

2.3. Derivation of the J-value for expenditure to reduce fatality risks

Suppose an annual expenditure ΔG is incurred by an individual to achieve an increase in discounted life expectancy of ΔX_d (and associated lifetime annual gain in utility) and $Q = G^q X_d$. Then, if ΔG and ΔX_d are small, and ΔG is expressed as a positive quantity, the proportional increase in the Life-Quality Index Q is given by $\Delta Q / Q = -q \Delta G / G + \Delta X_d / X_d$.

Thus, if the change in life-quality, ΔQ, is to be non-negative, $\Delta X_d / X_d \geq q \Delta G / G$ and so $\Delta G \leq G \Delta X_d / qX_d$. This maximum acceptable value of ΔG, multiplied by the size of the relevant population, N, defines *the maximum that the population should be willing-to-pay* (i.e. the value of reduction in fatality risks) and is designated by Thomas et al (2006a) as a_{pop}.

The J-value is then defined by the ratio $\hat{a}_{pop}/a_{pop}$, where $\hat{a}_{pop}$ is the *actual* amount spent on the safety measure. Thus $J = \hat{a}_{pop} / a_{pop} = \hat{a}_{pop} q X_d / NG\Delta X_d$ and is a measure of the cost/benefit ratio of the safety measure.

In practice the costs of a safety expenditure often approximate more closely to a single capital sum than an annual expenditure as assumed in the J-value derivation. Such a capital sum could be annuitised at some positive interest rate over the average population life span and Thomas et al (2006.1) suggest the pure time preference rate r_d for this purpose. However they recommend instead that, when a single initial spend is dominant, a zero rate should be assumed. This is computationally simpler, but it reduces the annuitised value of $\hat{a}_{pop}$. It therefore, for a given initial spend, favours expenditure on safety equipment. This is described as conservative.

2.4. Practical application

Thomas el (2006.1) applies the J-value to the Department for Transport conventions for road and rail, to the valuation of the Quality Adjusted Life Year (QALY), in the context of the NHS, by NICE, to the conventions of the National Radiological Protection Board, as it then was, and to the HSE convention for the Offshore Oil Industry.

Thomas et al (2006.2) applies the J-value to a range of case studies, some of which appear to have entailed substantial data collection. The cases are the Train Protection and Warning System (TPWS), the European Rail Traffic Management System (ERTMS), four drugs considered by NICE, public health countermeasures considered in the context of BSE / vCJD, BNFL plants for the removal of Technetium-99 and Krypton-85, and measures to control VOC (volatile organic compounds) in petrol service stations.

The papers do not make explicit the value of a prevented fatality (VPF) derived by their application of the Life Quality Index. However the J-value application to Department of Transport conventions shows that the LQI derivation, with a time preference rate of 0%, gives a VPF higher by factor of about three than the £1.22m VPF[2] then used by that Department, or a factor of two with a time preference rate of 2.5%.

Calculation of J-values for other regulators' conventions confirms, with one exception, that they have differing implied or explicit VPFs. A special case is the NICE criterion for the approval of NHS treatments, where the current policy criterion is expressed in terms of Quality Adjusted Life Years (QALYs), not a VPF. The QALY value can easily be compared with the J-value arithmetic, but as discussed in chapter 3 below it is a markedly different concept from and not necessarily comparable with the VPF.

The checks against the practical case studies tell a vivid story, albeit one largely familiar to those working in general safety regulation. In terms solely of expected lives saved the TPWS gave very poor value for money (VFM) even with the LQI-based VPF. The ERTMS would have given appalling VFM (as was recognised by HSC at the time that it recommended against it). The drugs, on an LQI basis, generally gave very good VFM. Early measures against BSE gave very good VFM and countermeasures after March 1996 very bad VFM.

[2] The DfT value includes lost net output and public service costs, which LQI estimate does not, but this is a minor issue as the willingness to pay element is around 90% of the DfT total for fatalities.

The BNFL Tc-99 plants gave appalling VFM, but this policy was implemented; the Krypton options gave worse VFM but were not implemented. The petrol station VOC measures gave VFM that was poor even with the LQI-based VPF, but only by a factor of about two.

Thomas et al (2006.2) suggest that it is "difficult to understand the logic" of disparities as severe as the spending of huge sums on the BNFL Technetium plant while not approving a drug that would save many more life-years. It suggests as noted earlier that the J-value has an "ability to translate a variety of cost-benefit formats into a common yardstick" and that "its adoption could lead to better targeting of health and safety expenditure in all areas of the economy".

3. Commentary on the J-value approach to expenditure on reducing fatality risk

Commentary on the City University work is preceded here by a brief description of the currently conventional analysis of fatality risks in UK government, to facilitate comparison with the J-value proposals.

3.1. The conventional approach to fatality risk

Since the 1990s it has been generally accepted by the relevant agencies in developed countries that small fatality risks to individuals can be valued in monetary terms. This value is often called in English a 'value of a statistical life' (VOSL, or sometimes VSL), although in the UK the term now most often used is 'value of a prevented fatality' (VPF).

The main component of this valuation, in countries such as the UK that have well developed methodologies, is the willingness-to-pay (WTP) of potential victims to reduce the risk to which they are exposed.[3] Other costs of a fatality, which in the UK are also explicitly valued, include the loss of net output following a death, and public service costs.

Public service (e.g. medical) costs, while large in the case of some illnesses or non-fatal injuries, are often small in the case of death. However costs specific to cancer risks include the medical costs of treatment (as well as the personal costs) before death. In practice this is not analytically well developed, but may provide a rationale for the HSE convention of multiplying the normal VPF by a factor of two for cancer risks.

In the case of risks faced by highly qualified staff, as noted in the MOD appraisal guidance (Ministry of Defence, 2006), the costs of specialised training of those that would take their place may be another, substantial, additional cost.[4]

In practice the VPF estimated by work commissioned and from time to time revisited by the Department for Transport has never been significantly contentious in Whitehall, nor indeed outside (except for those who do not accept the principle of valuing fatality risks). It is widely seen as a satisfactory figure for a "context free" VPF. The current difficulties in this field are generally perceived to be in two other areas:

§ The valuation of illness and non-fatal injuries: this is a complex and evolving story; but beyond the scope of the J-value literature and hence of this Report.

---

3 The value generally now used in the UK is derived by seeking people's stated preferences about small fatality risks relative to consumption. The methodology has evolved over the years to improve the handling the many difficulties that this presents. This is the approach adopted in most countries that establish a VPF. The US puts more reliance on the preferences revealed by wage differentials between high and low risk occupations, but this presents problems that are no less serious.

4 The MOD Guide notes that (paragraph 4.5.41) "Whilst the baseline WTP [willingness to pay] incorporates a societal element related to the lost output, this element will clearly be far higher amongst highly trained MOD personnel. It is legitimate to include training investment as an addition to the WTP figure. In some cases, this could increase the baseline VPF from around one million to as high as several million." The Guide also notes that equipment replacement or repair cost following an accident can be substantial, but this is a cost better handled on its own, rather than being embedded in a VPF. (In Department for Transport convention such costs are described as "accident costs", as distinct from and additional to "casualty costs".)

- The handling of special factors that might justify higher or lower numbers in special circumstances. Many such, overlapping factors might be considered, included the nature of those at risk (e.g. railway staff versus passengers versus trespassers, or age, or health state), the nature of the hazard (e.g. "dread" as some might apply to cancer, or the potential size of the accident, or the base level of risk), or the distribution of responsibilities (e.g. the duty of care that might differentiate public from private transport). These are the kinds of factor discussed in Appendix D to this Report.

One factor considered in Appendix D as *not* relevant to the VPF is the impact of an accident on the subsequent political or managerial climate. A nuclear accident on anything approaching the seriousness of Three Mile Island, or even a technically much less serious incident exposing members of the public to significant radiation, might have profound political and economic consequences. However such considerations are only very tenuously related to numbers of people killed. (The Three Mile Island incident for example killed no one, but had huge political and commercial consequences.) They are factors that need to be considered in the context of *accident* risks rather than fatality risks.

3.2. The J-value approach to fatality risks

3.2.1. Identifying what is new

The J-value is an ingenious new term for what would conventionally be called a benefit / cost ratio where, in this case, the denominator is the value of the expected prevented fatalities and numerator is the cost of achieving that improvement in safety.

However the claim that this particular ratio has an ability to translate a variety of cost-benefit formats into a common yardstick is misleading. All that is needed to provide a common yardstick for most comparisons of the kind to which the J-value is applied is either a standard value of a prevented fatality (VPF), or, even more simply, the comparison of conventions or applications in terms of *their own* "cost per prevented fatality". In the safety literature the cost per prevented fatality is often preferred, as its meaning is clear and the reader can compare the numbers instantly with the VPF generally used in that or other contexts.[5]

The main new feature of the J-value package is that it embeds an unconventional approach to the valuation of fatality risks, which implies a VPF higher than the values generally applied in the UK or elsewhere in Europe. The merits of this approach are discussed in section 3.2.2 below.

The Canadian literature, from which the LQI is drawn, sets out the VPFs implied by their calculation and how these values compare with those derived by other methods in the Canadian context.[6] Indeed the motivation of Pandey and Nathwani (2003) appears to be to promote the principle of assessing (civil) nuclear safety expenditures in cost benefit terms

---

5 A classic report of this kind, delivered in th early 1990s to a UK audience by a US OMB official, is Belzer (1992). Another US study, written largely by authors from a health economics background, where the value of a life-year (VOLY) is often more appropriate than the VPF, and presented in terms of implied VOLYs, is Tengs et al (1995).

6 Their VPF estimate in Pandey and Nathwani (2003) for a time preference rate of 1% (which is not wholly unreasonable in this context) was C$3.4 million (1996 C$), with values about C$1 million higher and lower for time preference rates of 0% and 2%. The 1996 £/C$ exchange rate was £0.47/C$. Their LQI value was therefore much higher than the UK value at that time, though not out of line with values then in use in Canada, derived mainly from wage risk models.

(where they believe there is sometimes excessive regulation) rather than to promote a new method of deriving what they describe as a VSL. They explain that "The point to make here is that the implied VSL estimates from the proposed model are in line with those obtained by wage differential and contingent-valuation surveys, thus providing an empirical validation of the LQI approach. Although the basis for selecting a VSL value is generally a controversial problem, the LQI model can still provide an alternative approach to this issue."

The J-value literature sells the LQI derivation more strongly. No explicit figures are given for the VPFs derived.[7] And the 2006 cross checks against institutional conventions and applications are presented as a test to see if current regulations are up to scratch rather than as a check on the (hidden) VPF derived via the LQI.

These checks nonetheless have the very considerable virtue of publicising prima facie inconsistencies between different regulators' conventions and between these conventions and their practical application. They add little to existing knowledge. Those working in the field have long been broadly familiar with the pictures portrayed. And as already noted (and setting aside the QALY issue noted below), the comparisons in the authors' 2006 papers could be made more simply and clearly, given the data they had collected, without the J-value arithmetic. However the J-value package may have made the presentation especially effective in the chosen journal, to a wide audience not previously familiar with the VPF concept and the apparent inconsistencies in its application.

The one case where translation is not straightforward with current conventions is that of comparing the VPF with the Quality Adjusted Life Year (QALY). Here, since the J-value / Life Quality Index arithmetic is itself based discounted on life-years, it can be tied directly to the QALY as well as the VPF. But as noted below this can be seen as a weakness rather than a strength of the Life-Quality Index. It is not clear that there should be a very close relationship between the QALY and the VPF.

3.2.2. The LQI approach to deriving the value of a prevented fatality

The Canadian Life Quality Index is an ingenious concept, but it presents problems as a basis for deriving people's willingness to pay for small reductions in fatality risks. Its main limitations are threefold:

1) It assumes that the value of a prevented fatality can be satisfactorily measured by discounting the value of the expected lifetime welfare that a person would enjoy if they lived to a normal lifespan. As a principle this is very questionable, as explained in Appendix A to this Report.

2) If the principle of valuing prevented fatalities by discounting future welfare were nonetheless to be accepted, the simplest approach in UK government, which is already discussed and sometimes applied in the context of illness and ***non-fatal*** injury, would be in terms of the Value of a Life-Year (VOLY), as developed in health economics.[8]

---

7 Perhaps this is partly because, in addition to the estimated VPF's dependence on the assumed time preference rate, it is also sometimes allowed to vary with life expectancy and hence age.

8 Long established in health economics, as a ***non***-monetised measure for comparing health states, is the Quality Adjusted Life Year (QALY). The question of valuing the QALY is contentious, but in some circumstances it is valued, usually as a way of valuing different health states. The full health value is often described as the VOLY (Value of a Life Year)

3) The assumptions about human preferences implied in the LQI derivation of people's WTP to reduce small fatality risks from measures of income, expected life span and the ratio of society's working hours to non-working hours, are manifold. They are not discussed, nor supported by any material empirical evidence.[9] This is not to criticise the making of such estimates. They can open up constructive questions. But such estimates cannot immediately displace other methods that, even if far from above criticism, have a strong empirical basis and a long history of development and review.

Lesser qualifications include the following.

- The LQI analysis does not, and makes no claims to contribute to the valuation of illness or non-fatal injuries, nor to the understanding of societal concerns, all of which are today more challenging issues than the valuation of fatality risks.
- The LQI approach does not lend itself very easily to the handling of front loaded expenditures on safety measures. The defence of the simplification of assuming a zero discount rate in allocating such expenditures over time – that this favours expenditure on safety equipment – is not very rigorous.
- One aspect of current UK conventions that is largely replicated in the J-value papers is the concept of indifference to the ***incomes*** of those at risk. The J-value comparisons do however sometimes discriminate by ***age***, which is not the current convention (except in the fairly extreme circumstances of people with very short natural life expectancies). However this could easily be changed and in any case it is not, in the examples considered, quantitatively significant.
- As noted above the LQI arithmetic, given its method of derivation, can easily be compared with QALY valuations as well as VPFs. However the comparisons in Thomas et al (2006.1 and 2006.2) overlook the nature of the NICE valuation. In contrast to the VPFs used in Whitehall, the NICE valuation of the QALY incorporates a public expenditure constraint. Thus while VPFs are generally valued in terms willingness to forego consumption, the NICE QALY is valued in the higher value numeraire of willingness to forgo public expenditure (on alternative uses). This reduces its monetary value significantly (because the distortionary effects of taxation give £1 of public expenditure a greater social cost than £1 of consumption), though by how much is debatable. This however is not itself criticism of the LQI derivation, but a marker on the need for some adjustment to its application to the NICE value of the QALY.
- The LQI takes no account of lost net output or public service costs associated with fatal accidents, but these could of course easily be added.

---

and, as discussed in Appendix A, it is sometimes suggested that there should be some equivalence between the value of a prevented fatality and the present value of the VOLYs expected from a person's natural remaining lifespan – in the same way that the LQI is the present value of a measure of future annual welfare.

9 One technical example is the unavoidable constraint of the elasticity of marginal utility, $-e$, in the LQI derivation. It is constrained to values between zero and -1, while the evidence (summarised in Appendix B to this Report) is now very strong that $e > 1$. More substantial concerns lie around the behavioural and social preference inferences that can be reliably drawn from such a "desk based" estimate. As examples, a large share of personal income is not wages (although an attempt is made to address this in Thomas, Kearns and Jones (2009)); the scope that people have for trading working time against income and leisure is usually very constrained; many or most of the trade offs between working and "leisure time" may be for untypical demands on non-remunerated time, such as caring for dependents; and many people may value their working time as a social benefit (to themselves) no less than as time given up in exchange for income.

- The use of the term "utility" in the context of the LQI is not quite right. The use of absolute measures of welfare is well established, as for example in what is often called the "happiness" literature, such as Layard et al (2007), which is cited elsewhere by the City University authors. But utility and utility functions are tightly defined concepts, which have meaning only in terms of comparative differences. However in this immediate context this is a matter of wording more than technical substance.

It appears that the Canadian work has been around for well over ten years, but it does not appear to have been taken up in any serious practical application. It may however have helped to achieve its possibly primary objective of promoting more informed debate about arguably excessive regulation of some nuclear hazards.

3.2.3. The J-value application of the LQI VPF

"J-value" is an ingenious term that might possibly encourage a wider interest in the concept of comparing costs and benefits of measures to reduce fatality risks. It has more buzz than dry terms such as benefit / cost ratio and value of a prevented fatality. Beyond this however it is hard to see advantages in it as an analytical instrument, relative to those of current conventions.

Apart from the new (LQI) approach to valuing fatality risks, discussed above, it brings nothing new to the table. Nor indeed does it lose any input data. But, from the user's perspective, much is hidden in the J-value equation, which produces a dimensionless number as if by magic. It is not instantly clear, for example, whether an LQI VPF applied in any particular case is a standard value or tailored to a case-specific age distribution. Nor indeed, discussed further below, is it clear that a monetary value of a prevented fatality is being applied at all.

The equation incorporates steps that would under conventional procedures be distinct, such as the identification of a VPF appropriate to the case in hand, the separate estimation of the number of people affected and the level of additional risk that they would face, and the separate identification of the cost of the risk reducing measure. All would be considered independently until, as a separate step, they were brought together to produce a number or numbers in the form considered most appropriate for the particular application.

Sometimes this might be (like the J-value) the ratio of the cost of a potential measure to the value of the expected fatalities prevented, or the inverse of that ratio, or sometimes the implicit VPF. The various inputs might often be combined with other costs, such as accident costs, or the costs of illness and non-fatal injury, before any final output is calculated. The J-value procedure precludes such flexibility as it is self-contained and concerned only with fatality risks.

Any mathematical process is "transparent", in the sense that a technically competent reviewer can dissect it and identify the inputs and how they are processed. However the J-value algebra is not readily transparent to users. Moreover the inputs incorporate many implicit assumptions about social behaviour and preferences that are not revealed. The current stated preference conventions, for all their weaknesses[10] are much more transparent to a critical reviewer.

In some politically sensitive contexts obscurity may be seen as a virtue. The concept of valuing fatality risks can be difficult for some audiences to accept, and the J-value conventions can apply such a value without its ever being explicit. But it is doubtful that analysts in any government department would regard this as good practice.

[10] A good assessment of the difficulties faced by almost any substantial stated preference studies was set out by Professor Graham Loomes (now of University of Warwick) in the final chapter, on "Cautions, caveats and future directions" in Bateman et al (2002).

4. J-values applied to non-human costs and risk aversion

Thomas and Jones (2009.1 and 2009.2)[11] and Thomas, Jones and Boyle (2009) develop a process to bring into a J-value framework all non-human costs (described collectively as "environmental costs") and incorporate a measure of "risk aversion". The analysis is lengthy, but summarised here briefly, drawing mainly on Thomas, Jones and Boyle (2009).

For any reader not familiar with the concept of a concave utility function, and the way in which it can be used to estimate a cost of risk, the outlines are explained, pictorially and with examples, in section C.2 of Appendix C to this present Report.

The City University analysis starts with the presumption that data is available on the probabilities of the non-human hazards, and on their monetary value, and, for a given hazard, the costs of protecting against it.

4.1. Preliminary analysis

The analysis assumes that the investing institution has a utility function with respect to the size of its assets, and that this function has a constant elasticity of marginal utility.

The authors note that this elasticity, apart from its sign, is algebraically identical to the "coefficient of relative risk aversion", which they denote (following common convention) by *e*. As *e* increases the curvature of the utility function, and hence its effect on the cost of risk, increases.

Four variables denoted by Thomas, Jones and Boyle are *A* (the organisation's assets, £), *B* (the cost of a prospective risk reducing measure, £), *C* (the cost of a potential accident, £) and *D* (the difference between the organisation's expected utility with and its expected utility without the risk reducing measure). The basic algebra presenting this difference, D, between those two expected utilities in terms of *A*, *B* and *C*, given accident probabilities of $1-p_1$ of the accident with the risk reducing measure and $1-p_2$ without it, is presented as the "ABCD model".

Account is then taken of the possibility of more than one accident within the time span considered, on the conventional assumption that the relative probabilities of 1, 2, 3, etc accidents follow a Poisson distribution.

4.2. The "risk multiplier"

The main focus at this stage of Thomas, Jones and Boyle (2009) is on how the monetary value of a risk reduction increases as *e* increases from zero. This leads to the valuation, over a range of values of *e*, of a potential accident risk, from the organisation's perspective, in terms of what maximum risk-reducing expenditure would be justified to eliminate it. This

11 Thomas and Jones (2009.1) is specific to MOD. Thomas and Jones (2009.2) is a text based on (2009.1), but adapted and submitted for publication.

analysis is followed through in terms of *e*, or in terms of an elasticity *q* (= 1 - *e*), with respect assets, of utility itself, as opposed to marginal utility.[12]

This leads to the derivation of a "risk multiplier", defined by Thomas, Jones and Boyle as "the maximum that the organisation should be prepared to spend on the environmental protection system at a risk aversion of *e* divided by the maximum it would be prepared to spend at $e = 0$." In other words the "risk multiplier" is the factor by which the curvature of the utility function increases the monetary value of eliminating the risk of a particular loss.

This ratio is illustrated in section C.2 in Appendix C to this Report, which explains and demonstrates a derivation of the risk multipliers described as exact figures in Thomas, Jones and Boyle Figure 7, which follows from their Table 1.

The derivation of the "risk multiplier" is followed by a short diversion into some actuarial literature, which uses an approximate method to estimate the effect of *e* on the monetary value of a risk that has not only a given probability of occurring but where the hazard itself, were it to materialise, has a probability distribution with a defined mean and variance. The authors test the effect of this, but find that the results are "relatively insensitive" to the variance and this does not reappear in the analysis. They do however present estimates of their "risk multiplier" using the approximate method, described as the Pratt approximation, to compare with their exact figures.

4.3. "Reluctance to invest", the "permission point" and the "point of indiscriminate decision"

At the end of their development of risk multipliers the authors say the "The question arises now as to the largest value the risk aversion, *e*, can take and still be reasonable." From here, in section 6 on page 19, beginning with equation (54), the Thomas, Jones and Boyle follow a new path, based on "the ratio of the utility difference, $D(u_1, u_2|e)$ at a given value of … *e* to the utility of the starting assets, *A*, at the same *e* … ".[13] This ratio is described as "Reluctance to invest" and is said to provide "a convenient, dimensionless measure of how reluctant the decision maker will be to invest in a protection system".

This algebra produces some results that at first sight look curious.

The authors in their Figures 8 and 9 plot " reluctance to invest" as *e* increases for the case previously used in Figure 7, to illustrate the risk multiplier. However as *e* increases the "Reluctance to invest" instead of decreasing as might be expected because of the increasing risk multiplier, increases steadily. Even more strikingly the (negative) "reluctance to invest" as *e* increases is not only asymptotic to zero but is shown as, for practical purposes, equal to zero when *e* is larger than about 1.2.

These figures can however be explained fairly easily.

---

[12] Although an elasticity of utility has no real meaning, as values of utility have meaning only as differences. Thus a utility value of zero does not mean "zero welfare" and equating zero utility with zero income, or consumption, or assets is arbitrary.

[13] u_1 and u_2 are utility without and with the protection system, the difference *D* being defined so that it is negative if the protection is worth its cost (i.e. $u_2 > u_1$).

When $e = 0$, utility, in the form explained above and in Appendix C to this Report, can be equated with wealth. Thus, taking the data from Thomas, Jones and Boyle Table 1 (and ignoring the very small expected cost of an accident even with the protective measure): the utility u_o of the starting assets can, for $e = 0$, be taken as 10,000 (since the starting assets are set at £10,000m); the expected utility u_1 without the protective system can be taken as 10,000 minus 5 (since the expected cost of the accident of is £5,000m x 10^{-3}); and u_2 as 10,000 minus 1 (since the cost of the protective measure is £1m). Thus "Reluctance to invest" for $e = 0$ is given by $(u_1 - u_2)/u_o = -4 \times 10^{-4}$, as shown in Thomas, Jones and Boyle Figures 8 and 9.

The calculation of how this ratio changes as e increases is algebraically tedious. However, as shown in section C.2.5 in Appendix C to this Report, for $e < 1$ and again the case set out in Thomas, Jones and Boyle Table 1, to a close approximation, $(u_1 - u_2)/u_o = (e - 1)(5M_R – 1) \times 10^{-4}$, where M_R is the (exact) "Risk Multiplier" plotted in Thomas, Jones and Boyle Figure 7. Thus, for example, for $e = 0.5$, "Reluctance to invest" is given approximately by $(0.5 – 1)(5 \times 1.17 - 1) \times 10^{-4} = -2.42 \times 10^{-4}$, as in Thomas, Jones and Boyle Figure 8.

The behaviour of this quotient, as e increases, follows from the changing structure of the conventional constant-elasticity utility functions used by Thomas, Jones and Boyle. For values of $e > 1$ the utility function, as the value of assets, A, increases, is asymptotic to some maximum value. This means that for numerically very large values of A (and Thomas, Jones and Boyle are using values of the order of 10^{10}) the utility curve becomes extremely flat. Thus as e increases the numerical value of the difference $(u_1 - u_2)$ for a given percentage change in assets becomes dramatically smaller. The numerical value of the "absolute utility" u_o of the starting assets also declines, but not as rapidly, as illustrated in section C.2.5 in Appendix C to this Report, which approximately reproduces the figures plotted in Thomas, Jones and Boyle Figure 9.

Thomas, Jones and Boyle present another case, in which the accident would cost 95 per cent of the organisation's assets. In this case, as e increases, the "reluctance to invest" does initially decrease, driven by the increasing risk multiplier, but the "reluctance to invest" reaches a minimum and, as in the previous case, is shown as rising to become asymptotic to and again for practical purposes equal to zero when e is larger than about 1.2.

In cases where, as e increases, the reluctance to invest passes through a minimum (so that the implied wish to invest passes through a maximum) the minimum is described as the "permission point". And the point at which the "reluctance to invest" becomes for practical purposes equal to zero is described as the "point of indiscriminate decision". This latter term is used because, by an obvious symmetry, while the algebra shows a negative "reluctance to invest" in a safety system, it also shows a positive "reluctance to invest" in a "danger system" that would reverse the action of the protection system. While always positive, the reluctance to invest in a "danger system" is also asymptotic to zero (from above the y axis) and is also shown to become for practical purposes equal to zero when e is larger than about 1.2.

Thomas, Jones and Boyle (2009, p21) suggest that in this situation "the decision maker will have to make a decision (for doing nothing is still a decision), but will be unable to discriminate between any of the options before him because to him there is now no measurable difference [between spending to move from the higher to the lower risk state and vice versa]. He will see all his options through a thick fog, and will be unable distinguish between good and bad."

In section 10 of the paper the algebra is applied to "the average UK adult", with assets, as derived from ONS data, of £180,000. This starts with the explanation that the previous sections had shown that "a risk averse decision maker will not have a unique value of risk aversion, *e*, but will vary it to find the optimum value, e_{opt}, known as the permission point, that will provide him with the greatest feeling of justification for spending on protection. This is the value of risk aversion actually applied by the rational decision maker in taking a decision to invest in a scheme offering protection."

This is followed by the development of a further concept of wishing to avoid ***entirely*** all cases where the risk of a hazard is more than some level that transforms the situation from a "risk" to a "gamble". It is suggested that this risk level, p_{10}, "might be set at 0.5 on the basis that probabilities below 0.5 would imply that the loss is more likely to happen than not" and "a strongly risk-averse person might regard the effective transition from risk to gamble as occurring at a higher value of p_1, perhaps as high as 0.8." This is developed with further algebra and graphical presentations.

At the end of this section (page 35) and in the paper's Introduction (page 5) the authors suggest that their conclusion that, for individuals, *e* lies broadly in the range of 0.5 to 1.1 is compatible with studies of this quantity by other authors.

At the very end of this section, just before the paper's Conclusions, the authors note that "setting the risk-gamble transition probability, p_{10}, to 0.5 would mean than the average value of applied risk aversion becomes $e_{optave} = 0.65$. Applying this value in J-value analysis produces an average value for human life that is roughly half the £2.5M normally calculated [in the J-value papers] and thus similar to current Department of Transport estimates. The implication that the J-value estimate has a disproportion factor of roughly two built into it requires further study."

4.4. The comprehensive J value

The application of a J-value formula to risk aversion is set out on Thomas and Jones (2009.2), with three new J values, J_2, J_{20} and J_T. J_2, the "Second Judgement Value", is the ratio of the actual proposed spending to protect against non-human cost damage to the estimated maximum 'rational' expenditure (i.e. the monetary value, after allowing for risk aversion, of the reduction in the risk of non-human costs). J_{20} is an aside. J_T is the ratio of proposed total spending divided by the maximum, 'rational' total expenditure when J_2 is constrained to be equal to the J-value adopted for the expenditure to protect against ***human*** harm, as described in chapter 2 above.

5. Commentary on the J-value approach to non-human costs and risk aversion

Commentary on the City University work is preceded here by a description of the currently conventional handling of risk aversion and non-human impacts.

5.1. The conventional approach to risk aversion

Appendix C to this Report summarises the current situation, in government and in private sector companies, with regard to variability risk and to the handling of large downside risks.

In nearly all corporate and government circumstances, especially in government, the case for "risk-neutrality" (i.e. optimising in the basis of expected values, or taking $e = 0$) is firmly established. The principle exception is the need to take account of the cost of systematic risk, which affects the cost of private equity capital, but is very rarely a significant issue in government.

Companies are (or should be) averse to financial risks that threaten their sustainability, but for government this is not a material issue for the generality of risks in public expenditure projects, such as those being considered on the current context by MOD.

However very large downside risks, especially major accidents, tend to be associated with costs, to those held responsible and sometimes to society more widely, that go beyond the valuation of the potential loss of life, other injuries, and direct material damage. Such costs, which are often categorised in government as societal concerns as set out in Appendix D to this note, are reflected in practical decision making, although the way in which this is done is often not well structured or clearly explicit.

5.2. The J-value approach to risk aversion

5.2.1. Institutional risk aversion

The J-value papers on risk aversion get off to an unfortunate start in supposing that utility functions, of the kind developed to describe individual preferences, apply also to non-sentient bodies such as government departments.

Even for people, a simple "constant-elasticity" utility curve serves well in only some very simple circumstances that are free from the psychological influences that often influence actual preferences.

For companies, and more so for government, there is nothing in the J-value papers that gives cause to question the normal convention of risk-neutrality. Very large hazards may often present potential social costs beyond the expected *direct* human and physical costs. But the algebra of risk aversion in the sense of constant-elasticity utility functions does not help to identify such potential costs or to value them.

However the "risk multipliers" derived in the J-value papers are correct for the utility functions that they assume.

5.2.2. "Absolute utility"

The authors show how the quotient of (utility difference between 'unprotected' and 'protected') divided by ("absolute utility" of the initial assets) approaches and is for practical purposes equal to zero for values of *e* more than about 1.2 and draw conclusions from this. The authors' choice of this quotient appears to have arisen from an eagerness (commendable of course in many circumstances) to normalise and to create dimensionless numbers. However, as the authors themselves recognise elsewhere, "absolute" utilities have no meaning.[14]

The concepts of "reluctance to invest", the "permission point" and the "point of indiscriminate decision" as derived in the J-value literature thus likewise have no practical meaning. The surprising behaviour of the numbers in the J-value papers is amplified by the very large numerical values of *A*. Had a larger numeraire been used for *A*, giving numerically smaller values, the concepts would still have no practical meaning but, as illustrated at the end of Appendix C to this Report, the "reluctance to invest" would not decay in such a dramatic way to a "point of indiscriminate decision".

It is surprising that Thomas, Jones and Boyle offer no intuitive explanation of the implication in their results that people with a utility function with a value of *e* more than about 1.2 (which for several decades in the literature, on the basis of much evidence as illustrated in Appendix B, has been considered quite plausible for people in general) could not make the decisions implied by the "risk multipliers" that are correctly represented in their Figure 7.

The suggestion near the end of Thomas, Jones and Boyle (2009) that individuals adjust their risk aversion to the point at which they are most satisfied with their choice is imaginative and has a flavour of insight into the ways that people self-rationalise. However the determinants of personal choice are remarkably complex, with many now well established consistent inconsistencies.[15] The constant-elasticity utility framework has little to contribute to this field.

5.2.3. Other issues

5.2.3.1. The value of *e*

The claim that the values of *e* derived by the J-value papers are consistent with other work does not fully reflect a very large and diverse body of economics literature. The main relevant passage (Thomas, Jones and Boyle, 2009, p5) is as follows.

"The first-stage, societal trade-off needed for J-value analysis, as described in Thomas and Kearns (2009.2) enables a central value of $\varepsilon \approx 0.82$ to be determined for individuals in the UK. This is similar to the estimate of 0.83 found using a diverse method by Blundell et al. (1994), and the recommendations of Pearce and Ulph (1995), who comment 'A value of 1 is defensible and a value relevant to UK conditions is ... 0.8 – 0.9'. Layard et al. (2007) use a number of surveys of subjective happiness to produce an estimate of 1.24. H. M. Treasury (2009) suggests that a reasonable range is between just below and just above unity, and it recommends a value of 1.0. Cowell and Gardiner (1999) perform a review of past estimates, and note that 'the estimates obtained for

[14] Thomas, Jones and Boyle (2009, p5) notes correctly that "only differences between utilities constitute a valid measure of preference". As they explain, " utilities are regarded as valid only to a positive linear (affine) transformation."

[15] As illustrated in part by the popular book by Professor Dan Ariely: ***Predictably Irrational: The Hidden Forces That Shape Our Decisions***, Harper Collins, 2008.

relative risk aversion [≡ elasticity modulus] by the indirect route of inference from the intertemporal substitution elasticity are fairly consistent. Most imply values for the [modulus of the] elasticity of marginal utility of just below or just above one.' However, they feel obliged to conclude that the elasticity modulus may lie within a rather wide range, from 0.5 to 4.0."

All of these references, and several others, are discussed in Appendix B to this Report. Taking each in turn of those cited above:

- Blundell et al (1994) is a first class paper, but not addressed to the valuation of *e*. The relationship between the elasticity of intertemporal substitution, on which the paper reports, and *e* depends upon several factors. The study was in any case based on 1970-1986 data that ended in the year in which UK retail financial markets were deregulated, which may have radically changed behaviour. The Thomas, Jones and Boyle quotation from Cowell and Gardiner (1999)[16] that "the estimates obtained for [*e*] by … inference from the intertemporal substitution elasticity … Most[ly] imply values for [*e*] of just below or just above one" is rather misleading. In the source document (Cowell and Gardiner, section 5.2, p 31) the first sentence begins with the words "In contrast, the estimates obtained …" Cowell and Gardiner are drawing a contrast between these (low but consistent) values and the (often very high, but also fairly diverse) values found by researchers measuring the intertemporal elasticity by experimental approaches, as opposed to analysis of economic statistics (see section B.2.2 in Appendix B to this Report).
- Pearce and Ulph (1995), while written by authors of deservedly high repute, is something of an outlier. Its origins were in consultancy and although a version was later submitted for publication in a refereed journal it was never accepted. It should not therefore be seen as very authoritative paper.
- Layard et al (2007) is a high quality paper and correctly quoted, although there may be reason to believe that this approach is subject to some downward bias in its estimation of *e*.
- The reference to HM Treasury (2009) is to the web version of the Treasury Green Book, as published in 2003. Its coverage of the discount rate is in places questionable and relied heavily on Pearce and Ulph.
- Cowell and Gardiner (1999) is a major reference and it is slightly misleading to say only that, looking across all the studies they considered, "they feel obliged to conclude that the elasticity modulus may lie within a rather wide range, from 0.5 to 4.0." These authors give substantial weight to derivation from the income tax structure and their final comment is that the range of 0.5 to 4.0 brackets "the values of 1.2 to 1.4 that are implied by .. the personal tax system [for the UK around 1999]".[17]

---

16 Cowell, F A and K. Gardiner. (2000) 'Welfare Weights', *Economic Research Paper* 20 (oft282.pdf), London, Office of Fair Trading

17 There is reason, furthermore, to expect that figures derived from the personal tax structure, at least since the 1980s, are biased downwards, as tax rates are set to strike a balance between equity (implying a fairly strongly progressive structure) and incentives (implying a less progressive structure.

Cowell and Gardiner overlooked the work, which has come more to life since their paper, on deriving *e* from income and price elasticities. This also consistently gives values of more than 1 (see section B.2.4 in Appendix B to this Report).

5.2.3.2. "Environmental impacts"

The use of the term "environmental", for wider impacts has a presentational appeal, and sometimes environmental costs may be a major potential cost (e.g. cleaning up the consequences of a large release over land of radioactive material). But it perhaps distracts from the many other wider impacts found in practice, as noted in Appendix D to this Report.

6. Conclusions

6.1. General observations

6.1.1. The value of interdisciplinary work on health and safety

It is good to see interdisciplinary work and new ideas in any field and in particular in applications such as analysis of the costs and benefits of potential spending to improve health and safety. There are important issues in the field that remain persistently resistant to formal analysis and no one discipline alone is equipped to handle them.

In the case of the J-value literature the authors have shown exceptional energy and enthusiasm and an interest in and willingness to delve into the economics literature. Their publication in a technological journal of comparisons of regulatory conventions and applications was a notable and valuable achievement.

It is encouraging also to see the willingness in principle of research councils in Canada and in the UK, including both the EPSRC and the ESRC, to support such work. And the MOD must be commended for seeking to develop formal analytical tools in an area of decision making that is not well served by current conventional approaches.

6.1.2. Clarity and analytical insight

The two J-value papers on “The Extent of Regulatory Consensus on Health and Safety Expenditure” are superbly clear, and a fine read, in the picture that they portray. But even there, and more so in the papers on risk, the algebraic structure is not always intuitively friendly and there is little intuitive explanation to support it. There is some tendency to perhaps overenthusiastic normalisation, to seek dimensionless numbers. The J-value ratios as presented tend to be less transparent and flexible for users than would be required for practical application in government.

The authors have done well to explore and exploit fields of literature that have been new to them, but perhaps have not fully appreciated the depth of current study and practice in the valuation of fatality risks, nor the limitations of utility functions. And they have not fully recognised the significance of the literature on systematic and non-systematic risk in companies and in government, and on personal decision making.

6.2. The J-value concept for the assessment of expenditure on fatality risks alone

Some of the J-value papers addressing fatality risks serve a very useful purpose in publicising some of the prima facie inconsistencies across regulators and between regulator’s formal positions and the regulations they implement. The differences presented offer no great surprises for experts familiar with the general field of safety regulation. They could also be presented more simply with the use of current government analytical conventions. But the J-value format may have helped in some contexts to publicise the issues to a wider audience.

However the potential of the Life Quality Index is limited and it is not a serious competitor to the methods already used in UK government and widely in other countries for valuing fatality risks. It is good that UK professionals should know of it, and it is a bold approach to estimating the value of a prevented fatality from published statistical data, far removed from the observation of individual stated or revealed preferences. But it is based on many strong assumptions about human preferences and behaviour that leave it much less robust against informed criticism than currently established methods..

One assumption is that people value small fatality risks in terms of "discounted lifetime welfare". It seems likely that people's preferences with respect to such risks are based mainly on other factors. But even if a lifetime welfare approach were to be adopted in the UK it would be based in practice on the well established 'Value of a Life-Year', originating in health economics.

"J-value" itself is a name used by the authors to describe one of several ratios that may be used to define value for money, but this by no means the best ratio for all circumstances. A significant problem with the methodology presented for calculating the value is that a great deal is compressed into a single formula, with little deference to transparency or flexibility in its practical application.

6.3. The J-value concept as applied to non-human costs and to risk aversion

The approach to risk aversion in the J-value papers starts from a misperception of the application of utility functions. There is no good reason to ascribe a personal utility function to a government department. Even if it were, the proposed method of analysis would be too complex and non-transparent for general application in government.

After correctly deriving the (fairly modest) "risk multipliers" that would be implied by an "constant-elasticity" utility function, the papers overlook the fact that "absolute" utility values have no useful meaning. Thus the derived concepts of "reluctance to invest", "permission point" and "point of indiscriminate decision", while demonstrating impressive energy, constructive imagination and algebraic dexterity, have no practical meaning.

6.4. Meeting the MOD requirement

The MOD requirement is for a better, quantified analytical rationale for spending more on some safety measures than can be justified solely by conventional valuations of risks of lives lost. This they sometimes describe in terms of rationalising the case, in some circumstances, for "gross disproportionality".[18]

This general problem arises in many fields, although the very different levels of expenditure per fatality prevented in different fields sometimes appear to reflect institutional incentives. More spending tends to be required, per expected fatality prevented, on preventing hazards that have painful repercussions for those politically or managerially responsible (e.g. rail

[18] "Gross disproportion" is usually used in the HSE context as a general requirement for safety expenditure. This general usage is challenged by outside experts as being now outdated, given the emergence of preference based valuations of health and safety impacts. It was condemned by the House of Lords Economic Affairs Committee in 2006. It would be better if it were quietly retired, but this appears to be institutionally difficult. However the MOD usage seems in contrast sensible, insofar as there are circumstances where the normal measures of proportionality are strongly supplemented by wider considerations. The J-value literature refers to gross disproportion in terms of especially large hazards.

accidents, other strongly newsworthy accidents) than on those that do not (e.g. road accidents; radon in the home).

However there are circumstances where some such “disproportionate spending” may be uncontentious, because there are obviously much wider potential impacts (e.g. on the political acceptability of nuclear power, or the diplomatic impacts of biologically trivial Technetium-99 releases into the sea).

Such wider impacts are often described as societal concerns, as discussed in Appendix D to this Report.

An attraction of the J-value papers in this context is that they offer an analysis that appears to support larger numbers for expenditure on preventing fatalities than current conventions. However the higher value for a prevented fatality emerges from a bold but very rough and ready method of estimation that cannot sensibly rival the stated preference methods now used. And the additional appeal to risk aversion is not technically well based.

There may however be other, more pragmatic avenues worth considering, given that there are often wider concerns that justify higher safety expenditure than is justified by conventional valuation of the expected lives saved. Such avenues might include the following, no combinations of which are mutually exclusive.

- Stronger identification of costs that are not societal concerns as normally defined, but simply extra costs specific to the MOD’s application. One such cost, which might be included in the value of a prevented fatality (VPF) when appropriate, is training costs as noted in JSP 507. There may be other identifiable costs, associated for example with equipment availability, that may not enter into a VPF but might routinely appear in value for money calculations and which are not at present fully developed.
- Development of a more formal, standardised, presentation of wider impacts in the form of an Appraisal Summary Table, as is used in DfT and is being developed in the context flood risk management by Defra.
- Engagement in this context with the economists in the HSE and MOD and the current Government Economic Service working group on the VPF, which could benefit from input from outside that discipline and from contact with more real operational problems. It might usefully be encouraged to develop the earlier progress on societal concerns outlined in Appendix D to this Report.
- Consideration of whether multi criteria analysis techniques, as have been used in parts of MOD, might be appropriate in some cases. This is typically a facilitated process in which senior experts allocate several hours to a process of prioritising options against well defined criteria.

The issues raised by the J-value papers might also usefully prompt some revision of JSP507 to produce a more user friendly set of perhaps web-based documents.

References

As noted in Chapter 1, references cited in the main text are listed here. References specific to the Appendices are listed at the end of, or in footnotes to the relevant Appendix.

Bateman, I.J., Carson, R.T., Day, B., Hanemann, M., Hanley, N., Hett, T., Jones-Lee, M., Loomes, G., Mourato, S., Özdemiroğlu., E, Pearce, D.W., Sugden, R. and Swanson, J. (2002) *Economic Valuation with Stated Preference Techniques: A Manual*, Edward Elgar for the UK Department for Transport

Belzer, R.B. (1992) *The Use of Risk Assessment and Benefit-Cost Analysis in U.S. Risk-Management Decision Making*, Paper to International Conference on Risk Assessment, Queen Elizabeth II Conference Centre, Westminster, London, 8 October

Ministry of Defence (2006) *MOD Guide to Investment Appraisal and Evaluation*, JSP 507, Version 3.0

Pandey, M.D. and J.S. Nathwani (2003) *A conceptual approach to the estimation of societal willingness to pay for nuclear safety programs*, Nuclear Engineering and Design, 224: 65-77

Pandey, M.D., Nathwani J.S. and Lind, N.C. (2006) *The derivation and calibration of the life-quality index (LQI) from economic principles*, Structural Safety, 28, 341-360

Tengs T.O., Adams, M.E., Pliskin, J.S., Safran, D.G., Siegel, J.E., Weinstein, M.C. and Graham, J.D. (1995) *Five-Hundred Life-Saving Interventions and Their Cost-Effectiveness*, Risk Analysis, Vol 15, No. 3

Thomas, P.J. and Jones, R.D. (2006) *J-value concept testing*, Report produced under contract by City University, April

Thomas, P.J. and Jones, R.D. (2009.1) *J-value Phase 3*, Restricted Report produced under contract by City University, July

Thomas, P.J. and Jones, R.D. (2009.2) *Extending the J-value framework for safety analysis to include the environmental costs of a large accident*, Paper for submission to Process Safety and Environmental Protection

Thomas, P.J. Jones, R.D. and Boyle, W. J. O. (2009) *The limits to risk aversion*, Paper for Process Safety and Environmental Protection, 13 July

Thomas, P.J., Kearns, J.O. and Jones, R.D. (2009) *The trade-offs embodied in J-value safety analysis*, unpublished but submitted for publication

Thomas, P.J., D.W. Stupples and M.A. Alghaffar (2006.1) *The Extent of Regulatory Consensus on Health and Safety Expenditure, Part 1: Development of the J-Value Technique and Evaluation of Regulators' Recommendations*, Trans IChemE, Part B, September, Process Safety and Environmental Protection, 84(B5): 329-336

Thomas, P.J., D.W. Stupples and M.A. Alghaffar (2006.2) *The Extent of Regulatory Consensus on Health and Safety Expenditure, Part 2: Applying the J-Value technique to Case Studies Across Industries*, Trans IChemE, Part B, September, Process Safety and Environmental Protection, 84(B5): 337-343

Appendix A. Objections to the valuation of fatality risks by discounting the value of future life years

One approach to estimating the Value of a Life Year (VOLY), is to regard VOLYs as an annuity, which, when discounted over the expected remaining life span of an individual at risk, would equal the value of a prevented fatality (VPF).[19] Conversely it sometimes proposed, as by Pandey and Nathwani (2003), that the valuation of future life years can be used to derive a VPF.

This has obvious appeal. It is closely analogous to how a physical asset such as a piece of machinery would be valued – that is, as the present value of its expected annual value over its remaining lifetime. It is also presumably how slaves are valued by owners who value only their labour. But free men and women differ from slaves or machines in ways that provide at least pause for thought about whether equating the VPF with discounted future individual utility in this way is reasonable.

One fundamental difference is that, while the destruction of a valuable machine imposes a subsequent loss of service to those who would have used it, the same is not true of the welfare that a person might have enjoyed had he or she not died.. That is not, after death, a loss to anyone.[20]

The most that can be said is that people can foresee the prospect of a natural lifespan ahead of them *ex ante*, and thoughts about how they would spend and hopefully enjoy it may well be one factor in their willingness-to pay (WTP) to reduce risks of premature death. This may be especially strong for those who wish to complete some special, say literary, artistic, business, academic, or sporting achievement. However there is little clear reason to suppose that this is even crudely proportional to years of life expectancy. In the absence of empirical evidence, it may be more reasonable to suppose that, on the contrary, for so long as they remain in robust health, people tend to accumulate more, and more strongly, missions that they wish to accomplish as life proceeds.

Furthermore, following a person's premature death, the consequences for those remaining bear little relationship to the welfare that the deceased might have enjoyed. The consequences include the net output that the deceased would have produced, but this can be, and usually is separately accounted for in valuing fatality risks and is uncontentious. More substantial, typically, is the grief from loss of love, companionship and / or care that the deceased would have provided to those close to him or her. There is no significant analogue for this following the loss of the inanimate machine.

---

[19] As clearly presented, but not necessarily advocated, by Pearce (2000).

[20] Unless the argument is carried into theology. One surreal model might be that of the person still being conscious in another form, in a miserable, zero VOLY state for the rest of his or her previously expected life span, grieving for the things that would otherwise have been enjoyed in life. Less absurdly, it might be argued that social welfare should be measured by aggregate utility, so that a population of one million with a certain quality of life is no more desirable than a population of two million with half the quality of life. A version of this view has some advocates in current moral philosophy (see for example Broome, 2004), but it is not an argument applied in mainstream welfare economics. (Provision of a given *per capita* gain to a large population is counted as proportionately more valuable than providing the same *per capita* gain to a small population; but simply increasing the population, with no increase in *per capita* welfare, is not of itself seen as a welfare gain.)

A major factor in aversion to the risk of death may for many people be perceptions of such grief and hardship that their death would bring. This might be expected to follow, over a lifetime, not a near-linear decline but an inverted U, as people acquire more dependants and in due course become themselves dependent. The objective *consequences* of death at various ages is a field in which more empirical research could be helpful. The welfare consequences (for those surviving) of the death of a typical 20-year old may be less (or more) than those of the death of a typical 40-year old.

Yet another motivation for reducing the individual risk of death, shared with other higher animals, may be a strong will to live, and so to mitigate risks such as those posed by dangerous carnivores. This is similar to the "love-of-life element" discussed briefly in Loomes (2002). It is not clearly at all age dependent, for people in robust health.[21]

The valuation of a life by discounting its future annual value is in any case some way removed from people's preference structures insofar as it implies that discounted future welfare measures the cost of certainty of death very soon. Few people, if any, would value certainty of death in this way. The valuation needed for policy purposes is one that reliably measures preferences about *small risks* of death.

All told, expecting or requiring a close relationship between the VPF and the present value of discounted VOLYs implies a very simple model of why people are averse to risks of premature death. This seems to be reflected in a general preference by governments for WTP values for risks of death, estimated by some means other than equating the VPF with the present value of an annuity of VOLYs, even if a VOLY value is used in the valuation of ***non-fatal*** impacts.

References for Appendix A

Broome, J. (2004) *Weighting Lives*, Oxford University Press

Loomes, G. (2002) Valuing life-years and QALYs: 'transferability' and 'convertibility' of values across the UK public sector. Chapter 5 in Towse, A., C Pritchard, and N Devlin (eds) *Cost-Effectiveness Thresholds: Economic and Ethical Issues*, King's Fund

Pearce, D (2000) Valuing Risks to Life and Health: Towards Consistent Transfer Estimates in the European Union and Accession States. Paper prepared for the European Commission (DGXI) Workshop on Valuing Mortality and Valuing Morbidity, November 13, 2000, Brussels. Revised, December. http://www.cserge.ucl.ac.uk/EU_Valuing_Lives.pdf

---

21 Although there is the question of whether the apparently lower stated preference value of young adults for reducing small fatality risks may be partly attributable to the 'recklessness of youth', or an immature sense of invulnerability. Whether it is society's ethical responsibility to impose 'mature' valuations is a political choice and not clear cut. It is also far from clear that considered public preferences would ascribe a declining VPF to children as they grow from infancy towards maturity.

Appendix B. Estimation of the elasticity of marginal utility

This Appendix summarises the approaches (other than the Life Quality Index approach) to valuing the elasticity of marginal utility (of consumption, with respect to marginal changes in consumption). The true sign of this elasticity is uncontroversially negative (i.e. the wealthier the individual, the less extra utility is gained or lost by the gain or loss of a few pounds or dollars). But by convention the sign is almost invariably reversed, as in this Appendix, in applied economics literature. In Thomas, Jones and Boyle (2009) and other related literature cited in the main text to this note the elasticity is denoted by $-e$ (where e denotes the coefficient of relative risk aversion, which is algebraically identical apart from its sign). In this Appendix numerical values for the elasticity are quoted as positive numbers and denoted by the symbol e.

B.1. The moral dimension

Whereas the valuation of "pure time preference" (i.e. time preference for marginal utility) is essentially ethical, the elasticity of the marginal utility of consumption is usually applied as a wholly utilitarian concept, ignoring any moral aspect of transfers from poorer to richer populations. Yet many might argue that "fairness" or "egalitarian" concerns would justify rather less redistribution of marginal utility (and certainly no more) from poorer to richer than implied by the simple utilitarian criterion of maximising the sum of individual utilities. This would imply for most policy purposes a higher value for e.[22]

B.2. Approaches to deriving the elasticity of marginal utility

B.2.1. Derivation from the personal tax regime

Derivation of the elasticity from personal tax regimes conventionally assumes that the schedule of income tax rates against income is based on the principle of "equal absolute sacrifice" (i.e. an equal loss of utility for each marginal dollar of tax paid). Combining this principle with a constant-elasticity utility function yields an implicit value for e.

Strengths of this approach are its conceptual simplicity and measurability, and that it may also include concern about fairness as well as marginal utility; but it has important limitations.[23] Social concern about *contemporary* inequality might differ from concerns about inequality *over time*; and personal taxation in many countries is influenced by concerns about incentives and personal freedom as well as fairness. Concern about incentives will give a downward bias to the estimation of e, relative to estimates based on the assumption of equal absolute sacrifice. Problems also surround the inclusion or exclusion of standard

---

[22] Dasgupta (2006) argues, in the context of discounting over time, for a value for an elasticity of about 3, incorporating a significant ethical weighting in favour of poorer relative to richer generations. This contrasts with a comment by Newbery (1992, p11) that"[HM Treasury's] preferred value … is 1.5, which is quite egalitarian, and one might quite reasonably defend a value of 1.0 or even less." However, despite the implication in the latter comment that the Treasury value incorporated an ethical judgment, this was not part of the Treasury's logic at that time.

[23] Many other qualifications apply to this and to all other methods of estimating this elasticity. One which applies to many methods is the assumption that the utility function is additively separable.

personal allowance or other tax allowances, and the significance of other policy instruments geared to income distribution. This is discussed by Evans (2005).

Cowell and Gardiner (2000) use this method to derive values for *e* for the UK in the late 1990s of 1.42 if applied just to income tax, and of 1.29 if applied to income tax and National Insurance Contributions (NIC).[24] Evans and Sezer (2005), Evans (2005) and Evans (2007) present results for large number of OECD countries, deriving an average value close to 1.4 with a perhaps surprisingly narrow spread of about $\pm$ 0.2.

The derivation by Stern (1977) for the UK income tax regime in 1973-74, before concerns about tax incentives became so prominent, may give a better measure of social judgments about the utility of marginal income marginal across the income distribution. Stern derived a value of 1.97, although Evans (2005) suggests that Stern's inclusion of the standard personal tax allowance gives a strong upward bias to the elasticity at relatively low levels of income. For US income tax from 1948 to 1965, Mera (1969, p469) found that "for a major portion of the income range" the rates implied a value of 1.5.

B.2.2. Derivation from personal savings behaviour

Many studies of household savings behaviour over the life cycle estimate the inter-temporal elasticity of substitution of household consumption θ. Under certain assumptions, as set out by Cowell and Gardiner (2000, section 4.2.2 and Appendix A3), household relative risk aversion (which is algebraically the same as the elasticity of marginal utility) is equal to $1/\theta$. Cowell and Gardiner consider some of this work, in particular Blundell et al (1994) on UK data. They note that the two principal models in Blundell et al imply, for the elasticity, values of 1.2 to 1.4, or of 0.34 to 1.0 (both sets of values increasing with income). However, as noted by Evans (2005), the sample period of 1970-1986 ends in the year in which UK retail financial markets were deregulated. It is unclear what either model would produce with data for the deregulated environment.

Barsky et al (1997) measure the inter-temporal elasticity of substitution directly by means of survey questions from the US Health and Retirement Survey. They obtain a mean value of 0.18, which implies a very high value of 5.6 for the elasticity of marginal utility.

The savings ratio, together with assumptions for pure time preference and for the long run rate of return, yields an implicit value for the elasticity of marginal utility. Stern (1977, p220) records that, taking plausible figures for the UK in the 1960s, and pure time preference of 2.5%, the implicit value of the elasticity was approximately 5 (or higher for lower values of pure time preference).

These studies suggest however that personal savings behaviour has little relevance in practice to the elasticity of marginal utility.

---

[24] Evans (2005) suggests that NICs should be ignored because the notionally insurance-based rationale for such deductions is "completely different" from that underlying income tax rates. In practice the contributions are widely seen to serve in effect as a (politically convenient) form of income tax, but, as Evans implies, it is hard to believe that the regime was ever designed other than to minimise disincentives for a given total NIC revenue. This however has for many years been an important objective of income tax regimes as well!

B.2.3. Derivation from direct evidence on personal risk aversion

Barsky et al (1997) report empirical measures of relative risk aversion of US respondents between the ages of 51 and 61. The authors give most weight to measures of the reciprocal, which they define as relative risk tolerance, for which the arithmetic average value was 0.24, implying a value for relative risk aversion of 4.2. This high value relates however to very significant risks to income, where factors such as the potential for regret would be expected to increase risk aversion.[25] They found no correlation across individuals between their risk aversion and their *inter-temporal* elasticity as described above.

B.2.4. Derivation from income and price elasticities

An approach with a long history (Fisher, 1927, Frisch, 1932, 1959) estimates the elasticity of marginal utility from the income and price elasticities of a preference-independent good such as food (i.e. a good that contributes an additively separable component to users' utility). Frisch (1959, equation 64) shows that $e = -E_i(1-a_iE_i)/(e_{ii}+a_iE_i)$, where E_i is the income elasticity of demand for the *i*th good, a_i the budget share and e_{ii} the elasticity with respect to its own price.[26]

Brown and Deaton (1972, p1206) report studies by other authors of data from several countries, together with work of their own on UK data for 1900-1970 that gave a value of for the elasticity of 2.8. They conclude that "though estimates obtained this way [from linear expenditure systems] fluctuate considerably and some are very large, an average value of -2 for [the true, negative value of the elasticity] seems consistent both with most such studies and with the results from fitting other models".

Kula (1984) reports values derived in this way for the US of 1.89 and for Canada of 1.56. For the UK, Kula (1985) derives a markedly lower value of 0.71. More recently Evans and Sezer (2002) derive for the UK a value of 1.6. Subsequently Evans (2004.1) has examined alternative specifications, deriving values for the UK of 1.6 (as above) by a CEM (constant elasticities model) and 1.2 by an AIDS (almost ideal demand system); and for France values of 1.8 and 1.3 (Evans 2004.2).

The validity of assumptions such as the constancy of the relevant demand functions over time and income is difficult to assess, and there are problems of data and definitions. The substantial effect of the model specification sounds a note of caution. This approach has nonetheless the merit of being a direct measure of the elasticity and it has been subject to empirical studies in many countries and over different goods.

B.2.5. Derivation from international happiness data

Recent years have seen the development of the a literature, now evolving into several strands, on the measurement and determinants of happiness. Much work has been done on international measurement. Layard et al (2007) use data from six surveys to derive a combined estimate for e of 1.26, with highest and lowest values of 1.34 and 1.19. The results

---

25 They may therefore be more relevant to the equity risk premium.

26 This is the "uncompensated" price elasticity of demand, as opposed to the compensated (or pure substitution) elasticity. In the latter case a price increase is compensated by an income increase, to maintain constant real income.

are similar for subgroups in the population. This method might tend to underestimate the elasticity insofar as the surveys will have measured individuals' preferences for relative consumption, which might decline less rapidly with relative consumption growth than preferences for absolute consumption, while the latter is more relevant for most analytical purposes.

B.2.6. Derivation from intuition

An elasticity of 1 implies that an extra $1 to someone with an income of $x gives twice the extra utility provided by an extra $1 to someone with an income of $2x. Values of 1.5 and 2 would imply factors of respectively 2.8 ($2^{1.5}$) and 4 (2^2). However although the judgement required is conceptually fairly simple, there are few points of reference by which to judge what is plausible.[27]

Scott and Dowley (1977) note the argument, maintained in Scott (1989), and which they report has the support of Little and Mirrlees and of Stern, that "it is reasonable to suppose that there is a maximum level of utility which anyone can derive from income", in which case the elasticity "must exceed one at least above some income level …".

B.3. Expert views

In the American literature there has long been some consensus around values, in the context of social time preference, of about 1.5. For example Eckstein (1958) considers a range of 0.5 to 2.0, and Feldstein (1965) a range of 1 to 2; Cline (1993) opts for 1.5; Boscolo et al (1998, p7) conclude that "the few available estimates suggest that the elasticity of marginal utility [ranges] from 1 to 2"; and Arrow (1995, p 6) suggests, on the basis of "rather thin evidence", 1.5 to 2.0.

In the UK literature, Stern's review of 1977 concludes that the evidence then pointed to the range of 1 to 10, with measurements based on consumer behaviour pointing to the middle of the range, and those based on government behaviour to around 2. Scott (1977, 1989), working back from market rates, estimates a value of 1.5. Little and Mirrlees (1974, p 240) suggest that "on admittedly extremely inadequate evidence, we guess that most people would put [the elasticity] in the range 1-3". Cowell and Gardiner (2000) conclude that the evidence supports a value in the range of 0.5 to 4, within which they give most weight to the range of 1.2-1.4 derived, as explained above, from the UK personal tax regime of the late 1990s. Evans (2005) regards a figure of 1.4, derived from the personal tax regimes of a large number of countries, and not inconsistent with derivations from food income and price elasticities, as plausible for many countries.

B.4. Conclusion

Of the several methods adopted for valuing the extent to which utility changes with marginal income, some seem to be too far removed from this aspect of personal preferences to be of much practical value. Thus the study of personal savings behaviour and direct evidence on personal risk aversion (at least from the limited data available) appear to offer little promise.

---

27 Although Stern's (1977) comment (p243) that a value "of around 5 does not seem ludicrously large" might for most applications be questioned.

And derivation from intuition – surprisingly given the simple way in which the issue can be formulated – appears to be too imprecise, beyond perhaps explaining the plausibility of a utility function that approaches a finite maximum utility. However derivation from tax structures, from income and price elasticities and from happiness data, all have strong features and give broadly consistent results. And the opinions of experts who have lived with and discussed the issue through distinguished careers deserve some weight. Given that there seems rather more reason to expect some downward rather than upward bias in the three most promising approaches, and the apparent sharing of this view among most leading experts, a value of 1.5, or perhaps a little higher, might be the most easily defensible.

References for Appendix B

Arrow, K. J. (1995) 'Discounting climate change planning for an uncertain future', Lecture given at Institut d'Economie Industrielles, Université des Sciences Sociales, Toulouse, 24 April

Barsky, R. B., F. T. Juster, M. S. Kimball. and M. D. Shapiro. (1997) 'Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Study', *The Quarterly Journal of Economics*, 112.2 (In memory of Amos Tversky), 537-579

Blundell, R., M. Browning and C. Meghir (1994) 'Consumer Demand and the Life-Cycle Allocation of Expenditure', *Review of Economic Studies*, 61, 57-80

Boscolo, M., R. V. Vincent and T. Panayotou (1998), "Discounting costs and benefits in carbon sequestration projects", *Development Discussion Paper* No. 638, Harvard Institute for International Development, Harvard University

Brown, J. A. C. and A. Deaton (1972) 'Models of Consumer Behaviour', *Economic Journal*, 82 (328), 1145-1236

Cline, W. R. (1993) 'Give Greenhouse Abatement a Fair Chance', *Finance and Development*, 30 (1)

Cowell, F A and K. Gardiner. (2000) 'Welfare Weights', *Economic Research Paper* 20 (oft282.pdf), London, Office of Fair Trading

Dasgupta, P. (2006) 'Comments on the Stern Review's Economics of Climate Change', comments for a seminar organised by the Foundation for Science and Technology at the Royal Society, London, on November 8, at http://www.econ.cam.ac.uk/faculty/dasgupta/STERN.pdf

Eckstein, Otto (1958) '*Water Resource Development: The Economics of Project Evaluation'*, Cambridge MA, Harvard University Press

Evans, D. J. (2004.1) 'The elevated status of the elasticity of marginal utility of consumption', *Applied Economics Letters*, 11 (7), 443 - 447

Evans, D. J. (2004.2) 'A social discount rate for France', *Applied Economics Letters*, 11 (13), 803-808,

Evans, D. J. (2005) 'The Elasticity of Marginal Utility of Consumption for 20 OECD Countries', *Fiscal Studies*, 26 (2), 197-224

Evans, D. (2007) 'Social discount rates for the European Union: New estimates', Chapter 12 in Florio, M. (ed.) (2007)

Evans, D. and Sezer, H. (2002) *A time preference measure of the social discount rate for the UK*, Applied Economics, 34, 1925-1934

Evans, D. J. and Sezer, H. (2005) 'Social Discount Rates for member countries of the European Union', *Journal of Economic Studies*, 32 (1), 47-59

Feldstein, M. S. (1965) '*The Derivation of Social Time Preference Rates'*, Kyklos, XVIII, (2), 277-287

Fisher, I. (1927) 'A statistical method for measuring marginal utility of justice of a progressive income tax', in Jacob H. Hollander (ed.) *Economic Essays Contributed in Honor of J Bates Clark*, New York, Macmillan, 157-193

Frisch, R. (1932) '*The New Methods of Measuring Marginal Utility'*, Tübingen, Mohr

Frisch, R. (1959) 'A Complete Scheme for Computing all Direct and Cross Demand Elasticities in a Model with Many Sectors', *Econometrica*, 27 (2), 177-196,

Kula, E. (1984) 'Derivation of Social Time Preference Rates for the United States and Canada'*, Quarterly Journal of Economics*, 99, 4, 873-82

Kula, E. (1985) 'An empirical investigation on the social time-preference rate for the United Kingdom'*, Environment and Planning A*, 17, 199-212

Layard, R., Mayraz, G. and Nickell, S. J. (2007) (revised July 2008) 'The Marginal Utility of Income', London School of Economics, *Centre for Economic Performance Discussion Paper 784*

Little, Ian M. D. and James A. Mirrlees (1974) '*Project appraisal and planning for developing countries'*, London, Heinemann Educational Books Ltd

Mera, K. (1969) 'Experimental Determination of Relative Marginal Utilities, *Quarterly Journal of Economics*, 83 (3), 464-477

Newbery, D. M. (1992) *Long term Discount Rates for the Forest Enterpris*e, Department of Applied Economics, Cambridge University, for the UK Forestry Commission, Edinburgh

Scott, M. FG. and Dowley, M. (1977) 'The Test Rate of Discount and Changes in Base Level Income in the United Kingdom', *Economic Journal*, 87, 219-241

Scott, Maurice FG. (1989) '*A New View of Economic Growth'*, Oxford, Oxford University Press

Stern, Nicholas (1977) 'The Marginal Valuation of Income', in Artis, Michael J. and A. Robert Nobay (1977) *Studies in Modern Economic Analysis*, Oxford, Basil Blackwell, 209-258

Appendix C. Some economics of risk

C.1. Corporations and government

C.1.1.Systematic and non-systematic risk

"Risk" in economics and finance literature, as generally in this section (but often not elsewhere in this Appendix or the man text of this Report) almost invariably means variability. The literature assumes that measures are taken to avoid bias in estimates of expected values, but that individuals may be averse to uncertainly about how outcomes will be distributed around these values.

The analysis of risk in this sense in economics and, especially, finance has changed beyond recognition over the past half century. Two sets of conventions are now firmly established, one relating to privately financed investment and the other mainly to public service benefits.

Since changes in utility are enjoyed by individuals, or households, all this analysis relates to the utility of investors, or the beneficiaries of public expenditure, or taxpayers, and the implications of these personal utilities for commercial or government bodies. The concept of an institution having a utility function is not generally helpful, although, as discussed in section C.1.2 below, it is sometimes rational for managers to be more than proportionately averse to large downside risks (in the more colloquial sense of the word risk).

Of the conventions developed from the 1960s the best known and most widely applied are those surrounding the capital asset pricing model (CAPM) to assess the cost of risk to equity investors and hence, to companies, of equity financed investment. In this model, which has stood the test of time, random (or 'non-systematic') risks (in the sense if variability about expected outturns) are for practical purposes costless, because investors can diversify their investments; but investors are averse to, and so need compensation for, expected variations of equity yields that are correlated with those of the market as a whole (i.e. 'systematic risks'). Thus the cost of equity for an investment will be equal to the risk-free rate (typically the government borrowing rate) plus a quantity denoted by b times the market average risk premium. b measures the covariance of the expected equity returns from the investment with the expected fluctuations in the equity market as a whole (generally the equity market).[28]

There is no correspondingly sophisticated model for the cost of private sector debt. It is generally assumed that the risk premium on such debt reflects the creditors' perception of the risk that the debt may not be fully repaid.

It is sometimes argued by financial economists that the equity risk premium reveals a cost that is not a function of equity markets but is inherent in the activity itself, so that the government should include such a cost in its investment appraisals. This is not however an argument accepted by welfare economists, or by the UK Treasury. The only analogy to equity market risk with public funding is the covariance of the costs or benefits with the fluctuations in the incomes of beneficiaries or taxpayers. In most cases the relevant

---

[28] Specifically $b = \text{cov}(r_a, r_p)$ / var r_p, where r_a is the expected return to the financial asset and r_p the expected return to a diversified portfolio, which is generally taken as the equity market average.

fluctuation is that of GDP, but, as compared with the equity market, GDP fluctuations are so small that the covariance with project costs or benefits could not justify a cost of capital premium more than a very small fraction of a percentage point.

Systematic risk is thus very rarely significant in public policy. An example sometimes quoted of where it may be material is that of an overseas aid project, where the benefits to a poor community from, say, a well, or better fertiliser or crop choice may be correlated with the cycle of famine that the community would otherwise face. The amount by which the expected monetary value of a varying cost or benefit to the individuals affected should be adjusted, to obtain its 'certainty equivalent' value, is given by e cov (X, Y) / E(Y), where Y is income, X is the project cost or benefit, and e, as in the J-value literature, is $-YU''/U'$, where $U=U(Y)$. This could help to quantify, for example, the extent to which, for a given average annual monetary benefit, a well that protected a community from periods of drought might provide more welfare than a better fertiliser that was most effective at times of relative plenty.

Thus finance textbooks and government guidance advise that, where costs and benefits can be expressed in monetary terms and risks (in the sense of variability) are not systematic, the objective should normally be to maximise the expected value.[29] The UK Treasury Green Book advice on the downside risks of overspends and overruns on capital projects is that objective estimates should be made of expected outturns.

C.1.2. The appraisal of major downside risks

A private sector company is placed differently from government in that it is plausible that it will sometimes be faced with potential risks that, if they materialised, could destroy the business.[30] Thus a medium sized business may forgo a business opportunity of a very large project that offered a very high expected value, but could bankrupt the company if some plausible risks materialised. Company directors may therefore regard increasing risks (in the sense here of impact times probability) beyond a certain point as more than proportionately unacceptable.

This is introducing a form of "risk aversion" for increases in extreme downside risks. It is unlikely that a company board would ever wish to build this managerial judgement into an algebraic formula (which would certainly not approximate to a constant elasticity marginal utility). A company board might however wish to see the potential outturns of a risky capital project presented in terms of the probability of the cost overrun exceeding some maximum acceptable level. Such an approach is illustrated, anomalously, in a report commissioned by the Department for Transport in 2004 and written apparently on the back of the authors' advice to private sector companies.[31]

---

29 This is described in economics terminology as risk-neutrality. The classic reference with respect to government is Arrow K. J. and Lind, R. C. (1970) *Uncertainty and the Evaluation of Public Investment Decisions,* American Economic Review, Vol. 60, 364-378. Some financial economists still quibble with its dismissal of systematic risk as a trivial in the government context, but the position of risk neutrality with respect to non-systematic risk is uncontentious.

30 As so starkly illustrated by Northern Rock, RBS, Lehman, AIG and other financial institutions, which, with their regulators, failed to adequately identify and quantify the risks they faced.

31 Flyvbjerg, B. in association with COWI (2004) *Procedures for dealing with optimism bias in transport planning,* Available, as issued by the consultancy group COWI, on the Treasury Green Book website. Flyvbjerg is known mainly for his lead authorship of an outstanding book on the history and causes of overspends on very large projects:

The government of a developed country does not in practice face such extreme financial risks from its public service expenditures. Even figures such as those considered in the J-curve literature, while serious and painful, do not threaten the state, nor any single part of it. Very large losses tend to attract attention that is painful for those responsible, and may be politically embarrassing and in some circumstances socially damaging. However the link between such impacts and the magnitude of the loss is very case specific and may, beyond a certain level, have little relationship with the magnitude of the loss. The media impact of say an NHS IT project failure may depend little upon whether the financial cost is estimated to be £200 million or £1 billion.

The principle of risk-neutrality in government appraisal is thus robust, with two important qualifications. One qualification is that the estimation of the expected outturns needs to be unbiased, with respect to issues such as the likelihood of capital project overspends (or failures).[32] The other is that assessment of the estimates must include all important potential impacts, including effects such as those that may arise from a major nuclear accident – even though some of these impacts do not lend themselves to monetisation. But this is a wider field, touched on briefly in the Conclusions in the main text of this Report.

C.2. Personal risk aversion and aide memoire for the 'J-value' analysis of risk

This section is provided mainly as an aide memoire for any reader of this Report not familiar with the conventional forms of personal utility functions, or with how a function that ascribes less utility to an extra pound as income increases implies that the individual is averse to uncertainty of income. It provides some illustrative examples of this effect and, in sub sections C.2.4 and C.2.5, illustrates the derivation of some of the numbers in Thomas, Jones and Boyle (2009).

C.2.1. Personal attitudes to risk

Personal attitudes to risk are extremely complex and have for forty years be subject to extensive academic study. It is nowadays an area that bridges economics and psychology and is studied by both disciplines, often working together.

It is rare for personal preferences to be well represented by anything so simple as a utility function with a constant elasticity of marginal utility, $-e$.[33] Such 'constant-elasticity' functions are however useful in cases where the issue is that of marginal changes in wealth to people with different incomes, where they are not directly aware of the change. This applies for example to general policy analysis on income distribution or of the weight that should be

---

Flyvbjerg,B., Bruzelius, N, and Rothengatte, W. (2003) *Megaprojects and Risk: An Anatomy of Ambition,* Cambridge University Press.

[32] Current Treasury guidance (repeated in Ministry of Defence, 2006), suggests some default factors, based on experience, to apply to initial cost estimates. But major health and safety risks such as those addressed in the J-value literature are of course more case specific.

[33] This does not for example apply to investors' aversion to the risks of equity investment, where the market premium is higher than can be explained by any plausible value of *e*. This is perhaps unsurprising given the nature of equity market fluctuations, which are sometimes far removed from a normal distribution.

given to marginal costs or benefits to richer future generations relative to a poorer current generation.

In such cases it is assumed that the utility enjoyed by individuals from their consumption (or income), increases with income, at a rate that continually declines, where the rate of decline follows a constant "elasticity of marginal utility with respect to consumption". If utility U and consumption, y, are related by $U = U(y)$, marginal utility is given by U' (which is positive), the rate of change of this marginal utility with consumption is U'' (which is negative). The elasticity of marginal utility is therefore given by yU''/U'.

As noted in Appendix B above this negative quantity is by convention usually described by its modulus (i.e. a positive number). This is often, as in the J-value literature, denoted by e, and is also equal to the 'coefficient of relative risk aversion'.

C.2.2. The constant-elasticity utility function

The following applies for any positive value of e.

$$U(y) = C_1\{y^{1-e} - 1)/(1-e)\} + C_2$$

C_1 and C_2 are arbitrary constants, for which the values taken or implied in the City University papers are, for $e < 1$, $C_1 = 1–e$ and $C_2 = 1$, so that $U = y^{1-e}$; and for $e > 1$, $C_1 = 1$ and $C_2 = 0$.[34]

As e increases, the function $(y^{1-e} - 1)/(1-e)$ changes at $e = 1$ from a function that increases indefinitely with y to one that is asymptotic to a finite ceiling. As e approaches 1 both numerator and denominator approach zero and the function approaches the natural logarithm $\ln y$.[35] It is thus usual to say that, for $e = 1$, $U(y) = \ln y$.

Such functions are unusual in not defining a unique scale or zero. Thus the numerical values for utility produced by the function for different values of e cannot be compared with each other, and for a single value of e it is only relative differences that have any meaning. The value of the ceiling to which U is asymptotic when $e > 1$ has no welfare significance.[36] Nor does the fact that $\ln y$ and $(y^{1-e} - 1)/(1-e)$ are negative when $y < 1$.

C.2.3. The effect of a concave utility function on the cost of risk

The following stylised diagram illustrates a constant-elasticity utility function, where utility, U, is a function of consumption Y. Such a function, with a declining gradient as consumption

---

[34] The form of $U = y^{1-e}/(1-e)$ is often attributed to Ramsey. The form with -1 in the numerator is conveniently positive for all $y > 1$ and is sometimes attributed to Atkinson, who used it in presenting the "Atkinson measure of inequality", used widely in the analysis of income distribution. Thomas et al describe this form as the Atkinson measure of utility.

[35] This perhaps surprising algebra follows from the standard expansions $(1+x)^n = 1 + nx + n(n-1)x^2/2! + n(n-1)(n-2)x^3/3! + \ldots$ and $\ln(1+x) = x - x^2/2 + x^3/3 - x^4/4 + \ldots$. Substitution, into the function $(y^{1-e} - 1)/(1-e)$, of $(1+x)$ for y and of n for $(1-e)$ and expanding $(1+x)^n$ gives $\{(1+x)^n - 1\}/n = x + (n-1)x^2/2! + (n-1)(n-2)x^3/3! + (n-1)(n-2)(n-3)x^4/4! + \ldots$. As $e \to 1$, $n \to 0$, and this series approaches the series for $\ln(1+x) = \ln y$, so that $U \to \ln y$. For practical calculations any logarithmic base can be used, since $\log_a y = \ln y / \ln a$, and multiplying a utility function of this kind by a constant (in this case $1/\ln a$) has no effect.

[36] For the Ramsey form of $U = y^{1-e}/(1-e)$ this ceiling is zero and U is therefore negative for ***all*** finite values of y. For the form with -1 in the numerator the ceiling is $1/(e - 1)$.

increases, is generally described as "concave". The rate of decline in the gradient is generally defined by *e* as explained above. If *e* = 0 there is no decline in the gradient, so ***U*** is a linear function of (and could be set numerically equal to) ***Y***.

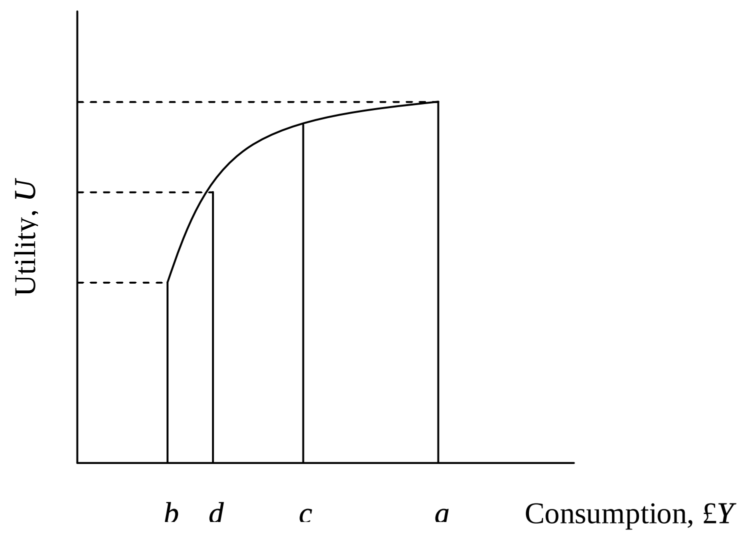

Suppose that this function describes the utility of someone enjoying consumption ***a***, who faces a 50 per cent chance of a loss that would reduce consumption ***Y*** from ***a*** to ***b***. If the loss occurred it would be equivalent in ***utility*** terms to a certain loss of 50 per cent of the ***utility*** loss in moving from ***Y*** = ***a*** to ***Y*** = ***b***. If utility increased linearly (i.e. *e* = 0), that utility loss would also be equivalent in monetary terms to a ***monetary*** loss of 50 per cent of the possible money loss – that is to ***ac***. However, with a concave utility curve as shown, that 50 per cent utility loss would be equivalent to a greater monetary loss ***ad***.

In practice, even if the preferences of people were well described in this way, the magnitudes are typically less than implied by this stylised diagram.

To take a still stylised example, suppose that someone has an annual income (or say minimal income and wealth of) £100,000 and faces the annual (or one off) 1 per cent probability of a loss of £10,000. Then the expected monetary loss (or the ***risk neutral*** premium that they would be willing to pay to insure against this risk) would be 1 per cent of £10,000 – i.e. £100. But suppose that they act to maximise their utility, with a utility function of the simple form *U*=log*Y* (i.e. *e* = 1). It can be shown as follows that the premium they would then be willing to pay would be £105, or in other words that the concave utility curve generates a "risk multiplier" of 1.05.

With *e* = 1 we can set $U=\log_{10}Y$

Thus *U*(100,000) = 5; and *U*(90,000) = 4.954,243. The difference = utility loss if the risk materialises = 0.045,757

Expected ***utility*** loss = 1% of this difference = 0.000,457,57

Thus expected utility, with this expected utility loss = 5 – 0.000,457,57 = 4.999,542,43; and the monetary value of this reduced utility = $10^{4.99954243}$ = £99,894.70

So "risk averse" annual "premium" = £100,000 - £99,894.70 = £105.30

C.2.4. Illustration of some numbers from Thomas, Jones and Boyle (2009)

C.2.4.1. The Risk multiplier

The following table is extends the risk multiplier calculation to a wider range of values for *e* and is applied to the example set out in Table 1 of Thomas, Jones and Boyle (2009). We here ignore for simplicity the residual risk after the safety measure has been installed. And the monetary figures are scaled down by a factor of 10^3 to bring the example closer to credible figures for a person or household. The assumptions are a utility function $U(A)$, with a constant elasticity of marginal utility of $-e$, where A is wealth and the following conditions apply.

Initial assets = £A_0 = £10,000,000

Cost of accident if it occurs = £5,000,000

Thus value of assets after accident = £5,000,000

Cost of safety measure = £1,000

Thus value of assets with safety measure = £A_2 = £9,999,999

Probability of accident without safety measure = 10^{-3}

Thus, **when** e **= 0**, value of assets without safety measure= £A_1 = £9,995,000

Given the strong assumption about the utility function, the maximum amount that it would be worth paying to eliminate this risk is as in column (3). The arithmetic calculation of U for a given A is of course different for the different forms of the utility function for e <1 and e >1. But the sequential steps are exactly the same as those in section C.2.3 above.

(1)	(2)	(3)	(4)
e	Form of utility function	Money value of risk	Multiplier (col(3)/5000)
0	$U = A$	£5,000	1.00
0.5	$U = A^{0.5}$	£5,857	1.17
1.0	$U = \log A$	£6,929	1.39
1.5	$U = (1 - 1/A^{0.5})$	£8,279	1.66
2.0	$U = (1 - 1/A)$	£9,980	2.00

The multipliers in column (4) correspond (except for the trivial effect of their omitting the residual risk after the safety measure has been installed) with those described as exact in Thomas, Jones and Boyle Figure 7.

C.2.5. The "Reluctance to invest"

As discussed in section 2.2.3 of the main text, Thomas, Jones and Boyle (2009) develop the concept of "Reluctance to invest", which they define as $(u_1 - u_2)/u_o$, where u_1 and u_2 are utility with and without the safety protection and, and u_0 is the "absolute utility" of the initial assets. This is a flawed concept because "absolute utility", in utility functions of this kind, has no meaning. However the ratio as shown by Thomas, Jones and Boyle approaches and is all but equal to zero for values of *e* more than about 1.2.

C.2.5.1. Derivation of the "reluctance to invest"

An illustration can be provided as follows of how "Reluctance to invest", as defined by Thomas, Jones and Boyle, changes with *e* when $e < 1$.

In this range, the utility *u* as defined by Thomas, Jones and Boyle can be set equal to $A^{(1-e)}$, where *A* is the monetary value of the organisation's assets. Thus, given that the cost of the safety measure (A_0-A_2) is very small relative to A_0: $(u_1 - u_2)/u_o = (A_1^{(1-\varepsilon)} - A_2^{(1-\varepsilon)})/A_0^{(1-\varepsilon)} \approx (A_1^{(1-\varepsilon)} - A_2^{(1-\varepsilon)})/A_2^{(1-\varepsilon)} = (A_1/A_2)^{(1-\varepsilon)} - 1 = \{1 - (A_2 - A_1)/A_2)\}^{(1-\varepsilon)} - 1$

It follows that, since $(A_2 - A_1)/A_2 >>1$, $(u_1 - u_2)/u_o \approx - (1 - e)\{(A_2 - A_1)/A_2\}$.

The ratio $(A_2 - A_1)/A_2$ for a given value of *e* can be derived from the value of this ratio when $e = 0$, together with the risk multiplier M_R. The monetary value of the risk when $e = 0$ is £5,000 and the cost of the safety measure for all values of *e* is £1,000. Thus, for this case, "Reluctance to invest" $\approx - (1 - e)(5M_R - 1) \times 10^{-4}$. Using the values of M_R from Thomas, Jones and Boyle Figure 7 this gives, to a close approximation, the figures for "Reluctance to invest" plotted in Thomas, Jones and Boyle Figure 8.

C.2.5.2. An illustration of the "point of indiscriminate decision"

Looking over a wider range of *e*, the table below illustrates the effects of increasing values of *e* on the constant-elasticity utility curve as used by Thomas, Jones and Boyle.

Rows 1 to 4 record figures for an example in which the values of *A* are numerically very high (of the order 10^7, which is however still three orders of magnitude lower than the values in the City University calculations – assuming that these take the numeraire for *A* as £1). At these numerically very high values the gradient of the utility curve, for values of *e* greater than 1, is extremely small, as it is asymptotic to an upper bound. This does not affect the validity of the curve if it is used conventionally, for comparing differences in utility values. But for such high numerical values, as *e* increases, the numerical value of the utility loss for a 50 per cent loss of wealth, shown in column (5), decreases very dramatically.

The numerical value of the "absolute utility", as shown in columns (3) and (4), also decreases, but less dramatically. And the "reluctance to invest", defined as the ratio of the utility loss to "absolute utility", falls to virtually zero at values for*e* of not much more than 1. Column (8) reproduces, to a rough approximation, the figures plotted in Thomas, Jones and Boyle Figure 9.

However if the numeraire of *A* is changed from £1 to say £1 million, as in rows 5 to 9, there are no such dramatic changes in columns (3) or (6). Columns (6) to (8), with "absolute utility" as a denominator, have no useful meaning whatever the numeraire. But rows 5 to 9 illustrate how the so called "point of indiscriminate decision", previously at some value of *e* a little over 1, fades away when the numerical values *A* are scaled down by the use of a larger monetary numeraire.

					(3) - (4)	(5) / (3)		- (7) x $4x10^4$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	e	Form of utility function	$U(A_0)$	$U(A_1)$	$U(A_0)-U(A_1)$	$\frac{U(A_0)-U(A_1)}{U(A_0)}$	(6) relative to value for *e = 0*	approximate "reluctance to invest"
Values if numeraire for *A* is taken as £1								
1	0	$U = A$	10,000,000	5,000,000	5,000,000	$500x10^{-3}$	1	- 4 x 10^{-4}
2	0.5	$U = A^{0.5}$	3,162	2,236	926	$293x10^{-3}$	0.586	- 2.34 x 10^{-4}
3	1.0	$U = \log_e A$	16.118	15.425	0.693	$43.0x10^{-3}$	0.086,0	- 0.34 x 10^{-4}
4	1.5	$U = 1 - 1/A^{0.5}$	0.999,683,8	0.999,552,8	0.000,1310	$0.131x10^{-3}$	0.000,262	- 0.001,0 x 10^{-4}
Values if numeraire for *A* is taken as £1 million								
5	0	$U = A$	10	5	5	0.5	1	- 4 x 10^{-4}
6	0.5	$U = A^{0.5}$	3.162,278	2.236,068	0.926,210	0.293	0.586	- 2.34 x 10^{-4}
7	1.0	$U = \log_e A$	2.302,585	1.609,438	0.693,147	0.301	0.602	- 2.41 x 10^{-4}
8	1.5	$U = 1 - 1/A^{0.5}$	0.683,772	0.552,786	0.130,986	0.192	0.383	- 1.53 x 10^{-4}
9	2.0	$U = 1 - 1/A$	0.9	0.8	0.1	0.111	0.222	- 0.89 x 10^{-4}

The changes with increasing *e* in columns (6) to (8) in rows 5 to 9 are not quite monotonic, as a consequence of the intersection of the utility curve with the horizontal income or assets axis flipping from the origin to $A = 1$ when *e* increases to unity. But there is no special significance in an absolute utility value of zero. It does not mean "zero welfare". It is only differences in utility that have meaning, whether they are between positive or negative absolute values. A zero value for $U(A_0)$ would of course take the "reluctance to invest" to $\pm \infty$.

Appendix D. Societal concerns

D.1. History and definition of societal concerns

The original presentation by the HSE of the "Tolerability of Risk" (TOR) concept, developed in the context of the Sizewell B Public Inquiry, included a picture of "individual risk" and "societal risk". The HSE later agreed to change this wording from "societal risk" to the more satisfactory "societal concerns".[37] Societal concerns are however defined in the HSE's R2P2 in two inconsistent ways.

One definition is that they are everything other than "individual concerns" (i.e. the concerns of those potentially at risk). This is not altogether satisfactory in that it includes some uncontroversial and relatively easily costed impacts such as ambulance costs and loss of net output. These are social costs, and sometimes referred to as societal effects, but are not really well described as "societal concerns".

The other definition is slightly garbled, as below, but defines societal concerns more explicitly as things that undermine trust in government institutions. The definition is honest in describing institutional concerns rather than considered public preferences, but some might question the objectives that it implies for health and safety regulatory regimes.

Societal concerns [are] the risks or threats from hazards which impact on society and which, if realised, could have adverse repercussions for the institutions responsible for putting in place the provisions and arrangements for protecting people, e.g. Parliament or the Government of the day (R2P2, paragraph 25). The text continues by noting that "*This type of concern is often associated with hazards that give rise to risks which, were they to materialise, could provoke a socio-political response, e.g. risk of events causing widespread or large scale detriment or the occurrence of multiple fatalities in a single event. Typical examples relate to nuclear power generation, railway travel, or the genetic modification of organisms.*

Societal concerns in this second sense – i.e. factors that bring political problems for regulators or governments – include those issues identified in the psychology literature, notably the work of Slovic and others in the 1970s and 1980s, on attitudes towards different types of hazard characteristic and on media amplification.

An excellent publication in the late 1990s by the Department of Health, generally known as "Pointers",[38] summarises that literature well, with illuminating tables of "Fright factors" and "Media triggers", in the context of "risk communication". The issues relate essentially to public indignation, or anxiety, or other emotions such as prurience, and are central to the HSE's concerns expressed in its definition of societal concerns as quoted above.

There are solid arguments, technical as well as political, for giving weight to the political consequences of specific types of incident. However the relationship between, on the one

---

[37] With the term societal risk retained for the allocation of more than proportionate weight to hazards that may lead to multiple fatality accidents, relative to those leading to smaller numbers of deaths in single incidents.

[38] Department of Health (1997) *Communicating about risks to public health: pointers to good practice,* http://www.dh.gov.uk/en/Publicationsandstatistics/Publications/PublicationsPolicyAndGuidance/DH_4006604 It was written by the OR Scientist Dr Peter Bennett.

hand, these political consequences and, on the other hand, considered public preferences and willingness-to-pay for risk reductions may be weak. A definition of societal concerns that might appeal more to analysts would relate also to people's preferences about how resources should be allocated to reduce their own and others' exposure to injury and health risks.

A substantial, three part study, commissioned by the HSE and published in 2007, provides a comprehensive and still relevant picture of where societal concerns stood in the mid 2000s.[39]

The first part examined attitudes to alternative causes of instant (or near-instant) death, such as road or rail accidents. Individual attitudes were viewed primarily from the perspective of people's 'self-focused' preferences concerning personal safety.

The second part considered – among other issues – causes of death typically preceded by protracted periods of pain and discomfort, and also investigated the public's attitudes to factors such as the victim's age and the question of blame or responsibility for the cause of death concerned. The focus here was directed more towards people's preferences as citizens, expressing their views on and attitudes towards general principles of social decision-making about life-saving interventions.

The third part reviewed literature of the early 2000s on societal concerns from a range of disciplines.

D.2. Impact of societal concerns

Societal concerns, however defined, often drive policy decisions away from the choices implied by the preferences of those at risk, and quite probably away from people's considered preferences as citizens. A few well known UK examples are as follows, but there appears to have been no general review across all sectors to collect evidence on actual practice. This enhances the value of studies such as the first two J-value papers – even if they might have been more usefully presented in terms of the implied values of a prevented fatality.

The HSE regulation of 'major hazards' puts more than proportionately extra weight on hazards that present risks of large numbers of deaths from a single incident, relative to hazards presenting the same total risks of deaths but in smaller incidents. This presumably reflects the more than proportionally greater political impact of larger incidents.

Nuclear risks to the public are given a very high weight, although this may be of little practical consequence, as explicit valuation of fatality risks in a cost-benefit analysis (CBA) usually has little to add to engineering analysis and other design judgements required in the management of very low risk, high consequence hazards.

In presenting the publication of R2P2, the HSE held an event that presented a few case studies that illustrated political policy preferences. One example was measures that implied exceptionally high values of a prevented fatality (VPF) for gas main hazards near residential property.

---

[39] Chilton, S., A. Covey, Jones-Lee, M., Loomes, G., Metcalf, H., Robinson, J., Spencer, A and Spackman, M (2007) *Valuation of health and safety benefits: Dread risks,* HSE Books www.hse.gov.uk/research/rrpdf/rr541.pdf

The HSE's long standing convention of applying an arbitrary factor of two to risks of death from cancer is being re-examined. Its basis is unclear, although it may reflect any or all of several factors. For example policymakers may believe that those at risk would be willing to pay more to reduce that risk relative to other fatality risks because it is conventionally a "dread" risk; or they may believe that the government might face exceptional criticism if deaths of that kind attracted media attention; or they may be making a broad brush adjustment for the fact that cancer deaths are often preceded by a period of illness (although on the other hand death itself is often long after the carcinogenic exposure, which should reduce its valuation at the time of exposure).

The road / rail contrast in transport is well known. For many years a higher VPF was applied formally to rail safety than to road safety. In the light of evidence that this did not reflect the preferences of those at risk this was changed in 2003 and the same values are now formally applied to both transport modes. However in practice the funds allocated to road safety do not allow this to be applied. The marginal VPF for roads implied by the allocated funds was in the early 2000s, and probably still is, about one tenth of the published value, whereas in rail safely it appears to be more than the published value.[40]

Another example of *low* safety standards in an area that, despite its even being a cancer risk, does not attract media interest, is radon in residential premises. Significant risks of death from cancer induced by radon in insufficiently protected buildings attract virtually no media attention, relative to many other risks of comparable magnitude, perhaps because radon is seen as 'natural', with no obvious institution to blame. But it would be surprising if people were less willing to pay to reduce the risks of premature death from this cause than from other carcinogens. Regulations and advice have in recent years been given a higher profile, but radiation risks to the public officially accepted in this context are far higher than those acceptable in other contexts.[41]

D.3. Potential for progress

D.3.1.The disaggregation of societal concerns

The term societal concerns is generally used by regulators to describe any factor that is not incorporated in empirical measures of people's WTP to reduce risks to the individual, but which the regulator may wish to consider. However there is no clear and generally accepted definition of what these factors are.

---

[40] The relevant empirical study and commentary on several of the wider issues in this note are well covered in Evans, A.W. (2005) ***Railway risks, safety values and safety costs***, Proceedings of the Institution of Civil Engineers, Transport 158, February 2005, Issue TR1, pages 3–9 http://www.thomastelford.com/journals/DocumentLibrary/TRAN1580102.pdf.PDF This (difficult) quantitative empirical work illustrates the difference between reality and the positions declared and probably often believed within the institutions concerned. A recent consultancy study for the NAO reports that local authorities stress the ***high*** priority given to road safety schemes. It is unsurprising that no official would publicly volunteer that they were underfunded.

[41] This effect may be enhanced by the extraordinarily widely dispersed and changing responsibilities for this low profile policy area. In England, building standards rest with CLG. Workplace issues, including schools, rest with the HSE, but school building policy rests with DCSF. The Department of Health leads on Health Protection issues and acts as the central focus for other departments on radon gas. In 2005, responsibility for the domestic radon roll-out programme in England was passed from DEFRA through DH to the Health Protection Agency. It is also a very low or negligible hazard in London and most of SE England.

Some progress on this was made in work published by the Railway Safety and Standards Board (RSSB) in 2006.[42] Subsequently the RSSB has established for itself a much clearer position, as follows, in its own guidance.[43]

… societal concern should not be taken into account by duty holders when deciding whether a measure is necessary to ensure safety so far as is reasonably practicable (SFAIRP). However, the impact of societal concern on a company's reputation might mean that the company takes account of it optionally for business reasons.
Societal concern refers to the concern and anxiety that the public feels about different types of risk. This concern might not reflect the true level of risk, is influenced by dread and other subjective or emotive feelings, and might change considerably following the occurrence of an accident. Societal concern about risk can result in pressure on the railway that is disproportionate to any objective evaluation of risk. It is not taken into account in the industry's determination of whether measures are reasonably practicable, although it can impact a company's profitability and performance, and therefore may be a factor in a commercial judgement. (Taking Safe Decisions, pp 3 and 6)

This removes societal concerns entirely from VPF (or VPI) and cost-benefit analysis, but defines it as a reputational factor to be considered at a later stage in the analysis. It is concerned with the same factors (insofar as they relate to railway safety) as those covered in the Department of Health publication "Pointers" cited above.

This is an advance on HSE's R2P2, but it is by no means clearly sufficient to place *all* the factors typically covered by the term societal concerns outside CBA. The case for more formal disaggregation of societal concerns was developed in that RSSB work, and carried a little further in a presentation to an HSE "Dread Risk Seminar" in September 2006. Section D.3.3 below provides a very simple outline of the structure that was presented at that seminar.

D.3.2.Personal considered preferences versus public outcry

The earlier psychology literature summarised in the Department of Health "Pointers" publication, as noted in section D.1 above, and indeed societal concerns as defined above by the RSSB, are concerned mainly with people's gut reactions. This also drives media coverage, with the consequent amplification of these concerns. From the perspective of corporate and personal reputations these aspects are important. They are bound to affect corporate decisions.

These attitudes may not however reflect people's considered preferences about the trade offs between expenditure and risk. Empirical evidence tends to suggest that people's aversion to premature death is generally not very dependent on the mode of death. Thus railway accidents tick more of the boxes of societal concerns in the RSSB sense than deaths from road accidents. They pose in particular much greater institutional risks. But focus group

---

42 Rail Safety & Standards Board (2006) *The Definition of VPF and the Impact of Societal Concerns*, Research Project T430, RSSB
http://www.rssb.co.uk/pdf/reports/research/t430 assessment of the value for preventing a fatality phase 1.pdf
Other important contributions at this time are reported under Project T571.
http://www.rssb.co.uk/pdf/reports/research/T517%20Modelling%20Societal%20Concerns.pdf
http://www.rssb.co.uk/pdf/reports/research/T517%20Anatomy%20of%20Risk%20model%20paper.pdf

43 Rail Safety & Standards Board (2008) *Taking Safe Decisions*, RSSB.

work suggests that there is little difference between individuals' willingness to pay to reduce risks from these two sources.

More empirical data will no doubt accumulate over time on how people's WTP varies for similar risk reductions from different hazards or for different potential victims (people's views as potential victims as opposed to citizens being another distinction, which in empirical work is not always easy to establish).

D.3.3. Elements of societal concern and how they might be handled

The outline below is a refinement of the categorisation of elements of societal concerns in RSSB Report cited above. That Report discusses each element in more detail. 'Risk' is used here in the HSE sense of a dimensionless probability.

A categorisation of this kind may provide a framework for identifying and handling the widely disparate elements that are explicitly or implicitly included in societal concerns.

It identifies some elements that might be incorporated (or are already incorporated) into the value of a prevented fatality (or non-fatal injury), separating those relating to those at risk from those relating to the hazard. It then identifies other elements that it may be best to handle at a later stage in the analysis.

Aspects of societal concerns that are or could be handled by adjustment of the value of a prevented fatality (VPF)

Personal characteristics

- Income/wealth *(equity value of VPF – i.e. same value for all regardless of wealth or age)*
- Age (adults) *(equity value of VPF)*
- Dependants *(equity value of VPF)*
- Health state unrelated to the hazard *(lower VPF?)*
- Disenfranchised/ disadvantaged consumers, including children *(higher VPF?)*

Hazard characteristics

- Baseline level of risk *(see section 4.4.2 above)*
- Fear/anxiety/dread *(adjusted VPF where there is evidence on considered preferences)*
- Health effects preceding death *(additional cost of health effect and adjustment for delay)*
- Distribution of responsibility (e.g. railway trespassers) *(adjusted VPF?)*

Aspects of societal concerns that seem best handled by adjustments to the analysis outside the VPF

- Fear/ anxiety uncorrelated with risk
- Distribution of risks and benefits
- Deterrence (e.g. crime)
- Blame/culpability
- Identifiably of victims
- Ambiguity (uncertainty about the nature of the hazard or the risk)

NERA
Economic Consulting

NERA Economic Consulting

15 Stratford Place
London W1C 1BE
United Kingdom

Tel: +44 20 7659 8500

NERA UK Limited, registered in England and Wales, No 3974527
Registered Office: 15 Stratford Place, London W1C 1BE