OPEN ACCESS

University of BRISTOL

Thomas, P., & Waddington, I. (2017). What is the value of life? A review of the value of a prevented fatality used by regulators and others in the UK. *Nuclear Future*, *13*(1), 32-39.

Publisher's PDF, also known as Version of record

Link to publication record in Explore Bristol Research
PDF-document

University of Bristol - Explore Bristol Research
General rights

What is the value of life?

A review of the value of a prevented fatality used by regulators and others in the UK, by **Philip Thomas** and **Ian Waddington**

The UK Department for Transport (DfT) values the prevention of a fatality on Britain's roads at £1.8 million (2016 £s). This value is used across government departments and agencies, including the Office for Nuclear Regulation, as a de facto standard for valuing the benefit of safety measures that preserve human life; however, there is no evidential basis for this valuation as it is derived from a statistical analysis of sparse survey data carried out 20 years ago that has now been found to be flawed. The methodology used to infer the VPF has been shown to be internally inconsistent, with the final recommended value being subjective. Members of the public whose opinions were rejected by the survey team actually gave entirely rational and understandable valuations based on human perceptions of utility. Another influential study – used to justify a significant reduction in spending to prevent multi-fatality rail accidents – has been found to be systematically biased against those very people who wanted more to be spent on preventing accidents causing multiple deaths. In contrast to the one-size-fits-all VPF, the J-value provides an objective, rational and statistically rigorous methodology that values the prevention of a premature death in terms of the amount of life that the potential victim would lose.

1. Philosophical problems

At the heart of UK health and safety policy is the demand that the risks involved in any venture should be reduced 'so far as is reasonably practicable', as specified by the Health and Safety at Work Act 1974. This broad requirement, interpreted as bringing residual risks 'as low as reasonably practicable' (ALARP), has been expounded in many sets of regulations and guidelines, notably the Health and Safety Executive's (HSE) document 'Reducing risks, protecting people' [1]. That document follows case law, in recommending that the ALARP requirement can be fulfilled only if the costs of reducing the risks any further can be demonstrated to be grossly disproportionate to any benefits received. Finding the balance between the benefit of preventing injury or death against the cost of doing so implies that a common metric should be used for both sides in the comparison, and this is conveniently taken as money. This entails ascribing a monetary value to healthy life.

The economic worth of a good can best be established from its market value – the price on which buyer and seller agree; however, health and safety benefits are, by and large, not free-market goods and so alternative evaluation methods need to be used. The next best method is to look at people's 'revealed preferences'. As the seventeenth-century philosopher John Locke said: "I have always thought the actions of men the best interpreters of their thoughts" [2]. Such methods are widely used in North America; for example, the US Department of Transportation uses labour market statistics to derive a 'value of a statistical life' of $9.1 million (2012 dollars) [3]. The method assumes that jobs involving a risk to life demand a wage premium, and this reveals the value the employee places on his

life. For a critique of these revealed preference methods see the review by Viscusi and Aldy [4].

When neither market values nor revealed preferences are available, an opinion survey of monetary preferences constitutes the final approach to estimating the financial value of averting a premature death. Such a 'stated preference' line of attack is sometimes called 'contingent valuation', since the estimates obtained are conditional or 'contingent' on the features of the scenario presented to the respondents taking part in the survey [5]. A variant of this technique is relied upon by the DfT for its published 'value of a prevented (statistical) fatality', generally abbreviated to VPF. The DfT VPF is held to be proportional to gross domestic product (GDP) per head and has a current value of £1.83 million (2016 £s) [6]. It is used across multiple government sectors and industries, including the Office for Nuclear Regulation (ONR).

Stated preference methods have inherent limitations, beyond any difficulties in particular methodologies [7]. For example a lack of information can lead to a reluctance to pay for something that would produce significant benefits or, conversely, people can be willing to pay for something with no discernible benefit. Moreover, people may be influenced by the perceived behaviour of others, and, for example, be willing to pay more if they believe that everyone is making an equal contribution rather than if their sacrifice is made on an individual basis. Beyond that, the maximum amount that a person will be prepared to pay will obviously depend on his or her ability to pay, which means that any survey must cover the full spectrum of wealths in the target population. (Later in this article, the procedure on which the UK VPF is based will be shown to experience particular problems over this requirement, see Section 3.)

But there is a further difficulty with the very principle of assigning a single value to 'preventing a fatality'. The wording itself embodies a confusion, since there is no way to 'prevent a fatality', as everyone is going to die at some point. In fact, the best that a safety measure can achieve is to restore the future years of life that would otherwise have been lost. It is self-evident that preventing the premature death of a young person is likely to have a greater value – in terms of the length of his or her future life – than preventing the premature death of an old person. To quote the Oxford philosopher John Broome: "Some people suggest that life has an infinite value but it does not: it has a finite value and it is a matter of quantity" [8]. This perspective leads to the notion of valuing the number of life-years lost to measure the impact of an injury or illness leading to premature death, rather than assigning a uniform value to every prevented fatality.

Some might argue that valuing loss of life expectancy discriminates in favour of the young and against the old. But it does not, as the next hour of life is valued equally for all, or, as Sunstein put it: "Everyone's life year counts as no less and no more than one" [7]. This approach to valuing life-years provides the basis recommended by NICE [9] for measuring the impact of illness and non-fatal injuries. The 'quality-adjusted life year' (QALY)

advocated by NICE assigns a monetary value to a year of life, reduced from its normal value in proportion to the severity of the ongoing illness or chronic injury.

The one-size-fits-all VPF has a limitation of particular significance to the nuclear industry, in that it fails completely to account for deaths that are delayed by many years. For example, the VPF approach would characterise death after, perhaps, two years of illness at the age of 70, resulting from an accident that happened 30 years earlier, in exactly the same way as an accident that cut off that same person's life immediately, at the age of 40. This is clearly a nonsense, as in the first case about a decade of life can be expected to be lost, while about 40 years can be expected to be lost in the second. The misconception is relevant to nuclear power because most radiation doses, whether as a result of an accident or coming from a continuing exposure, are low level, and any health effects, should they occur, will be delayed by tens of years on average. A similar effect applies to exposure to toxic materials in other industries, as evidenced by the delayed mortality from asbestosis among asbestos workers, and from 'black lung' (pneumoconiosis) among coal miners [10].

The UK VPF is based on a series of small-scale opinion surveys carried out 20 years ago by essentially the same team. That team dismissed its first attempt [11] in favour of its second [12]. It is a matter of some concern that, despite its significant philosophic shortcomings, the VPF is widely used by UK government departments and agencies, including the ONR, as discussed in Section 2. It is rather more troubling that the study [12] on which the current VPF rests has been shown to possess serious inconsistencies [13].

The debate in the literature over the validity of the VPF will be reviewed in detail in Section 3, which will highlight the logical inconsistencies in the reasoning that led to the accepted figure of £1 million (1997 £s). As will be shown in Section 4, a realistic model of people's attitude to risk demonstrated that a VPF of between £9 million and £12 million (1997 £s) seems actually to have been preferred by survey respondents. Furthermore, it will be shown in Section 5 that the government's decision to remove the premium on spending to prevent multiple-fatality rail accidents has not been justified, since the apparent 'justification' relies on analysing opinions using a method now mathematically discredited.

Section 6 will deal briefly with the J-value method, which conforms to the precepts of Locke, Sunstein and Broome, and offers a way of valuing life that is objective, rational and statistically rigorous.

Finally, Section 7 presents the conclusions of this review, including comments on the possible legal position.

2. The UK's VPF and its use

The closing years of the 1990s saw the Department of the Environment, Transport and the Regions commission the research that underpins the value of a prevented fatality currently used in the UK. An independent review of the VPF was conducted for the government by the accountants Deloitte LLP in 2009 [14], and this provides useful background and references for this section.

The current VPF is based on a series of opinion surveys reported by Beattie et al. [11] and Carthy et al. [12]. Survey respondents were asked about their willingness to pay for small reductions in the number of deaths or injuries from road accidents. From these responses, a value of £1 million per prevented fatality was reported (1997 £s), and this value was subsequently adopted by the DfT, and is updated annually by indexing to GDP per head. This VPF is recommended for use in cost–benefit analyses forming part of the appraisal process for new road schemes that are funded

by the DfT. The VPF is also used widely by local authorities for road safety [14].

The probability of a road fatality is higher than the risk of death in many other contexts (in the rail or nuclear industries, for example) and it was felt by the authors of the VPF research that survey respondents would comprehend a reduction in these 'large' risks more easily [11]. On the face of it, there appears indeed to be a much higher public acceptance of road fatalities than for other risks – it is hard to imagine an industrial or commercial setting in which 1900 deaths per year would be acceptable without protest (yet this is the mean of yearly road deaths that occurred from 2009 to 2013 [15]). A factor behind the de facto acceptance of this large annual incidence of fatalities may be the perception that the individual bears responsibility in a road accident, whereas for other forms of transport 'someone else' is responsible and has a duty to afford people greater protection.

Given the development of the DfT VPF figure in the context of the relatively high-risk area of road safety, it is an open question how far the value is appropriate for use in other contexts. Nevertheless, the DfT VPF is used widely across the UK, not only for transport, but in many other areas also.

Within the purview of the DfT, the VPF is used by the Rail Standards and Safety Board (RSSB). Following privatisation and a series of rail accidents in the late 1990s/early 2000s, the RSSB sponsored a number of studies about people's attitude to rail safety. One outcome of that programme was to remove a disproportion factor of three that had been applied to accidents involving multiple fatalities (see Section 5). The RSSB now values preventing a fatality on the railways with the DfT VPF, irrespective of the number of passengers at risk, with non-fatal injury prevention valued as a fraction of the VPF [14].

While regulators of the other major forms of transport, air and sea, do not recommend using the VPF in their safety assessments, six further government departments or agencies make use of the DfT VPF directly, according to the Deloitte review [14]. The Food Standards Agency uses the VPF to assess interventions to reduce food-borne illness, while the Department of Communities and Local Government uses the VPF for fire safety, flood prevention and domestic hot water safety schemes. The Home Office values the life of a homicide victim at the VPF and the HSE uses the VPF to value fatalities in workplace accidents. The latter include cancer-induced deaths, but here the value of the DfT VPF is doubled before application. Both the Environment Agency and the Health Protection Agency adopt the VPF. Meanwhile the Department of Health uses the DfT VPF in its derivation of the value of a quality-adjusted life-year (QALY) [9,15].

The Office for Nuclear Regulation (ONR), which devolved from the HSE in 2011, uses the VPF as its yardstick for valuing human life [17], and follows the HSE which, in appendix 3, paragraph 13 of its document 'Reducing risk, protecting people' [1], states:

"Currently, HSE takes the view that it is only in the case where death is caused by cancer that people are prepared to pay a premium for the benefit of preventing a fatality and has accordingly adopted a VPF twice that of the roads benchmark figure."

It may be observed, in passing, that it is not logical to double the VPF for a radiation cancer, which will deprive the victim of less than half the life expectancy lost through an immediately fatal road accident [18]. In any case, if the 'roads benchmark' VPF is arbitrary, as will be shown in Section 3, doubling its value will produce a similarly arbitrary figure.

In fact, the ONR suggests that the factor of two will not normally be

important (paragraph 6.16 of [17]), thus appearing to downgrade the importance of a cost–benefit analysis using the VPF. This is curious, given the organisation's expressed faith in the DfT VPF.

It is still, of course, incumbent on nuclear licensees to demonstrate that their residual risks are ALARP, meaning that further spending on safety has to be grossly disproportionate to the benefit it would generate. A licensee might have expected to apply the ONR-approved VPF to translate into pounds the benefit of any safety improvement deemed excessive by the licensee, since only in this way can an unambiguous comparison be made between benefit and cost. By the same token, the licensee might have expected the result of such a comparison to be accepted by the regulator. Reducing the importance assigned to this numerical comparison shifts the emphasis onto the subjective judgement of the ONR inspector and opens the door to regulatory inconsistency. The knock-on effect will be increased uncertainty in the costs faced by the licensee.

3. Flaws in the UK VPF
The UK VPF is based on face-to-face interviews with 167 people, as described in a paper by Carthy et al. [12]. In an attempt to avoid the difficulty of asking people directly to put a value on preventing a fatality, the investigators devised a two-part process to elucidate people's attitude to injuries and death, the two-injury chained method.

The theory of the two-injury chained method is summarised in the Appendix, including the concept of the 'value of a prevented injury' (VPI), which will morph into the VPF when the injury is fatal. But how does the theory fare in practice? Some indication of troubles ahead may be gleaned from the very wide range of values deduced for the average value of the VPF. Depending on the method used to filter and combine the results, the survey's authors found this to lie between £700,000 and £33 million (1997 £s; see Tables 2–7 in reference 12), with the lower limit reduced to £120,000 after the authors allowed themselves the freedom to choose the median instead of the mean. Despite this spread of two orders of magnitude, the authors were confident to recommend a value of £1 million (1997 £s) to their sponsor. Moreover, a report for the DfT in 2011, co-authored by some of the original investigators, concluded that there was no need to re-evaluate this recommendation [19].

But a comprehensive re-analysis of the methods and data used in the Carthy et al. paper was presented in 2015 by Thomas and Vaughan [13], who concluded that "there is no evidential base for the VPF that has been used for many years in the UK and is still in standard use today." This conclusion was dictated by the discovery that the original work contained several important shortcomings. Most strikingly, it was found that the two-injury chained method that lies at the heart of the Carthy et al. study contained inherent flaws.

Thomas and Vaughan were able to compare different estimates of the same individual's VPI for injury X using two methods put forward by Carthy et al. [12], which they applied to Carthy's own data. The first estimate could be calculated after the individual had stipulated two payments he considered fair for the injury: the maximum acceptable price (MAP, £) he would pay to avert the injury and the minimum acceptable compensation (MAC, £) he would be willing to receive to make up for enduring the injury. The second estimate could be computed using the two-injury chained method to link injury X to a lesser injury, W.

For the chained method to be valid, the two estimates would need to be at least approximately equal, but in fact the estimate coming from the chained method was a factor of 3–8 higher on average than the direct

estimate. Moreover, there was a scatter of many orders of magnitude between individuals' estimates under each method, and the estimates emerging from the two methods were found to be barely linearly correlated. See Figure 1, where the large scatter between points has dictated logarithmic scales so as to encompass all the results on one sheet of paper. The figure suggests that not only do individuals disagree strongly with the other members of the cohort, they tend to disagree strongly with themselves! Thomas and Vaughan concluded that the chained method was fundamentally flawed and should not be used.

In responding to these criticisms, Chilton et al. [20] said that they had themselves observed this inconsistency in their earlier paper [12]. They had noted the large discrepancy between their VPF estimates based on 'direct' calculation, £2–3 million, as listed in their Table 2, and their 'chained' estimates, £14–33 million, as given in their Table 6 (all 1997 £s). Chilton et al. suggested this disparity was due to a tendency for individuals facing a specified severe injury to require nearly equal, low probabilities of failure from different treatments, even when one would give a much quicker and better recovery if successful. Since a difference between failure probabilities occurs in the denominator in the two-injury chained method, small variations can lead to large differences in the VPI and hence VPF. The authors' admission ("This is undoubtedly a problem for the use of the chained approach") is tantamount to conceding that the method underlying the UK VPF is unstable with regard to its input data. Thomas and Vaughan's conclusion that the results of the original study cannot be relied upon is thus strengthened rather than diminished.

The subjective nature of the Carthy et al. [12] conclusions starts to become clear in the statement contained in Chilton et al. [20]: "If we *had* to judge, we would think it rather less plausible that this [£62,258] is the appropriate money equivalent for an injury that is completely healed with no further adverse effects after 18 months," preferring a figure of £8959 (equivalent to about six months' mean gross income in 1997). Any doubts
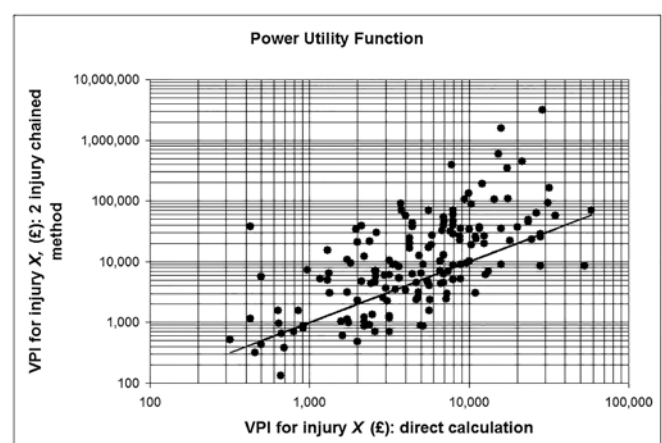

Figure 1: VPI for injury X, found for each individual by direct calculation and via the two-injury chained method (based on Figure 2 of [13]).
The very great scatter between points dictates the use of a log-log scale for display purposes. The straight line shown would have been followed if the two estimates had been equal: *y = x*. The best linear fit, *y = 8.24x*, reveals that the two estimates differ by an order of magnitude and have a very low linear correlation: $R^2 = 0.072$. Only 7.2 per cent of the variation in the second estimate is explicable in terms of the variation in the first.

about the dependence of the final figure on the authors' judgement are removed by the last sentence of [20], where the "blending of evidence with judgement" to help consolidate the VPF is cited as the first contribution of the Carthy study. It is worth remembering at this point that stated preference surveys are normally undertaken to find an answer based on the preferences of the respondents, not the preferences of the authors.

Thomas and Vaughan [13] went on to discuss a further inconsistency in the VPF results, which failed Carthy et al.'s own test to compare the two-part and three-part chained method. Thomas and Vaughan were able to reproduce this inconsistency by processing a random dataset through a simulation of the chained method that included a randomly chosen multiplier to represent the ratio of the three-part and two-part chained estimates. Chilton et al. [20] did not respond to this evidence that their results were no better than random.

Thomas and Vaughan [13] pointed out, in addition, that the wealth of the respondents deduced from the survey data was unreasonably small – below 10 per cent of the average adult wealth at the time of the survey. Chilton et al. [20] argued that wealth could not be inferred from the utility functions they used in their study for technical reasons. This argument was resisted by Thomas and Vaughan [21] on the grounds of consistency, whereupon two of the Carthy authors suggested, in a 'Final response' [22], that auxiliary utility functions could be used to cover low wealths. Thomas and Vaughan [23] were able to show, however, that this new suggestion would fail: it was a mathematical impossibility in two out of the three cases and economically unsound for the third. Even if the questionable economic validity were put aside for the sake of argument, it was still possible to calculate the implied average wealth. This remained the same as previously found – either £5252 or £1730 (1997 £s). Such low values, acknowledged by Jones-Lee and Loomes [22] to be "absurdly low", would have been unrepresentative of the national population – the average UK adult wealth stood at £78,300 at the time of the survey in 1997 [24].

The discussion then moved on to the validity, or otherwise, of censoring data by using a 'trimmed mean' rather than the sample mean as the estimator of the VPF. Carthy et al. [12] removed two outliers with apparently high VPFs from the sample before taking the mean (now unilaterally trimmed) of the remaining sample as their best estimate, finding a VPF of £0.9 million to £1.6 million (1997 £s), depending on the utility model used. But Thomas and Vaughan [13] used a model simulation to show that there was no basis to reject even the highest outlier, so that the trimming process biased the results illegitimately low. Restoring the arbitrarily censored values produced a sample mean of £3.3 million with a 90 per cent confidence interval of £61,000 to £6.6 million (all 1997 £s), although Thomas and Vaughan commented that the serious flaws affecting the basic method meant that the figures were of academic interest only.

Thomas and Vaughan [21] concluded they could find no reason for changing their earlier conclusions that the VPF currently in widespread use "has been shown to be based on a study that does not stand up to scrutiny."

4. Caveat investigator re-investigated

Prior to the Carthy et al. [12] surveys discussed above, essentially the same group of investigators (Beattie et al. [11]) published the results of two earlier surveys aimed at determining people's willingness to pay to prevent a road fatality. The first survey involved focus groups and one-to-one interviews with 83 people who were asked what they would pay for a car safety feature that would reduce their risk of injury or death in five different scenarios. The results showed that people were willing to pay £98 on average to reduce the risk of death by one in 100,000 and £139

for a reduction in risk that was three times greater. The authors expected people's responses to be linear in risk-reduction, at least roughly, but in fact respondents seemed to be willing to pay only 1.4 times more for a three-fold reduction in risk. On the basis of this result, the methodology of this survey was rejected and a further survey commissioned.

The second survey asked how much respondents were willing to pay to reduce the number of road deaths in "the area in which they lived", taken to contain one million people. A total of 52 individuals were interviewed. The questions involved reducing the number of deaths by five or 15 per year over a period of either one or five years. After providing their answers, the respondents were prompted to consider whether their valuations were linear with regard to the number of lives saved and were given the opportunity to revise their answers [11]; however, once again, the results suggested that in order to prevent multiple deaths, people were willing to pay only 1.3 times more to prevent three deaths as they were to prevent one.

The results of both these studies were rejected by the authors and their sponsors due to their perceived "serious doubt on the reliability and validity" of the study methods [11]. They went so far as to give their published paper the label 'Caveat investigator' – Let the investigator beware.

But were the authors correct to reject the opinions of the 135 survey respondents? Thomas and Vaughan [25] argued that they were not, and demonstrated that the survey responses could be accounted for by the same reasoning that explains the prevalence of offers such as 'Buy one, get one free' and '3 for the price of 2' in supermarkets. Thomas and Chrystal [26] had introduced the theory of relative utility pricing (RUP), using utility functions to describe how consumers value a multi-pack of a commodity relative to a smaller pack. The theory was able to explain the pricing of different 'packs' of milk, eggs and USB memory sticks, for example [27].

The RUP model applies to any commodity that is offered simultaneously in different pack sizes, and Thomas and Vaughan [25] were able to argue that the safety benefits offered to participants in the Beattie et al. survey could be considered as commodities. A necessary feature in RUP theory is that a person should assign a positive utility to a positive benefit, with a larger benefit attracting a higher utility. This requirement on discrimination eliminated approximately 45 per cent of the Beattie et al. samples from further analysis, leaving 74 respondents in total who were able to discriminate between the different safety measures or 'packs'. Thomas and Vaughan calculated the mean ratio of the MAPs that these respondents were willing to pay for the two safety packs and found similar values for the two surveys: 1.98 for the first and 1.93 for the second. Recalling that the second safety pack provided three times the benefit of the first in both the surveys, the MAP ratio of 2.0 deduced from RUP theory lay comfortably within the 90 per cent confidence intervals of the observed means.

Having shown that the respondents' valuations could be explained by the RUP model, Thomas and Vaughan then calculated the VPF implied by their results [25]. After removing the respondents who were unable to distinguish between the safety packs, a VPF in the range £8.7 million<VPF1<£11.9 million was found based on the first pack, and a VPF of £4.6million<VPF2<£5.9 million (all 1997 £s) was found from the second pack, which produced three times the benefit of the first. They noted that when the second pack is a much higher multiple of the first in terms of risks reduced or premature deaths averted, the RUP model predicts that VPF2 will converge towards VPF1. Thus they concluded that VPF1 has "more general applicability".

The point is that, if one puts one's faith in stated preference valuations, then the Carthy et al. surveys [12] have to be rejected because of the major

flaws exposed and the Beattie et al. surveys [11] are the ones to be believed. But the spending on avoiding a premature death, as derived from the Beattie et al. surveys (between £8.7 million and £11.9 million), will need to be an order of magnitude larger than that adopted by the investigation's sponsors (£1 million, all 1997 £s). Upgrading these figures to accommodate the increase in GDP per head since the Beattie study was carried out, this suggests that the current VPF should lie between £16 million and £22 million (2016 £s).

5. Multiple fatalities: RSSB's reliance on a biased study

Do people put a premium on preventing multiple fatalities in a single accident? Prior to 2003, the rail industry had mandated the use of two different VPFs when making a rail safety case: one for single-fatality accidents and a second, larger VPF for multi-fatality accidents [28]. Following its establishment in 2003, the RSSB adopted the DfT's VPF for all accident scenarios. This had the practical effect of reducing by a factor of three the amount that rail operators were obliged to spend to avert a fatality in an accident expected to cause many deaths if it happened.

The RSSB commissioned research in 2005 to investigate whether people would in fact prioritise safety spending based on the type of accident, in particular multi-fatality accidents, those involving children and fatal occurrences resulting from deliberate acts: trespass and suicide. The results were published in an RSSB technical report [29] and subsequently in Covey et al. [30]. The study, with a number of authors common to both of the VPF studies discussed above, concluded that the prevention of a death in a multi-fatality accident was not valued any higher than the prevention of an isolated death. On this basis, the RSSB continued and continues to recommend a single VPF when assessing all accident scenarios.

Thomas [31] reviewed the methodology reported in [30], which had been used in a number of influential studies on the VPF carried out in the UK for different government departments and agencies, including the HSE and the DfT, as well as in the Covey et al. report for the RSSB [29]. These studies all relied upon interpreting the results of one or more stated preference surveys, where respondents were questioned indirectly about their willingness to pay to avoid death in two scenarios, A and B. The respondent is asked to specify how many deaths in scenario A he considers to be as undesirable as a possibly different number of deaths in scenario B. Then, since the individual regards the two outcomes as equally bad, it is assumed that he would be prepared to authorise equal expenditure to avert each set of deaths. The ratio, R_{ABi}, of his expenditure per death averted in scenario A to his corresponding figure in scenario B will then follow the inverse ratio of the numbers of deaths he considered equivalently bad:

$$R_{ABi} = V_{Ai} / V_{Bi} = N_{Bi} / N_{Ai}$$

where V stands for value and N for number.

The question is then how to combine these ratios from all respondents into a single measure for the acceptable ratio of the spend against a death in scenario A compared with the expenditure against a fatality in scenario B. The sample mean of R_{ABi}, known to be an unbiased estimator, is easy to calculate and is the only statistic recommended for consolidating views measured by opinion surveys [32]; however, Covey et al. chose to employ a non-standard and more complicated statistic, the Valuation Index.

The Valuation Index possesses the property of reciprocality, in the sense that the index of the reciprocal is equal to the reciprocal of the index. This reciprocality comes at a price, however: the Valuation Index fails the test of "structural view independence", the principle that there should be no inbuilt bias in a consolidation algorithm that would render the views of

some people less important than the views of others [32].

Far from giving equal weight to each respondent's view, the Valuation Index is biased against anyone wanting more spent per death averted in a multi-fatality accident than in a single-fatality accident. The bias becomes more marked the more averse the person is to multi-fatality accidents (see Figure 2).

The built-in bias towards low valuations makes the Valuation Index particularly unsuitable for investigating whether or not society wants more to be spent to prevent deaths in a multiple-fatality rail accident. To the extent that they rely on such a study for support, any safety case, regulation or recommendation citing or based on the results generated using the Valuation Index has to be regarded as unsafe.

6. The J-value

As discussed above, the UK VPF has been shown to lack credibility in the last two years; however, there remains a continuing need for nuclear licensees to demonstrate ALARP if they are to comply with UK law. In fact, ALARP arguments may now be made objectively and rigorously using the Judgement- or J-value technique. Developed over the last decade, the J-value is a revealed preference method that is able to place an objective value on the increase in life expectancy that the safety measure brings about, summed over all its beneficiaries.

The philosophical and mathematical basis for the J-value has been laid out in this journal before [33], and a full set of papers is available at www. jvalue.org.uk. Rather than being reliant on the subjective opinions of a small group of people, the J-value is instead grounded in objective actuarial and economic statistics that characterise the lives and behaviours of millions of citizens. The parameter is thus suitable for assessing health and safety measures across all industries, from oil and gas, chemical and nuclear, through transport to the National Health Service in the UK. Moreover, unlike other approaches, the J-value allows immediate fatalities and loss of life in the longer term (e.g. after exposure, either of workers or of the general public, to a carcinogen such as nuclear radiation) to be differentiated but measured on the same scale.

The principal author recommending continued reliance on the UK VPF [19] composed a negative review [34] in 2009 of the J-value literature for the Nuclear Division of the HSE, now ONR. Despite being prepared for what is now ONR and stating that its author had worked as a "physicist/
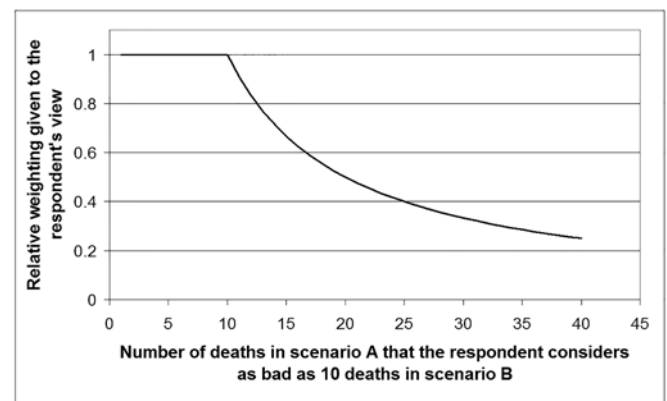
Figure 2: Some people's opinions count for less: relative weightings of respondents' views when the Valuation Index is used (based on Figure 1 of [31]).
Assign single-fatality accidents to scenario A and multiple-fatality accidents to scenario B.

engineer" in the nuclear industry, the report made no attempt to analyse the J-value papers dealing with nuclear risks (roughly half the literature available at that time). When the report was eventually made public, its objections were fully answered by Lind [35] and Thomas and Jones [36]. A 'final comment' was supplied by the author two and a half years later [37] but he made no attempt at a reasoned rebuttal of the points made in [35] and [36], seemingly content merely to restate his faith in the stated preference methods discussed above.

The J-value method has recently been validated against pan-national data [38], and its objective nature means that the method can offer consistency and certainty in regulation, particularly for cases where an ALARP argument needs to be made.

7. Conclusions

The VPF is used widely in the UK by government departments and agencies in the assessment of safety schemes to protect the public. Based on stated preference research carried out in the late 1990s for the now Department for Transport, it is updated annually in line with GDP per head and has a current value of £1.83 million (2016 £s).

But recent investigations have revealed severe methodological problems with the research upon which this VPF is based. The analysis of the stated preferences of the 167 people in the survey that underpins the VPF figure has been shown to be seriously flawed. By contrast, a new understanding of retail price offers – particularly '3 for the price of 2' – explains the findings of the previous 'Caveat investigator' study that was rejected by its authors and sponsors but was actually valid as a stated-preference survey. If stated preferences are to be believed, the UK public wants a VPF that is 10 times higher than the current value, at between £16 million and £22 million (2016 £s).

The clear findings on the lack of evidence for the UK VPF mean it is questionable whether any safety case presented from now on that relies in any way upon the UK VPF, whether on the roads, the railways or in the nuclear industry, could stand up to test in court.

Meanwhile, a survey cited as support for a three-fold reduction in safety spending on UK railways against big train accidents has been found to be biased against all those who wanted more spent on preventing big train accidents. To the extent that they are reliant on such a study for support, it is open to doubt whether any safety case, regulation or recommendation citing or based on results generated using the Valuation Index could resist legal challenge.

At a more fundamental level, the value of a prevented fatality – a uniform valuation for every possible premature death – is a very blunt instrument with which to assess the impact of safety measures. In particular, this approach cannot accommodate properly a delayed death that may occur many years after an accident, such as death from a radiation-induced cancer. A method of valuing human safety is required in the nuclear industry that treats everyone equally and fairly and is based on solid evidence – the VPF is not that method.

Appendix

The two-injury chained method [12] is based on the application of utility functions, economic constructs that are designed to model the benefit or satisfaction that an individual feels he is getting from his wealth. It boils down to the following. The respondent is asked to imagine he has been injured in a road accident and is offered two hospital operations, A or B, each of which will leave him with a serious, 'type 2' injury should it fail. Operation A will still leave the

patient with a less serious, 'type 1' injury even if it succeeds. But, provided it is successful, operation B will return the patient to full health almost immediately. Other things being equal, operation B looks to be the one to choose.

But what happens if the probabilities of success are different for the two surgical procedures? The Carthy authors thought that Operation B's potentially better outcome should cause the patient to accept that it could fail more often. They then theorised that information on acceptable failure probabilities, elicited via a 'standard gamble' procedure, should determine the ratio, m_{2i}/m_{1i}, of individual i's value of a prevented injury (VPI) of type 2, m_{2i}, to his VPI for a type 1 injury, m_{1i}.

Further theory expounded by Carthy et al. [12] suggests that an individual's VPI for a type 1 injury can be calculated after drawing out from the respondent both the maximum acceptable price (MAP, £) he would pay to avert the injury and the minimum acceptable compensation (MAC, £) he would be willing to receive to make up for enduring the injury. Multiplying the individual's VPI for a type 1 injury by the VPI ratio, m_{2i}/m_{1i}, found from the standard gamble model, then gives his VPI, m_{2i}, for a type 2 injury.

If the type 2 injury is taken to be fatal, then the type 2 VPI will be equal to the individual's personal VPF, m_{Di}. If the sample has been chosen to cover the full spectrum of wealths in the UK, averaging over all the people in the sample will then yield the VPF for the UK as a whole, m_D.

References

1. HSE (2001) Reducing risks, protecting people: HSE's decision making process, HMSO, London
2. Locke, J. (1690) An essay concerning human understanding, London, available at www.earlymoderntexts.com/assets/pdfs/locke1690book1.pdf (accessed November 2016)
3. US Department of Transportation (2015) Guidance on Treatment of the Economic Value of a Statistical Life in U.S. Department of Transportation Analyses www.transportation.gov/regulations/economic-values-used-in-analysis (accessed September 2016)
4. Viscusi, W.K. & Aldy, J.E. (2003) The value of a statistical life: a critical review of market estimates throughout the world. J. Risk Uncertain. 27, 5–76, doi: 10.1023/A:1025598106257
5. Carson, R.T. (2001) Resources and environment: contingent valuation. In International Encyclopedia of the Social & Behavorial Sciences (edited by N.J. Smelser & P.B. Baltes), Elsevier Science, London www.sciencedirect.com/science/article/pii/B0080430767041966
6. DfT (2016) Transport Analysis Guidance data book www.gov.uk/government/publications/webtag-tag-data-book-july-2016 (accessed August 2016)
7. Sunstein, C.R. (2003) Lives, Life-Years, and Willingness to Pay, University of Chicago John M. Olin Law & Economics Working Paper No. 191 http://ssrn.com/abstract=421341 (accessed October 2016)
8. Broome, J. (2006) Safety philosophy & morality. In The Economics and Morality of Safety, Royal Academy of Engineering, London
9. Glover, D. & Henderson, J. (2010. Quantifying health impacts of government policies: a how-to guide to quantifying the health impacts of government policies, Department of Health, available at: www.gov.uk/government/uploads/system/uploads/attachment data/file/216003/dh 120108.pdf (accessed October 2016)
10. Thomas, P. & Jones, R. (2011) Response to 'Review of the J-value literature – Final Report [to] The Health and Safety Executive'", Office for Nuclear Regulation, London, available at www.onr.org.uk/j-value-report.htm (accessed October 2016)
11. Beattie, J., Covey, J., Dolan, P., et al. (1998) On the contingent valuation of safety and the safety of contingent valuation: Part 1 – Caveat investigator. J. Risk Uncertain. 17, 5–25
12. Carthy, T., Chilton, S., Covey, J., et al. (1999) On the contingent valuation of

safety and the safety of contingent valuation: Part 2 – The CV/SG chained approach. J. Risk Uncertain. 17, 187–213

13. Thomas, P.J. & Vaughan, G.J. (2015) Testing the validity of the 'value of a prevented fatality' (VPF) used to assess UK safety measures. Process Safety and Environmental Protection 94, 239–261

14. Deloitte LLP (2009) Review of the Highways Agency Value of Life Estimates for the Purposes of Project Appraisal: A Report to the NAO www.gov.uk/government/publications/report-about-the-value-of-prevented-injuries-and-fatalities (accessed August 2016)

15. DfT (2014) Number of fatalities resulting from road accidents in Great Britain, https://www.gov.uk/government/publications/annual-road-fatalities (accessed October 2016)

16. NICE (2013) Judging whether public health interventions offer value for money. Local government briefing [LGB10], London, September www.nice.org.uk/advice/lgb10/chapter/judging-the-cost-effectiveness-of-public-health-activities

17. ONR (2015) Guidance on the Demonstration of ALARP, Nuclear Safety Technical Assessment Guide, NS-TAST-GD-005 Revision 7, London www.onr.org.uk/operational/tech_asst_guides/ns-tast-gd-005.pdf (accessed August 2016)

18. Thomas, P.J. (2017) Age at death from a radiation-induced cancer based on the Marshall model for mortality period, submitted to Process Safety and Environmental Protection

19. Spackman, M., Evans, A., Jones-Lee, M., et al. (2011, Updating the VPF and VPIs: Phase 1: Final Report [to] DfT, NERA Economic Consulting, London, available at http://assets.dft.gov.uk/publications/pgr-economics-rdg-updatingvpfvpi-pdf/vpivpfreport.pdf (accessed May 2014; no longer available as of October 2016)

20. Chilton, S., Covey, J., Jones-Lee, M., et al. (2015) Response to Thomas and Vaughan, Process Safety and Environmental Protection 93, 293–298

21. Thomas, P.J. & Vaughan, G.J. (2015) Testing the validity of the 'value of a prevented fatality' (VPF) used to assess UK safety measures: Reply to the comments of Chilton, Covey, Jones-Lee, Loomes, Pidgeon and Spencer, Process Safety and Environmental Protection 93, 299–306

22. Jones-Lee, M. & Loomes, G. (2015) Final response to Thomas and Vaughan, Process Safety and Environmental Protection 94, 542–544

23. Thomas, P.J. & Vaughan, G.J. (2015) Pitfalls in the application of utility functions to the valuation of human life, Process Safety and Environmental Protection 98, 148–169

24. Matheson, J. & Summerfield, C. (eds) (2000) Social Trends 30, 2000 Edition, Office of National Statistics, The Stationery Office, London

25. Thomas, P. & Vaughan, G. (2014) Explaining Perceived Inconsistencies in 'Stated Preference' Valuations of Human Life. Am. J. Industrial Bus. Management 4(9), 442–473, doi: 10.4236/ajibm.2014.49052

26. Thomas, P. & Chrystal, A. (2013) Explaining the 'Buy One Get One Free' promotion: the Golden Ratio as a Marketing Tool. Am. J. Industrial Bus. Management 3(8), 655–673, doi: 10.4236/ajibm.2013.38075

27. Thomas, P., Chrystal, A. (2013) Using Relative Utility Pricing to Explain Multibuy Prices in Supermarkets and on the Internet. Am. J. Industrial Bus. Management 3(8), 687–699, doi: 10.4236/ajibm.2013.38078

28. Bearfield, G. (2006) How safe is safe enough? International Rail Safety Council Conference, Dublin, 2006 www.intlrailsafety.com/Dublin/25_October_2006_Papers/01_George_Bearfield.pdf (accessed August 2016)

29. Covey, J., Robinson, A., Jones-Lee, M., et al. (2008) RSSB T616: Assessment of the Value of Preventing a Fatality (VPF). Report and Results, Rail Safety Standards Board, London

30. Covey, J., Robinson, A., Jones-Lee, M. & Loomes, G. (2010) Responsibility, scale and the valuation of rail safety. J. Risk Uncertain. 40, 85–108

31. Thomas, P.J. (2015) Testing the impartiality of surveys to measure differential risk perception. Measurement 60, 155–173

32. Thomas, P. J. (2014) Structural view independence: A criterion for judging the objectivity of economic parameters measured by opinion survey. Measurement 47, 161–177

33. Thomas, P., Vaughan, G. (2013) All in the balance: assessing schemes to protect humans and the environment. Nuclear Future 9 (3), 41–51

34. Spackman, M. (2009) Review of the J-value literature – Final Report, NERA Economic Consulting, London, dated 31 December 2009, published on HSE website February 2011. www.hse.gov.uk/nuclear/j-value-report.htm

35. Lind, N.C. (2012) Comments on 'Review of the J-value literature: final report', completed for HSE by M. Spackman. NERA Economic Consulting Report, 31 December 2009 www.hse.gov.uk/nuclear/j-value-report/nera-comments.pdf (accessed October 2016)

36. Thomas, P. and Jones, R. (2011) Response to "Review of the J-value literature – Final Report [to] The Health and Safety Executive' by Michael Spackman, NERA Economic Consulting. Available at: www.onr.org.uk/j-value-report/thomas-jones.pdf

37. Spackman, M. (2014) Final comment on the J-value approach by Michael Spackman www.onr.org.uk/j-value-report/final-comment-june-2014.htm (accessed October 2016)

38. Thomas, P.J., Waddington, I. (2017) Validating the J-value safety assessment tool against pan-national data, submitted to Process Safety and Environmental Protection

Abbreviations

ALARP	as low as reasonably practicable
DfT	Department for Transport
GDP	Gross Domestic Product
HSE	Health and Safety Executive
J-value	Judgement-value
MAC	minimum acceptable compensation
MAP	maximum acceptable price
NICE	National Institute for Health and Care Excellence
ONR	Office for Nuclear Regulation
QALY	quality-adjusted life year
RSSB	Rail Safety and Standards Board
RUP	relative utility pricing
VPF	value of a prevented fatality
VPI	value of a prevented injury

Authors

Philip Thomas

Philip is professor of risk management within the Safety Systems Research Centre in the Queen's School of Engineering, University of Bristol. He has extensive nuclear industry experience as head of several departments within the UK Atomic Energy Authority and as customer project manager for the £100 million project to decommission the Windscale AGR. He and his team have been responsible for developing a set of techniques, the J-value framework, that provide recommendations to decision-makers in high-hazard industries on how much ought to be spent to protect both humans and the environment.

Ian Waddington

Ian is a senior software engineer at Ross Technologies Ltd, Bristol. He has worked with Philip Thomas on the J-value for several research projects including NREFS, a project to examine measures put in place to protect power plant workers and the public after the nuclear accidents at Chernobyl and Fukushima Daiichi. He has provided mathematical support to the development of the J-value framework and led on the computing aspects of the J-value.