Contents lists available at ScienceDirect

Process Safety and Environmental Protection

journal homepage: www.elsevier.com/locate/psep

Pitfalls in the application of utility functions to the valuation of human life

P.J. Thomas*, G.J. Vaughan

Safety Systems Research Centre, Queen's School of Engineering, Bristol University, Queen's Building, University Walk, Clifton BS8 1TR England

ARTICLE INFO

Article history:
Received 27 May 2015
Accepted 2 July 2015
Available online 14 July 2015

Keywords:
Safety
Utility functions
Value of a Prevented Fatality
VPF
Health and Safety Executive
Human life

ABSTRACT

Safety strategies in the process and other industries depend ultimately on how much the owners and operators decide should be spent on protection systems to protect workers and the public from potential plant hazards. An important input to decisions of this sort is the value of life, which needs to be assessed in a valid manner so that safety decisions can be made properly. A key reference point for decisions on safety investment decisions in the UK is a 1999 study on the "value of a prevented fatality" (VPF), which employs a two-injury chained model that has been shown previously by the present authors to possess internal inconsistencies. The 1999 study made extensive use of utility functions to interpret survey data, and it is this feature that is explored in this paper. It will be explained here how different forms of utility function of the Exponential family can produce the same figure for an intermediate parameter in the calculation of the VPF from the two-injury chained model. Exponential utility functions are, however, unlikely to provide a realistic representation if their calculated risk-aversions need to be negative or zero in order to match survey data, which would imply an incautious attitude amongst those taking decisions on safety. The use of an incompletely specified wealth threshold in the utility modelling is explored in the light of a proposal by the authors of the 1999 study that a second utility function can be used to determine the individual's utility when his wealth lies below the threshold, which constitutes the lower limit of validity of the first utility function. The proposition is shown to be untenable. The results presented in this paper raise further concerns about the lack of validity of the 1999 study on which the UK VPF is based and hence on the safety decisions that have been made in consequence.

1. Introduction

Safety decisions in many industrial situations, particularly in high-hazard industries such as nuclear, oil and gas and chemicals, depend ultimately on a consideration of how much should be spent on safety measures to protect workers and the public from potential hazards. In the UK the legal requirement is to do all that is reasonably practicable to ensure the health and safety of workers and the public. This requirement means that it is necessary to compare the sacrifice (cost, time and trouble) of implementing measures to improve safety with the reduction in risks to those that might suffer harm (HSE, 2014). Thus the amount it is reasonable to spend on safety measures may be judged as a trade-off between the benefit that the system confers in terms of improved safety and the loss of benefit brought about by the costs of paying for those measures. In line with the Kaldor–Hicks compensation principle (Kaldor, 1939; Hicks, 1939), it is customary to assign the cost and hence a notional reduction in wealth to those being protected, even though they will rarely have to meet the bill in

* *Corresponding author*. Tel.: +44 117 331 5830; fax: +44 117 954 5208.
E-mail addresses: pjt3.michaelmas@gmail.com, philip.thomas@bristol.ac.uk (P.J. Thomas).
http://dx.doi.org/10.1016/j.psep.2015.07.002

practice. Utility functions are then used to characterise the fall in benefit these people experience as a result of their assumed reduction in wealth.

To compare the outlay on implementation with the costs in terms of loss of life and other detriments through a cost-benefit analysis (CBA), it is fundamentally necessary to place a monetary value on human life. The "value of a prevented fatality" (VPF) derives from one such valuation exercise. A figure for the VPF is published annually by the UK Government's Department for Transport (2013) and this is used also by the Health and Safety Executive (HSE) and the Office for Nuclear Regulation (ONR) in developing new safety regulations and determining whether safety measures meet the legal requirements—though it should be emphasised that the latter are the responsibility of those carrying out the work activity.

It is clear that ensuring the VPF is a true reflection of what should be spent is very important and to this end various methods of eliciting the VPF from social and economic surveys have been attempted. The VPF that is currently used in the UK is based on a study by Carthy et al. (1999), which uses a two-injury chained method, whereby individuals are asked to consider two serious injuries, with injury X more severe than injury W. The individual is asked to estimate the maximum acceptable price (MAP) he would pay to avert the specified injury and the minimum acceptable compensation (MAC) he would take as compensation for enduring the injury. The analysis of the data makes very extensive use of utility functions in an attempt to find the amount it is notionally reasonable to pay for a safety measure that will reduce by one the expected number of premature deaths in a large population, given that those deaths are associated with a specified hazard. This sum is deemed to be the VPF.

This approach to deriving a figure for the VPF originated in the transport sector but the concept is now applied much more widely. After tracing the history of the development of the UK VPF figure, Wolff and Orr (2009) concluded that:

"it appears that the Carthy study (Carthy et al., 1999) is now the primary source of VPF figures, adjusted for inflation and changes in GDP."

A 2011 report for the Department for Transport (DfT), with authors in common with the Carthy study, recommended "against any early new full scale WTP [willingness to pay] study" (Spackman et al., 2011). Thus the survey conducted by Carthy et al. of 167 people in 1997 (Carthy et al., 1999) remains the evidential base for the VPF used by the Government, regulators and many industries in the UK today, including the process and nuclear industries. It is obviously of crucial importance, therefore, that the Carthy study should be soundly based.

Updated for increases in GDP per head, the VPF is assumed to be the same for all people in the UK, irrespective of age or gender. While this might be a dubious assumption (see Nathwani et al., 1997, 2008; Pandey and Nathwani, 2003; Pandey et al., 2006; Sunstein, 2004a,b; Thomas et al., 2006a, 2006b, 2010a, Thomas and Vaughan, 2013), the VPF is used extensively in the UK.

We have questioned in earlier papers (Thomas and Vaughan, 2015a,b) the methodology used by Carthy et al, showing that their work contained serious flaws. Using the Carthy authors' own data it was demonstrated that the method was invalid in that it failed to estimate consistently a key parameter determining the size of the VPF. This parameter, m_{Xi}, is the individual's marginal rate of substitution between his wealth and his probability of not suffering injury X. The Carthy study's data allow estimation of m_{Xi} in two different ways, giving $m_{Xi}^{(1)}$ and $m_{Xi}^{(2)}$ respectively. These values should, of course, be equal for the method to be sound, but they are, in fact, very different and, indeed, barely correlated. Thus the two-injury chained method has been falsified in the sense used by Popper (1934). Indeed, whilst still defending their methodology, the Carthy co-authors have admitted to methodological problems:

> "there is a definite and seemingly systematic divergence between direct and indirect estimates which is illustrated by the comparison between $m_{Xi}^{(2)}$ and $m_{Xi}^{(1)}$" (Chilton et al., 2015).

Moreover, Thomas and Vaughan (2015b) found that many of the defences put forward by the Carthy authors in an effort to justify their methodology were flawed or mistaken and so concluded that their attempt to support the use of the methodology was not tenable.

A sub-group of those involved in the first defence published a second attempt at a justification of the Carthy study (Jones-Lee and Loomes, 2015). While not disputing the failure of the two-injury chained method in the fundamental validity test just mentioned, they raised two main points as a follow on to Thomas and Vaughan (2015b), concerning

(1) the form the utility function for wealth should take if the individual puts his maximum acceptable price (MAP) for averting the injury as high or higher than the minimum acceptable compensation (MAC) he would countenance to endure the injury;
(2) whether it is legitimate to deduce a wealth for the respondent from his stated MAP and MAC.

In considering these issues, it should be emphasised at the outset that they are both subsidiary to the previously accepted objection of systematic divergence detailed above, which is sufficient on its own to invalidate the two-injury chained method on which the UK VPF rests. We may conclude immediately, therefore, that the points now raised by Jones-Lee and Loomes are insufficiently important to affect let alone overrule the major criticism we put forward in our first paper. Therefore we reiterate our belief that the VPF derived by Carthy et al. and used for the past 16 years in the UK is not satisfactory for use in making safety decisions as it is not based on a sound analysis.

Nevertheless it must be recognised that the VPF figure used in the UK since 1999 rests squarely on the value produced by the Carthy study. It is thus very important that the attempted second defence of that study by Jones-Lee and Loomes should not be regarded as cover for continuing to use an invalid figure.

Much of the argument in Jones-Lee and Loomes (2015) concerned the use of utility functions and it is the purpose of this paper to investigate how the various utility functions may be used legitimately. Indeed the issues raised by Jones-Lee and Loomes (2015) are of both theoretical and practical interest. It is intriguing that different utility functions from the same Exponential family can give the same result for a key intermediate parameter in the calculation of the VPF when the two-injury chained method is used, namely the marginal rate of substitution, m_{ki}, of non-injury probability in place of wealth under injury k. The paper will explain why this is so.

Theoretically interesting results are also produced through considering the proposal of Jones-Lee and Loomes that a