What is the value of preventing a fatality?

Abstract: In deciding whether or not to introduce a safety measure a quantitative approach generally requires the decision maker to consider its financial cost, as well as its benefits in terms of saving of lives. To determine whether the measure is worth introducing a financial value is commonly placed on preventing a fatality (VPF - the value of preventing a fatality). Such a VPF is said not be the price society puts on a life, but rather the aggregated value of reducing a small probability of death for a very large number of people. This raises a number of questions. First, are the 'willingness to pay' methods used to determine this figure consistent with such an understanding? Second, is there any reason to think that the VPF should be a constant for a single person for different types or magnitudes of risk, or for different people over the same risk? Third, does current terminology stand in the way of a proper appreciation of the issues, and can this be improved?

1. INTRODUCTION

The idea of the ‘value of preventing a fatality’ (VPF) or, equivalently the ‘value of a statistical life’, is a technical term used in safety decision making. It gets its life from risk cost-benefit analysis (RCBA). Where a safety measure is designed to reduce the risk of death there will always be a question of whether it is worthwhile, whether on economic grounds, or some other.[i] As soon as it is admitted that there is a conceptual possibility that a potential safety measure could be too expensive to be worth implementing it is necessary to find a way of giving some meaning to this idea, and, if possible, some criterion to settle such questions.

At this point risk cost-benefit analysis offers a surprisingly simple formula for this purpose. To decide whether, according to RCBA, it is appropriate to introduce the innovation it is necessary to know three things. First, how much will the innovation cost? Second, how many lives, statistically, or in probabilistic terms, can it be expected to save? And, third, what, according to the relevant regulations, is the (monetary) value of preventing a fatality? The first two give the cost of preventing a fatality (CPF) for the project, which is to say how much the project costs for each statistical life saved. If installing a widget is going to save two statistical lives over the widget's lifetime, and costs £1m, then the CPF is £0.5m. This can then be compared with the VPF as set out in the regulatory framework governing the area under consideration. For example, the UK, for ordinary workplace risks, currently operates with a VPF of a little over one million pounds (Dept of Transport, 2004, Health and Safety Executive, 2001: 36).[ii] The US, I understand, uses a figure of around 6 million dollars (Sunstein, 2002: 165). Once these figures are known it is possible to give an answer to the question of whether the innovation should be introduced: it is necessary only to compare the CPF with the appropriate VPF. Current UK case law suggests that there is a requirement to introduce the safety measure unless the CPF is 'grossly disproportionate' to the VPF. In other words, in the UK if the innovation values saving a statistical life at much more than 1 million pounds there is no obligation to introduce it. Whether this gross disproportion test still applies is a matter of controversy, but in any case if the innovation costs less than the VPF then the regulations suggest that it should be introduced.

There are many questions that can be asked about whether and when RCBA is an appropriate way of approaching safety decision making. However, for the purposes of this paper I want to put most of these questions to one side. (For broader discussion see Wolff 2002 and Wolff forthcoming.) My topic here is the narrow one of simply trying to understand what 'the value of preventing a fatality' really means, and whether it can be morally defensible to operate with different VPFs for different circumstances.

2. LIFE-SAVING AND RISK-REDUCTION

For many people RCBA is an astonishingly cold and heartless way of making decisions about safety, as it appears to put a financial value on life. Defenders of the method deny that it has this implication. The Health and Safety Executive, in their major work on this topic, *Reducing Risk, Protecting People,* say the following:

> VPF is often misunderstood to mean that a value is being placed on a life. This is not the case. It is simply another way of saying what people are prepared to pay to secure a certain averaged risk reduction. A VPF of £1,000,000 corresponds to a reduction in risk of one in a hundred thousand being worth about £10 to an average individual. VPF therefore, is not to be confused with the value society, or the courts, might put on the life of a real person or the compensation appropriate to its loss (Health and Safety Executive, 2001: 65n).

This is a very interesting distinction. It is clear that purchasing – or declining to purchase – devices or services that make small differences to one's safety, or that of one's family, is an ordinary part of life, taking its place alongside other consumer decisions. Purchasing a smoke alarm, or deciding that it is not worth paying extra for a modification to a car to make it a somewhat safer, are decisions traded off against other consumption options.

Consequently we have to understand that the concept of saving of a statistical life is a different idea to the saving the life of a particular named individual. Cases where we are faced with saving the life of known individuals, such as miners trapped underground, or sailors in a stricken submarine, are, thankfully, rare and when they happen authorities rarely even raise issues of the cost of rescue. A statistical life is quite different: it is an accounting fiction pieced together from a patchwork of mitigations of very many low probability risks, each of which is equivalent to a perfectly ordinary economic transaction, and pervasive throughout life.

It has also been argued that in addition to the distinction between saving a statistical life and saving the life of an identified person, we need to bring in a third category: saving an anonymous life. This, it has been argued, is also quite different from saving a statistical life. (Jones-Lee, 1989: 174) Compare the following two situations involving a population of one million. In the first, everyone has an independent one in a million chance of dying from a particular hazard. In the second, one person – a person we cannot identify in advance – faces certain death from some other hazard. It is not unreasonable to say that for each individual in this population of one million there is an equal statistical chance of death in each case. Yet the situations are quite different. In the first any number of deaths from the hazard – from zero to one million – are possible, although one death is the most likely outcome, whereas in the second there can only be one death from the hazard.

However, the distinction between a statistical death and an anonymous death can become problematic, in at least two ways. The first problem is that for a particular individual there is no difference in practice between the two situations: in each case for all I know I am personally faced with a risk of death of one in a million. So although for the policy maker, or outsider, there is a clear difference, for those affected there is none, when considering only their own plight. Second, if the statistics are reliably based on regular past frequencies, then the distinction comes close to collapse. Consider road deaths in the UK. For some time around 3,500 people have died on the roads each year. How many will die next year? Theoretically it is possible that no one will, or millions will, but it is a very safe bet that the number will be between 3,000 and 4,000, and probably this could be narrowed down quite a bit more with a high degree of confidence. We might, then, think of this situation as closer to one in which we know how many people will die: we just do not know who they are yet. If so the distinction between anonymous deaths and statistical deaths is becoming blurred, even though, strictly these are statistical deaths. In policy terms, then, it becomes unclear how much difference this distinction makes.

However there is a further important distinction. In some cases, the chances of different individuals dying are not independent. Consider an incident such as a nuclear explosion from a power plant. I do not pretend that the following figures bear any relation to reality, but imagine that the risk assessment suggests that such an incident is likely to happen once every hundred thousand years, and if so it is most likely to kill one hundred thousand people. Once again, statistically, the expectation is one death per year, but it is far from obvious that the same framework of evaluation should apply to what we might call ‘catastrophic’ events, as to cases where risks to lives are independent and thus statistically one can expect a regular stream of smaller numbers of deaths.

We may, then, feel that there is a clear policy distinction to be made between ‘steady stream’ cases and ‘catastrophic’ cases and a different approach is needed. And it is does seem to be true that we have different intuitions about road safety and nuclear safety. Unfortunately, however, cases rarely differ in only one respect, and it can be a subtle question to work out what is driving intuitions in different cases, or even to detect what those intuitions are. Should we be more cautious about major nuclear accidents because they involve so many people, and have an impact beyond loss of life, or less cautious because they are likely to be so rare?

We will briefly return to catastrophic events later in the paper. At this stage I merely wished to raise some complications as a way of bringing out the central use of the concept of the value of preventing a fatality. As understood by regulators it is really a rather misleading way of talking about the value of reducing the aggregate probabilities – probably independent probabilities - of the risk of death for a group of people, rather than a way of talking about saving the lives of particular individuals whether known of anonymous.

3. DERIVATION OF VPF FIGURES

As should be clear, the RCBA process needs a value – or possibly a set of values – for the VPF. I mentioned that the UK currently operates with a figure of around one million pounds and the US about 6 million dollars. This variation could reflect different social priorities and values, different methodologies, or be a matter of pure contingency. Yet the figures on their own are mysterious, and a pressing question is how they are to be derived. In the early days of such calculations, the standard methodology seems to have been to base valuations on lost potential economic contribution. That is to say, human beings were regarded as 'human capital' and thus a source of potential income. The value of preventing a fatality is therefore equivalent, on this view, to the cost of losing that potential income. This, clearly, has some unfortunate effects: on this method anyone who is economically dependent, such as the old and the unemployed could have a value so low it could be negative, and if they are ill and in need of health care it could be significantly so (Mishan, 1971).

This idea that the 'value of preventing a fatality' or, equivalently, the 'value of a statistical life' is somehow based on 'society's investment' in that individual, or the opportunity cost of lost production, or some calculation based on both, is still common. Consider the following from Jared Diamond's otherwise excellent recent book *Collapse*

> The value of 'one statistical life' in the U.S. – i.e. the cost to the U.S. economy resulting from the death of an average American whom society has gone to the expense of rearing and educating but who dies before a lifetime of contributing to the national economy- is usually estimated at around $5 million. (Diamond, 2006: 504).

This, however, is to confuse the older methodologies , which typically yield a much lower average figure, with new methodologies which yield figures in the range Diamond mentions, but in which investment or lost production costs figure only, at most, as a minor element. Indeed, the idea that we should think of the value of a statistical life in terms of the investment in that life or its possible contribution to production has long been regarded as unsatisfactory. (Hammerton et al 1982: 182) One reason, of course, is the peculiarity of the results it achieves, including negative value, which we have already noted. A deeper reason is that it conflates the distinction insisted upon by the HSE between the reduction of a small risk for a large number of people, and the economic consequences of saving of a life. Even someone who is old, and has no productive contribution left to make to the economy, may still be prepared to pay good money to reduce a small risk of their own death, and even if they are not prepared to spend money this way 'we' as a society may consider doing so on their behalf. Reducing their risks of death is a benefit to them, even if doing so potentially increases total social costs in terms, say of healthcare.

When we understand the value of preventing a fatality in terms of risk reduction, rather than life saving, and we also realise that paying to reduce a small risk is a perfectly ordinary commercial transaction, then the possibility of new methodologies of valuation open up. The promise is that valuation of VPF can be conducted in terms of willingness to pay: the payments people do, or would, make in the market. The payments people *do make* yields the methodology of revealed preferences, and those people *would make* yields the methodology of expressed preference. The former looks at actual market behaviour to as a guide to valuation, the second uses purely hypothetical 'willingness to pay' methods, normally known as 'contingent valuation'.

Both methods have their attractions and drawbacks. Consider, first, revealed preference methodology. This has the advantage of looking at actual market behaviour: the decisions people have actually made in real markets, spending real money that they could spend on something else. The disadvantage is something it shares with any attempt to deduce people's underlying attitudes from their behaviour: what in Philosophy of Mind is called 'the holism of the mental' (Davidson, 1963). Given that both desire and belief play a role in the explanation of actions, it has been argued that every action is compatible with every desire if the surrounding beliefs are adjusted. A person's drinking hemlock is consistent with that person's desire to die, but also to live a long life, if they believe that hemlock is some sort of vitamin boosting drink. In the present context, people might choose to buy a dangerous product, believing it to be safe. Nothing, though, is revealed about that person's attitude to risk from this decision. Furthermore, it far from always the case for an action to be performed for the sake of a single desired goal. So, for example, it is very unlikely that the only difference between two complex products – two cars, say – is that one is safer than the other. The fact, then, that I have chosen to purchase a safer car does not show that I am prepared to the pay that premium purely for the risk reduction. It may be, for example, that I also prefer the image associated with one make of car rather than the other.

Consequently valuations based on revealed preferences are highly uncertain, and there is a huge variation in the figures such studies do in fact reveal. Sunstein, for example, sets out a list of valuations based on labour-market studies, which essentially explore the wage premium needed to recruit people into more dangerous jobs. These studies produce a VPF which varies from $0.7 million to $16.3 million, in 1997 dollars. (Sunstein, 2002:174) Although averages of many studies may be helpful, it is hard to know when a sufficiently broad or representative range of studies has been conducted. For all these reasons some theorists prefer to follow the method of expressed preferences or contingent valuation. The basic idea behind this methodology is that subjects are asked for their views of how much they would pay for a safety improvement, if one were available.

There are at least two main advantages for this methodology. First, the experimenter can set up questions in such a way to ensure that the subject must focus on the safety element alone in the choice. In a hypothetical example all other parameters can be fixed in a way in which this will rarely, if ever, happen in a real market case. Second, a given subject can, in theory, be asked many questions, and so a great deal more data can be generated.

In practice, however, such benefits are rarely seen. There are several limitations. First, there are framing issues. It is well known that if people say that they would pay a particular amount of money to avoid a risk, in general they will also say that they would need a higher sum of money in compensation if the feared event happened. This requires explanation, as many versions of decision theory predict that a given individual's 'willingness to pay' and 'willingness to accept compensation' should be the same. To get a sense of this imagine the highest price you would pay for a pair of tickets to some highly attractive event: the world cup final or a star-studded opera. Imagine you have just paid for the tickets and have them in your hand. Now, what is the lowest price you would sell them for? If there is a non-trivial gap between the two sums – as many people report there is – it shows that there is a difference between your willingness to pay and your willingness to accept compensation, and raises the question of what we should say is your 'price' for these tickets.

Second, given that in contingent valuation no money actually passes hands, it is unclear how seriously we can take the figures offered. There is a legitimate worry that some people are simply plucking numbers out of the air, rather than revealing willingness to pay.

Thirdly, and most seriously of all, human beings appear to be very poor at rational decision making involving very small probabilities. Subjects can very easily be led to make inconsistent decisions: for example being willing to pay more if a risk reduction is broken down into two steps rather than one, even if the resulting outcome is the same, or paying the same for a larger reduction and a smaller one. In sum, when subjects are asked to express their preferences concerning paying to avoid risks with small probabilities, very little reliable data is generated. Consider again the HSE example where a VPF of £1,000,000 is said to be equivalent to a payment of £10 to avoid a one in a hundred thousand chance of death. If asked 'how much would you pay to avoid a one in a hundred thousand chance of death?' most people, I think, would not feel that they could give a robust or reliable answer. Furthermore, when asked a number of different questions it seems very unlikely that many people would give a consistent set of answers.[iii]

To avoid these problems some experimenters have adopted modified techniques. One type of study, for example, asks subjects about their attitudes to risks of a greater probability. Michael Jones-Lee and associates developed one extremely inventive approach. (Carthy et. al. 1999) To simplify the description of a rather more complex method, the essence seems to come to the following. The subject is first told about a possible motor accident which will leave them hospitalised for a certain period, but from which they would make a full recovery. They are then asked how much they would pay to avoid such an accident, and how much compensation they would need to be prepared to accept that the accident had not left them worse off overall. These sums are then combined, by means of a standard calculation, to derive a single sum for the value of avoiding that accident. In the next phase of the experiment the subjects are asked, in effect, what risk of death would be less preferred than the accident. That is, if they choose option A they will have the accident so described, but if they choose option B there will be a risk of dying, but otherwise one survives unharmed. The question then is when does the risk of dying become unacceptably high in the sense that one would prefer the certainty of the accident to the risk of death. Once this question is answered it then becomes possible to calculate, for that subject, the value of preventing a fatality. The two inputs are the value of preventing the accident, together with the maximum risk of death one would 'trade' the accident for.

This method does, then, yield a determinate answer for each individual. The range of answers is very diverse, However, rather than discuss the range, which I have done elsewhere (Wolff, 2002: 57-69), I want to point out the way in which this methodology, at first sight, blurs the two questions separated by the Health and Safety Executive, and in doing so makes a particular assumption about 'linearity'. Indeed, the point may be obvious. In order to obtain meaningful data about the risks people are prepared to run, or to pay to avoid, subjects have to be asked about risks of fairly high probability. Yet the VPF figure is intended to represent the value of reducing a very large number of low probability risks to many people, which together add up to the saving of a life. This, then, presupposes linearity in attitudes to risk: that, say, that I would pay no more or no less than a thousand times as much to reduce a risk that was a thousand times as great, at least if I am rational. Yet this seems problematic.

To see the issue more clearly, consider the highly artificial example of being invited to take part in a game of Russian Roulette, for a payment. The gun, however, has a variable number of barrels, and let us also assume that there is no reason to doubt that it is a fair gun in the sense that the probability of any barrel being engaged is equal, whether or not it contains the bullet.

Now, I assume that for any person there is some number of barrels and some amount of money that would be sufficient to entice the person to accept the risk. So, for example, imagine that the gun has two million barrels, and hence one's chance of dying was two million to one. I imagine most people would be prepared to accept a finite sum of money – even a rather small amount - to run this risk, at least if it was presented to them in the right way. At the other extreme, imagine you are being asked to play the game with a gun with only two barrels, and hence there is only a fifty percent chance of survival. If attitudes to risk were linear than you should be prepared to accept a price for this of one million times the smallest price you would be prepared to accept for the one in two million chance. And, indeed, it may well be. But it seems hard to say that it is irrational for someone who would have taken the small risk for one pound to decline to accept one million pounds for a fifty percent chance of death. Of course some people may be happy to kill themselves, and even for those with a preference for life an exceptionally high sum could tempt them to risk their own life so they if they do survive they will have a life of luxury. Nevertheless, it would be very strange to argue that someone is irrational if they refuse to accept this price. For many people, I assume, there is simply no price which would tempt them to take a fifty percent chance of death.

Consider now something in between. Suppose the gun has twenty barrels. If the price was high enough more people might be tempted to accept the offer. But must this price be exactly a hundred thousand times as much as they price accepted for taking a one in two million risk? It may be an axiom of some forms of rational choice theory – although by no means all - that attitudes to risk should be linear in this sense, but it seems to be no part of common sense.

Whether these arguments put the Jones-Lee ‘chained’ methodology into question depend on whether the risks people are invited to consider in these experiments fall into a non-linear part of the curve. If so, then the studies would put too high a price on safety, assuming that people are prepared to pay proportionally more to avoid risks with very high probabilities. However, the values derived from this work may suggest that, actually, there is little to worry about on this score, although further investigation would be of help.

We need also to consider a suggestion from Richard Posner that linearity also fails at the other end of the scale, when we consider catastrophic risks with very low probabilities. (Posner, 2004). He may well be right about this, although it is very hard to gain firm evidence. One type of revealed preferences study could conceivably demonstrate a lack of preparedness to pay to reduce risks with very small probabilities. These would be studies of what governments are prepared to pay on their citizens’ behalf. Although he does not use it to demonstrate failure of linearity at miniscule probabilities, Posner discusses the example of possible asteroid collision. He quotes a UK task force which estimated that the chances of an asteroid collisions that would kill 1.5 billion people is 1 in 250,000 in any year. Using a VPF of $2 million, he calculates that RCBA would justify spending $12 billion a year to prevent an asteroid collision (Posner 2004: 180). At present the US government spends $3.9 million tracking large objects in near space: a mere eight hundred dollars for each statistical life. If we think it would be inappropriate to spend much more than this, then that appears to imply that either we think the probabilities are not as stated, or that we should use a very low VPF for such a risk, or that, somehow, the methodology of RCBA is not appropriate for risks of this nature. Perhaps we think all three of these things.

Matters, inevitably, are complicated. The risk of a natural disaster is quite different from the risk created by a project undertaken for human gain or benefit. Imagine a new, much more powerful, form of nuclear power. Suppose we are told that such a technology risks killing 1.5 billion people every 250,000 years: the same probability as an asteroid collision. It seems unlikely that we would be so phlegmatic about accepting the risks. Possibly the reason for this is that a nuclear power station is a risk we feel we do not have to take. Furthermore it creates profits which can be used to pay for safety improvements. By contrast, an asteroid collision is a risk that has not been created by human action, and does not generate a stream of revenue which can be used to pay for safety measures. Accordingly we may decide not to divert spending away from other valuable projects, and simply run the risk rather than reducing our present standard of living to avoid something that may actually never happen. We might also judge that the risk for the next 100 years is really too small to worry about, and that future generations will be in a position to defend themselves at lower costs. Because of such complications it seems clear that we might need to take into account a whole range of factors before we can decide what it is our social preferences reveal about our attitudes to risks with small probabilities. Incidentally, this is a powerful illustration of the problem of revealed preference methodology: any actual behaviour could have innumerable explanations.

While it is hard to decide what would count as evidence that our attitudes to risk are non-linear, except in the case of risks of very high probabilities, this is not evidence that they are, in fact, linear. We have to be prepared to consider the possibility that attitudes to risk are non-linear. If they are, in fact, non-linear, and we accept that our regulatory regime should be based on attitudes to risk, then it follows that we would have to accept different VPFs for different types of situation. Those which involve very small probabilities would demand very small VPFs, but where we get close to certain death the VPFs should be high. This, in fact, is one way of understanding what we actually do in 'rescue cases' where a child has fallen down a well, or miners are trapped underground, or sailors in a submarine. As mentioned above, often we seem to be prepared to spend without thought of whether it can be justified by RCBA with a standard VPF.

In effect the Health and Safety Executive itself has accepted non-linearity, by adopting what it calls a 'Tolerability of Risk Framework', distinguishing 'unacceptable', 'tolerable' and 'broadly acceptable' regions of risk (Health and Safety Executive, 2001: 42-46). A risk of death is unacceptable if it is above 1 in 1,000 per year for the workforce and 1 in 10,000 per year for the public who have a risk imposed on them, and is 'broadly acceptable' and thus requiring no special reduction measures if it falls below 1 in 1,000,000 per year. This is a type of 'limit' case of non-linearity, in which risks higher than 1 in 10,000 (for the public) are given infinite VPFs and those below 1 in 1,000,000 are given no VPF at all. Although crude, this may be the most appropriate approach for public policy for non-catastrophic risks, although it does have rather startling implications. A hazard to which the whole UK population is exposed might cause 50 deaths each year but would fall below the 1 in 1,000,000 threshold, and so, strictly, would not require any money at all to be spent on further mitigation steps.

4. THE VARIABLE VALUE OF PREVENTING A FATALITY

Once it is accepted that, in principle, there can be variable VPFs for different probabilities of death, this raises the question of whether there can be variation for other sorts of reasons. Consider the risks involved in railway safety. The ‘headline’ accident figures concern the number of passengers who are killed in train accidents. However this is only a small proportion – probably less than 3% on average - of the number of people who die on the railways. About two thirds of all deaths on the railways are suicides. Deaths of trespassers – adult and children – constitute the second largest group, with passengers and then workforce, other than in circumstances of train accidents, making up the rest. Although the industry, along with its government funder and regulator, has not officially adopted a variable VPF to cover these different cases, one may conjecture that unofficially it does, in that it is prepared to spend much more to prevent some categories of death than others. Passengers, workforce and child trespassers probably have higher priority, and adult trespassers and suicides lower.

One can understand safety decision makers being rather nervous about announcing different VPFs for these different categories. One can imagine the response: ‘are you saying that the life of a trespasser is less important than the life of a passenger?’ And, of course, no one wants to say that, even though this is how different VPFs for different groups would inevitably be interpreted. Yet it is important to see that variable VPFs do not have this implication when we understand that VPFs are simply a shorthand for accumulation of many small risk reductions. The question is whether it is right to pay more to reduce the risk to one group than it is for another. How can we answer this?

It may be helpful to reflect on more ordinary safety decision making. It seems plausible that people are prepared to spend more to reduce risks to their children than they are for themselves. One reason is that children are generally in more danger, being less able to take precautions for themselves. But putting that to one side, there is also the fact that parents feel themselves to be in a special position of responsibility with respect to their children. Finally people simply care a great deal about children, and particularly their own. Similar types of reasons can operate in the case of railway safety. The industry may rightly feel that it is in a special position of responsibility towards its passengers and its workforce, and much less so to adult trespassers and would-be suicides. This is not to deny all responsibility, but to say that, at the least, responsibility is shared. Child trespassers are in a different category, as they are far less able to exercise the right sort of judgement to take responsibility for themselves.

As soon as it is said that the industry has a greater responsibility to reduce risks to one group than to another, and should pay more to do so, then by the definition of VPF it follows that they are placing a higher VPF on one group than another. Yet one can understand a reluctance to admit this, for the reasons given: how can we value one person's life more than that of another? However this is largely a terminological problem, based on the confused idea that VPF puts a value on life. This, it seems to me, is a good reason for reconsidering terminology. The misleading content of the idea of the value of preventing a fatality is getting in the way of clear thinking. Perhaps it is time that this term is abandoned and replaced with another. We shall look at one such proposal shortly.

Before doing so, we might ask what the warrant would be for using different VPFs for different groups. One reason why this topic is something of a taboo is that if such differences are based on willingness to pay methods, we will find different VPFs for the rich than for the poor, as the rich are typically prepared to pay more for risk reduction, as for almost everything, than the poor. This outcome is avoided by averaging strategies, which are generally treated as justified.[iv] In the cases discussed in this paper the rationale for different VPFs is likely to be derived from judgements about higher and lower responsibility, rather than willingness to pay for different types of risk reduction, although there is room for conducting research, using willingness to pay methods. Clearly there is more work to be done here.

5. THE VALUE OF REDUCING A STANDARD RISK OF DEATH

In the last section I suggested that we might wish to rid ourselves of the confusing terminology of the value of preventing a fatality. Consider a new term: the Value of Preventing a Standard Risk of Death (VPSRD). This, clearly, is parasitic on the idea of a Standard Risk of Death. There are various ways in which this, in turn, could be approached. One suggestion would be to operate by orders of magnitude. SRD1 would be certain death. SRD2 would be a one in ten chance of death, SRD3 one in a hundred, and so on. However this is probably more complex than is needed. For regulatory purposes risks between one in ten thousand and one in a million are most relevant. It would be very convenient if we could also assume linearity in this area, which I will do for present purposes, while acknowledging that if it fails modifications to what follows would be necessary. This would then suggest that we could take a risk of one in a million per year as the base unit for a Standard Risk of Death. Translating current UK values into this terminology would suggest that, if we take a VPF of £1,000,000 then VPSRD would be £1. This would then be multiplied by the number of people involved and the probability of death as a multiple of one in a million to yield the financial value of the safety benefit under consideration. So far this would be identical in its workings and consequences to standard RCBA using a VPF figure. The only benefit so far would be transparency: that reduction of risk rather than saving of lives is at made more obviously at issue. However, once this step is made, it becomes possible for an industry to decide that reducing risks in one category – passengers, say – is a higher priority than another – adult trespassers. Hence it might decide to use standard figures for reducing risks to adult trespassers, but put a premium on reducing risks to passengers and to child trespassers. And once again, this could be defended not on the grounds that one category of life is more important that another, but because the industry has a greater responsibility to guard against some types of risk rather than another.

6. CONCLUSION

My purpose in this paper has been to consider the beneficial consequences of taking seriously the distinction between saving lives and reducing risks of death. Some of the objections to risk cost benefit analysis seem less pressing once this distinction is made clear, for it allows us to see safely legislation as more continuous with ordinary safety decisions made by consumers. It has the further benefit of it paving the way to opening up a clear debate on whether more money should be spent to reduce risks which fall into one category rather than another, even if they are of the same magnitude. Whether or not it is sensible to attempt to change standard terminology, there seems little doubt that the existing language makes the methodology of RCBA seem more vulnerable than it need be.[v]

Beattie, J., Covey, J., Dolan, P., Hopkins, L., Jones-Lee, M., Loomes, G., Pidgeon, N., Robinson, A., and Spencer, A. (1998). 'On the Contingent Valuation of Safety and the Safety of Contingent Valuation: Part 1-*Caveat Investigator*', *Journal of Risk and Uncertainty*, 17:5-25.

Carthy, T.S., Chilton, J. Covey, J., Hopkins, L., Jones-Lee, M., Loomes, G., Pidgeon, N., and Spencer, A. (1999) 'On the Contingent Valuation of Safety and the Safety of Contingent Valuation: Part 2 – The CV/SG "Chained" Approach'. *Journal of Risk and Uncertainty*, 17: 187-213.

Davidson, Donald (1963) 'Actions, Reasons, and Causes'. *Journal of Philosophy*, 60: 685-700.

Department of Transport (2004) *Highways Economic Note 1*. Online. Available HTTP: <http://www.dft.gov.uk/stellent/groups/dft_rdsafety/documents/page/dft_rdsafety_610642-01.hcsp> (accessed April 18, 2006).

Diamond, Jared (2006) *Collapse*, London: Penguin.

Hammerton, M., Jones-Lee, M.W. and Abbot, V. (1982) 'Consistency and Coherence of Attitudes to Physical Risk', *Journal of Transport Economics and Policy* 16: 181-199.

Health and Safety Executive (2001) *Reducing Risks, Protecting People* London: HMSO. Online. Available HTTP <http://www.hse.gov.uk/risk/theory/r2p2.htm> (accessed April 18, 2006).

Jones-Lee, Michael (1989) *The Economics of Safety and Physical Risk* Oxford: Basil Blackwell.

Mishan, E.J. (1971) 'Evaluation of Life and Limb: A Theoretical Approach', *Journal of Political Economy*, 79: 687-705.

Posner, Richard (2004) *Catastrophe* New York: Oxford University Press.

Sunstein, Cass (2002) *Risk and Reason* Cambridge: Cambridge University Press.

Wolff, Jonathan (2002) 'Railway Safety and the Ethics of the Tolerability of Risk'. Rail Safety and Standards Board. Online. Available HTTP: http://www.rssb.co.uk/pdf/policy_risk.pdf (accessed April 18, 2006).

Wolff, Jonathan (forthcoming) 'Risk, Fear, Blame, Shame and the Regulation of Public Safety', *Economics and Philosophy*.

Wolff, Jonathan, and Haubrich, Dirk (forthcoming), 'Economism' *Oxford Handbook of Public Policy* Oxford: Oxford University Press.

---

[i] Safety measures come into conflict, for example, with aesthetic values, such as in the modification of existing buildings. It is not obvious that this conflict can be translated into economic terms without loss. For some related discussion see Wolff and Haubrich (forthcoming).

[ii] For some purposes a further 'lost output' figure of around £500,000 is also added.

[iii] Such difficulties were observed in Beattie et al (1998), in which advocates of the contingent valuation model face up to some of the difficulties of the methodology as a

preliminary to exploring a way forward. They note: ‘Amongst the more worrying … anomalies and inconsistencies are so-called “embedding”, “scope” and “sequencing” effects. Essentially, embedding and scope effects refer to the tendency of many CV respondents to report much the same willingness to pay for a comprehensive bundle of safety or environmental “goods” as for a proper subset of that bundle. In turn, sequencing effects reflect a tendency for the order in which a sequence of CV questions are put to respondents to have a significant impact on the values that are implied by the responses to such questions.’ (p. 8)

[iv] It is important, however, to consider how the costs for safety measures are to be met. If safety costs come out of a public budget, funded by progressive taxation, then there is every reasoning for using averaging strategies as a form of distributive justice. However, if the price of safety goes on to the price of a ticket, then averaging will lead to the poor being forced to buy a level of safety they would not have chosen to. My conjecture is that at least some people will then try to return their level of risk to their preferred level by, say, every now and again evading the fare. Hence the breaking the law can function as a type of risk ‘black market’. I hope to make this the topic of future work.

[v] I am very grateful to Shepley Orr, Donald Franklin, Dirk Haubrich, Jesse Norman, Mike Jones Lee, Graham Loomes, Mike Robertson and Andrew Sharpe for discussions which have helped shape this paper. Written comments from Shepley Orr, Andrew Sharpe and Tim Lewens on the penultimate draft have been especially helpful.