

RISK, FEAR, BLAME, SHAME AND THE REGULATION OF PUBLIC SAFETY

JONATHAN WOLFF*

University College London

The question of when people may impose risks on each other is of fundamental moral importance. Forms of “quantified risk assessment,” especially risk cost-benefit analysis, provide one powerful approach to providing a systematic answer. It is also well known that such techniques can show that existing resources could be used more effectively to reduce risk overall. Thus it is often argued that some current practices are irrational. On the other hand critics of quantified risk assessment argue that it cannot adequately capture all relevant features, such as “societal concern” and so should be abandoned. In this paper I argue that current forms of quantified risk assessment are inadequate, and in themselves, therefore, insufficient to demonstrate that current practices are irrational. In particular, I will argue that insufficient attention has been given to the cause of a hazard, which needs to be treated as a primary variable in its own right. However rather than reject quantified risk assessment I wish to supplement it by proposing a framework to make explicit the role causation plays in the understanding of risk, and how it interacts with factors which influence perception of risks and other attitudes to risk control. Once an improved description of risk perception is available

* This paper was drafted as part of the project “Philosophical Foundations of Public Policy: Rethinking Cost-Benefit Analysis” funded by the Arts and Humanities Research Board under its Innovations Award Scheme. The final version was written up as part of an extended leave grant also from the AHRB. I am very grateful for the Board’s assistance. Earlier drafts of this paper were presented to the Society of Applied Philosophy in Manchester, and at Lancaster University, and I am very grateful to the audiences at both occasions for their discussion and comments. I am also pleased to thank Andrew Sharpe at the Rail Safety and Standards Board for bringing this topic to my attention, and his encouragement and useful comments on earlier drafts of this work. Finally I am very grateful to the referees and editors of *Economics and Philosophy*, and especially Luc Bovens, for suggestions which have led, I hope, to significant improvements in this paper.

it will become possible to have a more informed debate about the normative question: how safety should be regulated.

1. INTRODUCTION

When can people knowingly impose risks on each other? This question is central to such issues as transport safety policy, environmental risk, and health and safety at work. In many cases decisions will be taken by politicians and company directors, advised by engineers and economists. Different understandings and different approaches may well be followed in different cases. This gives rise to at least two causes for concern. First, as is well known, an uncoordinated approach may give rise to serious inefficiencies, in that it may be possible to use existing resources in far more effective ways. The second is also well known, if less often discussed. The question of what risks we can impose on each other is fundamentally a moral one: the proper aims of safety regulation and the proper means of achieving these aims are moral issues. It is unclear that between them politicians, company directors, economists, and engineers will be properly attuned to all the relevant moral issues. These two concerns come together in the following way: unless we have a firm understanding of the morally justified aims of safety regulation we cannot say whether any particular measure is rational or irrational, efficient or inefficient.

The purpose of this paper is to lay the groundwork so that the moral questions appear in clear focus. I will not, however, attempt to settle these questions here. Rather, in order to understand the proper aims of safety regulation we need to have a good understanding of its possible aims, or, better, the plausible aims. Eventually this should allow an evaluation of such possibilities. My strategy will be to look at the way in which safety decisions are influenced, made, and perceived in practice, in order to bring out the diversity of factors which come into play. Accordingly, the remainder of this paper falls into three main parts. Section 2 lays out the background to the issue, looking at the standard paradigm of quantified risk assessment, and the argument that some current practices can be seen to be irrational in such terms. Section 3 introduces and explores a number of complicating factors – summarized as fear, blame, and shame – while Section 4 produces an overall framework in which it can be seen how these differing factors relate to each other. The main theoretical proposal is to add the concept of “cause” to that of hazard and probability, as primary variables in the analysis, and to show how this inter-acts with fear, blame, and shame which are here treated as “secondary” variables. One conclusion is, perhaps, a predictable one: that greater clarification of the proper aims of safety regulation is necessary before we can argue that any particular safety measure should be denounced as irrational or inefficient. But the main pay-off of the analysis is the beginnings of a

model of the “anatomy of risk” which sets the groundwork for informed normative discussion of the proper aims of safety control and regulation.

2. QUANTIFIED RISK ASSESSMENT

2.1 Safety has a price

Life is a risky business. We all face threats to life and safety every day. Some circumstances, and especially some working environments, seem especially risky. What should be done about this? What should be done, for example, about hazardous working conditions?

As a first thought, it might be proposed that it is always wrong knowingly to inflict risks on others, and so there should be an absolute duty on factory owners and others to eliminate all known risks to their workers. But a moment’s reflection shows that this is an impossible aim. Virtually any human activity involves some risk. Even a perfectly maintained and serviced machine might malfunction with unpredictable effects. Even very sensible workers can trip on an even floor. These are risks we know about, yet cannot eliminate entirely.

However, even though we cannot eliminate all risks, we might be able to reduce many of them. So perhaps the goal should be to reduce risks in so far as this is technically possible. But this again seems to have some absurd consequences. We might virtually eliminate fatal road accidents by lowering the speed limit to 10 miles an hour. We can end injuries to coal miners by closing the mines. Neither suggestion would be treated as a serious contribution to safety policy debate. Safety – and therefore life and limb – is not the only thing we value, nor, it seems, is it always the highest value.

The lesson is that safety has a price, in terms of its impact on other things we want or value, and there are limits to what we are prepared to pay. It seems that in generating policy we are forced to put a value on life – and this inevitably seems to mean a financial value – which helps us generate rules about how much firms, and in some cases the government, can reasonably be expected to pay for safety improvements.

This may seem callous or inhuman. Don’t we know that life has infinite value? But what is the alternative? Not putting a price on safety? Allowing companies to operate with dangerous machinery, because we can’t put an infinite value on life and any finite value is arbitrary and demeaning? This hardly seems an improvement.

2.2 The standard paradigm

Safety is regulated in somewhat different ways in different jurisdictions, although the differences in detail need not detain us. As it is the example I know best I shall take the UK as my main example. Work-related safety

is regulated by the Health and Safety Executive, and its general approach is explained in a publication called *Reducing Risk, Protecting People* (Health and Safety Executive 2001). Here, for simplicity, we will concentrate on risks of death, although other risks are covered too. In standard cases the basic approach is to divide risks into three categories. Some risks of death are too high, in probabilistic terms, to allow and must be reduced (unless there are special circumstances). Some are so low (in the sense that the probabilities are minute) that they do not require any special measures. In the large middle ground are risks which although in some sense are “broadly tolerable,” should be reduced “as far as is reasonably practicable” (44ff.).

If a risk falls within the broadly tolerable region those in control of the risk are required to perform a risk cost-benefit analysis. To carry this out risk assessors must calculate the probability of death for the risk under consideration. Suppose, for example, a piece of machinery could trap and kill a careless and negligent worker and that machines of this type kill 1 in 10,000 of their operators every year. And, as ease of calculation would have it, in your large factory you have 1,000 operators. Hence you should expect a death every 10 years or 0.1 death a year, assuming that there is nothing special about your factory.¹

Let us suppose there is a possible modification to the machine which could reasonably be predicted to eliminate half the deaths in your factory, thus saving 0.05 lives a year. Should you introduce the modification or not?

For simplicity let us assume that you believe that the machines will be in use for another 10 years, and let us also apply no discount rate for future deaths. To know whether to introduce the modification two further pieces of information are required. First, how much the modification will cost, and second, what financial value should be placed on preventing a fatality (VPF). Currently the UK operates with a figure of a little over £1 million (£1 million at 2001 prices, and hence a few thousand more now, allowing for inflation). How to calculate such a figure is a matter of some controversy, which I shall not enter into here, but for the purposes of this paper it makes no difference what figure is selected.² So let us round down to 1 million. Consequently as the modification will save 0.5 of a life over ten years you would be required to introduce it if (and only if) it will cost less than £500,000.

¹ Where there are frequencies of this nature probabilities are relatively easy to estimate. Of course this is an unusual case, and especially in the case of new risks there can be great controversy about actual probabilities.

² For an illustration of one important method for determining such a value see Jones-Lee *et al.* (1999). In the US different figures are used. According to Richard Posner, current estimates range from \$4 million to \$9 million with a mean of \$7 million. Posner (2002: 166). For further discussion of the US approach see Sunstein (2002: 153–90).

This general approach is intended to apply to all work-related risks in the UK. It was first devised as a response to the risks of nuclear power stations. It also, for some reason, applies to the risks of railway travel, and the location of airports, although not to road travel, and not for product safety. However, it could be easily be adapted to these cases too.

2.3 The irrationality argument

What I have described is one application of risk cost-benefit analysis, which in turn is a form of quantified risk analysis. Many people find these types of approaches rather chilling, at least at first. However risk cost-benefit analysis can be used as a very powerful tool for examining current practice. One familiar debate concerns the contrasting situations of rail and road safety.

In recent years the UK public has become very concerned about railway safety, and in particular about train crashes. How many passengers die in railway accidents in the UK each year? In a recent, as yet unpublished, study around 1000 people were asked this question. Their answers lay in the range 10 to 2,000, with a mean, excluding outliers, of 99. In the last decade the actual average number was about six deaths of passengers per year. Passenger deaths, in fact, are only a small proportion of deaths on the railways in the UK. Over recent years the annual average number of deaths on the railways as a whole is about 275, with the vast majority being suicides and trespassers. (Others include members of the workforce, people at stations, and occupants of vehicles on the line.)

Recently technology has been introduced, at a cost of about £585 million, to make it less likely that a train will run into another train if it runs through a red light. (Commission for Integrated Transport 2004: section 3) Arguably this is already saving lives, albeit at a cost of somewhere between £5 million and £20 million a life saved, depending on what is counted (whether or not injuries are counted as fractions of death) and how long a time period is taken. Further technological innovation – Automatic Train Protection – is being discussed which would make it theoretically impossible – i.e. impossible if the system works – for a train to run a red light. The system currently under active discussion, on the lowest estimate I have seen, would cost £3.6 billion. Even its defenders admit that it will cost close to £100 million for each rail passenger life saved.³

³ Commission for Integrated Transport, 2004. Some non-passenger lives, it has been claimed, would also be saved, although some workers may die installing it. There are further arguments on both sides claiming that the cost per life saved would be higher or lower. It is worth noting that currently most of those who defend the introduction of ATP do so on the basis of the commercial and performance benefits it is calculated to bring, and not on its contribution to safety.

Turning now to the roads, in recent years about 3,500 people die in the UK each year (the figures are remarkably stable), of whom around 2,000 are occupants of cars. By every measure travelling by road is much more dangerous than travelling by rail. While there are always reasons to be skeptical about any particular calculation, there are indications that the hundreds of millions recently spent on the railways could have saved perhaps ten times as many people on the roads. And for £3.6 billion miracles could happen, at least if road safety campaigners are to be believed.

To bring out the disparity between our attitude to road and rail safety consider the arguments of commentators who have discussed the aftermath of the Hatfield rail accidents. At Hatfield a high-speed train was derailed when the track it was travelling on shattered. Four people died, and others were injured. A very cautious response followed, and speed restrictions were enforced throughout the network so that all relevant track could be checked for similar faults. The resulting memorable chaos meant that train travel was unreliable to an unprecedented degree; it was as if there was no timetable. Frustrated passengers took to their cars. It has been estimated that there may have been as many as five extra road deaths in the first month as a result.⁴ Although the comparison may seem rather mischievous, it appears that in some sense we would have been better off with no speed restrictions and a Hatfield sized crash every second month, compared to what, it is claimed, actually happened.

These examples bring out a stark general message. We can easily save more lives by spending our resources in different ways. Indeed some analysts are raising the possibility (behind closed doors) that we should significantly reduce the amount of money we already spend on railway safety, diverting the resources to road safety, public health, or even foreign aid. Essentially the same argument is made in the US concerning consumer protection and environmental protection. Huge sums are being spent to mitigate tiny risks, while much larger risks go ignored (Sunstein 2002). It seems hard to avoid the conclusion that this is an irrational state of affairs. For this reason I shall call it the irrationality argument, and it is growing in popularity. So, for example, Bjorn Lomborg has notoriously argued that instead of spending trillions of dollars slowing down global warming by a few years, we would do better to spend a fraction of that money helping developing countries build the level of infrastructure that will

⁴ This figure is reported as an estimate, but without attribution, by Sunstein (2002: 2). However Sunstein's diagnosis of the change in behavior as resulting from individual over-reaction to small risks seems mistaken. Rather it was a perfectly rational response to massive disruption to the service, caused by the industry's highly cautious response to the incident. Further, his comment that these five road deaths is "nearly equal the total number of deaths from train accidents in the previous thirty years" is extraordinary. The true figure, according to Evans (2005), is close to 200.

allow them permanently to cope with effects of global warming (Lomborg 2003, drawing on Lomborg 2001). Although, no doubt, the science and economics are contestable, this line of argument – another application of the irrationality argument – can be made to seem quite compelling. Those who oppose it are portrayed as supporters of the politics of gesture and a dangerous menace to rational thought and even to life on earth.

I hold no brief for current practices, and I accept that thinking through the consequences of the irrationality argument can be liberating. Yet I will argue that we should not be quite so quick to think that the irrationality argument settles anything. There are subtleties which we need to investigate first. This is another way of saying that it is possible that risk cost-benefit analysis, as currently used, does not capture all the information that is needed in order to make the best decisions.

3. COMPLICATING FACTORS

3.1 Fear reduction

Each of us is afraid of some risks, but less so of others. What is the relation between risk and fear? To make progress we must distinguish objective risk and subjective risk, or in other words belief in risk. For there need be no relation between objective risk and fear. How can you fear an unknown risk? Well, of course you can fear the unknown, but there is no reason to think that anyone's fear will be related to the actual risk. Furthermore, it is well known that fear may be out of all proportion to the objective risk; this is the central finding of those who argue for the "social amplification of risk": essentially the idea that there are numerous social mechanisms which can make people feel that risks are much higher, or lower, than in fact they are (Pidgeon *et al.* 2003). But, more pertinently, is there a clear correlation between subjective risk and fear? Obviously they are not the same thing as one is a belief and the other an emotion. But there could be a causal, or even a partially constitutive, relationship. Let us, in the first instance, assume there is some sort of direct connection, although we will examine this shortly.

It seems clear that people may pay their own money, or agree to spend taxpayers' money, to reduce risks in order increase their sense of safety.⁵ Or, to put this another way round, it is reasonable to believe that people are prepared to pay more to reduce those risks they fear most. But notice if the point of risk regulation is to reduce each individual's subjective sense of being at risk then already we can see that the irrationality argument may not go through in such straightforward fashion. For the irrationality argument concerned objective risk. It argues that particular risk reduction

⁵ The distinction between death reducing policies and anxiety reducing policies is explained well in Schelling (1968).

policies are irrational because they are inefficient means of achieving the goal of reducing objective risk. The point about inefficiency relative to the goal of risk reduction can be granted, but irrationality need not follow. For if the point of safety policies is, or includes, the reduction of subjective risk – to increase feelings of security – then the failure to reduce objective risk is no longer decisive. Indeed, we can generalize this point. If the goal of safety policy includes anything other than objective risk reduction, then it is moot whether the irrationality argument goes through. Everything needs to be recalculated in the light of the new goals of policy.

Here, though, defenders of the irrationality argument may well change tack. The point, they will say, of safety regulation *ought to be* reduction of objective risk. How plausible is this? Note that this response need not downplay the importance of fear, anxiety, and insecurity in people's lives. Such emotions, it can be conceded, are terrible things to suffer. Perhaps they are much worse than the presence of small risks in one's life. After all, small risks rarely lead to actual harm, whereas fear and so on can have a constant dampening effect on one's spirits. But, so the argument goes, the way to respond to this is not to introduce expensive means of reducing what may already be barely significant risks. Instead, education is needed so that public fears track the real risks, and people worry about only what they ought to be worried about. False fears should be calmed by good information and the same means should be used to ensure that people come to fear the objective risks they face.

While this appears very attractive it nevertheless relies on some assumptions which may well be false. In particular, it relies on an intuitive assessment of the effectiveness – and hence the costs and benefits – of alternative policies. Changing public attitudes and emotions through provision of information is very difficult. Or rather, it is difficult to change public attitudes in a positive direction. It is expensive to attempt, and rarely more than marginally effective. Who can we rely on to provide accurate information? In the current climate people profess to distrust scientists, doctors, the government, bankers, big business, the police, the media, civil servants, lawyers, educationalists, anyone in the employ of the government and, indeed, anyone on a decent salary. In the light of this it is rather hard to see how anyone comes to any beliefs about anything. More pertinently, the prospects for a public education strategy which would bring subjective and objective risk into step seem pretty bleak. I am not proud of humanity for this, but it may turn out that once we do the sums, the most cost-effective way of reducing public anxiety could be to spend huge amounts of money on almost useless safety devices. Certainly anyone who has traveled by air lately, and seen what is being done in the name of reducing risks of terrorism may well have had the thought: obviously not much better than useless, but nevertheless somehow strangely reassuring, at least for some people. But in fact the practice of symbolic safety measures

to reduce fears is much older. Has, in recent times, anyone's life been saved on a standard commercial aircraft by a life-jacket? Or by that little whistle?

The distinction between objective risk, and fear, is hardly news. It is becoming common in the UK to make a distinction between reduction in crime, and reduction in the fear of crime. This is clearly inspired by the recognition that fear of crime can have a deeper impact on people's lives than crime itself, coupled with the thought that one way of reducing crime is to make people hyper-vigilant, which may make them hyper-scared too. So the two goals have a complex relation. Yet was it right to assume that fear and anxiety are so directly correlated with subjective risk, which I understand as belief in risk? This is not entirely clear. Studies show that women are more fearful of crime than men, even though they know full well that men are more often the victims of crime than women (Burgess-Jackson 1994). However, this too is complex. Consider an example from John Adams. Are roads safer now for children pedestrians than they were in the 1950s? Statistically the result is surprising. Fewer child pedestrians are killed on the roads now than for decades. But this, he argues, is because we believe that roads are so dangerous that we keep our children away from them (Adams 2001: 10–14). What this shows is that we need to be very careful in how we collect and present our statistics. If women do not go out on their own late at night it isn't surprising that men are more likely to be victims of street crime at night. But if we were to measure "crime per risk taken" then the figures may be very different. Or they may not. We need careful studies by people who are not setting out to confirm a particular hypothesis.

But nevertheless although for a given individual whether or not there may be good reason to believe that there is a positive correlation between subjective belief in risk and fear, there is no reason for postulating an interpersonal correlation. People just have different personalities. Some are neurotic, some are oblivious. And there are many shades in between. On an aggregative ethic of fear reduction we may have to reduce small subjective risks for one group before addressing larger subjective risks for another. And all of this is independent of objective risk. But the main lesson is that this reinforces the claim that the irrationality argument fails to take into account that saving lives is not the only possible point of risk reduction policy.

3.2 Causation and blame

We have sketched out so far two main theories of risk regulation: risk reduction and fear reduction. This distinction is well known. Yet we are far from finished. There is another dimension to which we should pay attention.

Standard risk analysis begins with two concepts: hazard and probability. The only hazard we are concerned with here is death, and so the risks we have been concerned with are probabilities of death from particular causes. However it is vital to recognize that essentially the same type of hazard can have more than one possible cause. Take the example of death in a house fire. Some people die in house fires caused by electrical faults, caused, in turn through freak accidents – rodents gnawing through wires, for example. Or the fault could have been caused by negligent workmanship. Or through deliberate arson. The death in each case is equally gruesome, although not equally morally culpable.

Imagine that as a society we take a decision to reduce the number of deaths in house fires. Policies are proposed to combat each of the three causes: by means of regular safety checks; better training of workers; and better policing to track down and punish arsonists. In deciding which policy to adopt one possibility would be to carry out a cost-effectiveness analysis, working out which policy saves most lives for a given budget. Yet an alternative approach would be to argue that it is more important to eliminate some causes of house fires than others. This would be to make the judgment that some processes by which risks are created and sustained are worse, in some non-statistical sense, than others, and so should be a priority to eliminate even if this does not lead to the most cost-effective way of eliminating risk. Thus the hazard/probability analysis is too superficial. We must also take into account the process by which the hazard comes into being, or is sustained, or perhaps, permitted.⁶

It seems that, in general, people worry about some processes more than others (Baier 1986). For example it is widely documented that “man-made” hazards are regarded as in some sense “worse” than “natural” hazards. However there is more than one sense in which a hazard can be worse than another of the same objective magnitude. One is that it makes people more fearful. Another is that it generates greater moral concern or outrage. While these may often go together they need not. For example, I could be outraged at the existence of a risk that I do not even face.

In the light of this it might not be surprising if a society chose to eliminate morally blameworthy culpable behavior first. Morally culpable behavior comes in various forms. Roughly we can distinguish malice, recklessness, negligence, and incompetence. Malice is to set out a course of action with the deliberate aim of imposing harm or risks to people. Recklessness is to act knowing that it could cause harm or risk, but not taking this properly into account in deciding whether to act. Negligence is to fail to consider whether or not your action carries risks to others, when such risks were reasonably foreseeable. Incompetence, in this context, is to carry out a proper risk assessment and decide to take appropriate

⁶ For another perspective on the importance of cause in risk analysis, see Hopkins (2004).

action, but fail to do so. We blame people and organizations where we feel they have violated some moral norm, and an extreme form of blame is outrage.

There are important distinctions here, and we may well feel differently about different types and levels of culpable behaviour, but to keep the discussion within manageable bounds I will consider only malice here. And it seems that we do have policies which give a priority to reducing hazards brought about by malice. After all, it is not obvious that the resources put into deterring, detecting, and punishing murderers, or preventing terrorist attacks, can be justified on a risk cost-benefit analysis valuing each saved life at £1 million.

Is it plausible that we should, as a priority, eliminate hazards caused by malice? I think that this would be a common view. But what explains it? One possibility is simply that we think very badly of malice, and take particular satisfaction in eliminating its effects, or, to put it differently, we find some actions moral outrageous and we find ourselves giving a high priority to removing sources, or potential sources, of outrage. I'm sure that this is at least part of the story, although not all of it. An alternative explanation appeals to the vital distinction between risk and uncertainty. Risk involves known hazards and probabilities, whereas uncertainty involves lack of knowledge, either of the precise nature of the hazard, or the probability of its occurrence, or both. For most people, in most of their life, they are faced with uncertainty, at least within a range, rather than risk in the technical sense. This puts us in quite a different situation, both practically and technically. For risk cost-benefit analysis assumes that we know the hazards and probabilities, or, at least, have a good basis for estimating them, or at least enough stability to apply some other methodological approach. Without this the analysis can't even get started.

The relevance of this distinction is that it is not implausible that once malicious human beings threaten, we are moved into a world of uncertainty, not risk. And perhaps what in part explains any belief that we should give the rooting out of bad behavior special attention is the further belief that bad behavior places us under conditions not of known probability but uncertainty, and eliminating this uncertainty is the priority. With a few arsonists running around we cannot predict what is going to happen. And it is the same with killer sharks lurking in shallow waters, even though we don't tend to hold them morally to account. This may indicate that part of the problem does indeed lie in uncertainty.

However, it seems highly likely that we find ourselves with two converging explanations in these cases: root out bad behavior and control uncertainty. Both are distinct from risk reduction, and may lead to irrational results in such terms. But, many will argue, so much the worse for risk reduction, as the sole aim of safety policy.

3.3 Reputation and shame

Alleged bad behavior, though, comes into the picture another way too. Often a firm will pay more for safety than is mandated by the regulations. Why? Sometimes a firm may want to take expensive safety measures because it thinks it is the right thing to do. Sometimes there is a commercial advantage in having a squeaky clean reputation for safety. But for either reason a firm may decide to spend more than the regulations require. Is this irrational? The irrationality argument is that the redistribution of resources across sectors may reduce overall risk. However, it is rare that it would be within the power of one firm to do this, as firms operate only within a restricted domain. All they can do is regulate their own area. Their own budgeting trades safety against other aspects of the quality of the product, together with prices and profits. Their choice is simply to spend more money or less on safety, knowing that they could well suffer commercial damage from any accident involving their own goods or services. Hence it is very likely that a “not on my watch” phenomenon will sometimes operate; we know that there are going to be accidents but we don’t want them here. Individual decisions may then lead to a very uneven provision of safety. Some firms may overspend while others underspend, depending on how prepared they are to risk harm and consequent reputational damage, as well as legal liability and possible bankruptcy. That is, different attitudes to being blamed for causing harm will lead firms to different attitudes to safety. And it could be that in a given industry no one can afford to be singled out as relatively dangerous, even when general standards are very high. Conceivably this is true of air transport, where, we already saw, it could be argued that far too much is spent on safety measures which have either only a negligible, or perhaps symbolic, effect. But an operator may feel compelled to match “best practice” since no operator can afford a reputation for being less safe than the competitors. This we could call the problem of clean hands – no one wants to be the site of where the harm takes place.

This problem goes all the way up. A safety regulator cannot, for example, tell firms to stop spending money on safety improvements but pass the money to the health service instead. And, understandably, it will want as few deaths as possible in the areas it regulates. Excessive media attention, and a reputation for poor safety, however undeserved, follows accidents. A concern for reputation may lead to apparently irrational over-provision and even over-regulation, in particular areas, relative to the goal of risk-reduction. Yet, once again, if the aimed-for goal is “reputational damage reduction” this turns out to be rational after all.

And it does not stop with the reputation of the safety regulator. In the aftermath of a train crash people often feel ashamed to be identified with a country where this sort of thing can happen. It is shocking to find

that even in a modern, industrialized economy at many junctions the only thing stopping our trains crashing is the driver giving the correct response to a trackside light, and a bell. Admittedly the system is set up so that if a driver loses concentration the train will stop automatically, but there is still room for driver error. And if a train crashes through driver error, our shame that this can still happen in our country may overwhelm the thought that these things happen very rarely, and that for the last 50 years in the UK there probably has not been a week where more people died on the railways than on the roads.

This “shame” perspective can be based on presumed international comparison and a thought about how we must look to others. Once more, when one has got “into the skin” of this approach it may seem quite reasonable to want to take steps to reduce the potential for shame, even if they are expensive. Yet from a risk reduction perspective it is an absurd waste of money.

4. MODIFYING THE RISK ASSESSMENT PARADIGM

4.1 Perception of risk and societal concern

This issues I have mentioned – anxiety, malice, recklessness, negligence, incompetence, reputation, and shame – are not unknown to those who theorize and regulate risk. However risk management has struggled to work out how to incorporate them. Two leading approaches are what we can call the “perception of risk” framework and the “societal concern” framework.

The perception of risk framework, drawing especially on the work of Paul Slovic and associates (Slovic 2000), pays attention to how individuals perceive the seriousness of risks. So, for example, it is commonly noted that some risks seem to give special concern. Particularly important categories are those that are “dreaded,” such as the fear of cancer, and those that are outside the control of individuals, either in the sense that individuals have no influence over whether they are exposed to the risk or that there are no strategies they can personally adopt to mitigate their risk. Here traveling by car and traveling by air are an interesting comparison. Of course, whether or not one is exposed to the risk at all is generally a matter involving choice. But once the journey is under way a car driver has a measure of influence over subsequent events in a way that no air passenger has. Such utter reliance on others seems to create special concern.

Risk management needs to decide what to do about risks that give rise to special concern. One possibility, of course, would be simply to ignore special concern, which implicitly is what the irrationality argument recommends. But if one is unhappy with this approach something else is needed. The UK regulations deal with dread risks of cancer by doubling the

value of preventing a fatality from cancer.⁷ We can see how this modifies the standard risk cost-benefit analysis (RCBA). As we saw, in its simplest form RCBA derives an appropriate spending figure for risk reduction by means of a two-value formula of hazard (number of statistical deaths, each valued at £1 million) multiplied by probability reduction. The “dread” factor can then be used as a multiplier, so, in effect, a third variable. Hence anything that the perception of risk framework wishes to include can be added as a multiplier (or indeed divider) of the result that would otherwise be derived from the RCBA. In this manner the standard paradigm can be used in a much more flexible way.

“Societal concern” is a different matter. This term answers to the need to generate a concept to capture the idea that society may have concerns which go beyond the sum of concern each individual has for his or her own life. How to specify this is fraught with difficulties, given that its key defining feature is a negative: not the sum of individual concern. Therefore what is to be included can be contested. The Health and Safety Executive say the following:

Societal concerns [are] the risks or threats from hazards which impact on society and which, if realised, could have adverse repercussions for the institutions responsible for putting in place the provisions and arrangements for protecting people, e.g. Parliament or the Government of the day. This type of concern is often associated with hazards that give rise to risks which, were they to materialise, could provoke a socio-political response, e.g. risk of events causing widespread or large scale detriment or the occurrence of multiple fatalities in a single event. Typical examples relate to nuclear power generation, railway travel, or the genetic modification of organisms. Societal concerns due to the occurrence of multiple fatalities in a single event is known as societal risk. Societal risk is therefore a subset of societal concerns. (12)

Clearly a number of different issues are brought in here. For example, societal concern seems to include both direct costs – “large-scale detriment” – and indirect costs such as loss of confidence in government and presumably the safety regulators too. This may have further direct costs – loss of business to better regulated countries, for example – or the costs may be less tangible, such as scorn or being the butt of jokes. While many analysts seem to agree that there is such a thing as societal concern, and that it should be taken into account, there seems a great puzzle about what it is, precisely, and how it could be taken into account. One attempt at least to do something has been the railway industry which has used a

⁷ “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” (Health and Safety Executive 2001: 65).

higher VPF for multiple fatality accidents than those that involve a single death. So, for example, one accident that kills four people is considered significantly worse than four that each kills one person. Again, this is to use societal concern as applying a multiplier to the standard VPF in the basic formula.⁸

Is it right or wrong to use an enhanced (or reduced) VPF to reflect varying individual concern (anxiety, dread) or varying societal concern? Here, I think, the correct thing to say is that it is right to wish to modify the standard “hazard multiplied by probability” paradigm, but this is a crude way of doing so, especially as we do not have an agreed understanding of what the nature of societal concern is, and how it is engaged. Can anything better be envisaged?

Before we can settle this normative question – which is not in any case the task of this paper – we need first to understand how all the various attempts to understand and modify the analysis of risk can be brought together. Until we have an accurate understanding of the various factors in play it is hard to know what should properly be taken into account and how. So first we have to provide a descriptive model before a complete and compelling set of normative recommendations can be made. Yet before that we need to be as clear as we can about what the model is a model of: what precisely are we attempting to model? Here the best thing I can say is that we need an outline model of all the factors that affect human attitudes to risk, in the sense of influencing beliefs about how much we should do, pay, or sacrifice to mitigate risks. Many of these factors affect public perception to risk. Others, such as the “not on my watch” phenomenon – influence the attitude of those subjecting others to risk, or those who regulate risk. These can also influence public perception when there is an identification of some sort between members of the public and the organization in question, and so any failure of that organization can cause shame or embarrassment for the public. In the next section I will introduce a model of (part of) public concern, showing how fear, blame, and shame operate as factors in the public perception of risk, and in calls for its mitigation. For shorthand, I will call this a model of the “anatomy of risk.”

4.2 The anatomy of risk

The main lesson so far is, I believe, that the standard paradigm that risk is to be understood as hazard multiplied by probability is inadequate as

⁸ There is also the issue that some groups are especially vulnerable and so some weighting needs to be given to this. This may or may not fall under the heading of societal concern. Clearly it is a serious and important point. However it is beyond the scope of the current enquiry to consider this, as is the further, fundamental question of the degree to which public consultation is integral to the risk assessment and management process. I hope to address these issues in future work.

a way of modeling the actual concerns people have about risk, and the way in which those concerns translate into pressure to mitigate one risk rather than another. The “perception of risk” literature provides a way of incorporating some further features by means of allowing that fear may be out of proportion to objective risk, and this can be included as a multiplier or divider of risk. This is already reasonably well understood in the literature.⁹ However, as noted above “outrage,” or alternatively “blame” is distinct from “fear” and so needs to be incorporated as a second perception factor.

This does not yet incorporate all the issues included under the heading of “societal risk”, which I have assumed stands for those factors which can influence attitudes to risk, but go beyond each individual’s concern for the risks that he or she faces. Drawing on the examples and the analysis of Section 3 of this paper it appears that for analytical purposes the concept ‘societal concern’ is too vague to do any work, encompassing too many factors of quite different types. We do better to approach the analysis a different way. Let us begin by contrasting two rail accidents, the first caused by faulty track maintenance, the second by a car which bursts a tyre at speed, skids, and breaks through a fence onto the track, into the path of a train. The hazard in the two cases is for the purposes of analysis identical: loss of lives through derailment of a train. The probabilities may be hard to assess, but there is no reason to treat one as more probable than the other. If they are believed equally probable then again there is no reason to think that one will be more feared than the other. Yet the blame that will attach to the rail industry is likely to be very much higher in the case of faulty maintenance than in the case of the careless car driver. Indeed in the latter case there might even be sympathy for the rail industry.

If this is so then it seems that in this case at least blame attaches itself not to the hazard or the probability but to the cause of the hazard. Hence, it appears, the cause of the hazard must appear as an independent variable if we are to model public concerns about risk. Cause concerns how a hazard is created or sustained, and in consequence whether it can be viewed as a matter of culpable human action or inaction, especially the culpable action of those supposed to have a special responsibility.

One convenient way of understanding the various factors in play is to divide them into “primary variables” and “secondary variables.” On the present analysis the primary variables are three: cause, hazard, and probability. The secondary variables so far introduced are fear/dread and blame/outrage. They are called secondary because they take the primary variables as their object. In the standard cases fear attaches to hazard and probability – the “bigger” the risk, the more it is feared – whereas

⁹ For a review of some of the relevant literature, and an initial sketch of some of the ideas developed here, see Wolff (2002).

blame or outrage attaches to cause, as illustrated. But note that each of fear and blame/outrage can take as their object each of cause, hazard, and probability. Outrage can attach to the hazard, independently of cause. If a hazard will involve many deaths, or deaths in a particularly frightening manner, this may create pressure to do more to mitigate the hazard, even from those who are not personally at risk and so have no fear for themselves or on behalf of family and friends. Here, then, it must be a sense of moral concern, rather than personal dread, that moves such people. Finally, outrage can attach to probability. Even if the cause and the hazard do not generate outrage, the fact that something is happening "too often" may do so. Note that it is not claimed here that everyone will react the same way, but rather that broad trends may be detectable.

It may be that the secondary variables are indeterminate in number, in that a wide variety of further responses may be possible. Importantly, though, they must include shame, as introduced in Section 3. Shame can only be present for those who feel either identified with, or partly responsible for, the regrettable event. This is the sense in which ordinary citizens can feel ashamed of the loutish behavior of the football hooligans who follow the national team, or ashamed at the racism of other members of their own family. Shame, then, seems to pre-suppose some form of identification. Shame, also, is likely to lead to the worry that reputation will also be damaged, which can then have further damaging effects.

Consequently, while those directly involved in the organization imposing risks may well feel shame if those risks are culpably imposed, it does appear that this shame can also spread to the public. It could do so because the organization at fault is something with which there is strong identification – perhaps if it is a nationalized industry. Or it may be because a regulator, who acts on our behalf, has not done a good job. Or it may simply be because we as citizens have voted in a government which has allowed certain things to happen, or prevented other things from happening. Hence even "natural disasters" can give rise to shame, where the thought may be "how did we allow things to get to a state where this could have happened?" Furthermore, like blame, shame can attach itself to each of cause, hazard, and probability.

Shame, then, may affect public attitudes to risk, and may intensify public calls that a scandalous situation should be redressed. But, as indicated, shame can be more private than this, affecting only those who are directly involved in allowing a risky process or event to take place. And, as also noted, there can be reputational and commercial factors also to include, and these will influence how an organization wishes to manage risk. It is easy to see how the model can be developed to include more secondary variables, such as these, and perhaps even tertiary variables, such as regret that an organization is the cause of shame in its citizens. The central point of this paper, however, is to argue for the inclusion of

cause, alongside hazard and probability, as a primary variable, if we are to understand how various attitudes to risk are formed and sustained.

One further remark about cause is needed. In so far as we are concerned with strategies to mitigate risk, cause appears relevant only when it engages the secondary variables of fear, and especially blame and shame. Where there are no such factors then there is no role for cause to play and the standard paradigm of hazard and probability appears sufficient to capture general concern, subject to issues about distribution, consultation, and radical uncertainty, which I have not discussed in this paper. In this respect cause differs from the other primary variables. However this is not an argument that cause is generally redundant, but rather an argument that other effects cannot be properly understood without paying attention to cause.

Some readers may be surprised that I have attempted this analysis without appeal to the idea of "altruistic preferences" or to the distinction between a "consumer perspective" and a "citizens perspective." Both ideas are a useful reminder that distinctions need to be made but neither, in the end, takes us far enough. The idea of altruistic preferences has been used to model the thought that people care about other things than their own self-interest (see, e.g. Jones-Lee 1991, 1992). However using a single concept to attempt to capture the whole range of issues that fall under the heading of "not-entirely-self-interested concern" is problematic given their variety. A similar problem afflicts the idea of a "citizen" perspective.¹⁰ Insofar as a citizen is someone with moral beliefs, this concern is better captured under the idea of outrage or blame. Insofar as a citizen is a member of a collective entity which in some, possibly very indirect, way is responsible for the policies, including safety policies, adopted by his or her country, the idea of "shame" captures the relevant aspects. Hence the model adopted here is intended to provide a deeper analysis.

The model, then, points to a range of factors that currently influence judgment. However it is not obvious that there are no others, or that they should all be taken account of in the risk management process. For example, as we have seen, it is sometimes thought that perception and reputational effects should be managed by better public relations rather than by taking "skewed" safety decisions. Clearly when it is appropriate to attempt this is a vital question and it could be that practice ought properly to vary. So I should make clear that by setting out this model I am not assuming that the only way of managing the secondary variables is by spending money on safety measures. My point is only that we need a recognition of these factors and that it may turn out that the best way of dealing with them is, indeed, by spending money on safety measures.

¹⁰ The distinction between a "citizen" perspective and a "consumer" perspective is discussed in Wolff (2002).

5. CONCLUSION

My argument, then, is that to arrive at a model of the anatomy of risk – an account of the factors we need to include in order to decide how to manage particular risks – attention must be given to cause, hazard, probability, fear, blame and shame, and we must acknowledge that this list is unlikely to be complete. However the model provides considerable insight into the factors which influence public perception. Once we have this model we might decide that normatively it is nonsense, and all that we should be concerned with is saving lives in a cost-effective fashion, which is what the irrationality argument assumes. Here I take no view on this question. The preliminary task is to understand the factors which affect how generally we think about risk management. With a descriptively more accurate picture of the factors we actually take into account we can sensibly debate how much should be included in the normative framework of risk management, and why.

REFERENCES

- Adams, J. 2001. *Risk*. Routledge
- Baier, A. 1986. Poisoning the wells. In *Values at Risk*, ed. D. MacLean. Rowman and Allenheld. 49–74
- Burgess-Jackson, K. 1994. Justice and the distribution of fear. *Southern Journal of Philosophy* 32: 367–91
- Commission for Integrated Transport. 2004. Rail safety: revision of Factsheet 10. URL=<http://www.cfit.gov.uk/research/railsafety/03.htm>
- Evans, A. W. 2005. Fatal train accidents on Britain's main line railways: end of 2004 analysis. URL=<http://www.cts.cv.ic.ac.uk/html/ResearchActivities/publicationDetails.asp?PublicationID=465>
- Health and Safety Executive. 2001. *Reducing risks protecting people*. HSE Books. URL=<http://www.hse.gov.uk/risk/theory/r2p2.htm>
- Hopkins, A. 2004. *Safety, culture and risk*. CCH Books
- Jones-Lee, M. W. 1991. Altruism and the value of other people's safety. *Journal of Risk and Uncertainty* 4: 213–19
- Jones-Lee, M. W. 1992. Paternalistic altruism and the value of statistical life. *Economic Journal* 102: 80–90
- Jones-Lee, M. *et al.* 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
- Lomborg, B. 2001. *The skeptical environmentalist*. Cambridge University Press
- Lomborg, B. 2003. Paper given to Spiked conference: panic attack, London, 9 May
- Pidgeon, N., Kasperson, R. E., and Slovic, P. (eds.) 2003. *The social amplification of risk*. Cambridge University Press
- Posner, R. 2004. *Catastrophe, risk and response*. Oxford University Press
- Schelling, T. C. 1984. The life you save may be your own. In his *Choice and consequence*, 113–46. Harvard University Press
- Slovic, P. 2000. *The perception of risk*. Earthscan
- Sunstein, C. 2002. *Risk and reason*. Cambridge University Press
- Wolff, J. 2002. Railway safety and the ethics of the tolerability of risk. Rail Standards and Safety Board. URL=http://www.rssb.co.uk/pdf/reports/research/railway_safety_and_the_ethics_of_the_tolerability_of_risk.pdf