Sci Eng Ethics (2010) 16:147–168
DOI 10.1007/s11948-009-9125-6

ORIGINAL PAPER

The Social Nature of Engineering and its Implications for Risk Taking

Allison Ross · Nafsika Athanassoulis

Received: 28 November 2008 / Accepted: 18 March 2009 / Published online: 7 April 2009

Abstract Making decisions with an, often significant, element of risk seems to be an integral part of many of the projects of the diverse profession of engineering. Whether it be decisions about the design of products, manufacturing processes, public works, or developing technological solutions to environmental, social and global problems, risk taking seems inherent to the profession. Despite this, little attention has been paid to the topic and specifically to how our understanding of engineering as a distinctive profession might affect how we should make decisions under risk. This paper seeks to remedy this, firstly by offering a nuanced account of risk and then by considering how specific claims about our understanding of engineering as a social profession, with corresponding social values and obligations, should inform how we make decisions about risk in this context.

Keywords Engineering ethics · Risk · Virtue ethics

Introduction

Making decisions with an, often significant, element of risk seems to be an integral part of many of the projects of the diverse profession of engineering. Whether it be decisions about the design of products, manufacturing processes, public works, or developing technological solutions to environmental, social and global problems, risk taking seems inherent to the profession. Despite this, little attention has been paid to the topic and specifically to how our understanding of engineering as a distinctive profession might affect how we should make decisions under risk. This

A. Ross · N. Athanassoulis (✉)
Centre for Professional Ethics, Keele University, Keele ST5 5BG, UK
e-mail: n.athanassoulis@keele.ac.uk; Nafsika@komselis.gr

A. Ross
e-mail: a.ross@keele.ac.uk

paper seeks to remedy this, firstly by offering a nuanced account of risk and then by considering how specific claims about our understanding of engineering as a social profession, with corresponding social values and obligations, should inform how we make decisions about risk in this context.

Engineering projects provide us with the technological means of overcoming some of the physical limitations that are a consequence of being human. For example, it is thanks to developments in engineering that we can now travel long distances in a relatively short period; irrigate crops and dispose of sewage and other waste; seek solutions to world hunger, etc. However such developments would not have been possible if those involved had not taken some risks. Engineering projects are often innovative, long-term and involve the co-ordination of so many different variables that it is impossible to predict absolutely accurately what their consequences will be. In addition, because of the scale, and infra-structural nature of these projects there is often significant potential to do harm should something go wrong. So engineering projects inevitably involve risk. Some risks are short-term and local, they are run for the duration of the project and those exposed to them are the people involved in carrying them out. Other risks are longer term and more general, they are run for as long as the product of the project exists and the people who are exposed to these risks will be users of the product and the general public whose daily activities bring them into its vicinity. It is not obvious, however, how we ought to respond to the prospect of having to take either sort of risk.

The standard models for evaluating the acceptability of risks discussed in the literature, adopt an openly consequentialist approach to the evaluation of the acceptability of taking the risk. However, these models are inappropriate, partly because they render decisions about the acceptability of risk hostage to subjective preferences in an undesirable way, and partly because of more generic concerns with this kind of reasoning, e.g. problems with aggregation, standard 'integrity' objections, concerns about the 'quantification' of some of the harms caused by risks actuating, etc.[1]

We will consider the nature of risk taking decisions and the different elements that feed into their moral evaluation and the assignment of responsibility for risk-taking. We will emphasise the moral significance of the 'situatedness' of risk—the fact that risks are taken, imposed and borne by specific actual persons or groups of persons under particular circumstances.[2] Moral evaluations of risk-taking decisions depend on a number of factors including the existence of a consensus about the value of risk-taking—e.g. routine as opposed to extraordinary—the distribution of the burden of risk exposure, the extent to which people exposed to risk are complicit with or culpable in the risk taking, and in general the circumstances in which risks are taken and risk decisions are made. We will then consider how these decisions

[1] For a more detailed account of these problems see, Hansson (2003).

[2] In this we will be in agreement with two general critiques of consequentialist approaches to ethics:

- The broadly deontological critique that it fails to pay sufficient attention to the way in which the context of action often evolves as a result of previous actions which ought to influence our decisions about what is and is not moral.
- The distinctly Rawlsian point that Consequentialism fails to adequately accommodate the moral significance of the separateness of persons (Rawls 1999, p. 25).

about risk are shaped within the context of engineering by the values and goals of the profession.

We will conclude by arguing that the core values of engineering in combination with the better model of risk-evaluation we endorse, must mean that the there are some sorts of risks that ought not to be taken and there are certain moral obligations that attend risk-taking in engineering projects, which, however, can only be understood within a wider conception of engineering as a social profession and which can only be assessed through a complex judgment incorporating the values and goals of the profession.

Risk

Perhaps the greatest difficulty we face in thinking about the ethical dimensions of risk is that we are unsure whether to treat it as part of the context in which we act or as an intentional action in its own right. If we were to treat risk, as we do luck, as a feature of the world over which we have little or no control then questions about responsibility for the consequences of risk would have to be answered in the same way that questions about responsibility for the consequences of luck are answered. Cases of luck involve an outcome that is outside of our control, or, possibly comprehension, and thus are at odds with judgements of responsibility, which fundamentally suppose the agent's control. Questions about luck and responsibility are notoriously intractable and should issues concerning risk and responsibility fall under the same heading they would face the same intractability.

Fortunately for those who seek to establish answers to questions about risk and responsibility there is a significant difference between luck and risk. Risk is not essentially a phenomenon in respect of which we need remain passive—we can actively manage risk whereas we cannot 'luck'. Because risk has this active dimension we might expect assigning of responsibilities for the consequences of risk to be somewhat less complex or fraught than is the parallel task of assigning responsibility for results that are partly due to luck (usually ill-luck).

However, it will still be useful to distinguish between risks that we are subjected to, risks whose occurrence is entirely out of our control, and risks the occurrence of which we actively bring about. Rescher makes a distinction between risks we face, e.g. natural events, and risks we take, e.g. those involving choices with respect to the risk.[3] An example of the first kind of risk are natural disasters. The risk of a natural disaster is out of our control, as we cannot manage the occurrence of such risks. However, unlike in the case of luck, we do have some control—we have control over the extent to which we manage these risks. For example, we have no control over the amount of rainfall in any given year, but we can control whether we build houses in flood plains that are likely to be severely affected in years when rainfall is above average. By contrast the risks we take involve a decision as to whether to expose ourselves to them, e.g. accepting the risks of driving cars has an element of control which is missing from being subject to the risks of heavy rainfall.

[3] Rescher (1983, pp. 5–6).

We are primarily concerned with aspects of risk that are related to judgments of responsibility, so either risks we take or the management of the risks we face. Since these aspects of risk are associated with choices and decision making, in what follows we want to look at the circumstances of risk, how these circumstances give rise to the action of risk-taking and at the features of such action which impact upon the way we ought to morally evaluate it.

Risk-Taking: A Preliminary Definition[4]

The notion of 'risk' seems to be defined by two parameters: the idea of uncertain outcomes (which may include a probabilistic calculation of the likelihood of occurrence of these outcomes) and the idea of loss. The latter involves the claim that cases of risk involve a chancing of negativity,[5] broadly construed to include both direct loss or harm, but also the inverted risk of giving up on something good.[6] Given our claims above about the nature of risks we take and their link to decision making, risk-taking is characterised by a very specific set of circumstances in which it occurs—it arises when a set of statements like the following are all true:

1. Someone (S) is considering performing act A
2. S is not in complete control of the outcome of A—there are a number of possible outcomes of doing A
3. However S is able to, roughly, estimate the likelihood that her action will have one or other of the range of possible outcomes.
4. Some of the outcomes of A are desirable (either to S or to others) and other potential outcomes are undesirable (either to S or to others)
5. S is strongly motivated by the prospect of achieving the good outcomes of A and this represents a reason for her to discount or weigh less heavily the potential bad outcomes of A.

Consider a gambler, for example, who wishes to place a bet, knowing that she may or may not win and that the probability that she will win is, say, 1%. She is strongly motivated by the prospect of winning to the extent that she considers this prospect—even though unlikely in this case—to be a reasonable justification for placing the bet. She finds herself in the circumstances of risk.

So far we have described an agent in a particular situation, which we might call this 'the circumstances of risk'. Now, when S, recognizing that statements 1–5 above are true, actually does A, then S's action becomes an instance of 'risk-taking'. Risk-taking occurs when there is a relationship between the circumstances of risk, a decision or action and an outcome. People get drawn into/take up positions in this relationship in three main ways. They can be:

[4] For the purposes of this paper we will seek to define risk from the perspective of the person who makes a decision about it rather than from the perspective of the person who is subjected to risk through no choice of her own. We do so because we are interested in responsibility which seems to us to be tightly linked to the decision-making process.

[5] Rescher (1983, p. 5).

[6] This seems to be Rescher's position, Rescher (1983).

1. the Decision Maker: the person(s) who has the information about the elements of the decision, weights up the alternatives and comes to a decision about whether the risk is worth entering into or not
2. the Harm Bearer: the person(s) who runs the risk of bearing (or bears) the cost of the decision, i.e. the person who is at 'harm's way' because of the decision
3. the Potential Beneficiary: the person(s) who runs the chance of gaining (or gains) from the decision, i.e. the person who may benefit from the decision.

Any one person may occupy more than one of these roles, but there may be situations where different individuals occupy these three roles (and even more than one person for each).

Some risk decisions have an identified risk/gains bearer whereas others may have an unknown target. That is, some decisions will be made in the knowledge that they run the risk of affecting an identifiable individual, e.g. the particular patient on the operating table runs the risk of the surgeon's clinical decisions during the surgery, whereas some decisions will be made with a statistical probability that someone may be harmed/benefited but without knowing who that individual will be, e.g. many people will buy a faulty product the company knowingly sold, but only some of them will experience the fault and only some of those will be harmed as a result. Finally, some risks may be negligible at lower levels, but a have a kind of 'threshold effect',[7] so that over a period of time the next act imposes a much more significant risk than the original one.

Risk and Responsibility

Other things being equal we are usually responsible for the actions we freely choose to do. So, if S is our gambler and she loses we are inclined to say that S is responsible for her loss—she has no basis for complaint against either her fellow gamblers or the practice or institution that made it possible for her to place the bet in the first place (presuming no one cheats). But is it always fair/reasonable to judge people responsible for the risks they take? It certainly seems that some risks are taken with impunity—everyone who drives a car takes the risk of being involved in an accident but the car-driver is not responsible for the harm that is suffered by victims of a crash in which she was involved *merely* because she took the risk of driving a car. To see this consider Williams' example of 'agent regret', which is underpinned, we argue, by just such thoughts.[8] In the example, a lorry driver of a well-serviced and carefully driven vehicle hits and kills a child. Williams argues that the driver will and ought to feel regret DESPITE not being responsible for the death of the child. The driver took a risk, he drove a heavy vehicle at a particular speed knowing that if a small child happened to run out into the street his vehicle would hit and kill the child. Driving is a risky activity, simply due to the nature of cars, the nature of driving and how they interact with pedestrians. We manage this risk by imposing speed limits, highway regulations, driving with due care and

[7] This point is nicely made by Thompson (1986, p. 175), using the example of smoking in the presence of another person.

[8] Williams (1993, p. 43)

attention, etc., and although all these parameters *limit and manage* risk they do not *eliminate* it entirely. Given that young children are likely to, sometimes, run out in the middle of the road without looking, any kind of driving posses a risk to them in such a situation, and only refraining from driving altogether would eliminate the risk entirely. However, once a driver has complied with all the requirements for managing the risk of driving, as the lorry driver did, he is not responsible for any residual risk actualizing.[9] Why then is the driver not responsible for killing the child? Perhaps this is because all pedestrians and all drivers accept the residual risks of driving as reasonable, in such a way that when they do actualize the agents involved in these situations are not morally responsible for them.

In order to know how to apportion responsibility we need to know what the significant differences are between the car-driver and those we do hold responsible for the consequences of the risks they take. Judith Jarvis Thompson's classic and graphic example of the morally culpable risk-taker is the agent who decides to play Russian roulette with the life of some unknowing victim.[10] The agent playing Russian roulette with another person exposes his potential victim to the risk of death and this seems to be significant regardless of whether the risk actualized or not. The victim had no knowledge of the risk he was under, he was not in any way complicit in the decision to undertake this risk or in any way culpable for being in this position; while the attacker exposed his victim to the risk for no particular reason other than his own gratification. This kind of analysis concentrates on the fact that the victim was exposed to the risk, rather than whether the risk actualized or not. To an extent, all this also holds true in the car driver case, since by driving we are exposing all other road users to the risks of driving, however there are significant differences which can allow us to distinguish between the two. Firstly, almost all road users, with the exception of the life-long, exclusive pedestrian, benefit from the possibilities afforded by driving. Secondly, there seems to be a social agreement that driving is necessary and indispensable to our way of life. We structure our lives, e.g. where we live, where we work, what we do for fun, etc., around the possibility of driving, making us all complicit, in a sense, in both the benefits and the risks of driving. Lastly, all road users share equally in the managed risk of driving and this seems fair.[11] We all benefit from driving, we all find it a necessary and indispensable part of our lives and we all share equally in the risk, none of which applies to the Russian roulette example.

So, to summarise, it looks like whether or not an agent is morally responsible for the exposure of another to risk is determined by:

- Whether the goal of the particular instance of risk-taking is worthy
- Whether the particular instance of risk-taking is (objectively or at least inter-subjectively) routine.

[9] Williams is not alone in arriving at this sort of conclusion about responsibility for risk-taking. Judith Jarvis Thomson arrives at similar conclusion when considering the case of neighbour who, in switching on her gas to make a cup of tea triggers an explosion which destroys the adjacent house. See Thompson (1986), pp. 182–183).

[10] Thompson (1986, pp. 181–182).

[11] With the exception of the life-long pedestrian.

- The distribution of the burden of risk-exposure—whether it is shared equally between all persons or at least between the risk-taker and the person exposed to risk or not.
- the extent to which the person(s) exposed to risk is (are) complicit in the risk taking or culpable for being caught up in someone else's risk taking.

It is our view that in ethical risk-taking, all of the above will feature in the deliberations of agents contemplating the prospect of taking a risk. That is to say that each of the above will function as distinctively moral reasons for taking/not-taking a particular risk. We further assert that it is the process of deliberation and particularly the reasons the agent has for acting that ought to be the focus of moral evaluation. In what follows we wish to make some general observations about the sorts of things these reasons are and their implications for responsibility attributions.

Reasons for Risking

Much is usually made of the 'objective' features of reasons for risking. Risk decision making is often thought of as a paradigm application of rational choice theory and judgements about the acceptability of risk are portrayed as almost entirely dependent upon probabilistic calculation. If the objectively determinable probabilities favour risk-taking, then it is legitimate and if they don't, it isn't.[12] It is a mistake, however, to conflate 'acceptability' as used in rational choice theory and moral acceptability. The former is always 'acceptability to a particular decision-maker' and thus conditional upon his/her desires. Moral acceptability is something much broader and can only be decided after rational scrutiny of the aims and desires that move the agent. This is not to say that probabilistic reasoning does not factor in assessing the morality of risk-taking but rather that it is not the sole determinant of the moral status of an action. Indeed more will be said about the role of probabilistic reasoning later. For the moment we want to focus attention upon the role of the decision-maker and his/her subjective states in determining the legitimacy of risk-taking.

The first generalisation we wish to make about the sort of things that reasons for risking are, is that they are practical reasons—reasons which satisfy what Christine Korsgaard likes to call 'the internalism requirement'.[13] The important thing about practical reasons is that they move agents to action.[14] So the reasons an agent has for taking risks will reflect the sorts of things she wants or values. But this generalisable fact should also lead us to recognise that the reasons people have for accepting and rejecting risks are not entirely objective reasons, they depend upon the conceptions of the good life or settled values of the reasoner and these are widely held to vary quite dramatically between persons. Evaluations of risk necessarily have a strong subjective component and decisions about which risks are so widely acceptable or reasonable that people cannot be blamed for taking them

[12] Lewens (2007 p. 1–7).

[13] Korsgaard (1996, chap. 1).

[14] They are to be contrasted with reasons for believing which do not have the same action directing function.

will be 'infected' by that subjectivity. The reasons people have for taking risks are subject or personal-value dependent and as such reflect the characters of the agents who are motivated by them. When we morally judge an incidence of risk-taking we are (at least in part) judging the character of the risk-taker.

This subject-dependency of the reasons for risking can lead to genuine and intractable controversy. Consider, for example, the debate about nuclear energy. We have shaped our lives in ways that generate great energy requirements. Currently there are a number of options for generating energy, of which nuclear power generation is one. The risk of malfunction at a nuclear power station is relatively low, suggesting that perhaps this is a risk worth taking as it is not likely to actualise. However, the consequences of malfunction (what it is that is being risked), are so significant that this fact seems to sway the argument the other way. There is wide support for the view that even a small likelihood of this harm occurring is unacceptable, making this a risk that a reasonable person or group of people would not accept as routine (unlike the driving case). The controversial status of nuclear power can, in part, be seen as a consequence of the fact that people react to the question of nuclear energy based upon a personal evaluation of how bad things would be (the seriousness of the harm that would be done) if the worst case scenario was realised combined with their particular inclination/disinclination towards caution. In this case there is no consensus and so no strong grounds for treating nuclear risk-taking as routine or as extraordinary. However there will be cases where we do get consensus. These will occur where the values to which deciders appeal are widely shared e.g. the kind of mobility that we have by using cars is widely recognised and valued.

So what sort of reasons for risk taking are we likely to be able to reach consensus about? It will be helpful here to classify the practical reasons agents have for taking risks into three categories according to the sort of motive with which they present the agent for whom they are reasons:

A. Reasons to do with the possible loss of good things, e.g. by seeking a new job, I risk losing my current employer's goodwill.
B. Reasons to do with avoiding harms or the risk of being exposed to harms, e.g. by not eating beef I avoid the risk of contracting CJD.
C. Reasons to do with gaining benefits by running a risk of harm, e.g. by taking on mountain climbing as a hobby I may gain the benefits of exercise, the joy of the great outdoors, and a sense of satisfaction from the great achievement, at the risk of a climbing accident.

First, let's compare cases of types A and B. We standardly assume that it is morally less bad to fail to benefit than it is to actively do harm. Similarly in the case of risking we might be able to say that, all other things being equal, risking a harm has more serious implications than risking not gaining a benefit, as the former involves exposure to a harm, whereas the latter merely loss of opportunity of a benefit. However, it is not obvious that this is right—it is possible that some losses of opportunity are so serious or that serious losses of opportunity for third-parties have such great moral significance that the previous thought does not always hold true. So it looks like we will have to decide the moral seriousness of individual cases

on their own merits. A more detailed look at cases of type three confirms this impression. Consider that such cases will involve an evaluation of the relative merits of the harm to which people are exposed to in order to gamble for the benefit. Again there might seem to be grounds for generalisation here, after all, we standardly assume that more severe harms are worse than less severe harms, as are harms of a longer duration, intensity and scope. So it may be with risk—it seems sensible to claim that risking great, enduring and wide-reaching harm (even if the likelihood of this manifesting is very small) is more morally questionable than risking minor, short-term harm to a small group. However this comparison ought to leave us somewhat uneasy. One reason for this is that consideration of the scope of the risk exposure seems to pull us toward the opposite conclusion—that the risks I take in car-driving are more serious than those I take in playing Russian roulette. After all, my one-off indulging of my violent inclinations in a game of Russian roulette affects one person whereas my driving a car puts a great number of people at risk.

A further level of complexity is added to the process of reasoning about risk and responsibility as there is a question about the characterisation of the outcomes should they actualise. We take this to be one of Thomson's points when she compares the case of the agent whose lottery numbers do not come up and the agent who risks going home via Unpleasant Way only to be mugged. There are judgements about the appropriateness of taking the risk in the first place, e.g. if the first agent spent the family's rent money on the lottery ticket then the decision to do so seems callous and irresponsible, as does the decision to take the Unpleasant walk on the day the second agent was responsible for the school run and had four young children with him. These judgements will, in themselves, be subject to debate, but there is a separate question as well which relates to the harm actualising. It is in this respect that Thompson concludes that the lottery loser has no grounds for complaint about losing whereas the victim of the mugging has been wronged.[15] Had the walker not been mugged then he would have not been wronged, but this is a separate question about the question of the reasonableness of his decisions to walk in an unsafe area in the first place.

Even further complexity results form the fact that, as we have said above, there are three possible positions one might occupy with respect to risk; that of the decision maker, that of the harm bearer and that of the potential beneficiary. If one person occupies all three roles, matters are less problematic as this goes to the heart of our conceptions of agency and responsibility, i.e. the person who makes the decision is the person to be affected by it. Making decisions on behalf of others is evidently more problematic as it involves the extra responsibility of how one person's decisions affect another. Situations where one person is both the decision maker and stands to benefit, whereas another person is risking being harmed, are clearly the most problematic. The distribution of the risk is significant, both in the relationships between the three roles an agent may occupy with respect to risk, but also with respect to justice issues regarding distribution of harms and benefits.

[15] Thompson (1986, pp. 189–190).

Another aspect of circumstances which will affect the legitimacy of risk-taking is the extent of the options open to the agent(s). Generally, although again not inevitably, where the options for achieving a particular good/avoiding a particular harm are wide and some less risky than others, then the legitimacy of risk-taking is more questionable than it is in cases where risk-taking is the only means to a particular desired end. Think again of the nuclear power example. In the past we were inclined to think that in the face of a nuclear disaster, fossil fuels seem like a very good alternative, however, given concerns about global warming the risks of using fossil fuels have changed. The decision becomes much more complex now: the nature of the harms themselves becomes very complex, with serious harms on either alternative, as well as some uncertainty as to the nature of these harms. In addition, the cohort of potential victims also presents difficulties to do with indeterminacy, future generations, etc. In this context our assessment of the legitimacy of treating nuclear risk-taking as routine may change.[16] This would suggest that the content of any consensus on whether particular risks are routine is context-dependent and liable to change as conditions change.

Finally, we said above that probabilistic calculation will affect judgements about responsibility for risk-taking. We think that this is a consequence of the conceptual link between what it is reasonable to expect and intention, and the fact that part of moral assessment is an evaluation of intentions. In general, the higher the probability of a bad outcome the more responsible the risk-taker for any resultant harms and this is linked with what it is reasonable to expect[17]—if the probability of an outcome is high it is reasonable to expect that it will occur and one's act is carried out under the expectation of that bad event. By contrast if the probability of the bad outcome occurring is low then it is reasonable to expect/hope that it will not occur and one's act need not be overshadowed by any negative expectation.

So, it is not only values and personal attitudes to caution that influence consensus about the status of particular incidences of risk-taking—the circumstances in which risks are taken can (and should) influence our judgements about whether or not these risks should be considered acceptable or even routine. Judgements about responsibility for risk depend upon the sort of risk decision this is, the reasons or motives that move agents to make such decisions, whether or not the person affected by the decision is the person who takes the decision and the probability of the risk being actualised, the nature of the possible consequences, the distribution of the burdens and benefits and the nature of the probabilistic calculation, as well as how these factors interact.

It is not surprising, then, that a fail-safe ground of consensus about the morality of risk-taking still eludes us. Perhaps what we should learn from the above discussion is that although generalisations based upon the duration, scope and intensity of the harm risked are often reliable guides to the appropriate moral evaluation of risk-taking they will firstly, not always be possible and secondly, not always give 'the right answer'. When it comes to arriving at a moral judgement about any particular incidence of risk-taking, the precise nature of the harms and

[16] As it seems to have done in the minds of the members of the current UK government.

[17] See Athanassoulis (2005, pp. 9–14).

benefits involved and the question of precisely who stands to benefit, who is exposed to risk and who takes the risk will all matter. Risk evaluation and the assignment of moral responsibility for risk-taking and for the results of risk-taking needs to be done on a case by case basis because the necessary evaluations are context-dependent. We now wish to turn our attention to the special reasons for acting that engineers have and the way in which these impact upon their decision-making in the circumstances of risk.

Engineering Professionals and Their Reasoning About Risk

We have, so far, argued that responsibility for risk-taking will be best handled by an account of responsibility that can accommodate the agent and context sensitive nature of the reasons that people have for acting under conditions of uncertainty. We have suggested that risk-taking can be either practically rational or irrational. Whether or not a particular decision to risk is rational will depend upon a wide variety of agent and context specific factors which all have a role to play in determining what reasons for action the decision-maker has.

In this section of the paper we want to argue that engineers, like other professionals, have distinctive reasons to take or refuse to take risks that they acquire by being members of their particular profession. That is to say that where the decision maker is an engineer, and significant determinant of the rationality of a decision to (or not to) risk is the fact that the decision-maker is an engineer. In particular, in the circumstances of risk-involving decisions, engineers have special reasons to be very concerned about safety. We will further argue that when these reasons are more closely examined they embody and balance degrees of caution and (social) ambition that are appropriate to the circumstances and nature of the decisions that engineers have to make and, more significantly, to the role which engineers occupy in society. It will not be possible to say how much caution/ambition is an appropriate motivator in general—it follows from the above discussion that there cannot be generalisable claims about which sorts of risk are appropriate and which are not. However it will be possible to demonstrate that the balance of caution and ambition that is appropriate will be affected by the sorts of factors discussed in the previous section and by the role which an engineer fulfils when taking these decisions. So that the decisions an engineer should make may very well be different to the decisions that a generally sensible and compassionate person ought to make. It will further follow that responsibility for bad outcomes of risky decisions may be attributed to engineers/the organizations for whom they act.

The Role of the Engineer and its Impact on Decision Making About Risk

Perhaps a good place to start any inquiry is to look at definitions that attempt to explain what is ‘engineering’. There are a variety of definitions of engineering offered by professional bodies, for example, the Engineering Council’s definition:

"Engineering is the practice of creating and sustaining services, systems, devices, machines, structures, processes and products to improve the quality of life; getting things done effectively and efficiently".[18]

or the American Engineers' Council for Professional Development definition:

"[T]he creative application of scientific principles to design or develop structures, machines, apparatus, or manufacturing processes, or works utilizing them singly or in combination; or to construct or operate the same with full congnizance of their design; or to forecast their behaviour under specific operating conditions; all as respects an intended function, economics of operation and safety to life and property."[19]

These definitions, and many like them, have certain shared elements, namely engineering involves a particular kind of knowledge, used to create particular practical applications, for a particular purpose. Interestingly enough the definitions all comment on the purpose of engineering, so that not any kind of application of engineering knowledge to practical issues will count as an instance of engineering, but rather only those that fulfill the purpose of engineering. We will go on to consider what this purpose might be, but it is worth pausing to consider the import of the very existence of a purpose. These definitions seem to suggest that engineering is a purposeful activity, carried out to fulfill a specific function, which implies that it is also possible to fall short of the mark, i.e. if one is applying engineering knowledge to purposes other than the true function of engineering, than one is not really an engineer.

To make sense of this point consider the following analogy: medical knowledge can be applied in various spheres, for example, someone who has knowledge of medical anatomy would make a more effective torturer, but we would not want to say that a torture's use of medical knowledge is in accordance with the function of medicine. It may be difficult to give a precise definition of the function of medicine in all instances, but in broad terms, medicine is concerned with the best interests of the patient, which torturing clearly is not. So although both activities, medicine and torture, make practical use of the same knowledge, only one qualifies as fulfilling what medicine is for, i.e. benefiting patients. We would like to make a similar claim about engineering, that is, that engineering is a similarly purposeful activity, it involves the application of a certain type of knowledge but this is constrained by the aims of engineering as an activity, so that not all kinds of applications of this type of knowledge will count as an instance of engineering.

Consider this characterization of engineering from the National Society of Professional Engineers:

"The highest type of engineer is proud of his profession. He is proud of his competence, proud of what he can do. He captures the forces of nature and puts them to work for the improvement of industry, and finally for a better life for the people. He is proud to be able to do things with materials and forces which astound

[18] From the Engineering Council's "Engineering 2000" quoted in the Royal Academy of Engineering's report "A Universe of Engineering", chair Sir Robert Malpas, http://www.engc.org.uk/documents/Malpas_report.pdf

[19] The Engineers' Council for Professional Development, p. 456.

lay people. But the highest type of engineer is proud of something else which is more important. He is proud of the good he does in the world. He is proud that he serves the people; he is proud that ministry in the community is him principle objective; that he is not merely a 'laborer for hire'."[20]

Regulatory bodies and codes of ethics define engineering in these terms because they see it as a purposeful activity in which excellence is defined as achieving the function of the activity. Engineering is distinctive as it is not merely an activity in which people engage in for profit. It is different from, say, selling vegetables or trading on the stock exchange. Unlike these things engineering, like other professions, is a practice in Alistaire MacIntyre's sense. For Macintyre a practice is:

> any coherent and complex form of socially established co-operative human activity through which goods internal to that form of activity are realized in the course of trying to achieve those standards of excellence, which are appropriate to, and partly definitive of, that form of activity, with the result that human powers to achieve excellence and human conceptions of the ends and goods involved, are systematically extended.[21]

The professions have distinctive profession-defining internal goods, which are realized through the activities of their practitioners. Moreover the activities themselves can be carried out in more or less excellent ways and standards of excellence extend to the way in which goals are achieved as well as to the end result. Usually, the goods internal to/definitive of a profession, reflect its function—a combination of the goals towards which it strives (preservation of health; technological progress and design etc.) and the role that the profession occupies in the wider society. It is our contention that the chief good internal to the practice of engineering is safe efficient innovation in the service of human wellbeing and that this good can only be achieved where highly accurate, rational decisions are made about how to balance the values of safety, efficiency and ambition in particular cases.

In order then to find out what engineering is and how engineers should conduct themselves qua engineers, we need to ask what engineering is for. In the same way that medicine is for the good of the patient, engineering is the innovative application of relevant scientific knowledge for the good of human beings. The question still remains what constitutes a good life, but this is a question that permeates every aspect of human existence and human endeavors and whose answer can preoccupy an entire life. Engineers will struggle with this question as much as the rest of us, and in their pursuit of engineering excellence they will struggle with aspects of the good life that are relevant to the appropriate application of engineering knowledge. We say the 'appropriate application' because this account of engineering makes it a moral pursuit, one which can be carried out in accordance with virtue. Determining what constitutes the good life for human beings is a communal project which will include individual ambitions and social values. In what follows we will try to give a

[20] National Society of Professional Engineers, Canons of Ethics for Engineers.

[21] MacIntyre (1981, p. 187)—emphasis ours.

bit more content to how we understand engineering's purpose within current social values and virtues.[22]

So the strain of ethical engineering falls squarely upon the process by which decisions are made and upon the values internalized by practitioners. Governing bodies, across the world, recognize this and have sought to give some very general guidance in their codes of conduct and statements of ethical principles. The Engineering Council UK/Royal Academy of Engineering begins its statement of ethical principles as follows:

> Professional Engineers work to enhance the welfare, health and safety of all whilst paying due regard to the environment and the sustainability of resources. They have made personal and professional commitments to enhance the wellbeing of society through the exploitation of knowledge and the management of creative teams.[23]

The prioritization of the values set out in the above statement suggests that engineers don't just strive to find technological solutions to human problems, they strive to do so in a manner fitting for the conduct of an engineer which involves consciously foregrounding the values of safety and sustainability. The engineering profession recognizes that its goals are social goals (they are shared by all members of society and the professional only gets to pursue them because society provides the infrastructure that makes this possible). The professional acts not just for him or herself but also for the society in which s/he practices. S/he has a public role and the successful execution of that role demands the internalizing the professional attitudes to safety and sustainability—so that decision-making which attributes appropriate weight to these values becomes 'second-nature'. In doing so s/he satisfies obligations that arise as a result of her role. It is distinctive of any profession that practitioners acquire 'role obligations' we mean moral obligations one has by virtue of one's institutional role. For example, the moral obligations of doctors are shaped by the goals and values of medicine, are understood within the context of health promotion, and apply to doctors in their capacity as doctors and in virtue of being doctors. Hardimon (1994) considers role obligations at length and argues that the sort of identification with a role that is so characteristic of professionals highlights three things about moral judgments that are often ignored.

Firstly, an important part of the moral life is this identification with one's role,[24] and secondly, role-identification has a transformative effect on identity and thus on the values which drive action: "One comes to conceive of oneself as an occupant of

[22] This characterization of social values and virtues are potentially subject to change is in no way a commitment to moral relativism, rather it stems from a recognition of the multiple aspects of what constitutes a good life, which can vary depending to the individual and the context. In this we see ourselves as following familiar Aristotelian thoughts, illustrated, for example, in the Doctrine of the Mean (Nicomachean Ethics II 1106b 36-1107a 2).

[23] Statement of ethical principles issued by the Engineering Council UK which can be found at the following web address: http://www.engc.org.uk/publications/default.aspx, accessed 2nd January 08)—emphasis ours.

[24] Hardimon (1994, p. 360).

a role, a member of the social world, and a participant in the moral life that is lived through institutions".[25] The institutional roles we occupy come to define who we are and the values we hold. Thirdly, and most importantly, role obligations provide a special source of motivation; to paraphrase Hardimon, if you are an engineer who identifies with the role of engineer, the fact that this is something engineers do will give you a reason for doing it.[26]

But is it really the case that there is 'an attitude to safety' which informs 'what engineers do', something that will be a reason for acting in a certain way for engineers but not for others? Our claim is that there is and we think that the fact that there is, is best demonstrated by what happens when the ethical role of engineers is undermined in the course of the production of technology. An infamous case that illustrates this is the case of the Ford Pinto. A design flaw in the original Ford Pinto meant that the manufacturers were faced with a choice between undertaking changes to the design of the Pinto or exposing customers to the risk of the vehicle going up in flames, if it was subjected to high-velocity impact from the rear. In order to make the decision the company did a simple cost-benefit analysis—they compared the cost of re-design, estimated at $137 million and found that it outweighed the cost of compensation for victims at $49.5 million.[27] On the grounds of this analysis they decided against re-design. This case is often used as a paradigmatically unethical behaviour in the field of engineering. Ford's judgement, and, more importantly, the grounds on which it was made, i.e. the blatant financial evaluation of the loss of human lives as one of many economic parameters to be factored into the decision, strikes us as a failure to consider as important something which is (a) extremely valuable and (b) something which Ford as a seller of engineering products has an obligation to consider. Details of the case suggest that at various stages of the production process engineers working on the project were over-ruled or forced to act against their better judgement.[28] So one way of understanding what went wrong in this case is to say that the role of the engineer in the production process was undermined and that engineers were coerced into failing to fulfil their professional obligation to make certain sorts of judgements about the acceptability of risks involved in the production of technologies.

When seeking to understand who has responsibility for what in this case, one might be tempted to say that the sales people and management 'were simply doing their job', which involves seeking to get the highest return for the least output. The persons with responsibility for safety were the engineers. If we think like this then we do think of whatever moral failure happened here a failure on the part of the engineers. So the fact that we think there was a failure on the part of the engineers suggests that we also think engineers have role obligations and that considerations of safety ought to represent reasons for acting for engineers even when they do not

[25] Hardimon (1994, p. 360).

[26] Hardimon (1994, p. 358).

[27] Case details cited in Engineering.com, http://www.engineering.com/Library/ArticlesPage/tabid/85/articleType/ArticleView/articleId/166/Ford-Pinto.aspx, accessed 16th April 2008.

[28] In this case engineers were placed in a position in which they would have most likely been dismissed for raising concerns about safety.

represent reasons for acting for others.[29] Furthermore it suggests that failure to act on such reasons should undermine one's ability to consistently identify with the profession. All of this is perhaps illustrated by the real-life case of Lou Tubben—an engineer who worked on the Ford Pinto. The pure cost/benefit approach to decision making clashes with the sort of conception of the engineering profession that engineers themselves subscribe to. An anonymous engineer who worked for Ford at the time that the above decision was taken relates the following story:

> Lou Tubben is one of the most popular engineers at Ford. He's a friendly, outgoing guy with a genuine concern for safety. By 1971 he had grown so concerned about gas-tank integrity that he asked his boss if he could prepare a presentation on safer tank design. Tubben and his boss had both worked on the Pinto and shared a concern for its safety. His boss gave him the go-ahead, scheduled a date for the presentation and invited all company engineers and key production planning personnel. When time came for the meeting, a grand total of two people showed up—Lou Tubben and his boss.[30]

This story reveals a number of things. Firstly, the engineer sees his role as involving taking the practical responsibility for, and arguing the case for, the safety-measures built into the project under design. Secondly, Lou Tubben is characterised as having a genuine concern for safety, i.e. not merely an instrumental concern or a desire to merely comply with professional requirements for fear of sanctions or litigation worries. His is an internalised concern for safety which implies a genuine understanding of the value of safe products and an appreciation of the crucial role for engineering in designing products that meet the needs of the customers but within safe parameters. Tubben doesn't merely pay lip service to the need for safe design, he actively takes steps to ensure that the results of his labour are safe for the end users. When he thinks safety may be threatened he takes a personal risk by setting up a meeting to air his worries, consult with colleagues and, presumably, effect change. Thirdly, Tubben exemplifies the sort of link between practising the profession of engineering and having or developing particular attitudes. That is not to say that an ordinary averagely compassionate person who is not an engineer would not be moved by concerns about safety in this case. Rather what we think it highlights is the fact that for an ordinary lay-person the appropriate balance between the conflicting goods of this case—seeking to ensure safety and seeking personal well-being—would be different to the appropriate balance of the same goods for an engineer practising his profession. The values internalised by Tubben in the course of his training would mean that decisions taken by engineers in the course of their work are decisions that exemplify the distinctive 'right balance of goods for engineers'. In the same way that we can make sense of the obligation of a doctor to treat victims in an epidemic at risk to his own well-being and see this as part of what

[29] It is not our contention that safety considerations may never operate as reasons for action for marketing professionals, or in this example they rightly did not, rather that they have a special force for engineers qua engineers.

[30] Quoted on the Engineering.com site, http://www.engineering.com/Library/ArticlesPage/tabid/85/articleType/ArticleView/articleId/166/Ford-Pinto.aspx, accessed 16th April 2008.

it is to be a doctor, certain considerations will act as reasons for action for engineers qua engineers.

The development of a reliable capacity to respond to risk with the appropriate attitude is crucial because the role obligations of engineers will not be the sorts of things that can be specifically and exhaustively legislated for in codes and guidelines. Engineers will always be involved in seeking a balance between safety and other goals, so, for example, we ought not to think of the role obligations of engineers as taking the form of a singular general obligation to always and rigidly 'put safety first'. The special reasons that engineers have for choosing to risk or avoid risking qua engineers are likely to reflect a 'balance' of many different considerations. It is also likely that the exact nature of this balance will vary from case to case and that the process of arriving at it will require, on the part of the engineer, the possession of 'good judgement'. It is perhaps also worth noting that the achievement of balance is not a simple matter of weighing consequences and their likelihood against one another. It is a process of being appropriately influenced by appropriate reasons. The capacity to achieve the balance which best reflects engineering values will be a skill that is initially put in place by training, then fully learned and refined by practice. Our claim is that professionals acquire, though training and thought, settled dispositions to judge in accordance with their distinctive professional values and thus can be said to exemplify a kind of professional practical wisdom. These are general points about the uncodifiability of ethics and the futility of attempting to use codes of ethics as a regulatory guide, rather than precise answers expected to answer every single possible moral dilemma an engineer might come across. What codes of ethics can do instead is capture the general values and virtues which inform the profession, and the very process of trying to articulate what these might be for the purposes of the code may turn out to influence the shape of the profession itself.

If what we have said above is correct then we think it is fair to make the general claim that decisions made by engineers are not made in a 'moral vacuum' or a setting which expresses solely one's personal preferences and values, rather they are made by engineers qua engineers, professionals who have a conception of what their professional role entails, the demands of their commitment to this profession, who have internalised these demands and see them as giving them reasons for action. However it is not only the internalised values and attitudes that will be distinctive of moral decision-making about risk by engineers. In addition to character we also emphasised the role of circumstances in our discussion about moral risk taking. We now wish to maintain that there are certain reasonably generalisable features of the circumstances in which engineers take decisions about risk that can shed some light on the ethical decision making about risk in this field.

The Context and Nature of Engineers Decisions About Risk

Consider the following routine sort of engineering decision. An engineer who is designing a railway bridge in a busy metropolitan area has to decide the level of stress that the bridge should be designed to withstand. If the bridge fails to withstand the stress to which it is exposed, the lives of train passengers and those living in the

vicinity of the bridge will potentially be put at risk. However building anything to withstand stress is expensive and to build in more 'safety' than is necessary is to spend resources in wasteful way. So the engineer needs to decide upon how much 'safety' is enough in conditions where it is uncertain whether the sorts of events that would represent stresses on the bridge will manifest and, further, it is also uncertain whether, should the worst circumstances manifest, the resultant disintegration will a) occur and b) really do any harm. This example tells us quite a lot about the features of risk-involving decisions made in engineering. It tells us, first of all that what counts as a 'routine' risk for an engineer in the course of her practice and what counts as a routine risk for anyone else in the course of his daily life are very different. This is the case despite the fact that the sorts of risk-involving decisions that are taken by engineers in the course of their work are paradigmatically decisions to develop technologies which subject third-parties to some level of risk.

Furthermore these third parties who will be exposed to risk are generally:

- persons who are unknown and unknowable to the decision-maker (there is no relationship between the engineer and the people who may be affected)
- groups of persons of indeterminate but probably large numbers
- a subset of the persons who will benefit from the fact that the risk is taken. So, on the whole, those at risk will be mildly complicit in the risk-taking although the risk will not be equally distributed across all those who make themselves complicit by interacting with the relevant technology (people who live in the immediate vicinity of a power cable enjoy the benefits of power as do others but the others are not exposed to the risk of the cable coming loose etc.)

A final feature of the situation which characterises engineering risks that we wish to highlight is that decisions about risk made by engineers require them to weigh their concerns about risk against economic considerations. Engineering projects are expensive and it is probably safe to say that funders operate under a requirement to maximise efficiency (either in order to produce profit or because they are under obligation to maximise the benefit that can be achieved with public funds). It is also usually the case that the demands of efficiency and safety/minimisation of risk tend to conflict. Lets look at each of these features of distinctively engineering decisions in more detail. First of all the expansion of the notion of routine risk. What justifies this? Presumably it must be the combination of the expertise (knowledge) and professionalism of the engineer that legitimizes this widening of the category of 'routine risk'. What then is it that the engineer knows which the lay-person does not and which is relevant to the legitimacy of risk? We might be able to claim that the engineer is in a position to better estimate the extent of the risk involved—the possible consequences and the probability that they will occur. Arguably lay-persons are inclined to under-estimate the significance of future risks in the face of immediate pay-offs. It is quite plausible that this is the case for major technologies—flying in an aeroplane must be incredibly risky but most of us do so without concern because we simply have no idea of the risks involved. An aeronautical engineer on the other hand would have to have a much clearer conception of the risks involved in order to do his work. But we might want to ask whether this sort of knowledge is enough to explain the difference.

In our earlier analysis of reasons for risking we said that there were two main components to the reasons for risking or not risking—one is a clear grasp of what is being risked, the other (which we emphasized) is a clear (and correct) understanding of the relative value of the risk-taking which is reflected in the extent to which the prospect inspires sentiments of something like caution and ambition and in the way that these are balanced against each other in a rational manner. Might we reasonably expect engineers to be better than the average person at responding to projected risks which the right amounts of caution and ambition or at achieving a situationally appropriate balance of the two? If we could expect this of engineers it would not be because of what they know but because of what they value and the way in which they make evaluative decisions. If engineers as a group or profession do more reliably make good value judgments in the context of their work then this would have be the result of their (uniform) experience or their education or both. We have argued above that part of being educated as an engineer is coming to practice, grasp and embody the norms associated with distinctive role obligations. If this is correct then it may well be that we can expect engineers to make particularly good evaluative calls when it comes to potentially risky actions within their field because they have internalized the goals and values associated with their role.

Now let us consider the question of the status of the persons who are exposed to risk as a result of decisions made by engineers. On a consequentialist account of the morality of risk-taking, the fact that they are likely to be 'unknown numerous others' will probably entail that the extent of the legitimate risks is proscribed quite significantly. However it would appear, as we have seen above, that this is not how things work. Large risks are taken in the course of engineering projects. A consequentialist might explain this in terms of the great value of the technological progress that will be balanced against concerns about putting people at risk. Alternatively we might appeal to the fact that the sorts of projects engineers work on enjoy a high level of public support—on the whole people just are prepared to take risks in order to enjoy the freedoms and conveniences offered by successful engineering projects. Whether it be the development of an iphone or the building of the Channel Tunnel people are very cavalier about the dangers associated with the use of or proximity to engineered products. But that does not seem quite right either. We can be legitimately concerned that such cavalier attitudes are the product of (wilful) ignorance about the risks and the non-generalisable conclusions of many personal cost/benefit assessments.[31] What all of this and in particular the 'unknown' nature of the risk others are exposed to suggests is that the engineer should be concerned with whether any reasonable person could object to being exposed to the risks associated with the choices she makes in the course of working on a project. The person exposed to the risk considered by the engineer has to be hypothetical and any good reason they could have for objecting to the decision ought to be influential—probably definitively so. If this is so then part of the education of the engineer—part of what enables him to fulfil his role obligations will be an ability to sort reasonable from unreasonable objections to risk exposure. It is our contention

[31] The latter would be avoidable if we considered attitudes 'behind a veil of ignorance' rather than actual attitudes.

that such an ability is a skill which cannot be taught explicitly but which is acquired through observation of those who do it well and through skill-building tasks which involve a gradual increase in responsibility for decision-making.[32]

Finally we turn to the balancing of economy and safety—this is perhaps the area in which engineers are most conscious of being caught up in a moral dilemma—as illustrated by the Ford Pinto case discussed above. We saw above that the company took a pure financial cost/benefit approach to decision making about the acceptability of risk. We also saw that such an approach clashes with the sort of conception of the engineering profession that engineers themselves subscribe to in the course of pursuing work with which they identify. We have argued above that the engineer will have acquired and will identify with a distinctive set of values which are associated with her role and that real-life examples suggest that the capacity to be moved by the need not to endanger the safety of strangers will be central to such values. It follows that her assessment of the potential costs—in human rather than financial terms—will and should be a powerful determinant of her choice of the degree of 'safety' she builds into the design of the car. She has reason to be cautious that non-engineers would not have. This is why we entrust the design of bridges to engineers rather than to developers and entrepreneurs.

Responsibility of Engineers for Risk

We began our deliberations on the concept of 'risk' by remarking that the element of control inherent in risk making decisions had implications about moral responsibility. The above discussion has suggested that judgements about responsibility with respect to risk have to be made within the context of how we perceive engineering as a distinct profession, instantiating particular values and goals, which reflect social contexts and are internalised by engineers so that they are motivated to act in accordance with specific reasons qua engineers.

More often than not, engineers will find themselves in the position of decision makers, people whose decisions will primarily affect others, which, in and of itself, imposes specific obligations with respect to risk making decisions. Furthermore, and crucially for engineering, we accept some risks, such as driving, as reasonable and seek merely to manage rather than eliminate them. Within the context of a profession that seeks to harness technological advancements to provide solutions for a wide range of social problems, from food supply issues, to manufacture, to solutions for global warming, etc., a social conception of which risks should be managed rather than eliminated is indispensable. Clearly this judgement, about social values and goals to which engineering skills can be put to use, is not one for individual engineers to make in isolation, rather one for the profession to consider and which will need to be revised in light of new possibilities and in the context of how we all understand engineering as a profession.

Such thoughts are not entirely novel and indeed shape our conceptions of other professions such as medicine. In the health care context, a socially shaped understanding of the appropriate aims of medical skills guides the decisions doctors

[32] For more on this see Athanassoulis and Ross (2009).

make. Shared professional values about what medicine is and what medical professionals should aim at, form the centre of every individual doctor's deliberations on difficult decisions such as those involving taking risks. A similar model applies to engineering, even if it is the case that some of these values are less clearly articulated, or instantiated in the education of engineers, or present in the consciousness of the public when thinking about engineering.[33]

Furthermore, weighing up different outcomes, with the complexities introduced by the nature of the harm/benefit risked, its duration, distribution, relation to complicity or culpability for bringing it about, availability of other options, and all the other factors identified in the first part of this paper during our analysis of risk, can only be undertaken within the context of the values of engineering as a profession. None of these factors can be considered in 'a moral vacuum', rather must be understood and weighed within the context of the engineer qua an engineer.

Conclusion

We began this paper by considering how one could go about deciding what would be good or ethical risks for engineers to take. We presented an analysis of risk and its ethical implications which suggested that there is no clear formula with which we can answer this question—the ethical implications of risks taken by any agent will depend upon the quality of the reasons the agent has for taking the risk as well as the prevailing social and physical circumstances in which risk-taking occurs. We then went on to suggest that engineers, like doctors, identify with a set of goals and standards of good-practice which has an important effect on their risk-taking decisions. In particular the subject-dependence of ordinary reasons for risk-taking is substituted for profession-dependence. So we can expect a level of uniformity in the decisions made by engineers about what is and is not an acceptable risk. In consequence there are public, profession-endorsed, standards by which decisions regarding risk in an engineering context can be judged. The difficult decisions of deciding exactly what these standards should be and how they are to be interpreted within particular contexts still remains for engineering as a social profession to determine.[34]

References

Athanassoulis, N. (2005). *Morality, moral luck and responsibility: Fortune's web*. Basingstoke: Palgrave.
Athanassoulis, N., & Ross, A. (2009). A virtue ethical account of making decisions about risk. *Journal of Risk Research*, (forthcoming).

[33] Considering such issues about education and training is well beyond the scope of this paper, but are clearly related to our deliberations here.

[34] We are grateful to three anonymous reviewers for their very helpful comments on an earlier version of this paper.

Hansson, S. O. (2003). Ethical criteria of risk acceptance. *Erkenntnis, 59*(3), 291–309. doi:10.1023/A:1026005915919.
Hardimon, M. O. (1994). Role obligations. *Journal of Philosophy*, *91*(7), 333–363
Korsgaard, C. (1996). *The sources of normativity*. England: Cambridge University Press.
Lewens, T. (2007). *Risk: Philosophical perspectives*. Oxon: Routledge.
MacIntyre, A. (1981). *After virtue*. London: Duckworth.
National Society of Professional Engineers. (1957, June). *Canons of ethics for engineers*. Available on line at: http://ethics.iit.edu/codes/coe/nat.soc.pro.engineers.June.1957.html, Accessed January 26, 2009.
Rawls, J. (1999). *A theory of justice: Revised edition*. Cambridge, MA: Harvard University Press.
Rescher, N. (1983). *Risk: A philosophical introduction*. Washington: University Press of America.
The Engineers' Council for Professional Development. (1941). *Science, 94*(2446).
The Royal Academy of Engineering's report "A Universe of Engineering", chair Sir Robert Malpas, http://www.engc.org.uk/documents/Malpas_report.pdf, Accessed January 26, 2009.
The Royal Academy of Engineering. *Statement of ethical principles* http://www.engc.org.uk/documents/Statement_of_Ethical_Principles.pdf. Accessed January 2, 2009.
Thompson, J. J. (1986). *Rights restitution and risk: Essays in moral philosophy*. Cambridge: Harvard University Press.
Williams, B. (1993). Moral luck. In D. Statman (Ed.), *Moral luck*. State University of New York Press: USA.