The Royal Academy of Engineering

PHILOSOPHY OF ENGINEERING

Monday, 27 March 2006

at 29 Great Peter Street, London SW1P 3LW

Chairman: Dr Keith Guy, FREng

PHILOSOPHY OF ENGINEERING

Monday, 27 March 2006

WELCOME & INTRODUCTION

Dr Keith Guy, FREng

Let me welcome you to this seminar. I sit on the Engineering Policy Committee but I have done none of the work for this occasion: it has all been done by primarily Natasha McCarthy and also by Tony Eades.

Natasha presented a paper at the Engineering Policy Committee, a version of which has since been published in *Ingenia*, outlining the opportunities for philosophy in engineering. It was the first thing that I had seen in any of the Academy's programmes that I actually felt excited about. The way the board works is that, if you grab it first, you get to look after it – and it has been a great privilege to do that. This is an extremely exciting challenge, not least because I truly believe that, if as engineers we can engage philosophers to say that we are worthwhile, then others will believe that far more than if we say it ourselves.

Let me hand over to Natasha, to outline the programme.

BACKGROUND TO THE MEETING

Natasha McCarthy

I would like to reiterate that welcome and thank you all very much for coming – it is pleasing to see a group of such esteemed philosophers and engineers here together and I hope we will have a very fruitful discussion.

My background is in philosophy – the philosophy of science and the philosophy of physics – but I came to the Academy as a policy adviser last summer. Coincidentally, at that time, Andrew Fox, who is here from Coventry University, approached the Academy, asking whether they ever did any work in the philosophy of engineering, or whether they would be interested in having some seminars on this topic. It was therefore given to me to scope out this project, to see whether there was any mileage in the philosophy of engineering.

I did some research, which was the background to the paper that I sent to all of you. I found that there was some work in the philosophy of engineering, primarily in the US, some of it in Europe and a little in the UK. However, I really thought there were some interesting topics in the philosophy of engineering – metaphysical topics, epistemological topics and moral and ethical topics. Hence I thought this was something in which the Academy should be involved.

The Academy is very well placed to be involved in this because it has a broad fellowship of engineers with different backgrounds who are very interested in topics outside of engineering.

The aim is to have a series of seminars on topics to do with the philosophy of engineering. The purpose of this particular meeting is to have a group of philosophers and engineers with an interest in this area, who will identify the questions and issues that they think are most interesting. The idea is that this will allow us to find those topics which are of mutual interest to philosophers and engineers, and on which philosophers and engineers can work together. Our objective for today is to arrive at a series fruitful topics on which to have future seminars.

During the afternoon, we will break out into discussion groups. You have been given some questions to discuss but these are just ways to start off discussion, and I would not want you to feel limited to discussing just what is included on that list. Please raise whatever is of interest to you. After that, we will have a session in which to report back to the main meeting. I hope those discussion will be very fruitful and that you will enjoy them.

Keith Guy: The first part of this seminar will include four presentations. I would like to thank our speakers for agreeing to talk. We will start with presentations on 'What is Engineering?', and we will hear from Professor Igor Aleksander first.

WHAT IS ENGINEERING?

Professor Igor Aleksander, FREng

First, let me thank Natasha, Tony and everyone else who has been kind enough to ask me to speak at this meeting. I would not have missed it for the world because I have always been interested in philosophy, although I am not a professional in the business.

Some of you might know that the work I do has to do with the cogs and pulleys of what it is that makes our brains conscious, and one immediately confronts the views of philosophers there. Consciousness has two ways in which it relates to this question of what is engineering. First of all, the question 'what is consciousness?' has parallels with the question, 'what is engineering?'   Most of us here do engineering, but if anybody asks us to explain what it is, we find it very difficult. However, that does not stop people from trying. If we ask the question, ‘what is it to be conscious?’, most of us are conscious, but we find it very hard to answer that question, but that too does not stop people from trying.

Another point is that, as an engineer, I have found an interesting relationship with philosophers. It is from this interesting relationship that I would like to look at the question of what one is engineering. Many people working in consciousness become very upset because they are constantly being criticised by philosophers. One typical example of that is Francis Crick, who I see not so much as a scientist as an engineer, because he created models of the DNA molecule – he recreated them until he got it right. He became very upset, however, about being criticised by philosophers for his work on consciousness, and he had a great down on philosophers. If you read his book, *Astonishing Hypothesis*, it is there in black and white. What did philosophers ever do for mankind?, is the question. When did they ever invent anything?

I happen not to take that line. It is right, and it is the job of philosophers to be critical – and for each and every one of them to be critical in a different way. That has helped me a great deal in developing my engineering.

[*Portrait of Georg Wilhelm Friedrich Hegel*]

Let us start with this fellow. Most engineers will recognise him as the chairman of the last interview panel they attended, whereas philosophers might recognise him as Georg Wilhelm Friedrich Hegel, who brought up the end of the Age of Enlightenment. He had an interesting way of developing a philosophical argument, which has an engineering overtone to it. He suggested that, whatever it is you want to argue about, a thesis should be accompanied by an antithesis – in other words, a criticism of the thesis. These would then be brought together in a synthesis, and that synthesis could become another thesis. Thus, this process could go on, until you had a good argument. Engineers do that with design: you have a design, you criticise it and then create a new design, which can then be criticised further. At the end of the day, you get fed up with criticism and creating new theses, and you have a product.

I would like to apply this sort of idea to what it is that people say engineering is.

Thesis

The subject of the thesis is, ‘what is engineering?’. We have some folk views and some less folksy views. You read things like those on this list. You can look it up on the internet, you can look up engineering in encyclopaedias. You can look at what great engineers have to say about it, but you end up with a list. It could be something like:

- creating something useful, using maths and science – and that is an American dictionary definition;
- turning ideas into reality;
- practical application of science to commerce and industry;
- making properties of matter useful in structures and machines – that is rather old-fashioned, and I will tell you why in a moment;
- and this is what I was told by my Dean of Engineering on my first day as an engineer at university – 'making for a dollar what any fool can make for 10.'

There is a more serious side to these definitions. The Royal Academy of Engineering brought out a publication called *The Universe of Engineering*, which tries to define engineering. This is quite an interesting approach to the problem because it is just blanket – it covers everything. You can design, build, make operate, sustain in each cycle or retire, for a specific concept or model – so that is quite good. I will look at that again in a moment.

Antithesis

Turning now to the antithesis, I would like to look at what would be a philosophical perspective on these things that I have just quoted. Here, I want to lean on another fellow, Alfred North Whitehead, who wrote *Principia Mathematica* with Bertrand Russell. He was very concerned with teaching philosophy to mathematicians. He felt that a mathematician was not fulfilled unless he could stand up in a philosophical argument. He defined as philosophers those who, first of all, develop their own premise for future inquiry. So the engineer who just goes for the last specially promoted programme of EPSRC is not a philosopher – he has got to go for the responsive mode.

There are those who are dissatisfied with the concurrence of eminent and sensible people, including themselves. We all sit there comfortably and say, 'This is how you do this, or this is how you do that', but perhaps one should always question these things.

There are those who assault the boundaries of finitude. Yes, you hear people saying, 'No, you cannot do it this way because it will cost too much', or whatever, but then the next day someone has done it.

There is then the 'endeavour to enlarge the human scope of application of every doctrine that constitutes a framework of thought'. That is very much what engineers actually try and do – they try to make sure that their frameworks of thought turn out to be something that is of human benefit. This is the antithesis which we should now apply to those things that we saw earlier.

What is a philosophical perspective?

‘Creating something useful using maths and science’. Useful is a relative term. If something is useful to side A in a war, it may not be so useful to side B, and in fact it may be destructive and, at the end of the day, it may not be useful to anybody. That points to a question of ethics in this sort of philosophical argument.

Using maths and science: the reason I have written here ‘using a toy train set’ is because one of the best experiences that I have had as an engineer was to visit Alec Reeves’ laboratory. He was the guy who invented pulse code modulation and he did not care for maths and science. In his office, he explained how pulse code modulation worked with a toy train set – that was his way of doing it. So you do not have to do maths and science – although most of us find it useful!

I will not say much about turning ideas into reality! [Slide reads: ‘having an affair with the neighbour’s spouse’]

The practical application of science to commerce and industry: yes, without us being interested in commerce and industry, there would not be all that much scope for engineering. But, one must remember that that leaves out engineering innovation in its own right. For example, I do not think that Claude Shannon, who worked for Bell Labs, saw what he did as a practical application of science to commerce and industry, but he did see that he was handing over an idea that would revolutionise what is done in communication.

Making the properties of matter useful in structures and machines – that again leaves out the informational side of things. When I was a student, everything had to do with matter, in one way or another, whereas, nowadays, it does not help to think of information as matter, though it helps in some cases.

‘Making for a dollar what any fool can make for 10’: most eminent engineers have been totally wrong about their assessment of costs over the long term. I remember someone once saying that the way you make large memories for computers is to use ferrite core, because anything else would cost too much – like transistors, perhaps.

Too-broad definitions self-destruct

On this, I have made some kind of criticisms. If a definition is too broad, it tends to self-destruct. What is left out of all of that is the questioning, critical dimension.

Synthesis – engineering from a philosophical perspective

Finally, there is the synthesis, which is a very simple message that is important to me. When I look at some engineers in the educational world, at Imperial College, being critical is not something that they feel that they are being asked to do. I would suggest that engineering from a philosophical perspective is a critical assessment and pursuit of method and processes. It challenges the performance claims for existing systems. I am saying these things not *instead* of the list of things we saw before, but something that might perhaps be additional if there is a philosophy of engineering, which all engineers might follow. It strives to cross the apparent boundaries of existing methods and continuously attempts to discover new ways of doing things.

The next speaker will probably give you the right answer to 'what is engineering?'

[*Portrait of Hegel*]

I will finish by thanking Hegel for leading us along. [*Applause*]

WHAT IS ENGINEERING?

Dr Robert Hawley, FREng

You could describe what I am about to say as random thoughts from an engineer: as a doctor of philosophy, I do not have the first idea of what philosophy is. However, I do know that engineering is the process that converts science into technology and then into wealth-creating products. That is a very simple definition.

This is a process that is dependent on the final product itself. For example, is it a complex product, such as a transistor laser? Is it a large product, such as a bridge or a turbine generator, which requires a knowledge of many branches of engineering? Or is it a simple product, such as a locking device, that requires basic mechanical engineering knowledge?

The process also has to take into account the operating conditions experienced by the product. Is it subject to high rotational speeds or extremes of temperature – hot or cold? Are there extremes of pressure – for example, must it work in outer space? Are there high vibrations, or extreme weather conditions such as snow, rain or sun? Or there could be high electrical or mechanical stresses? And so on.

It can be a process that follows very carefully previously developed design rules and procedures. For example, the use of the correct radius at the changes of physical angles on a rotating or flexible shaft: if you do not have the right radius, you can guarantee that it will fatigue.

It can be a process that has to develop the design rules to cope with situations not previously experienced, such as designing to cope with dramatically increased electromagnetic forces or high voltages. The development of these new rules can depend on very advanced mathematical simulation, or finite analysis, or results from experimental rigs, or operating experiences and resultant feedback where failures have occurred. With Duncan here, dare I mention the Millennium Bridge and its wobble: that is experience, and that is feedback. Whilst, for an experienced engineer, intuition is important, it cannot be solely relied upon.

I have wrestled with describing the philosophy of engineering that has enabled me to become an experienced engineer in many fields, so I look forward to hearing the next two speakers. However, let me first describe briefly the process by which I became an engineer.

At school, I enjoyed maths, physics and art and so, when I left school at 16, I applied to become a draftsman. After a day of interviews, BICC, to whom I applied, decided that I would become an apprentice in the high voltage laboratory, and that was a decision that led to the first part of my career. At 19, I went to Durham to read electrical engineering and then did a PhD which led me into physics.

I was then recruited by Parsons to set up a high voltage dielectric research laboratory. Those were exciting times, when I learned how to discover scientific facts and turn them into design rules. For example, there are the limits of electrical stress that would not cause breakdowns in large transformers.

I was then told to join the turbine generator design department as a designer. Here, I learned the design rules for conventional machines but, at that time, we were also developing the first 500MW turbine generators that were to operate at previously unused levels of electromagnetic stresses and mechanical forces. Those early 500MW generators experienced many problems that had to be solved by mathematical and analytical analysis. There were large test rigs and detailed measurements on prototype generators, and painstaking examination of all failures. All of this involved all areas of engineering. Here, I learned the need for clear examination of the facts, and careful analysis to propose the solutions which, whenever possible, benefited by being kept as simple as possible.

I was born thinking in straight lines and circles. Professor Peter Lawrenson, who many of you will know, of motor fame, taught me to think in other ways. Why should a shaft that revolves always be circular and not oval? Is that philosophy, or is that lateral thinking? And what is the difference? However, a logical thought process and an ability to think laterally, and to postulate potential unforeseen circumstances, was vital. At the same time, one had to be proficient in electrical engineering, mechanical engineering and many other areas – and experience is vital.

For years, I have been concerned about young graduates who go straight into a computer, without appreciating the interaction of parameters and they then believe the computer results. It is rather like a person who goes straight from school to university and then, with little experience, goes on to do an MBA and is then absolutely convinced that they can run a major company. Experience is important.

Let me explain how I came to my views on the wrong uses of a computer. In the early days at Parsons, we were putting in place a computer programme to design turbo-generator equipment. Prior to that, we were using slide rules and log tables – and some of you in this room will not even know what a log table is, but we were using log tables and slide rules. One day, in the drawing office, I asked the draftsman for the length of a rotor shaft, and the reply was, '67.3721 ft'. What? 'It must be true', he said, 'that's what the computer says.' So I asked him what temperature the computer said the shaft was at, and he did not understand that. Belief in computers can be very risky.

We then looked at the possibility of computer design and transfer directly to the machine tools, to manufacture the components – in other words, without drawings. I dreaded the day that the works would ring up and say, 'This turbine generator, which is a mile long and one inch in diameter – who, exactly, is it for?' Sir Terry Farrell tells me that they give young graduates a wall to design, without specifying the foundation conditions. Once complete, they put in those parameters, and the wall falls down. Is that experience, or philosophy?

Was it young people without experience who did the programmes for the beams that fell down in the Scottish Parliament, and at Wembley? The question to ask is, where does experience fit into the philosophical approach to engineering?

To be fair, in my career I was too busy to realise that I was becoming a broader engineer. I can no better than quote the president of the Rensselaer Polytechnic in the US, Dr Shirley-Ann Jackson – and Andrew Ramsey gave me this quote and so he takes full responsibility for it. I had the pleasure of meeting her in Washington and she said:

> "Since, increasingly, engineers create the settings for and means of human interaction, a grounding in the liberal arts is especially essential and necessary for engineering today."

Here are some of the reasons. There is a growing need for engineers to communicate, in order to be effective in engineering itself. There is a greater than ever need for broadly educated engineers to heighten respect for technological solutions and to alleviate a cultural fear that occasionally challenges progress. Engineers must be sensitive to the social consequences of their work, and that surely translates into ethics and ethical questions: i.e., ‘it’ – whatever it is – should it be done? Ethics education is critical.

Our society needs technologically knowledgeable individuals in its highest councils. A liberal education ultimately makes engineers more creative by expanding their minds and exercising their imaginations. This could all be regarded as what makes an engineer and hence defines engineering and the philosophy of engineering. Thank you. [*Applause*]

Keith Guy: Those were the presentations from the engineering side. Peter will now speak on the philosophy side.

WHERE IS THE PHILOSOPHY IN ENGINEERING?

Professor Peter Simons, FBA,
University of Leeds

Thank you for the invitation. I am very pleased to be here, although I feel slightly uneasy – what is a philosopher doing here in the Royal Academy of Engineering? There is a cartoon by my favourite American cartoonist, Gary Larson, which shows an elephant at a piano on a stage. The elephant is looking very alarmed and is saying to himself, ‘What am I doing here? I am a violinist!’

Engineering philosophy

My ambiguous title is meant to signal that I am not quite sure what it is that engineering and philosophy have to do with one another, and I am not sure that anyone else is, so I do not feel too bad about it.

How can engineers and philosophers help one another?

I will start by asking the general question, how can engineers and philosophers help one another? The fact that you are here suggests that you think that the answer is not, ‘In no way, whatsoever’, so we will press on. The 99.9 per cent of our respective professions who are not here may disagree with us, but we are not talking to them today.

One obvious area in which philosophy and engineering overlap – and we have heard the words more than once already today – is in the issue of engineering ethics. It is part of the duty of an engineer as a citizen to consider the ethical consequences of the projects in which he or she is involved. Similarly, it is part of the academic duty of the ethicist in philosophy to consider all possible applications of their theory. Since engineering applications are among those, then clearly there is an area where the two of them will overlap in their interests and concerns. Since my colleague, Jo Wolff, will talk about the ethical side, however, I will press on. I am not an ethicist – although I believe I am a good person!

One area in which philosophers claim that they can help anyone who comes to them is in conceptual clarification. There are definitely some areas of engineering that are probably ripe for conceptual clarification – one of which sees to be the concept of engineering itself, as we have already heard today. I was very pleased, incidentally, that Professor Aleksander mentioned one of my favourite philosophers, and not the one beginning with H, but the one whose name beings with W – Alfred North Whitehead – who was an under-rated genius of British philosophy. I will come on to what kind of clarifications there may be in a moment.

In terms of the help that engineers might offer philosophers – engineers can at least offer a great many real life examples to philosophers, either in the ethical field or in the cognitive field. I will talk about some of those later on. Philosophers looking for unusual and different examples might well look away from the history of science, for once, and look at the history of engineering.

This is the area I shall spend most of my time talking about today – it is the notion of ontological engineering. This is a relatively new concept which is not yet completely formed. Ontological engineering has two sides to it – the ontology in engineering and the engineering in ontology. I will tell you what ontology is in a moment – I do not want to get ahead of myself.

Ontological engineering can take place either directly – that is, engineering applications using ontology and ontological concepts or, more frequently, as happens, it takes place via data models which, for large scale engineering projects, are computer-based models of the processes and objects involved. Where philosophy and engineering rub shoulders, there is usually a computer scientist somewhere between the two shoulders, and that is where things can go extremely wrong. I will say a little about that in a moment.

Conceptual clarification areas

I will briefly mention one or two possible clarification areas. What you are interested in as a philosopher or an engineer is a personal matter, and these are areas in which I am not so interested, so I will go through them fairly quickly.

There is a lot of background in philosophy of the theory of action, intention and realisation of intentions. This can be anything from building a bridge to seducing your neighbour's spouse. There is a great deal of philosophy about that, and a good deal of background going back two and a half thousand years. I would be surprised if there were not theories and concepts in there which could be of use for analysing engineering intent and action.

Typically, engineering is engaged in in groups – and sometimes by quite large groups. There is also a certain amount of analysis of social intention, action and responsibility in the philosophical literature, which could possibly be drawn on.

Something on which I have worked with engineers is the notion of strategy, which obviously applies to enterprises. Another area I have worked on, together with some other people such as the University of Delft, is the concept of an artefact. These things tend to be Sunday occupations for engineers, I imagine – if that.

Here, however, is something which it is well worth stressing. Philosophers, being cognitive types, tend to stress the importance of cognitive knowledge – propositional knowledge in the head, what you know. Engineers, however, are not wholly cognitive types and they have to make things that do things, which requires a certain amount of technique. I have philosopher colleagues who only have to walk past a computer for it to crash – that is how good their non-propositional knowledge is.

There is a very considerable tendency in the academic community in general and in the philosophical community in particular – and even more particularly in the *English* philosophical community – to de-stress the importance of practical knowledge, or non-propositional knowledge, or what Gilbert Ryle, philosopher, called 'knowledge *how*'. Especially important is the way in which it is transmitted from one person to another, or from one group to another, because this is typically done by word of mouth and by example, rather than by writing manuals or giving written instructions. It would be very important from the point of view both of understanding the notion of knowledge from the philosophical point of view, and from understanding the transmission of what might call the culture of doing things, as distinct from the technique of doing things, if we focused a little more on the nature and transmission of non-propositional knowledge or know-how.

Engineering as a source of problems, examples and tests

From a philosophical point of view, one of the reasons for being interested in engineering is that it gets you away from pocket-sized examples. For instance, it is one thing to build a Lego aeroplane, but it is another thing to build a real aeroplane.

It can be of use in testing the categories that you apply. Most philosophers have a very narrow focus these days and are only interested in a few problems – their reading tends not to go beyond the latest issue of the journal that specialises in their specialism. However, there are important jobs for them to do in, for instance, the area of categories and classifications. This has applications in engineering, as I know from talking to engineers who have to apply it in a very wide range of contexts.

There is a requirement when you are an engineer to see that all aspects of the system work together, and philosophers can and should learn from that. Hegel was probably the supreme systematist in modern philosophy as regards the philosophical system although, personally, I would not trust an aeroplane if it had been designed according to Hegel's philosophy. Nevertheless, there is a requirement on philosophers to be systematic – although, unfortunately, current academic conditions tend to militate against that.

There is the prospect, which philosophers very rarely have, of getting some kind of genuine empirical feedback on those theories. Build a plane and see whether it falls down out of the sky, using your philosophy – obviously, the feedback is a little more indirect than that, but the general idea is that there might be something other than just a nasty article from somebody who disagrees with your view in another journal, which could make you think about whether your theory is correct or not.

Ontology and engineering

Here are a couple of ways in which ontology and engineering can get together. Ontology means two things in the modern world, unfortunately. There is the good, old-fashioned sense in which it is part of philosophy, meaning the theory of what there is, but nowadays it also means – in cognitive science and computer engineering – 'platform independent conceptual schemes'. Unfortunately, both of those are relevant because, if you are to do a good platform independent conceptual scheme, it had better line up nicely with reality, which means that you have to use ontology in the philosophical sense. Currently, the line-up is not especially good in many areas.

Philosophers can draw on, and engineers can benefit from, long-standing accounts of various very general concepts, such as property, object, part, quantity, process, relation and action – all of which enter any adequate description of what engineers do. Probably, the best thing that an ontologist could offer an engineer is some kind of system which allows the engineer to check, directly or indirectly via people who understand what the point of the systematic checklist is, whether they have taken all possible points of view into account. This would obviously be something that should be taught to engineers, but it would come within the realm of the know-how that is imparted to them, rather than the know-that which is laid down in manuals.

Philosophers are very adept, and in fact probably too adept, at talking about language and nomenclature, which has been their stock in trade for 100 years, since Bertrand Russell. Interestingly, his principal pupil in philosophy was also an engineer, called Ludwig Wittgenstein. Whether he was a good engineer, or indeed a good philosopher, opinions are slightly divided.

Philosophers are not scared by abstraction and breadth of topic; engineers are not scared, at least in their waking hours, by complexity and detail. They should therefore be able to help each other in that way.

Two examples

Here are two examples, both from my own experience of working with engineers in various ways. The first is on the notion of mereology, the theory of part and whole – something I wrote about from a conceptual point of view nearly 20 years ago now – as applied to what I call the multiple Bill of Materials (BoM) problem, and I shall explain what that is in a moment.

The other area is enterprise engineering, which is using engineering, especially systems engineering principles, to design companies or parts of companies. In other words, it is to treat an enterprise the way you would treat a product, from a theoretical point of view. There are reasons why you would need to do that.

Multiple BoM problem

Let u s look at the multiple Bill of Materials problem. Bill of Material is a mereological or part/whole breakdown of a complex artefact and it turns out that there is not just one for any sufficiently complex object but there will be several – at least three and as many as 10 in different cases. There are good reasons why there are different breakdowns, having to do with the different jobs that the different kinds of engineers do. The discrepancy can be extremely large and 40 per cent is a case that I have come across, and that is an extremely ripe and mature enterprising project, namely the Lockheed Hercules, which first flew in 1954 – so this is an aeroplane that is over 50 years old and yet, between the different Bills of Materials, there is up to a 40 per cent discrepancy, even in the latest version. I have that from the horse’s mouth.

MEREOS

Multiple Bills of Materials are extremely disadvantageous for various reasons. As a result of this, the United States Air Force Manufacturing Technology Division set up a project in the 1990s called MEREOS, the analogy with the notion of mereologies, which was intended to do some conceptual modelling to see how to tackle that problem. I was involved with that project indirectly.

MEREOS was meant to be a system for tackling the integration of multiple Bills of Materials and here you have two schematic outlines of Bills of Materials on the left and on the right, as produced by different engineers – and there are differences between them, as indicated by the yellow linking lines.

Structure definition

One kind of difference is when you have to make one part out of several parts, for engineering reasons. In this case, the rocket nozzle cannot be engineered as a single piece and therefore has to be made in three parts which you join together – called articulation.

Another is factoring, which is where you have some very complex part which has repetitive features in different places. In this case, we have three helical gears with slightly different sections and different widths and it turns out that the easiest way to make those is to make one long piece of helical gearing and then cut it up and do further milling. This is called divergence fire factoring.

Process definition

You need integrated declarative process representation if you are going to take account of this. This is meant to be the computer representation of what is going on in this case.

[*Slide*]

What you end up with in the end is the poor guy in the middle – in this case, the designer – who has to cope with the discrepancies between the two different kinds of representation of the artefact, for instance, as designed and as manufactured. He deals with it via a database which enables you to do proper costing. The idea is that you make all sorts of efficiency savings here. The fact that this guy is sitting on a wall is absolutely intentional: these are two different kinds of engineers, and there is a wall between them and they do not talk to each other readily.

Enterprise engineering

The other way in which philosophy has at least had a hand at trying to help engineers is through enterprise engineering. This was because of problems in adapting an old industry to new circumstances. The industry in question is the aerospace airframe manufacturing industry in the United States. They were having all sorts of difficulties in adapting to the post-

Cold War situation in the 1990s and these were some of the issues that they were facing at the time. The fact that these were real is shown in that, in the 1990s, the number of major airframe manufacturers in the United States decreased to the number two – basically, Aerospace East and Aerospace West.

Example

The example is Lockheed Martin at Marietta Aeronautical Systems who make the F22 and the Hercules. They were having trouble in the 1990s. Their methods were expensive, and they still worked on Cold War lines when they were supposed to be working along modern, market-oriented lines. The government was no longer prepared to pay them an open-ended cheque for the planes they produced and so they needed things to help them re-design their enterprise.

LMAS strategic intent

This is one of the things that we produced to help them re-design their enterprise. I will not go into detail but this was essentially a graphical representation of how to check when you are trying to change the enterprise from how it is to how it should be, and whether you have got everything sorted out. It is essentially a large, graphical check-box list. This is based on a philosophical analysis of intention, action and realisation of intentions, in conceptual and physical surroundings, with labels.

Formal architecture

Here is part of the formal stuff that we were using to design this. This is based on database design ideas, going back to IBM in the 1960s, and it is called a three-schema database. We applied that to the architecture of LMAS and came up with that list of checks that we ought to take account of.

LMAS enterprise architecture

Here is another piece of graphics which shows essentially the same thing, iterated and applied in the context of the model of intent realisation. In this case, the intent was to govern the enterprise by a systemic strategy process. It failed because that part of Lockheed Marietta was actually taken over by another part, which ran on completely different lines. So Georgia became part of Texas, as far as Lockheed Marietta was concerned.

As you have probably realised, my own experience of co-operation with engineers is that, whenever I come into the picture, then within two or three years, the whole system has folded – which may not be a very good advert!

What philosophers would need

What would philosophers want from engineers, in order to pursue the idea of philosophy of engineering, and engineering in philosophy and vice versa? We need open-mindedness. We need a willingness to co-operate, by which I mean from both sides. I actually have a much worse reception from my philosophical colleagues than I do from engineers – at least I have novelty value to engineers!

We need a little funding, including pump-priming support. We need somewhere where we can publish. We philosophers have one product: it is called publications. That is how they measure our achievements. We also need a little respect from both sides. And that is it. [*Applause*]

WHERE'S THE PHILOSOPHY IN ENGINEERING?

Professor Jonathan Wolff

University College, London

After what you have heard, what I am about to say may seem rather mundane. The issue that I shall be talking about – Peter mentioned engineering ethics, but it is not even engineering ethics as a whole topic, but rather one area where I, as a moral and political philosopher, have been involved in a small corner of engineering. This was on the topic of railway safety.

The issue came up very strikingly to me, because a certain amount of engineering deliberately addresses questions of safety. Probably, virtually every project that anyone carries out from an engineering point of view involves some sort of risk and issues of safety. Particularly if you are in a large-scale project, you will be imposing risks on people – whether it is workers or members of the public, or perhaps the resulting project will have a dangerous aspect in some way. Every engineer has to take into account the question of 'what risks should I allow to go forward?'

From the point of view of moral philosophy, this is a moral question. It is not an engineering question, or it is not *only* an engineering question. What is the tolerable level of risk? Under what circumstances is it right to impose risks on people, and under what circumstances is it wrong to impose risks on people? This is an area that all of you probably have to deal with in your professional lives. Sometimes, you will look at a manual which will tell you the answer, so that you do not have to think very deeply about it, although it is custom and practice and so on. However, that custom and practice became established somehow, and sometimes it should be criticised or re-assessed. So how do you go about doing that?

As far as I know, this is a topic that has not been discussed by philosophers in the UK at all. I must be wrong about that, and there must be some philosophers who have discussed this question about when is it right to impose risks, and when is it not right? Some American philosophers have discussed this, as have some in Continental Europe. There is a very interesting volume called *Values at Risk* which was published in the US in the 1980s. This came out of a symposium sponsored by the nuclear power industry in the US. They gathered a dozen excellent philosophers together and paid them to think about risk. They did so and came up with a very good volume and then they went back to their day jobs afterwards. Thus, we have one volume of philosophers paid to think about this area for a few months – they did the job, and then went away again. All the papers from philosophers in that volume were in a state of first draft, although they were published.

We have an interesting situation in that you, as engineers, have a fundamental moral problem to deal with. We, as philosophers, have not really stepped in to give you any help with this. Maybe you think that is just how it should be, and that engineers do not need any help in things like this, and there is no need for a philosopher to come in. The reason I became involved in this had nothing to do with the love of the issue of railways or safety. A former student, who had been working for the Health and Safety Executive, and then started working for what was then called Railway Safety, which is now the Rail Safety and Standards Board, got in touch with me. He said that his boss had come up with a moral problem and they wanted a moral philosopher to take a look at it – a rather intriguing idea. They therefore sent me a paper, and I was just amazed to see what it said.

The point of the paper was to set out a moral dilemma. Something that is probably very familiar to all if not most of you here, or at least the engineers, but which was completely unfamiliar to me, was the way in which safety decisions are theoretically taken now, or at least the way in which they are regulated by the Health and Safety Executive. There was this concept of VPF (Value of Preventing a Fatality) being using in a type of risk cost/benefit analysis.

For the philosophers in the room who do not know about this, imagine you are carrying out a project which may impose risks, or perhaps you have a factory which has a machine which is somewhat dangerous, or machines of that type occasionally kill people. There is a safety innovation that you could introduce, but it is expensive. Under what conditions should you introduce it? Naively, most people would think that if you have a dangerous machine and you can make it safer, then you have an absolute obligation to make it safer. But of course, as you all know, that is no good, because if you had to make everything safer, then you would just close down the factory and then you would not be causing any risk to anyone at all. There must therefore be some sense of when a risk is worth taking. We need to look at the benefits of taking a risk, and the costs of doing so.

The regulatory regime that we have from the Health and Safety Executive, tells you that you have to put a value on a statistical life saved, and you have to work out how much that particular project would value a life. So, if the project would cost you £10 million and it will save one life, then that project values a life at £10 million. If it is going to cost you £10 million to save a life, should you do it?

The regulations say that you do not have to do it, because if the VPF is more than £1 million, or grossly disproportionate to £1 million, then you do not have to introduce that innovation. If it is less than £1 million you have a statutory duty to introduce it, a duty to save a life. It was amazing to me, to read this, but this is apparently what we do in this country. We put a value on saving a life and we apply this to decide whether or not safety innovations should be introduced.

Who asked me about that? Is this what we should be doing? Who came up with this idea? Why £1 million and why not £10 million, or why not £5? Where did we get that figure from? Is this the right way to think about it in any case? Should we be putting a value on a life in that way, or in any way?

This is set out particularly starkly in the document that was sent to me in about 2001/02, at a time when the railway industry was under a lot of pressure. There had been a number of high profile crashes – although not many people were dying. On average, even during those difficult times, if you take years as a whole, the average number of deaths per year – actual deaths rather than serious injuries – was probably no more than 10 per year for passengers. This was considered scandalous, compared to the 2,500 people dying on the roads and so on. Ten fatalities on the railways was considered a major national scandal. Safety innovations had been brought in, valuing each life saved at around £8 million – already eight times as much as the statutory requirement.

Under discussion was a new safety system, ATP. It is very hard to know what this would value a life at, because so few people were dying. It might save two lives a year, probably, and it was going to cost about £6 billion. That was the cheapest estimate, and it was probably on the low side. The system that was under discussion would also reduce capacity by 10 per cent on the railways, people might die putting it in and so on. There would be massive disruption, and so there would have been huge negative costs for bringing it in. *The Observer* was running a story saying that the industry was negligent about safety, in not spending £6 billion to save two lives a year, and have all these other effects.

Suppose that you valued each life at about £100 million. If you have £100 million and your goal is to save lives, then probably the most stupid thing to do is to spend it on railway safety. You could spend it on road safety, you could spend it on the NHS, or you could give it to the Department for International Development – you could do whatever you wanted with it, but the last thing you should do with it is to give it to the railways. Yet that was supposedly what public pressure was for.

If you think about things in those terms, it was absolutely absurd to think that we should be spending more money on railway safety. However, you can have that thought and then, next morning, you switch on Radio 4, and there is someone from a bereavement group – the mother of someone who lost their life on the railway – saying, 'In this day and age, it is a scandal that we have trains that go through red lights! One death is too many. How can we tolerate a system in which anyone can die at all? This is the 21st century, and we still have people dying on the railways.' And you think, yes, that is right, actually.

There are two ways of looking at this. One the one hand, it all looks far too expensive and we should not be doing it. On the other hand, how can we condone any death that is avoidable? We know how to avoid it but we just have to spend the money. I thought this was a very interesting problem. You have one group of people who will say that you should not be spending this money, and we have another group of people saying that you should not even be calculating in terms of how much it will cost – because there should not be any financial consideration, and we just have to do it. You have those two groups. In fact, however, for most people, you do not have two groups – you have one group, which is all of us, who can switch backwards and forwards between thinking about it in these terms. It looks far too expensive to spend all this money for rail safety but, on the other hand, look at those people who are dying and we could have done something about it but we chose not to do so. Most of us are flipping backwards and forwards between these two moral perspectives.

How do we go on from there, and where do we move forward from that? One of the problems is that these ethical questions have been considered in a purely formal framework. That is, engineers and economists are very good at dealing with things that they can put numbers on. Thus, if you can have a way of approaching an ethical problem, which allows you to encode it into a numerical framework, that will go down much better with engineers, with economists, with lawyers and regulators, and with everyone who is involved with the actual process of having to put these things into practice. It gives a consistent and an accountable approach, and it gives you something with which you can be comfortable.

The problem is that the standard approach that is used does not capture all the issues that normally trouble people when they are thinking about moral questions. We have a situation where the railway decision-makers feel that they are doing everything that is required of them. They are valuing a life at £1 million – 'Well, we got that years ago, and we are on £8 million, or £20 million.' Railway engineers are far in excess of statutory requirements and yet, when there is an accident, everyone says that they are terribly negligent of safety.

They feel rather put upon and brow-beaten, and rather defensive about this. They are spending all this money, but the public still hate them. Well, how can we move on from that? The issue here is one that moral philosophers can help engineers to think about, and social psychologists can help too. I am not sure how much is required of the philosopher, and how much other forms of humanities and social science. In this case, however, we have a very interesting issue. If we go back to the statistical safety record of the railway industry, suppose that it is true – as it is true – that it is a rare year in which more than 10 people die on the railways. I do not think that, since during the war, there has been a week when more people have died on the railways than on the roads – although I have not been able to check that. Even in the case where there have been the worst railway accidents, there have been more people dying on the roads that week than on the railways.

It therefore cannot be the statistical number of deaths that is worrying people about railway safety, because the statistics are so small. It cannot be that they think that 10 is an outrageous figure. In fact, most people have no idea. There was a recent MORI poll, in which 1,000 members of the public were asked how many passengers were dying on the railways each year. In the MORI poll, respondents gave answers in a range of 10 to 2,000. The average, even excluding outliers, was 99, in this study – but the true answer is five. No one knows how many people are dying on the railways. They all thought that it was many more than is in fact the case. *Perhaps* it could be that people are worried about the numbers dying – but they do not *know* the numbers who are dying.

What they *are* worried about, however, is the particular way in which some accidents happen. They are not indifferently concerned about all accidents on the railways. If you consider the accidents we have recently, there was the one at Great Heck, where a trailer ran onto the railways lines. In some research that we did, the public expressed their sympathy for the railway industry for this, because it was a member of the public who had damaged the railway industry. In other cases, if a rail breaks or a signal malfunctions, then the railway industry gets it in the neck. The same number of people could be dying, it could even be that you calculate it to be the same probability, and yet there is a completely different public reaction.

In the Health and Safety Executive guidance, there is no way at the moment for factoring in, in any formal way, the different causes of the hazards. You have hazard and probability, and hazard and probability gives you the sum. Thus, this type of consequentialist framework is very congenial to engineering calculations, and very congenial to economists' calculations, but it misses out the most important variable – or at least a highly important variable, which is the cause of the accident and how that relates to the trust in the organisation and so on.

My message here is only one about how one may be able to improve not necessarily engineering, but the type of interface between the public and the engineering, by having a better understanding of what moves people's moral beliefs, and trying to think about how to incorporate a wider range of moral issues into a framework of ethical decision-making. I have been talking to the railway industry about this for a few years now. It is a very delicate issue, and they do not like it when I tell them that the problem with the railway industry is that there are too many engineers running it. I am sure that will not go down very well here! It is actually one of very few places in the world where they could do with a few more management consultants, rather than engineers – they need some engineers, but not only engineers.

Peter was talking about know-how earlier. They need people with moral know-how, and not just people who know the formulas. It is not just propositional knowledge, but a type of moral experience. That goes against the type of formal system building approach that engineers have from day one. The railway industry, to the degree that they have absorbed the lesson that they may need a more complicated way of moral thinking, what they want is a more complicated rule book, which has more things in it. It will take us a long time to come up with that but maybe we will get there.

Just as a closing thought, one of the reasons why the railway industry had not approached philosophers in the past is because they have had a bad run-in with sociologists. They have occasionally thought that they have this notion of societal concern, and they are trying to formalise that notion – and they think sociologists ought to know about this. Sociologists write these very long papers, which are extremely interesting, but the take-home message is that it is all too complicated to do anything about it and that we have to give up quantitative risk assessment and do some further research.

We are hoping that, perhaps, we will be able to come a view that quantitative risk assessment is a good thing, in its place, but that you have to understand its place and its limitations, and what other types of factors also need to be taken into account. [*Applause*]

Questions & Panel Discussion

Keith Guy: We have a short time now for questions and panel discussion. I was intrigued by your last comment Jo. I am old enough to remember when we actually tried to understand, on a philosophical basis, what was an acceptable risk to the public. They decided that involuntary risk, being hit by a meteor, was considered to be something that the public would accept. However, I believe that, if a large asteroid fell on a conurbation, many questions *would* be asked about why something had not been done about it, so maybe this is *not* an acceptable risk. We have to understand the reasons for different perceptions of risk too. When you get into your car you are king and you think you are in control, whereas when you get on an aeroplane or a train, somebody else is in charge and we have a different perception of that risk. What you say is absolutely true, and there are many reasons behind it.

Are there any questions or comments?

Nigel Gilbert (University of Surrey): I wanted just to defend both sociologists and the Royal Academy of Engineering by referring you to an Academy publication called *The Societal Aspects of Risk*, which looks at exactly these issues. This was published in 2004 and in fact raises – although not in the context of the railways – a much older incident, a case with the Ford Pinto, where a similar kind of issue arose about the valuation of human life. This report suggests that there is something called a multi-criterion decision analysis which, if you like, is a much more sophisticated version of the cost/benefit analysis, and which has some quantitative and qualitative aspects, which is an appropriate way to go forward. So perhaps you should just tell the Health and Safety Executive to read the Academy's publication.

While I have the floor, perhaps I could make a quick response to Professor Aleksander, one of whose heroes is clearly Hegel. I note that Karl Marx had something to say about Hegel – Marx was the antithesis to Hegel's thesis! One of the things he said, rather more generally than this of course, was that one should not be looking at some essentialist aspect of what engineering is, in some kind of platonic, conceptual way. If we want to define what engineering is then we should go and look at what they *do*. We should define engineering in this very practical and pragmatic kind of way.

This is where I come back to the sociology. If we want to do that, then we need to study how engineers work in reality, rather than how they ought to work, or how we might think philosophically that it might be best that they should work. There is a body of work in sociology, which takes that view and does that, and which could be extremely relevant to the discussions of this group.

Sir Duncan Michael (Trustee, Arup Group): In listening to the speakers, I remember that I studied natural philosophy, which was actually physics, while my mates studied moral philosophy, which I thought was posh economics. So it is a matter of getting to know you again.

A serious question from me is, do philosophers exist outside academia? The other point is, if it is roads 3000, railways 10, can philosophers deal with the irrational?

Jo Wolff: Can philosophers deal with the irrational? They do not *deal* with the irrational, because that is perhaps a little ambitious, but they can explain the irrational. One of the difficulties with much of the discourse about here is that there is a very quick rush to label things as irrational. Quite often, some other consideration has to be taken into account in the thinking, which just is not being accommodated in the model, which I think is what you are talking about in this case. It is perfectly reasonable to think that, if you are travelling by train and other people are responsible for your safety, then at least they ought to be able to make it stop at a red light.

It is not actually the statistics that are worrying people, but it is the way in which accidents are happening that is worrying people, and that puts it into a different level.

David Jenkins (BHP): There is of course a science of what we have just been talking about, and it is called the science of outrage. We, as engineers, talk about hazard and we refer to it as risk – but actually risk is a function of hazard and outrage. Outrage is all about the way the public perceives the nature of the hazard. You mentioned control and there is a particular issue about that. A fair amount of work has been done on this. Engineers like to enumerate, and you can indeed enumerate and look at how many orders of magnitude is the difference between the way people perceive both risks, which they think they control, and risks which they know that they do not.

That is just one point. The other point that we have not talked much about, but which intrigues me is that, as humans – and this becomes a philosophical question, we are the only beings on the planet who insist on changing the natural environment. Engineering is all about the way we change it, and engineering is all about how much resource we use. Engineering is all about how we handle the waste from the resources that we use, and we have not talked about that but it is a philosophical question.

Peter Simons: There are animals, apart from humans, that change the environment – from birds making nests to beavers building dams. However, they have not had the same kind of large scale impact as we have. On the other hand, it is arguable that large numbers of organisms have had a huge impact on the way the climate is, because it was blue-green algae and so forth which produced the oxygen in the atmosphere. They did not do it intentionally but, boy, did they make a difference!

Jo Wolff: I am aware of the theory of outrage. The literature I have seen has tended to use factors like outrage as a modifier of other variables, as a multiple or divider. That is why I object to it, and it has to be treated outside of that framework.

Robert Hawley: Could I just tease Jonathan slightly? He used the words 'people are too quick to label', and then he came out with the idea that there are too many engineers on the Rail Board. I remember – and you will not, because you were not born – Sir Peter Parker, who took over the railways when they were in a hell of a state. If you read his book, he says it was because there was not one engineer on the board. At the top, you have to have balanced skills, and I should hate to think of a railway that was just a mix of sociologists and philosophers.

Jo Wolff: I agree. Yes, I did say that engineers are needed, but engineering is not the *only* skill.

Ibo van de Poel (Delft University of Technology, the Netherlands): I was quite interested to hear all the themes and topics about where engineering and philosophy might fully co-operate, but I would like to add a further topic – that of design, or engineering design. This is a topic where there could be a fairly fruitful connection between the two.

Design is one of the things that distinguishes engineering from science. Design is about making things, and it is there that there is a new distinction between science and engineering. Actually, from a philosophical point of view and, it might be that in design, philosophy of action is more important or relevant than the philosophy of science. That is therefore one topic where there could be a fruitful combination.

Then there is the aspect of ethics. In design, many of the ethical problems come out in a fairly practical way. Sometimes, engineers do not recognise that they are ethical. On the other hand, you might say that philosophers talk about conflicting values, for example, but engineering designers solve conflicting values by making technology. Here, there is something interesting that philosophers might learn from engineers. Not that engineers always do it the right way, but there is a practical way of dealing with a quite different sort of problem here. The topic of design could be really interesting to look into, if you want to delve into the relationship between philosophy and engineering.

Robert Hawley: That is an extremely good point. The laws of science, if you like, are precise – for example, gravity is just about the same everywhere. However, if we were designing a product for the UK market – *designing* a product – or designing it for your country, it is highly likely that it would be different. You may use different design rules and so on, so it would be different – a bridge here would be different from a bridge in your country. That is a fascinating subject.

Peter Simons: I mentioned intent realisation under the philosophy of action. The example that I gave was designing and enterprise – although I did not use the word 'design'.

Graeme Gooday (University of Leeds): Linking together two of the papers, it struck me that one of the things we have not talked about is the ontology of risk. We have been talking about risk as if it is a thing, it has been somewhat over-reified. The difference between a car risk and a train risk case is that it involves very much whether the driver of the car has the agency, and whether the passenger has no role to play in enhancing the risk. Should we think about perhaps looking at this much more carefully, to understand this complex constitution, its parts, and thereby to educate the public and ourselves to talk rationally about this?

Jo Wolff: That is correct, for the simple models that at least the Health and Safety Executive are putting forward at the moment – though they may well have more complex models shortly if they read The Royal Academy of Engineering's booklet! On the one hand they have a very simple model and then they tell you about all these other factors that have to be taken into account, without actually telling you how to do it.

Greg Hunt: I agree entirely about design being central to engineering, but we have to avoid going through the agonies of the social scientist of the last 50 or 60 years by saying, what are we if we are not a natural science? There is always a danger of that occurring. It seems to me, in fact, that if you look at the philosophy of science in the last 10 years, it has been particularly influenced by followers of Hegel. In some ways, you can characterise natural science as a design of theories. Of course, it is not producing artefacts, but there is a problem – there are evidential problems and more general problems. You are trying to develop a product – namely a theory, and you want to develop it in ways that are very much akin to the ways that one might develop a rather more tangible product in the case of engineers.

To avoid falling into the perils social scientists have for 50 years, we need simply to ask whether we have got science right? Is science not just like engineering, in all sorts of very important ways? Decisions have to be made, markets have to be served, there is the international community, and things like that. The amount of decision making that goes on with the sciences is systematically underplayed, to make a contrasting case with engineering. I do not think that there is all that much difference but then, of course, the immediate concerns and constraints are very particular.

Robert Hawley: I understand what you are saying. I was chairman of the Particle Physics and Astronomy Research Council and I went to CERN, and I could not tell the difference between science and engineering. The whole thing was totally linked together, with the engineers making big rigs for the scientists. In the whole thing, you could not tell the difference.

Shepley Orr (University College, London): I am presently working in a department of engineering, but I have worked in economics and sociology. I have a question about intellectual and academic history. This concerns an anomaly. Very often, engineering departments in the United States will have a unit within them called Operations Research, but you do not see that in United Kingdom departments, so far as I know. You have very few departments of operations research. As far as I can tell, these departments engage in two kinds of research that would be relevant to the issue of risk appraisal and how to deal with risk.

First, they very often work in the theory of decision making under risk, rather relevantly, and in theories of rational decision. They have these very obvious connections to people who work in elicitation of risk attitudes in psychology and economics. They do the same kind of work in these departments of operations research that you see economists and psychologists doing.

The other thing that operations research departments very often do is project appraisal. They are doing cost benefit analysis and they are asking, what are the variables that we might want to put a cost on? In fact, cost benefit analysis is becoming much more attuned to the idea of saying, how do we value societal concern about the cause of risk, and about so-called dread risk, and so on? That would seem to me to be a very helpful way to get some of the philosophical issues that we are discussing more centre-stage in engineering. So I am curious as to why these units within engineering departments are very common in America, (and I suppose this question is for any engineer who might know the answer) and why they have not happened here?

John Uff (King's College, London): I am a part-time lawyer as well as an engineer. One answer to the question - what is engineering? – is that it is what engineers do. Assuming that as a definition for a moment, I wonder what the panel had in mind as to what type of engineer they are talking about. Very few engineers are masters of what they do. They all work for companies, partnerships, universities, government and so on, and they have primary legal responsibilities, statutory duties and so on.

When we come to engineering decisions, which is what much of engineering is all about – and to take the example of the Health and Safety Executive – they are largely directed by the amounts of funding that government makes available. The HSE is controlled in terms of which accidents it investigates by the resources it has available – which are always scandalously inadequate, as the HSE will freely tell you.

Many of these questions then come down to matters of ethics. Should individual engineers put up with it? Just to take an example of the engineer working for the company and the railways, we have seen over the last 10 years various increasingly hopeless attempts by the prosecution authorities to hold *someone* liable for manslaughter, whenever these relatively statistically insignificant accidents happen and 30 or so people are killed. Once again, that figure is completely swamped by the numbers killed on the roads.

The government prosecution authorities decide to spend many tens of millions of pounds in an attempt to put somebody in prison for about six months, and they always failed. They failed recently with the attempted Hatfield prosecution. If you try to prosecute the individual, they invariably get off because the jury is persuaded that the company is responsible; if you try to prosecute the company, then the courts decide that you cannot prosecute the company and, in any event, there would be an individual somewhere who had not yet been identified.

Is there not a problem running through all of this, as to who we are talking about? What does the panel have in mind? Is there a notional, philosophical engineer somewhere that they are addressing? Or who is it?

Robert Hawley: Sitting next to market forces, you mean. It is a very difficult situation. Having been in the nuclear industry, we were very concerned on the board that we had a director of engineering who was, within his terms of reference, responsible for the nuclear safety of the reactors. He had the authority to do things that needed to be done. Whether or not, if something went wrong, you could have him up and cross-examine him in court and make him guilty, I am not really sure, but it was his job to make sure that the systems were in place, the people were competent and properly trained, and all the other good things. That was his responsibility and, after that, it was then the board's responsibility.

Peter Simons: There is a genuinely difficult area of conceptual analysis to be done around the notion of collective and corporate responsibility. I am not sure whether anyone has come up with a stable solution to that yet. It is certainly something on which public money could be well spent, to get practitioners and theoreticians together, to see whether they could hammer out something useful which could then be enshrined in law.

John Uff: Could I just add that the Royal Academy has addressed this in terms of ethics, through issues like whistle-blowing, which from times to time comes into the news. That is one answer, but I wonder whether the panel had any other.

Jo Wolff: The problem is that the moral issue is one of whether you are taking an acceptable or unacceptable risk, and not whether a bad thing happens. Even if you are taking an acceptable risk, every now and then a bad thing will happen to someone. From a moral point of view, we need to make sure that safe and acceptable practices are in place but, from a moral point of view it does not really matter where the accident falls, if you have already decided that you will tolerate a certain level of risk. In the public mind, however, everything switches onto the actual accident as a focal point and so it becomes a much bigger political issue to prosecute someone when the accident has happened. But the moral issue may say that that person is no more to blame than a thousand other people who did exactly the same thing but had no accidents.

Igor Aleksander: I have just a few comments, although I have felt like responding to every one of the points that has been made.

First, I must strenuously deny that Hegel is one of my heroes! I only mentioned him because he seemed to have used the design process in philosophy, which is an example of a way in which engineering might help philosophy.

One of my heroes is certainly Shannon, because he did all of the things that Whitehead was suggesting that people with a philosophical outlook on life would do – although he never described himself as a philosopher.

A key point that comes across, even when I hear people speak about looking at what engineers do – that, to me, does not help with the evolution of engineering. I would be hoping that philosophy, in conjunction with engineering, would give engineering the impetus to look *beyond* what people do, to see what they should be doing next.

Keith Guy: At that point, we will break for refreshments, which will be followed by the break-out session.

[*Tea*]