Don't Just Reduce Risk—Transform It!

dave rosenfeld

In order to ensure a specially high degree of safety, the job was so highly compartmentalised that those in different departments often could not or were not allowed to communicate with each other.

—Robert Jungk, on the Cap de la Hague reprocessing plant, in The Nuclear State

. It would be possible to write a history of the inventions made since 1830 for the sole purpose of supplying capital with the weapons against the revolt of the working class.

— Karl Marx, Capital, *Vol 1.*

THE arguments for nuclear safety have become familiar: 'The risk from a radiation does of 1 rem is the same as the risk of being obliged to smoke 1/20th of a cigarette every Sunday', claims Walter Marshall, chief of the Central Electricity Generating Board (*Atom* October 1982). Independent studies conclude that nuclear electricity in the whole fuel cycle gives fewer deaths per unit of electricity than does coal (see Table in Appendix). The argument appears absurdly simple, yet anti-nuclear objectors get locked into a debate on the exact level of risk from reactors, reprocessing plant, transport bottles, dumped nuclear cargoes and nuclear disposal sites. Hiring experts to challenge the well-funded orthodoxy is expensive. Counter-experts are marginalised from the 'scientific community' and are denied even minor funding. Some, like Professor Sternglass, have been crudely attacked as 'fearmongers'. Some scientists find themselves moderating the objectors' 'overreaction' as they see the risk as numerically low, even if it is really at double or treble the industry's risk figures. For the anti-nuclear side the debate is demoralising and endless.

But there is more at stake than simply challenging the true numbers of deaths due to each technology. The very idea of quantitative 'risk' contains hidden assumptions that influence not only the style and outcome of the debate, but also the detailed choice of technology by industry's managers. This article will draw out some hidden assumptions and challenge them. I end by showing that this is not merely word-play, as the analysis gives useful perspectives for a number of radical science issues, not least nuclear safety.

What is This 'Risk'?

Risk determination is used by the nuclear industry in all its safety cases to marshall safety information, from a number of disciplines, into figures of expected deaths, accidents or disease per reactor, per disposal site, per unit of electricity, etc. For the initial breakdown of equipment, elaborate engineering models ('fault trees' and 'event trees') are used in what's called 'probabilistic risk analysis'. For the spread of contamination, environmental modelling is used. For the effects of radiation on the exposed workers and residents, toxicological models (of 'dose-response') are used. It is not claimed that the models yield causal predictions. Rather they give a figure of risk that is said to give the chance of harm from nuclear power, or from coal when applied to coal. The calculation of risk figures use disciplined methods and are arguably worked through with a degree of scientific rigor often lacking in safety analyses. Yet the actual notion of risk, defined as 'probability of harm', has so far gone unchallenged. And it is precisely this notion, when it is analysed, that turns out to be the vehicle of prevailing prejudices about technology.

Hidden meanings can be seen once we ask why 'risks'—seen from presentations of tables of deaths, in say coal, nuclear or other industries—always seem so fixed. These figures, we are told, represent the risk of an industry, or rather, it is implied—and this is the crucial unspoken step—the hazards *due to the technology* alone. Thus the 'risk' of nuclear power is 1 death in a million people per year. When risk analysts look at average accident rates over decade periods and claim that, on current improvements since 1940 in the coal industry, mining deaths will be down to 0.3 per 10,000 per year by 2000, it reinforces the inevitability and asocial character of risk. It is simply the 'risk of a technology'. We can only wait for the technology to become safer.

This kind of fatalism is reinforced by a risk analysis that excludes the social relations of hazards and so makes risk into a thing, as an inherent, technical property of a particular technology. This definition of the problem serves to relegate 'decision-making' to elites acting on behalf of potential victims. In particular it uses 'risk reduction' to justify supplanting older skilled technologies with supposedly 'safer' ones (coal vs nuclear), or adding on extra 'safety devices to dubious designs for nuclear technologies (British-style PWR for Sizewell).

The prevalence of risk analysis (or 'risk determination' as they call it these days) in the Great Nuclear Debate raises some important issues for the politics of science. Its use is based on the idea that risk is due to a technology and measurable as probable deaths per year. This idea rests on several definite assumptions, rarely spelled out: 'risks' are *in* technology, asocial, randomly striking and quantifiable. Once having analysed these assumptions we will see the dangers of being led to tilt at technology, and of acutally inviting greater managerial control of work via the choice of the 'least risky' technology.

The assumptions that lie behind risk as the 'probability of harm' are as follows:

(1) *The calculated risk is due to the technology alone* (in the sense of 'hardware'). Thus the risk figure is applicable with the same technology to Britain and Brazil, to 1983 and 2003, to managers and workers.

(2) *We are 'exposed' to risks that strike from within the technology; the probability of harm is the probability of being hit.*

(3) *The risk is randomly striking.* For the use of probability:

a) The specific hazardous encounter with technology is identical every time.

b) The risk *exists* in every encounter, but it only *strikes* with harm as the 'outcome' in a few cases, entirely at random.

c) This randomness gives rise to the fluctuations around an average number of people struck per year, as a hand might draw differing numbers of peas from a black bag, but always approximately a handful. It is the average number of deaths that is measured by probability.

(4) *The quantifying of risks makes all activities comparable.* The essential part of this 'risk' is its numerical level. Once that is known, the essence of the risk is grasped. Otherwise incomparable activities can then be compared by ranking them on a scale.

The current use of risk analysis, then, leads us to a particular view of the structure of hazards—risks are asocial, they are external, they strike randomly. If the construct of the risk analysts is to be believed then we are surrounded by randomly-hitting techno death-threats. Is it some kind of macabre prediction when we are given the retort, 'Well, *everything* has a risk?, Is it a wonder that analysts speak of people as being risk-averse?

Safety Science Is Dangerous

How valid is this description of hazards given by risk-analysis? We shall look at its assumption. Some are easily debunked, while others are more subtle.

(1) *Is risk due to technology alone?* For workers from a workforce to die year in, year out, in similar numbers, they must be born, fed, clothed, transported to work, trained, ordered, paid, put in hazardous situations, kept healthy or replaced. In short such statistics require that the whole system that gives rise to the hazard is reproduced. Mortality and morbidity figures for industries, then, do not simply measure some 'technological risk'. Rather, they measure the overall *social reproduction of industrial harm*. That is, they measure social and economic forces that bring people into contact with hazards as much as they indicate any intrinsic hazard of working with the hardware. Accepting the false assumption leads to a false strategy. Attacking technology as the cause of the intolerable hazard will only solve half the problem—you must attack the forces bringing the technology into being, or else a substitute for the same purpose will be found.

(2) *Is risk a thing that 'exposes' and strikes us from within technology? Do risks always strike at random?* These questions are linked by the idea of an alien threat hitting in an unknowable way from outside our experience. Are hazards in essence randomly striking alienated threats? The answer is not straightforward: yes or no. At one level, some hazards do strike randomly. For example pipe breaks in complex plants carrying toxic materials are due to metal fatigue or mechanical failure that is basically a random process. Likewise the effects of exposure to agents like radiation, viruses, invisible asbestos fibres occur in individuals in a way which is governed by random biophysical events.

However, the assumption is false at another level—that is, at the level of individual or collective control over a hazard. Take the case of detailed control by operators over a task. To speak of the 'exposure' to the risk of cutting your own finger with a chisel (assuming the blade is perfect) would sound odd to the skilled craft worker. The ordinary risk of a cut finger is not random to the woodworker since it arises at well-defined times when care can be taken to avoid it. It is the term 'skilled' that we associate with remedial actions being on a human scale, where the 'exposed/at risk' categories would misrepresent that control.

In the case of collective control, information on workplace organisation can transform the nature of risk. For example, ask a worker at the Windscale nuclear fuel reprocessing plant to repair pipework in a high-radiation area unfamiliar to him. Even if there are only a couple of lethal 'hotspots' where doses are high, the whole area appears hazardous. To him (women are not employed in high-radiation zones) a walk in a straight line is like crossing the road blindfold. As the management gives him a chart of hotspots and a pocket alarm meter, he feels sure to avoid deadly spots, confidently and consistently—as long as experience tells him the management or union safety committe have assured the chart meter are reliable.

In this example, if two areas had the same death rate, it would be preferable to work in an area A (most of which had lethal hotspots but where you're given accurate, trusted information) than in an area B (with very few lethal areas but no clue where they are). If some workers have died in areas A from carelessness, it seems less of a problem to negotiate than if the same number died in area B from a withering beam by bad luck. *The randomness of the risk depends on workplace relations.* So at one level there are indeed randomly striking hazards e g, cancer risks from radiation, viruses, chemicals. At another level hazards are subject to detailed control, so that harm is due to bad design, bad procedures, poor training, poor information or carelessness, e g, woodcrafting and work in 'hot' radiation areas.

(4) Quantification adds no further assumption that is not already made in the earlier points. However, the acceptance of such numbers as inherently comparable serves to hide the assumptions that we've identified.

By assuming that all hazards are randomly probabilistic, risk analysis treats people as passive objects of technological hazards: individual and collective consciousness are ignored. The immediate consequence is that, in general, *strategies to reduce 'risk' evade the issue of control over safety instead encouraging a 'lower risk' technology that is alienated from worker and community control.* The problem is not too little management attention to safety analysis, but too much. The quantitative comparison of different risks encourages deference to the intrinsic 'risk of a technology' and thus to management control over safety.

Historical Origins of 'Acceptable Risks'

How does 'safety' get reduced to a technical choice among different technologies or different devices within a technology? This is done so compellingly by risk analysis by quantifying risk as the probability of harm, yet while hiding the assumptions involved. The 'nuclear is safer than coal' argument, then, involves more problems that the simple, hopeful predictions about nuclear accidents and health damage from radioactivity. More insidious than that, the whole framework treats the harm from each industry as technical, fixed, while at the same time pretending not to make value-judgements. Ironically, this type of risk-analysis arose from an historic admission of 'no safe level'. Let us see how.

1966 was a turning point for risk analysis. The International Commission on Radiological Protection (a self-appointing, non-governmental body) accepted that radiation may have no safe level of exposure for cancers and genetic defects. Secondly, in that same year a fast-breeder reactor near Detroit burst a fuel pin and partially melted down, thereby exceeding its official 'Maximum Credible Accident'. The safety engineers quickly realised that automatic safety devices can and do fail, so that there is no accident that could be made strictly impossible by engineering safety features. So by 1967, both in toxicology and in engineering, it could no longer be assumed that risks had a 'safe level' that could be found scientifically.

Since the acknowledgement that some hazards have no safety level, new theories of how to predict risks from drugs, rays, bugs, chemicals machines and industries have arisen together with social theories of how to resolve conflicts over safety. One model, popularised by W W Lowrence in 1976, particularly captured the minds of researchers. It appeals to the administrator role of professionals in the scientific/industrial/managerial world. In this view:

i) Everything has a risk—most agents, technologies, occupations, leisure. The magnitude of this is a technical judgement for scientists.

ii) Scientists should rightly avoid statements that something is 'acceptably safe', which amounts to a political judgement.

The view is articulated as follows. No longer can scientists simply show that risks are non-existent and then reassure the union, the patient, the plaintiff or the tenants association. As risks always exist on a numerical scale, a safety standard represents some level of hazard. So someone (society, not the scientist) must either ban the technology or set an acceptable level of hazard, given the benefits of the technology. In effect scientists purport to present simply the facts and then allow others to set 'acceptable risks' within that rigged framework.

Faced with a new technology, or a newly contested one, scientists must now find the objective level of risk from a whole range of exposures. A new science, *risk analysis*, was painstakingly developed, supposedly to give a predicted sliding scale of risks from different levels of technology. This scale was to be input to *decision-making*, where conflicts are resolved by choosing to expose people to a level of risk which is outweighed by the benefits from the risky technology. The decision-making could operate 'rationally'—that is, by using economic arithmetic to spend the limited safety and health resources in proportion to the level of risk.

However, decision-making will be bedevilled by the need to make expedient decisions because, so the model goes, the level of risk people tolerate and the vigor of opponents to risky projects is not in proportion the 'objective risk', as ascertained by risk analysis. People (not the numbers) must be wrong, ignorant, vindictive or irrational. Decision-makers therefore need a second input, apart from risk analysis, namely a socio-psychological study of why people's reactions do not correspond to the objective level of risk.

This whole view, originating in the nuclear and aerospace industries, was seen as the proper way for scientists to set about denying adjudicating on 'acceptable risk'. And, so it appears, scientists could voluntarily abdicate the politically-laden role of reassuring on safety (if only they would!) and could instead retire to the rigorous and objective world of toxicology and safety engineering methods in the new science of risk-analysis. Appealing though this scheme is, however, it has been openly admitted that risk analysis is not entirely objective. And this admission has jeopardised the credibility of 'rational' decision-making and socio-psychological studies of people's irrationality toward risk.

For example, in developing measures of the risk of an industry, risk analysis has used various ways of combining the incidence of different accidents and diseases into an index that would measure the total amount of harm. However the conclusion became inescapable that, where these grand indices were used, risk analysis was making moral judgements. For example, if we simply add days of work lost from different diseases per year, we make judgement about the relative amount of harm from physical versus mental suffering, suffering long versus dying early, frequent isolated deaths versus infrequent mass-killings.

To avoid making such obviously moral judgements, practitioners decided not to aggregate diferent forms of harm. The disciplines of safety engineering and toxicology can systematically analyse failure rates of machines and arrive at does-response curves for toxins. Mathematical formulae then come out with probability of harm per year for *each type* of injury or disease separately. It is these well-defined methods, usable by anyone, together with the experimental data and industrial accident data, that supposedly make risk analysis figures objective. It is only by a detailed look at the definitions of risk (as in the previous section) that we were able to get to grips with this last formulation and understand precisely how risk analysis is value-laden.

Practical Consequences

Risk analysis lies at the heart of the prevailing idea of acceptably risky technologies. Debunking it (i) is useful in general radical science perspectives, (ii) informs some perplexing questions for anti-nuclear campaigns, (iii) opens up an approach for opposing safety analyses that use probability and (iv) facilitates intervention in the 'risk debate'.

Firstly, we have worked through an example of how science clearly acts as ideology. The scientific content—here the very definition of 'risk'—is the vehicle of the manager's right to dominate workers' activities. New questions arise: Should socialists abandon entirely the elaborate and rigorous methods of engineering risk analysis? Is there a socialist science of risk? Will the simple expedient of stating the social relations of (quantified) risk overcome the objection?

Secondly, we can see how the dominant managerial definition of nuclear safety fits well into the wider nuclear project. Although nuclear energy will not deliver any net *energy* until after the turn of the century—due to a net consumption during the construction part of the programme—it could generate 30% of *electricity* by 2000. That is much more than nothing at times when transport workers or miners block the remainder.

> . . .a nuclear programme would have the advantage of removing a substantial proportion of electricity production from the dangers of disruption by industrial action by coal miners or transport workers.
> —leaked minutes of Thatcher's Cabinet Sub-Committee on Economic Strategy, October 23, 1979 (originally published in *Time Out*)

Nuclear safety systems

- —require often intense workplace discipline, e g, the restrictions on internal communication at Cap de la Hague;
- —require no-strike agreements to achieve their assumed performance; and
- —rely on highly qualified scientists whose work and attitudes are highly integrated into management perspectives.

Risk analysis, portraying risk as a neutral quantitative matter, paves the way for work organised under nuclear managers' discipline, even blackmail, all in the name of *more* safety. Is it not preferable to back coal, where a knowledge and control of the system by workers give higher chances of winning a better safety deal than with the polutonium state? What guarantee is there that current nuclear safety standards will not be eroded when the Tory government has further cowed workers and other opponents with nukes? Already Reagan has been declaring the nuclear industry to be 'over-regulated', and promising side-steps to licencing hearings.

As scientists are involved in designing and assessing technologies, they should be aware that their work influences the relations of control. When people reject 'dread hazards', this is not simply a matter of their ignorance of 'the true risks', but a sign also that no trust has been built between people facing hazards and those in control. Sometimes distrust comes from consistently negative experiences, as many shop stewards in industry will testify. But in other cases it is a sign that there is no consistent experience that people can rely on. In short, there is a gulf between people's experience of technological change and scientists mostly aligned to managerial priorities, as defined especially by risk analysis. It is this class nature of scientists role, and not some fundamental irrationality of people, that generates hostility to new hazards.

Thirdly, is it possible to challenge nuclear science as science? The technical literature clearly paints the picture of a systematic safety programme that's the envy of chemical safety campaigns. Time and again the literature asks whether this effort is 'enough'. It seems a logical definition of the problem. Yet once we are in the thick of the numbers game, it is hard to challenge the expensive and extensive studies that show nuclear power as 'relatively safe' or 'about as safe as coal'. Most critics stick to the general issues of the politics of plutonium, rather than wade through the swamp of reactor science, radioecology and radiotoxicology. There has been no analysis that criticises nuclear science as a starting point for identifying the problem. Our analysis removes this quandary. By analysing the precise use of probabilistic risk, we can argue specifically on the science of safety and win! That is, we can expose the class base of the technology by dissecting a key scientific concept, quantitative 'risk'.

Lastly, radical research on safety is given a new focus by this analysis. We've discovered that the key debate is not

Low-Level Radiation – Out of Control

Workers at Windscale experience, by and large, a comprehensive programme for avoiding the 'hot-spots'. This programme has engendered a basic confidence that risks are under control. At the plant, workers are understandably quite defensive when 'outsiders' allege that their work is dangerous, as they find themselves identifying with management on the safety issue. This attitude affects relations over the low-level exposures. With low-level radiation, in contrast to the high-level situation, there is no alternative to an intrinsically random. untouchable, unknowable hazard. You do not even know whether you have been affected until twenty or so years after the event. Routine, random low-level exposure is tolerated in the knowledge that the chance of harm per person is very remote. But this 'knowledge' is chimerical and entirely different from daily witnessing a competent hot-spot management operation. For low-level work, badges are issued and exposures tallied, but the knowledge of the hazard is available only from scientific studies. BNFL's interpretation of the studies is accepted. This acceptance might be seen as a spin-off from the confidence won from their track record in hot areas, their vindication by 'independent' Government bodies and their success in dealing with scientific challenges in the press and at inquiries. As BNFL win this confidence game, the scale of low-level radiation risk remains uncontested as a workplace issue. (What has been contentious is the administrative practice of 'burning out', or exposing unruly workers to the yearly maximum allowed radiation, thus disqualifying them from further work in the year—but this is not directly linked to the 'randomness' issue.)

This historical aquiescence by most workers, on the lower-level exposures, should not distract from the very different relations of control. With high-level exposures, concern is highlighted when for example a direct injury occurs, with visible effects within hours or days. The response is obvious: press for effective measures to prevnt any repetition, with whatever muscle you have. With low-level exposures, your concern over getting a cancer or deformed kids is heightened only by claims that the old risk-estimates are wrong, that it's more dangerous now. Workers don't control science; the pronouncement could appear to come out of the blue unless you have reliable information from scientific allies. And you have already been exposed; you cannot change that fact. Without those scientific allies with a proven track record, the choice is between management's reassurance and abject worry. Thus the class role of scientists is crucial.

about the level of risk from Sizewell B's proposed PWR. The numbers used in risk analysis talk only about the hardware failing. When operators are mentioned, the chance of 'operator error' is used entirely as if operators were hardware, albeit defective. At nuclear plants we must expect that the real failure rates will be dependent on the social system that organises it:

Real responses of thinking, waged operators

—Workers may be on strike in a minor accident sequence.

—Management commitment to a training programme may lapse.

—Complex but infrequent modes of failure may happen in too short a time for anyone to cope.

Smoothly running quality assurance programme

—Subcontractors may falsify quality inspections (and have done so).

Smoothly running maintenance programmes

—Non-unionised temporary maintenance workers may rebel against 'burning out' practices.

Effective inspection programmes

—Inspectors close to the industry may be lax, believing risks low. (Hendrie, chief US regulator, was sacked after Harrisburg.)

'The techbology'

—Like each car, each reactor has its own unique history of construction and maintenance.

—When politicians like Reagan (or an 'over-regulated' industry) change standards, the PWRs built and maintained may be different.

In non-nuclear safety there is an increasing tendency to follow the assessment methods developed in the nuclear energy debate. In fire safety, asbestos control and chemical plant safety we see use of 'cost-benefit analysis', 'reasonably practicable reductions', 'engineering risk assessments'. These all rest on the idea of balancing costs of reducing 'the risk' against supposed benefit of using the technology.

There are many obvious questions to raise about these schemes: Who benefits from the product? How do you measure the cost of life? Our approach goes further by challenging the scientific definition of 'the risks' that were measured in the first place and that were assumed to exist in the technology as a thing, not as an organising system. Through our approach, the question of control over the risk can be made central to the debate even before monetary costs are raised. Indeed, the particular social construction of nuclear risks turns out to be less a cost than itself a benefit to nuclear management.

The numbers game has often led environmentalists and hazards campaigners into a blind alley of demanding 'zero risk'—an idealistic and unrealistic focus. This quandary points to the real difficulties with either rejecting or accepting a (supposedly apolitical) 'balance' between health 'costs' versus industrial 'benefits'. That kind of choice usually confronts us as utterly compelling, universal rational. For example, could socialist societies delay reconstruction programmes until all industry is conclusively proven 'safe'? Our approach offers a way out of the quandary: while defending the primacy of health, we can assert the issue of control as central to any 'acceptability' of hazards in the name of wider benefits.

In conclusion: Beware the 'low-risk' technology of safety science, which serves to usurp control over hazards and thus guarantee management's safety from workers. The important safety question facing workers and communities is not some precise, numerical level of safety. Rather it is how we can gain detailed control over deciding which risks we take, so that we are confident, at all times, they're worth the benefit. How do we transform alien hazards?

[Reprinted from Science Radical Journal 1989.]

(Continued from p 118)

the international atomic energy authorities as a small hazard.

Safe management of radioactive waste is an unanswerable problem of the age because toxic products are not only highly lethal but remain radioactive for several million years. The more we go nuclear, the more we are adding to the problem of survival of our future generations.

There have been some romantic suggestions disposing of toxic waste from the earth by rocket into space or deep burial under the Antaractic ice but no adequate solution has yet been devised. Uptil now only high level radioactive wastes are stored in carbon or steel-concrete tanks which last 30-50 years; and low and intermediate level wastes are either dumped into the sea or buried underground in concrete silos. Proposals bave been made to solidify the highly toxic waste in glass blocks to be stored in shafts drilled in the seabed or under hard rock.

All the attempts and plans are far from reaching any real solution. There have been leaks from the storage sites contaminating the surface and ground water and the atmosphere and causing serious health hazards. We know very little about India's waste management programme.

The operation of a nuclear reactor generates astronomical quantities of radioactive waste of different types and of varying half-lifes ranging from a few seconds to a few thousand years. The amount of radioactivity produced from these elements is in direct proportion to the operation of the reactors. Even after Chernobyl which has put a big question mark on the future of nuclear power, India's nuclear policy is unchanged. We have an optimistic plan of 10,000 MW electricity from nuclear plants by 2000 AD! It is estimated that one year's operation of a 1000 MW nuclear plant generates fission products equal to that of a 23 megaton fission bomb; that is more than 1,000 bombs of the Hiroshima size.

Safe, permanent and absolute isolation of these radioactive poisons from the environment is the only condition for nuclear power to be acceptable. And this is simply not realistic. There is no disagreement today about how much radioactive poison is produced by the nuclear power plants. There is little or no disagreement about how lethal these poisons are. The disagreement lies in the quality and quantity of routine release of radioactive elements during all steps of the nuclear fuel cycle. Will the nuclear advocates give a satisfactory answer to this? No, they cannot and will not. The only answer is:

STOP NUCLEAR POWER