Risk and the Regulation of New Technologies

Jonathan Wolff[1]

Blavatnik School of Government, University of Oxford, Oxford OX2 6GG

Jonathan.wolff@bsg.ox.ac.uk

Sept 20, 2018

Abstract: New technologies can bring tremendous benefits. But they also have costs, or risks, some known, some unknown. How should authorities regulate new technologies in the light of the possible costs and benefits? A standard approach to decision making under risk is to use formal risk cost-benefit analysis. Yet there are clear limits to this approach where risks and probabilities are unknown. Furthermore, simple cost-benefit analysis ignores questions of moral hazard – where benefits and costs fall – and the political dimensions of the introduction of new technologies. In this paper, I discuss how to frame a reasonable precautionary attitude to the risks of new technology, setting out a series of questions that need to be taken into account before a technology should be approved.

Key Words

Risk, Regulation, New Technologies, Precautionary Principle, Precaution

1. Introduction

I once read an article by an American college professor who said that he had asked his class two questions. The first was 'what are the greatest technological advances of the last hundred years?' The students answered: nuclear power and plastics. His second question was 'what are the greatest technological challenges we now face?'. The students, so he said, responded: dealing with nuclear waste; and the disposal of unwanted plastic.

---

[1] My thanks to audiences in Kobe, Oxford and Reading for their very helpful comments, and to Tom Simpson, Karthik Ramanna, and Henry Shue for discussing the themes of the paper. I owe a special thanks to Hélène Hermansson and Sven Ove Hanson for giving me the opportunity to respond to Hélène's PhD thesis, which contained their joint work and set me on the path explored here.

New technologies can bring tremendous benefits. But they also have costs, or risks, some known, but some, such as damage to the ozone layer and asbestosis, unpredicted or even entirely unpredictable. How should authorities regulate new technologies in the light of their possible costs and benefits? Our standard formal methods for risk analysis, most notably risk cost-benefit analysis, cannot be applied in cases of radical uncertainty. Some theorists have tried to extend those formal methods or replace them with other quasi-formal principles – such as 'the precautionary principle' – with limited success. Here I argue that a number of factors, including moral hazard and the political dimension of risk, can be overlooked, with damaging effect, in the attempt to apply a formal method or principle. Instead of hoping to formulate a precise 'precautionary principle' we do better to use a 'precautionary attitude' (Wolff 2014) or perhaps even a 'precautionary checklist' (Myers 2006) which reminds us which values and concerns should be at the forefront in the political debate that will ultimately make the decision (see also Manson 2002, Munthe 2011).

2. The Introduction of a New Technology

While there is no single process by which all new technologies are introduced, often a development can be traced through a number of stages. A breakthrough in basic science uncovers a new mechanism, or discovers an unexpected property of a known object, or creates a new material, or makes some other surprising finding. Innovators suspect that there may be an opportunity for commercial exploitation, and engage in experiments to put the new discovery or invention to a use that will have commercial value. That stage could be carried out by the original scientists, or by a team of technologists, or, increasingly in recent decades, by amateur enthusiasts working out of a basement or garage laboratory. Ultimately a product is created, and typically it will need to go through a sequence of tests before it is approved for general use, to show that it meets a certain safety standard.

In many areas common standards have developed and have been operationalised into a series of tests that can be applied in a straightforward fashion. This will be particularly so when the proposed innovation is little more than a development of existing technologies, the risks of which are well known. But underlying those tests is likely to be a more complex pattern of reasoning that has been instrumentalised in this particular way. In designing those deeper tests it is appropriate to consider all stakeholders in the decision, and the benefits and costs that could accrue if the technology is approved for use. It is rare, however, that costs and benefits can be known with certainty. In other words, technologies have risks and those risks also need to be understood if an informed decision is made about whether to authorise their use, and if so under what conditions. Such analyses have been carried out, at least informally, for many decades. For example, the introduction of the motor car in the UK was debated in parliament, and was allowed in 1865 only if preceded by a person waving a red warning flag (Setright 2004).

Stakeholders in any such decision will include those who wish to introduce the technology and would expect to benefit commercially from doing so. This will normally be a company, and for ease I shall call this the corporate interest. There will also, naturally enough, be a potential consumer group. Consumers would expect to benefit, but with a new product it is always possible there are unintended effects, either known or unknown. There can also be unknown costs or benefits for the corporate interest (for example finding itself liable to clean-up costs, or reaping secondary benefits from an unexpected source). And very often there will be a series of third parties (including future generations), who will either benefit or lose from the introduction of the technology, through, for example, pollution, or increase in property values, or the creation of new employment opportunities. However, very little of this will be known for certain in advance, and therefore estimates of probabilities will need to be made. Hence, we have the basic question facing any regulator of a new technology: how can the balance of risk of costs and benefits to the various stakeholders be evaluated in a way that allows a reasonable decision about whether a new technology should be made commercially available, and what restrictions, if any, should be applied? I will use as a running example through this paper genetically modified (GM) crops. How should the regulator of agricultural products assess whether they should have been permitted?

Some new technologies have little impact beyond their immediate sphere of application. Others can be revolutionary and reshape society beyond the immediate technical context. The printing press, the spinning jenny, the railway, the computer, the internet and the mobile phone are all examples of technologies that have had profound effects, not only on individual well-being and livelihoods, but throughout society. No one would deny the social effects of the railway or the internet, but this entails that they have also enabled shifts in power and influence. Little of this can be known at the time of the first introduction of the fledgling technology that led down this route. But still, as I will claim later, technologies can redistribute power and influence in ways that can be included in a form of risk assessment.

3. Cost-Benefit Analysis and Its Limits

A standard approach to decision making under risk is to use risk cost-benefit analysis. On this familiar approach, a range of options is considered and evaluated against each other. The simplest case is where a decision is to be made whether or not to make a single change. The two options are either to remain as we are, or to make that change. The many possible contexts include ordinary business decision-making, safety planning, a change to public policy, and the introduction of a new technology.

To use cost benefit analysis it is necessary to come to an understanding of all known possible costs and benefits, and their probabilities. It is also standard practice to convert costs and benefits into a common currency, most notably money, so that they can be weighed against each other. The probabilities of things working out well or badly for the different stakeholders can be combined with the positive and negative value of the different possible outcomes so that the expected value of making the change can be calculated. In the simple case 'business as usual' can be given a value of zero, and the question then is whether the change yields a positive or negative expected value. Of course, there will also be cases where more than one alternative to the status quo is being considered, and others where the status quo is not sustainable and hence not an option or very expensive to maintain. But similar considerations apply and the technique generally recommends the option with the highest expected value.

This approach has been used, for example, to decide whether to introduce safety improvements in transport systems. A new signalling system, for example, could be predicted to save a number of lives, and the saving of a life will be given a particular financial value, as will other costs of an accident. Preventing an accident therefore has a value. The cost of the technology is also calculated, and if the benefits outweigh the costs, this is a strong consideration in its favour, as it will yield higher value, according to the calculation, than doing nothing. In other cases, the cost of the new technology might be considered to be too high, and the method would recommend against introducing the safety measure, even if, by all probability, it would save lives. (For discussion see Wolff 2011, pp . 83-108.)

It should not be completely obvious, however, that cost-benefit analysis is acceptable for areas such as transport safety. I have argued that cost-benefit analysis can be the right approach to take in economic decision-making where the following conditions are met (Wolff 2006):

1. The range of possible outcomes, and their probabilities, are known (to a reasonable degree of accuracy).
2. It is reasonable to put a monetary valuation on all relevant costs and benefits.
3. The decision is one of a repeating series of similar decisions.
4. The probabilities in the repeating series are true, in the sense that they are not biased towards or against one group.
5. The loss to any individual or group is not catastrophic.

The basic idea is that in a long series of repeating decisions under risk, where all costs and benefits are comparable, maximising average expected value is a rational procedure, for, under the conditions mentioned, the great majority of people can expect an individual outcome that is close to the expected value. You will be on the losing side in some cases, and the winning side on others, but few people will do badly over a long series. Nevertheless, even if the probabilities are fair, in a big group it is likely that some people will have an unlucky run, coming out on the wrong side time after time, and a form of secondary scheme will be needed to compensate those who lose from those who have a run of lucky gains. With such a secondary mechanism in place, and under the assumptions set out above, cost-benefit analysis has much in its favour.

The conditions set out above make evident when cost-benefit analysis is particularly useful, and, arguably, it would be irrational to follow any other procedure. But it can immediately be seen that there is a limit to the application of *risk* cost benefit analysis. For many of the cases where we are tempted to use risk cost benefit analysis involve possible loss of life, which is a catastrophic outcome for an individual, and cannot be made up for by benefits issuing from future cases. This appears to violate condition 5 above. For this reason, some have thought that cost-benefit analysis should never be used in life and death situations.

Conceptually, however, the decisive move was made by Thomas Schelling in the US (Schelling 1968) and Jacques Drèze in France (1962). If the railway company introduces a safety measure that can be expected to save a life, we will, nevertheless, not know whose life it has saved. Indeed, we will not know whether it really has saved one life, two lives or even more, or none at all. We have no way of knowing what precisely would have happened had the safety development not been introduced. But what we do know is that the introduction of the safety measure has made every traveller's life a tiny bit safer. Therefore, although it is convenient to talk about lives saved, and the value of a statistical life, the reality is that a safety improvement does not save a known particular life but is an aggregation of many small risk reductions. It is argued that when the risk to any individual is very small, we can still reasonably use risk cost benefit analysis, for it is continuous with ordinary economic life. We buy smoke alarms or improved bicycle helmets, and these reduce our risk of death slightly. A statistical life simply aggregates this. Of course, the fundamental issue – that if you die you cannot be compensation by a later run of good luck – holds. But if we were to treat all the normal risks of life as requiring some sort of special treatment, as if we had a right not to be subject to risk of death, then nothing would be possible. (Elsewhere I have called this the 'problem of paralysis', Hayenhjelm and Wolff 2012.)

When risks rise, and especially if they become concentrated on particular groups, risk cost-benefit analysis becomes much more problematic. We can see this concern reflected in current UK policy. Very low risks are simply accepted as part of life and no special measures are needed. At a particular threshold a higher, though still low, level of risk is taken as acceptable if it would be disproportionately expensive to eliminate. This is the level at which risk cost-benefit analysis applies. At higher levels, special steps should be taken to reduce the risk, and the highest levels should simply not be permitted unless there are special circumstances. In other words, cost-benefit analysis is seen as the appropriate methodology for dealing with risk in the moderate range – not so low that it just fades into background noise, or so high that it rings alarm bells (HSE 2001).

Something else that can ring alarm bells is high variation from year to year. In the case of transport safety, statistics from one year tend to be very similar to the years before and after. But compare this with the safety of a nuclear power station. If we are told that a nuclear power station will cause three statistical deaths a year, while supplying vast quantities of energy, we might accept that as a cost worth paying. But consider two different scenarios. In the first leaking radiation causes three deaths from cancer every year. In the second, most years no-one dies, but we expect a serious incident every thirty years in which ninety people die. Risk cost benefit analysis suggests that (ignoring discount factors) the two cases are identical. But at least in terms of how they are perceived they may feel very different. Hence where there is the possibility of a single incident causing a large number of deaths some will wish to be more cautious in the use of cost benefit analysis.

However, the argument that a 'statistical life' is really the aggregation of many small risk reductions also allows us to answer a second, related criticism. Standard cost benefit analysis requires all values to be reduced to monetary terms, as reflected in condition 2 above. But how can we put a price on life? The answer is that the method does not require us to put a price on life. Rather we are putting a price on a small safety improvement. And this is a perfectly ordinary economic transaction, as we noted in the case of smoke alarms or improved bicycle helmets. Putting a price on risk-reduction is familiar to us as consumers. Techniques building from this practice – either revealed preference or 'contingent valuation'

– allow a value for a statistical life to be estimated (which is not to say that the resulting figures are uncontested). (See, for example, Carthy et al 1998)

I will only discuss the third and fourth conditions briefly at this point as they will be important later on. The third was that the decision should be part of a repeating series of similar decisions. Of course, the decision to introduce a new technology is always in one sense a one-off. However, new technologies are being introduced all the time, and although some turn out to have enormous costs, others do not and are extraordinarily beneficial. As opponents of precaution point out, if we have a restrictive attitude to new technology we risk losing much that is of great benefit and low cost. Hence, we must not jump to unnecessary precaution. The trick will be to decide when precaution is necessary and when it is unnecessary, or, in other words, when the introduction of a new technology should be treated in a different way. I will return to this later, as it is the key theme of this paper.

Finally, I mentioned that the probabilities must play true. The threat is that either immediately or over time the benefits will accrue to one person or group and the costs will fall elsewhere. This is one of the most sensitive issues of all concerning risk regulation, and again I will return to it later in detail.

4. Risk, Uncertainty and Radical Uncertainty

I have so far passed over the first condition above, that the situation is one in which the probabilities and outcomes are known. It is this that allows the calculation of expected utility. My own detailed work on risk has largely been on railway safety, where there are large statistical databases, and past frequencies are a good guide to present probabilities. The mechanical systems are relatively straightforward, with very few complete unknowns, and it is reasonable to present options and their probabilities. Yet this is the exception, rather than the rule. In other cases, it is a fiction to pretend we have a good basis for calculating probabilities. Hence it is common to distinguish decision under risk, where the probabilities are known, and decision under uncertainty, where probabilities are not known.

Decision under uncertainty is very common. Consider the offshore oil and gas industry. There have been a small number of very spectacular accidents in oil and gas rigs. Clearly it is a high risk business. But it would be very foolish to try to conduct a risk analysis purely on the basis of statistics collected so far, given the limited data available. How, then, is risk analysis to be done? One way is by 'critical factor' analysis: how likely is it that a critical factor will fail, and what would happen? In some cases, frequency analysis is available at that level, but in others the problem repeats. There is also an issue of interaction between different critical factors that have as yet unknown effects in combination. Hence different analysts, looking at the same situation, can make wildly different judgements about the probability of different types of failure. In these cases we are dealing not with risk, but uncertainty.

In the case of oil and gas I have assumed that we know the possible hazards, but lack solid information that allows an estimate of probabilities. Nevertheless, it is possible to talk about upper and lower bounds, and use other methods that allow the application of a modified form of risk cost-benefit analysis. But for very new technologies we often did not have an understanding of what problems we might have in store when they were introduced. Examples I have given so far are the disposal of nuclear waste and of unwanted plastics, the depletion of the ozone layer, asbestosis and GM crops. Examples can be multiplied. Some detrimental effects of technologies must have been anticipated but others are wholly unexpected. Standard cost-benefit analysis in effect treats the probability of unpredicted events as zero, which is clearly problematic. But if you do not know what the possible outcomes are, it is impossible to make a fair assessment of their probabilities, even in terms of an upper and lower bound. This is the problem of decision-making under radical uncertainty. What can be done?

From the perspective of existing theory, the natural approach is to try to derive a formal method to help with such cases. Can risk cost-benefit analysis be extended through some clever techniques for estimating unknown possibilities? Is there some way of using a 'maximin' principle to help us steer clear of the most damaging possible outcomes, or could a version of a 'precautionary' principle help us? Or do we need to take a different starting point?

It may be that there is a type of path dependency which has led from simple and useful cost-benefit analysis to a potentially very damaging possibility of using formal methodologies in contexts to which they cannot apply. This, so opponents claim, has the consequence of advancing the interests of a political and economic elite. I will try to explain this suggestion later in the paper. First, I need to say more about the claimed path dependency and the alternatives.

Recall simple cost-benefit analysis. Under the assumptions set out above (almost) all people will benefit in the medium term if social and economic policy attempts to maximise average expectations in every decision. Those who lose out now will, if the probabilities run true, be more than compensated in the future (ignoring the case of death). But in order to apply that technique it is necessary to reduce all values to a form in which they can fed into the necessary formulae. We need to value all inputs – including life and death – in financial terms, and this can be done by a variety of valuation methods. Activists and philosophers have been worried about the reduction of life, and other values, such as the natural environment, or the preservation of species, into monetary terms. However, sometimes it is accepted that it is necessary to do this in order to have a seat at the table when cost-benefit analysis is being discussed.[2] With some reservations, therefore, the reduction of everything to monetary values can be helpful as, for example, it allows environmentalist to discuss the matters that concern them in terms that economists recognise and policy makers need if they are to take everything into account using formal methods. As long as it is understood why this is being done, the damage of giving everything a cash value can be limited.

There can be, therefore, a pragmatic reason for reducing all value to the monetary when considering decisions under risk, and, perhaps, under some level of uncertainty. But consider the case of radical uncertainty, where there is incomplete knowledge of the range of possible outcomes and hence very little understanding of any probabilities. The case for reducing everything to a monetary value in order to be fed into a formula recedes dramatically. Before, it was a pragmatic accommodation; the reduction is needed if we are

---

[2] I owe this point to John O'Neill

to use the method. Now, though, things looks different. The reduction seems as much ideological as technical. If there is no decision formula there is no particular reason to render all inputs into a form that can used in the formula. In fact, doing so can do more harm than good, as I will endeavour to explain in the next sections.

5. The Risk Triangle and Moral Hazard

It has been pointed out by Hermansson and Hansson (2007, and see also Hansson 2018) that in any decision regarding risk there are generally three roles to consider: the agent deciding that the risk should be taken; the agent that will benefit if the risk pays off; and the agent that will lose if it fails. This simple insight is very powerful, for it allows us to understand that the number of possible structures of risk decision-making is limited. Suppose that agents A, B, and C are different people. If only one agent can occupy each role the only possible structures as laid out below:

	Decision Maker	Benefits go to	Costs go to
Individualism	A	A	A
Paternalism	A	B	B
Moral Hazard	A	A	B
Moral Sacrifice	A	B	A
Adjudication	A	B	C

Of course, there are variations on these, where, for example, the benefits are shared between someone who makes the decision and someone who does not, but in terms of pure structures only these five are possible.

The first is morally speaking the most straightforward. All costs and benefits are absorbed by the decision maker. In practice, however, for any significant decisions this situation is rare. The government is normally a supervisory decision maker, providing rules for what is and is not permitted. Where the government is the only decision maker, regulating what will happen to a single individual we move to the second case, paternalism, which also applies whenever one agent takes a decision for another. Again, some level of shared decision making, with the consent of both, is the much more common model in practice. Paternalism does raise moral questions: why should people not be permitted to take the risks they want to? This is particularly relevant regarding the regulation of drugs, for example.

The third case, however, is the most troubling: where one agent is in a position to decide whether a risk should be taken, and would benefit if it works out, but another would pay the cost if it fails. This is the source of many of our difficulties in policy, including, according to some analysts, the financial crisis of 2008. It is highly relevant to the introduction of risky new technologies and I will come back to it in the next section. For completeness, I should finish the picture first. The fourth case appears to be an unusual one, for one agent can decide to take a risk, but while benefits would go to others any cost will fall on the agent who decides. If it does go ahead it is an act of self-sacrifice. Elsewhere I called it ‘maternalism’ as it is a behaviour that some mothers will engage in for the sake of their children (Wolff 2011b). However, we should understand that when faced with this situation the most likely outcome is that the risk will not be taken: why should I decide to take a risk if others will benefit but I would suffer the costs? It may seem fair enough to decline to take a risk from which you cannot benefit, but consider the issue from the point of view of cost-benefit analysis. If the potential costs to me are smaller than the potential benefits to you, then it is, in a sense, inefficient, even if perfectly understandable, for me not to take the risk. In policy terms this is very important, and could be a significant barrier to growth. The refusal to make this ‘moral sacrifice’ could be called ‘moral cowardice’. However, it is not central to the current analysis and so I will leave it aside for the purposes of this paper.

Finally, there is the case one where party takes the decision, another would get the benefit and a third suffer the loss. This is, perhaps, standard for government when regulating behaviour that impacts different groups in different ways. This is why I have called it 'adjudication'. The difficulties arise if all the benefits begin to pile up on one side, over time, case after case.

Understanding these different structures allows us to see some of the potential difficulties with cost-benefit analysis. Simple models aggregate costs and benefits without consideration of where those costs and benefits fall. In principle no distinction is made between the five structures outlined. But morally we can see that they are different, and we can also understand why I have insisted that cost-benefit analysis is relatively unproblematic only as part of a repeating series in which the probabilities of receiving costs and benefits run true. That will not be so in repeated plays of the several of the structures outlined.

It is also worth noting that the issues of moral hazard and moral sacrifice are another way of identifying 'negative externalities' and 'positive externalities' in economic theory. It is a common-place in economics that the free market, left to itself, will 'over-supply' goods with negative externalities, as this is a way of dumping costs on others, and 'under-supply' goods with positive externalities, as individuals will not want to take on costs for the sake of others. If we think of cost-benefit analysis as providing a benchmark of efficiency, and therefore a normative account of what ought to be supplied, then standard economic theory maps onto the table above very well. For the standard policy response is to find ways of 'internalising the externalities' which, in my terms, is also a way of attempting to move from a problematic structure to a less problematic one, either individualism or adjudication.

To have a fuller understanding of any case we need to know not only the structure, but the types of costs and benefits that are likely to be engaged. Risk analysis begins with an attempt to understand all factors in play, and to formulate a list of possible costs and benefits in qualitative, realistic, terms. For example, in deciding whether to open a new runway at an airport, it will be important to consider such things as air pollution to those who live nearby, as well as noise pollution, disturbed sleep and traffic congestion. Benefits would include reduced travel time for those able to make use of the airport, additional flights, and the commercial stimulus. Some of these, such as increased profits, are straightforwardly stated in monetary terms, but others, such as loss of sleep, are much more difficult to translate. As mentioned previously, where the decision is being made by means of a formula that requires inputs in monetary terms it makes great sense to convert everything into monetary terms, uncomfortable though that may be. But if the decision is being made without the use of a formula, then this appears to be an unnecessary, and, I will argue, potentially highly damaging step.

For let us consider the types of costs and benefits that can be associated with the decision to introduce a new technology. As an example, let us return to the introduction of GM crops, which has been accepted in the United States of America, but was much more controversial in Europe. Those in favour of introducing GM crops regarded the issue as essentially one where those understood science favoured GM crops whereas those who were opposed were scientifically illiterate. Defenders of GM claimed that the technique is not very much different from the widely accepted practices of selective breeding or of creating hybrids. Taboo, suspicion or sheer ignorance was alleged to be the reason for opposition. (For one discussion, see Nuffield Council 2004.)

There is no denying that some opponents did not understand the science well. But at the same time other, more subtle, lines of opposition also existed. One was purely at the level of risk analysis. It may be true that the new techniques were similar to selective breeding. But that does not mean that they are safe as selective breeding. Perhaps they would have risks that we are simply not aware of. Hence, while defenders framed the issue of one of straightforward cost benefit analysis, with the benefits far outweighing the costs, opponents thought that the situation was one of radical uncertainty, at least in part. And there is a cost to living with risk or uncertainty, even if nothing bad ever happens. (Wolff and de-Shalit 2007, Roeser 2014)

Second, it was also noted that once GM crops became widely used, all farmers would have to use them if they wanted a competitive yield of their crops. But the seeds would be patented, and the company that held the patent would have near monopoly power, which would make farmers very vulnerable. Even if the seeds were initially introduced at low cost, the manufacturers could abuse their monopoly position, giving them not only economic power, but also political power, as the productivity of an entire region could be in the hands of one manufacturer. This is particularly worrying in a developing country where farmers may have few resources to fall back on if things go wrong for them.

Hence in enumerating the costs and benefits engaged, it is necessary to take a very wide view. The possible benefits obviously include the well-being of those who have better nutrition and the farmers who might make more money, as well as the profits of the corporate interests who invent and market the new seeds. But it also includes the power and prestige of those who introduce and profit from the new technology. Being early to market with GM crops will attract attention from industry and government. It will lead the company executives to receive honours, to be invited to lecture at Business Schools, to sit in government advisory roles and help shape future legislation, which in turn will be more likely to be moulded round their interests. Prestige and power will follow from the roll-out of a new technology, provided it isn't an instant disaster. But prestige and power are zero sum and hence one party's gain is another's loss.

How much of this can and should be reduced to purely monetary terms in order to allow formal risk analysis to proceed? In terms of what is possible, the company involved will reap economic benefits, and these can be included unproblematically. A heroic attempt at a reduction of the intangibles of power and prestige to financial terms is no doubt also possible. But to repeat an earlier point, if there is no formula to enter the values into, it is unclear why it is worth going to the trouble, especially as estimates would have a vast margin of error. And to amplify a point made earlier but not yet explained, converting everything into a financial quantity bleaches out the nuance of the detail, and takes off the table what so alarms activists. It is not so much that if the innovation goes ahead some people will get very rich, but rather that those people will become powerful in diverse ways, none of which are beneficial for those who now become beholden to the company for their livelihood. These effects can last long in to the future, even when the technology is outdated and replaced by others. Hence economic modelling hides the politics. Cynics will say that this is precisely why economic modelling is so popular. But the question of how to regulate a new technology with unknown effects is a political question, not just a question for technocrats.

6. Decision-Making Under Radical Uncertainty

I have suggested that if we start from a position in which risk-cost benefit analysis is the default position for making decisions about risk and uncertainty, then as we move to situations of radical uncertainty, to which it does not apply in straightforward terms, the temptation is to try to modify the approach to deal with the more difficult cases. So far in these paper I have introduced four obstacles. First, because we do not know what the possible risks are, or their probabilities, we do not know how to maximize expected benefits. Second, some of the options may include devastating risks which would be masked by an averaging process. Third, in a one-off case we cannot assume that those who lose out will be compensated by gains elsewhere. And finally, attempting to reduce all costs and benefits, such as long term political influence, to a monetary value is problematic. To overcome those obstacles it is tempting to try to move to a principle that side-steps the problems.

Sometimes it is suggested that there is a 'precautionary principle' that we can appeal to in such cases. Quite what the precautionary principle is supposed to be is a matter of debate, but it seems fairly clear that it is not a principle in any decision-theoretic sense. Rather, as I have suggested elsewhere, it is better thought of as an attitude (Wolff 2014, see also Nuffield Council 2004). Others have used the idea of a 'precautionary checklist' (Myers 2005). In my understanding, we should regard the precautionary attitude as a series of linked questions, as follows. The preliminary stage asks two questions about immediate risks and benefits.

1. Are the costs and risks of the new technology acceptable?
2. Does it have significant benefits?

A second stage considers the social need for the technology:

3. Do these benefits solve important problems?
4. Could these problems be solved in some other, less risky, way?

The precautionary approach suggests that we should proceed only if the answer to these questions is: yes, the risks are acceptable; yes, it has real benefits; yes, it solves real problems; and no, we cannot solve the problem in other ways.

Consider, for example, some opposition to the introduction of GM foods. First, there was a disagreement about whether the technology was known to have unacceptable risks, which is, of course, itself a vague notion. Defenders were keen to normalize the risks by pointing out the similarities between genetic modification and selective breeding, which has been common for centuries; critics pointed out the discontinuities with earlier approaches. Second, there was broad, if not universal, agreement that GM crops could bring real benefits in terms of increased crop yields. Third, there was a somewhat less broad agreement that the benefits would solve real problems by helping to feed the developing world or, less significantly, reducing food prices in the developed world. However, only the opponents were interested in the fourth question: could we end hunger in the developing world in other ways? Here they had many suggestions, such as changing the protectionist trade policies of the wealthy countries, or investing in irrigation, infrastructure and other known development techniques. The critics said that the developing world goes hungry not through lack of technology but through lack of political will. We know, so they say, how to solve the problem in a way that does not put the environment and even human life at risk. So why introduce a very risky new technology?

Now, the opponents of GM could be wrong. Perhaps we can’t feed the world without GM crops. In other cases, it may be that the new technology does address a problem for which we do not have other solutions. We are beginning to see, for example, such claims being made about geo-engineering to reduce the greenhouse effect (for discussion cf RS/RAE 2017, and Lenzi 2018). It is possible, therefore, that a new technology will pass this initial, technical, stage, at least in some circumstances. But even so, we need to move to a more difficult political stage in which we explore the longer term effects of introducing technologies, in which we need to consider two difficult questions:

5. What are the possible longer term economic consequences of introducing the technology?
6. What are the possible longer term political consequences of introducing the technology?

Here is might be argued that there is little point in raising these questions, as we have already acknowledged that we are in a position of radical uncertainty; indeed this is the motivation for the discussion. Trying to answer these questions, it might be said, is to fall into the realm of fortune-telling. But in response, it is very likely that refusing to discuss them plays into the hands of those who wish to introduce the technologies, as it will make some of the significant benefits accruing to them, and costs to others, invisible.

For these questions return us to the discussion of moral hazard, power and prestige. In the case of GM foods critics argued that an enthusiasm for innovation plus corporate profit-seeking was the real motivation for pushing the technology forward. Corporate profit-seeking constitutes the most obvious moral hazard element. If GM crops work, and producers the world over are tied into the manufacturer's patented products, profits will soar to the great benefit of shareholders, and those individuals who work for the corporation. Of course, there are risks to these same agents. In the extreme, the company will go bust and people will lose their jobs. But shareholders typically hold a diversified portfolio and accept some losses as a price of investing in innovation. Those developing and marketing the technology will have skills and experience that make them highly employable elsewhere, even if the company collapses and everyone loses their jobs. On the other side, if GM crops fail, or become too expensive, farmers, especially in developing countries, may lose their livelihoods without being able to replace it, particularly if their land is damaged by the failure of GM crops. And consumers of the GM crops could also greatly suffer if the worst fears are realized.

Hence the issue of moral hazard is real: those promoting the technologies and putting pressure on governments to accept them have a lot to gain and relatively little, in comparison, to lose. Those who will be using the technologies have something to gain, but an awful lot to lose. The economics are unbalanced between corporate and individual interests and this should give regulators reason to pause. At the very least they need to take a very hard look at any evidence presented to them by the corporate interest.

At this point the issue shades into the questions of power and prestige. I would not wish to enter into anti-business conspiracy theory in which companies develop products in order to strengthen their political power. Any political benefits could be an unintended consequence. But, and to simplify, let us suppose there are two types of unintended consequence: those that benefit the interests of those who run the corporation and those that are detrimental to their interests. It seems very likely that agents who are able to do so will take steps to mitigate or overturn practices that are detrimental to their interests, while ignoring those consequences that further their interests. They may even take for granted as 'just how things are'. In this case, the unintended consequences may well be that those who are responsible for introducing successful new products will receive accolades and prizes, and will be invited to join committees to shape future regulation of their area, where they will be able to consolidate their success and magnify its effects.

In sum, we need to be particularly watchful of two types of problematic aspects regarding the introduction of a new technology. First, those who have developed the technology and are pushing for its introduction, and who would benefit from it, are very likely to focus much more attention on its benefits (to others) than on the costs and risks. However hard they try (and they might not try at all) it is very difficult to put self-interest entirely to one side and perform a properly balanced assessment.

The second level concerns the political consequences, in terms of the consolidation of corporate power, and increasing vulnerability of those who are already vulnerable. The basic point is very simple. As more people become dependent on the products of a smaller number of suppliers, the power of the supplier increases, even to the point of 'regulatory capture' where they help compose the regulations that they must follow (Ramanna 2015).

This process is self-fueling unless checked. Therefore, even if there is an agreement that the technology should go ahead, safeguards need to be put in place to provide a check on growing corporate power, if possible.

To put this point in terms of the risk triangle introduced earlier, we worry, initially, that corporations producing goods such as GM crops are in a position of moral hazard. Unchecked they will hope to make gains for themselves, pushing risks on to the vulnerable consumers. Government is brought in to regulate, thereby transforming the situation into something closer to what I called adjudication in which the government acts as a neutral broker. However, if the technology is allowed, a sphere of regulation will develop, which in turn will draw on those who know the industry. If it turns out that the regulators are sympathetic to corporate interests then we have an element of regulatory capture, where the industry, in effect, regulates itself, in part at least. This, then, converts what was meant to be adjudication back into moral hazard. Arguably this mechanism is the governance failure of our age, repeated over and over again, and incredibly hard to guard against. Nevertheless, even when we have observed such risks, it may be the case that going ahead is the right thing to do, even in the face of likely adverse long-term political consequences. My argument is that the listed questions raise considerations to be taken into account, not a set of necessary and sufficient conditions.

7. Conclusion

My conclusion is a fairly commonplace one. Although many areas of risk can be approached and dealt with by technical tools of economic assessment, doing so in the case of risky new technologies is highly problematic, as, first, the tools are not appropriate for cases of radical uncertainty, and second, attempting to use these tools hides the deeply political dimensions of new technologies. There is no real alternative to having what will be fraught and contested debates about the costs and benefits of the technologies, where those costs and benefits need to be understood in their complexity rather than reduced to a colourless monetary sum. All the questions raised in the previous section need to be addressed. In particular, however, the absolute key question is whether the new technology has a real chance of being the best response to a genuine problem. If not, it can be hard to see why it is worth taking a step into the unknown, with the variety of problems to which it can leave us open.

Carthy T, Chilton S, Covey D, Hopkins L, Jones-Lee MW, Loomes G, Pidgeon N, Spencer A. 1998. On the contingent valuation of safety and the safety of contingent valuation: Part 2 - The CV/SG "chained" approach. Journal of Risk and Uncertainty 17: 187-213.

Drèze, J. 1962. L'Utilitè Sociale d'une Vie Humaine. Revue Française de Recherche Opèrationelle 6: 93-118.

Hannson, S.O. How to Perform an Ethical Risk Analysis (eRA). Risk Analysis 38: 1820-1829.

Hayenhjelm, M, Wolff, J. 2012. The Moral Problem of Risk Imposition. European Journal of Philosophy 20, S1: E26–E51.

Hermansson, H. Hansson, S. O. 2007. A Three-Party Model Tool for Ethical Risk Analysis Risk Management 9: 129-144.

HSE. 2001. Reducing Risk, Protecting People. London: HMSO.

Lenzi, D. 2018. The ethics of negative emissions. Global Sustainability 1, e7, 1–8.

Manson, N. A. 2002. Formulating the Precautionary Principle. Environmental Ethics 24: 263 – 274.

Munthe, C. 2011. The Price of Precaution and the Ethics of Risk. Dordrecht: Springer.

Myers, N. J. 2005. A Checklist for Precautionary Decisions. In her (ed) Precautionary Tools for Reshaping Environmental Policy. Cambridge Ma.: M.I.T. Press. 93-106.

Nuffield Council of Bioethics. 2004. The Use of Genetically Modified Crops in Developing Countries: A Follow-Up Discussion Paper. London: Nuffield Council of Bioethics.

Ramanna, K. 2015. Thin Political Markets: The Soft Underbelly of Capitalism. California Management Review 57: 5–19.

Roeser, S. 2014. The Unbearable Uncertainty Paradox. Metaphilosophy 45: 640-653.

RS/RAE. 2017. Greenhouse gas removal. London: Royal Society/Royal Academy of Engineering.

Schelling, T. 1968. The life you save may be your own. In: Problems in Public Expenditure Analysis, ed. S. Chase. Washington, DC: Brookings Institution. Pp. 127-62.

Setright, L.J.K 2004. Drive On!: A Social History of the Motor Car. London: Granta Books.

Wolff, J. 2006, Making the World Safe for Utilitarianism, Royal Institute of Philosophy Supplement, 58: 1-22.

Wolff, J. 2011a, Ethics and Public Policy. London: Routledge.

Wolff, J. 2011b, Five Types of Risky Situation. Law Technology and Innovation 2: 151-163.

Wolff, J. 2014. The Precautionary Attitude: Asking Preliminary Questions. In: Synthetic Future: Can We Create What We Want Out of Synthetic Biology?, special report, Hastings Center Report 44, no. 6: S27-S28.

Wolff, J. and de-Shalit, A. 2007. Disadvantage. Oxford: Oxford University Press.