Precautionary Principle

The economic perspective

SUMMARY

How should society deal with risks when there is scientific uncertainty about the size of these risks? There has been much recent discussion of the Precautionary Principle, which states that lack of full scientific knowledge should not be used as a reason to postpone cost-effective preventive measures. We show in this paper that the Precautionary Principle contradicts one important intuition about the right way to act in the face of risk, namely the principle of 'looking before you leap'. When we expect to learn more about the future, the effectiveness of our preventive measures will be greater if we learn before we act. However, a number of other ways of taking uncertainty into account are consistent with a reasonable interpretation of the Precautionary Principle. First, postponing preventive measures may increase our vulnerability to damage, which induces a precautionary motive for risk-prevention, similar to the precautionary savings motive. Secondly, stronger preventive actions often yield more flexibility for the future, so that acting early has an option value. Thirdly, when better information comes from a process of learning-by-doing, the risk associated with early events is amplified by the information they yield about the future. This plays a role analogous to that of an increase in risk aversion, making us more cautious. Fourthly, because imperfect knowledge of the risk makes it difficult to insure, the social cost of risk should include a risk premium. Finally, uncertainty about the economic environment enjoyed by future generations should be taken into account. This raises the benefit of acting early to prevent long-term risks.

If the Precautionary Principle sometimes gives good and sometimes gives bad advice, there is no escape from the need to undertake a careful cost-benefit analysis. We show that standard cost-benefit analysis can be refined to take account of scientific uncertainty, in ways that balance the Precautionary Principle against the benefits of waiting to learn before we act. Furthermore, it is important that they be used to do so, for instinct is an unreliable guide in such circumstances. Abandoning cost-benefit analysis in favour of simple maxims can result in some seriously misleading conclusions.

— *Christian Gollier*

Economic Policy October 2001 Printed in Great Britain

Should we beware of the Precautionary Principle?

Christian Gollier

Université des Sciences Sociales de Toulouse

1. INTRODUCTION

Precaution is certainly a virtue, but too much precaution can be as damaging as too little. In the early 15th century, China led the world in scientific knowledge, technology and the welfare of its people. Along the East African coast, Chinese flotillas far surpassed in grandeur the small Portuguese fleets that were to come later. In 1405, one of these Chinese fleets consisted of 317 vessels carrying 28 000 men. But, as David Landes (1999: p. 96) describes, 'after some decades of tugging and hauling ... the decision was taken not only to cease from maritime exploration but to erase the very memory of what had gone before lest later generations be tempted to renew the folly. ... The abandonment of the program of great voyages was part of a larger policy of closure, of retreat from the hazards'. This precautionary aversion to change struck many generations of visitors to the Celestial Empire. Europeans, in contrast, accepted many novelties brought home by their explorers, with what must have seemed like recklessness to the Chinese. Europe lost some of its plants and biomass, but in return came corn, tomatoes and turkey. In later centuries, this openness to innovation triggered the industrial revolution, as a consequence of which Europe left China economically far behind. In an interesting reversal of history, the Club the Rome recommended in 1972 a halt to economic growth because of pollution hazards and

In the process of writing this paper, I benefitted from helpful discussions with Giuseppe Bertola, Jacques Drèze, Tore Ellingsen, Olivier Godard, Jean-Jacques Laffont, Benny Moldovanu, Paul Seabright, Jean Tirole and Nicolas Treich. The technical assistance of Edwige Sauvé for collecting information on the web is acknowledged.

uncertain reserves of natural resources. Half a dozen years later, China was to commence one of the most remarkable bursts of economic growth in human history.

This paper asks how society should manage hazards whose characteristics are not perfectly known. The Precautionary Principle (PP) has been proposed as offering guidance in these circumstances. This principle was enshrined at the 1992 Rio Conference in its principle 15, which states that: 'where there are threats of serious or irreversible damage, lack of full scientific certainty shall not be used as a reason for postponing cost-effective measures to prevent environmental degradation'. The European Community officially endorsed the PP in Article 174 of the EC Treaty. It is clear that the applicability of the PP is not limited to environmental risks. However, the PP does not stipulate how the phrase 'cost-effective measures' should be interpreted when full scientific certainty is lacking. Thus the PP is subject to divergent interpretations.

Many experts and politicians, particularly in France after the infected blood scandal of the 1990s, now favour a restrictive interpretation in which everything should be done to reduce risk to its lowest possible level, whatever the cost. Even if we could identify a lowest possible level of risk, though, such a policy is incompatible with any kind of economic efficiency, as it would virtually eliminate innovation in our society. The extraordinary innovation-driven growth rate of our economies over the last two centuries clearly calls for the rejection of this extremist approach. The extremely low variability of the growth rate over this period is also an indication that the riskiness of the innovation process has been under control, at least at the macro level.[1] But this does not mean that risks have been adequately managed at the micro level, nor that we are sure to be able to avoid catastrophic dangers in the future. Some risks taken in our economies are socially undesirable. But detecting them is difficult, not only before but even after the event. In particular, even when avoidable damage takes place it may not be right to conclude that unreasonable risks were incurred. Sometimes it is right to take risks even when we know that such risk-taking will occasionally go wrong.

How can we decide whether a risk is acceptable to society? Using the language of cost-benefit analysis, we can say that the risk is acceptable if its benefits to society exceed its costs. But to say this is merely to re-state the problem, for by assumption the benefits and costs are uncertain. Where these benefits and costs have known probabilities, and where individuals can diversify away their own risk through insurance and other markets, we know from the work of Arrow and Lind (1970) that such a risk will be acceptable if its expected net present value is positive. This criterion is a standard rule used by public decision-makers in a wide variety of fields from road safety to long-term investments in the energy sector. Where we cannot measure risks precisely, however, we cannot simply apply this technique mechanically. But this does not mean we should abandon the spirit

[1] Lucas (1987) came to a similar conclusion. He showed that the observed variability of consumption per capita around a secular trend growth rate of 2% per year has the same effect on the well-being of ordinarily risk-averse consumers as a reduction of less than 0.01% of the growth rate! Risk and risk aversion are a second order phenomenon compared to the effects of compounding on growth rates.

and general methods of cost-benefit analysis. The central message of this paper is that we can assess risks systematically even when we lack full scientific knowledge about the size of the risk. Some of the ways of taking such lack of knowledge into account are compatible with an intuitive interpretation of the Precautionary Principle, though others are not. As scientific knowledge develops, it will turn out sometimes that large sums of money have been spent in fighting phantom risks, whereas not enough attention has been devoted to other risks that happened to be catastrophic.[2] We shall need to evaluate such developments without hindsight, judging the reasonableness of risk-prevention measures against the state of scientific knowledge available at the time.

Whatever its other merits, the Precautionary Principle has one implausible implication in circumstances where actions we may take to avoid risk in the future are a good substitute for actions we might take today. Because we can reasonably expect to have better information in the future about the consequences of our actions, future actions are likely to be more effective in avoiding any given risk. This should make us more willing to wait and see than to take preventive action immediately, to 'look before we leap'. This implies the desirability of a 'learn-then-act' principle, whose message is the opposite of the PP. But we show in this paper that the 'learn-then-act' principle is only one of several ways in which we need to take scientific uncertainty into account, and most of the other ways of doing so are in agreement with a reasonable interpretation of the PP.

Two particular issues are important. The first is irreversibility. Some present choices reduce our freedom of action in the future. This loss of flexibility has a cost. This type of irreversibility may characterize carbon dioxide emissions and genetically modified organisms (GMOs), which are hard to remove from the environment once introduced. Taking the irreversibility into account militates in favour of earlier prevention, in line with the PP.

A second issue is the fact that we often learn about future risks not just by engaging in research, but also by observing current developments. The worse the current state of things the more pessimistic we are likely to become about the future. Here the literature on hedging demand in finance suggests that we need to take more preventive efforts against bad outcomes than if we had full scientific knowledge of the risks. Again, this is compatible with a reasonable interpretation of the PP.

The recurrent message of the paper is that the occurrence of 'new risks' in our society should not be used as a pretext for abandoning cost-benefit analysis as a guide to public decision-making.[3] Taking careful account of scientific uncertainty can indeed influence our decisions in ways that make sense of some of the intuitions embodied in the

[2] The history of innovation is full of examples. Asbestos is an example of a new product that happened to have adverse health effects. In contrast, many innovations have been delayed because of insufficient knowledge about their effects. In the 18th century, Parmentier had to protect fields of potatoes imported from America because of the common wisdom that they were poisonous. Similarly, a large fraction of the population was against the building of railroads, because it was believed that human beings could not survive acceleration to high speeds.

[3] See Godard (1997) for an exposition of these alternative views of the PP where the comparison of the distribution of costs and benefits becomes inessential for the decision. Godard (2000) compares the kind of economic approach to the PP considered in this paper to an interpretation of the PP taken as a social norm.

Precautionary Principle. But simply following intuitive reasoning without the discipline of cost-benefit analysis can lead us to make many costly mistakes.

The structure of the paper is as follows. In Section 2 we set out some basic facts about two problems where there has been the most public and widespread appeal to the Precautionary Principle: namely, global warming and the BSE ('mad cow') crisis. In Section 3 we examine the assumption made in standard cost-benefit analysis that individual decision-makers are expected-utility maximizers. While ignoring the effect of changing knowledge over time, we ask what difference it makes that the probabilities attached to various risks are themselves uncertain. Section 4 looks directly at the effect of changing knowledge over time. How does the fact of progress in our understanding of risk affect the optimal timing of efforts to reduce risk? Insurability problems and long-term effects are examined in Sections 5 and 6. Up to that point, we assume that, although there is uncertainty, that uncertainty is common to everyone in society; we ignore the particular problems of asymmetric information about the nature of the risks involved. So Section 7 considers a range of issues arising from asymmetric information: the right of citizens to take their own decisions about risk, the traceability of products and foodstuffs, and the question whether experts and politicians have the right incentives to reveal their information to the public in an unbiased way. The appendix sets out a reasonably general formal model, special cases of which provide a number of the applications discussed in earlier sections.

2. ILLUSTRATIONS

Examples abound of risks whose potential cost to us is only imperfectly known: hazards from various wastes, species extinction, global warming, low doses of radiation, electromagnetic fields, cellular phones, genetically modified (GM) food; and the list could go on. This section focuses on two cases: bovine spongiform encephalopathy (BSE) and global warming. It summarizes those aspects of existing knowledge that are useful for public decision-making.

2.1. The mad cow crisis

The three mad cow crises of 1996, 1999 and 2000 raise the difficult question of whether to impose common European prevention rules in the face of BSE. Should we impose a complete ban on the use of recycled cattle bones in animal feed, given evidence of the transmission of the disease through animal proteins, and of fraud in several countries resulting in the remains of cattle being fed to pigs and poultry? Should we, as in France, kill all the animals in a herd where a case of BSE has occurred, in spite of the evidence that among the thousands of animals killed, only one extra cow has been found to have BSE? Should we require testing for BSE of all cows that enter into the food chain?

Because of scientific uncertainty about the transmission mechanism of new variant Creutzfeld-Jakob disease (nvCJD), the current best estimate of the cumulated number of victims in the UK over the next 20 years is 6000, but with a minimum in the hundreds and a maximum at 250 000. For a population of British citizens at 60 million, this puts the individual probability of getting nvCJD within the next 20 years between 2×10^{-5} and 4×10^{-3}, with a mean at 10^{-4}.[4] From the start of the BSE epidemic in 1988 to March 1996, the British government's expenditure on fighting BSE had already reached £247 million, just for compensating farmers, destroying cattle, administering controls and research. Due to measures taken in 1996 in the face of the first human cases of nvCJD, the cost of preventive action went up to around £700 million per year, or around 0.1% of GDP. Whereas the costs of the preventive efforts are relatively easy to measure, the central question is to determine how to take account of the ambiguity of the health risk in evaluating the benefit of these efforts.

How did consumers react to the risk? In March 1996, French consumers reduced their purchase of beef by 35%, but the reduction was only 10% by the end of the year. The price fell by 20% during the crisis, and recovered to 10% below its original level by the end of the year. The crisis was much stronger in November 2000 after potentially tainted meat was discovered on supermarket shelves in France. During the last quarter of 2000, beef consumption slumped by 27% across the 15 countries of the European Union, while the wholesale price of beef in France went down by almost 25%. The effect of the crisis on the consumption is depicted in Figure 1 for each of the 15 countries.

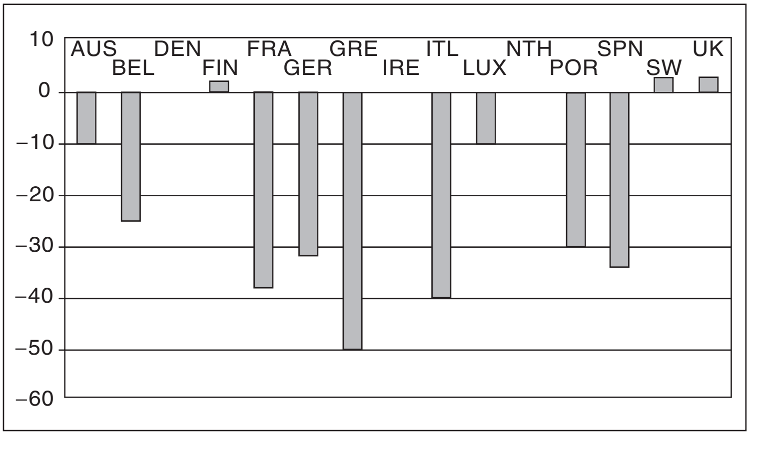

Figure 1. Percentage change in EU beef consumption, 1 October–31 December 2000

Source: European Union

[4] For the sake of comparison, the annual death toll from car accidents and cancer are 3421 (in 1998) and 156 000 (in 1996) respectively. This puts the probability of death over the next 20 years at respectively 0.001 and 0.06 for these two hazards.

2.2. Global warming

The risk linked to global warming is affected by a similar degree of scientific uncertainty. For example, the Intergovernmental Panel on Climate Change (IPCC, 1995) has published estimates ranging between $5 and $125 for the net present value of future damage generated per tonne of carbon emitted today. Because each litre of petrol consumed yields 2.36 kg of carbon dioxide, this implies it would be desirable to impose a tax somewhere between 1.2 and 30 cents per litre of petrol – or issue consumption permits whose equilibrium price would lie within this range. The wide range of damage estimates does not reflect natural risks due to random shocks to the biosphere. Instead, it comes from our limited scientific knowledge of the mechanisms at various stages of the process linking emissions to damages: imperfect knowledge of the impact of greenhouse gases on the climate, imperfect knowledge of the effect of climate change on the environment (including its possible feedback loops), and imperfect knowledge of the impact of environmental change on human welfare. This uncertainty, combined with the prospect of rapid scientific progress, was among the arguments used to justify the recent decision of the Bush administration not to ratify the Kyoto Protocol.

3. DECISIONS WHEN RISK PROBABILITIES ARE UNCERTAIN

How should a decision-maker behave when there is such a large degree of uncertainty about the distribution of the risk? In the next two sections, we ignore doubts about the trustworthiness of public decision-makers, assuming for the time being that they are benevolent, and capable of aggregating any relevant information in order to measure the risk.

3.1. Subjective expected utility

Suppose you are contemplating two uncertain prospects: a picnic in the country (with the attendant risk of rain), versus an afternoon at the cinema (with the risk that the movie may turn out to be terrible). You decide that, all things considered, you prefer the picnic.

Now you learn that your car has developed a fault which means that, with some probability, it will break down after a few miles and prevent you from reaching either destination. It will be equally unpleasant having to repair the car whether you break down on the road to the picnic or the road to the cinema. Is your preference between the picnic and the cinema affected by the fact that your choice is no longer a definite one, but a choice conditional on your not having to spend the afternoon repairing the car? If your choice of whether to set out for the picnic or the cinema is unaffected, whatever the probability of not actually getting to your chosen destination, then your preferences satisfy the so-called 'independence axiom'. This axiom, introduced by von Neumann and Morgenstern (1944) and extended by Savage (1954), is a founding stone of (subjective) expected utility theory. Under very general conditions, it implies that preferences are linear in probabilities. For our purposes, this in turn implies that when the probability of an event is itself uncertain decision-makers should act as if it were certain, and equal to the mean of the possible values it could take: choices depend only on the compounded probability of the various outcomes that affect individual utility directly.

To illustrate, consider the case of nvCJD in the UK. We assume that being a victim of the disease generates a damage that is equivalent to a financial loss of 50 times per capita GDP per capita.[5] We also assume that victims are fully compensated for the reduction of their life expectancy. This implies that the risk is well diversified across society (an unrealistic assumption in practice, and one which we shall relax in Section 5). In spite of the high uncertainty of being a victim of the disease, we know that, under the independence axiom, we can just focus on the average probability $\bar{p} = 10^{-4}$ (see Section 2). Under this theory, the risk for British citizens over the next 20 years is equivalent to a reduction in their wealth amounting to $50\bar{p} = 0.5\%$ of GDP per capita. In other words, if there existed a method to eliminate nvCJD risk for human beings in one shot, it would be efficient to implement it only if it cost less than 0.5% of GDP. Using this method to deal with scientific uncertainties makes precaution not very different from standard approaches to protection (where probabilities are objective): just take expected probabilities and apply standard cost-benefit analysis with these probabilities to determine the optimal risk-reduction effort.

3.2. Aversion to ambiguity

Not everyone agrees that focusing on expected probabilities is the appropriate criterion when dealing with scientific uncertainty. For example, most environmental organizations like Greenpeace favour an alternative wording of the PP: 'When an activity raises threats of harm to human health or the environment, precautionary measures should be taken even if some cause and effect relationships are not fully established scientifically'. This is interpreted as requiring society to behave as if the worst theory were true, even if not fully established scientifically. This strong interpretation of the PP is even clearer in the conclusion of the French Conseil d'Etat about the infected blood affair (1995): 'In an uncertain situation, an hypothesis that cannot be rejected should be taken as temporarily valid, even if it cannot be formally proven'. Can we justify such an interpretation of the PP?[6]

In fact, some people do not behave in a way consistent with maximizing subjective expected utility. They do not behave in the same way in the face of two uncertain environments with the same average probabilities, but with different weights of

[5] There is much to say about the value of life. The debate about the PP can also be examined from the viewpoint of a crisis about the value of life and of environmental assets. We ignore these issues here in order to focus on the role of scientific uncertainty. If we think it appropriate to impute higher values to lives lost, the calculations in the text will need to be adjusted upwards in proportion.

[6] Dupuy (2001) provides various philosophical arguments in favour of the strong interpretation of the PP.

evidence. More precisely, they are ready to pay more to get rid of a more uncertain risk. Ellsberg's paradox (see Box 1) states that most people's choices are affected by the weight of evidence about probabilities. The European Commission (2000: p. 4) says exactly this when it states that 'decision-makers need to be aware of the degree of uncertainty attached to the results of the evaluation of the available scientific information'.

Gilboa and Schmeidler (1989) have proposed an alternative decision criterion to explain the Ellsberg paradox. Under their criterion, individuals perform a sequence of two operations. First, for each possible theory, they compute the corresponding expected utility that they would get by assuming that this theory is the true one. Next, they behave as if the true theory is the one that yields the lowest expected utility. Notice that the advantage of this criterion is that we don't need to assign any probability to each possible theory. This maximin criterion obviously generates a strong aversion to ambiguity that can be interpreted as maximum pessimism.

In the nvCJD case, this would mean that we should act as though the theory predicting a cumulative death toll of 250 000 were the true one. Under Gilboa and Schmeidler's criterion, the UK government should be ready to pay as much as 20% of GDP immediately in order to get rid of BSE.

Gilboa and Schmeidler's criterion has its own problems, however, notably because of its extreme pessimism. People behaving in this way would be pathologically averse to risk, including the daily risk of crossing the road. As far as large medical risks are

Box 1. Ellsberg's paradox

Daniel Ellsberg's (1961) experiment is based on the Keynes–Ellsberg 'two-color' problem. There are two urns each containing red and black balls. Urn 1 contains 50 red balls and 50 black balls, whereas urn 2 contains 100 red and black balls in an unknown proportion. A ball is drawn at random from an urn and the player receives 100 euros if she correctly predicts the colour of the ball. Whichever urn is used, most subjects are indifferent between betting on red or black, a fact which is interpreted to show that their subjective probability for each color is 0.5, whether they have been told the proportions or not. If they were subjective expected utility maximizers, they would therefore be indifferent between using urn 1 and urn 2. However, most people prefer to gamble with urn 1, where the 'weight of evidence' is felt to be greater. Betting on urn 2 is like facing scientific uncertainty. There are 101 possible 'theories' about the actual risk taken by those who bet on urn 2, with theory i assigning a value i to the number of red balls in the urn, $i = 0, 1, ..., 100$. In this example, with no additional information, it seems reasonable to assign the same probability to each theory, so that the probability of winning with urn 2 is exactly equal to 1/2.

concerned, it is always possible to find some expert favouring an extreme and catastrophic scenario. Even if it is extremely unlikely to occur, ambiguity-averse agents will do everything they can to avoid it. They will never invest, and they will never innovate. Moreover, the optimal decision will be extremely sensitive to the set of 'plausible' distributions of risk, which raises an important problem of implementation.

There is good reason to doubt just how averse to ambiguity people really are. Savage himself argued that although people might make mistakes, they would change their minds after a more detailed evaluation of the situation. Notice that if consumers were as ambiguity-averse as Gilboa and Schmeidler's criterion requires, they would all have stopped eating beef entirely because of the ambiguity surrounding the risk of nvCJD. This is not what has been observed. In a different context, Viscusi and Chesson (1999), using a sample of 266 business owners facing risks from climate change, show evidence of both ambiguity-seeking behaviour and ambiguity-averse behaviour. More precisely, people seem to exhibit an aversion to ambiguity for small probabilities of suffering a loss, and an attraction towards it for large probabilities. Fox and Tversky (1995) also showed in a series of experimental studies that ambiguity aversion, present in comparative contexts in which a person is confronted with both clear and ambiguous prospects, seems to disappear in noncomparative contexts in which a person evaluates only one of these prospects in isolation.

4. UNCERTAINTY THAT IS RESOLVED OVER TIME

The previous section did not take account of the intrinsically dynamic nature of scientific uncertainty, which can be expected to be resolved over time through learning. The Precautionary Principle makes a very clear claim about this: scientific uncertainty should not be used as a reason for *postponing* preventative efforts. In order to evaluate this claim, we need to consider how learning will affect the nature of the actions that can be taken in the future. A generic model of dynamic risk taking with Bayesian updating is presented in the Appendix. All applications discussed in this section are special cases of it.

4.1. The learn-then-act principle

Consider a new good whose consumption generates some immediate added value, but poses a future risk. When the future risk depends upon the cumulative total of past consumption, we can think of a reduction in current consumption as representing a form of precautionary saving. It allows for consumption to be increased in the future for the same level of future risk. In this sense, preventative efforts tomorrow are perfect substitutes for preventative efforts today. The problem is thus to determine the effect of scientific progress on our willingness to postpone the consumption that results from innovation. To illustrate, it is not just the current emission of carbon dioxide that determines the current climate. Rather, it is the current *concentration* of this pollutant which determines it, and this concentration is the result of all past emissions. It is also the accumulation of past exposures to nuclear radiation, cellular phones and smoking that seems to determine the health risk.

How does the expectation of scientific progress affect the desirability of reducing consumption of a pollutant good today? Better information improves the efficiency of actions taken in the future. It has an effect similar to an increase in future revenues. Because of the willingness to smooth consumption over time, that should induce us to enjoy some of the benefits of that improved efficiency in the present, in other words to raise the immediate consumption of the good. This is a very intuitive effect, which tends in the opposite direction from the Precautionary Principle. When future efforts are good substitutes for current efforts, it may be optimal to wait for information before taking potentially drastic actions. This is the learn-then-act principle.

But postponing preventative actions may also *increase* future risk. Leland (1968) and Kimball (1990), among others, have observed that an increase in future risk raises consumers' willingness to make provision for the future if they display *prudence*, that is if the marginal utility of their consumption is convex. What does this mean? In principle people could react in one of two ways to the news that the future is likely to be less predictable than they had previously thought. They could save more, because of the chance of being poorer in the future. Or they could save less, because of the chance that they may be much richer, and their saving will have been unnecessary. If the first motive dominates the second, as seems plausible, we say that people display prudence. This prudential effect is an argument for the PP. Whether it outweighs the learn-then-act principle depends on how large is the degree of absolute prudence. For utility function displaying constant relative risk aversion, Gollier *et al.* (2000) show that prudence will outweigh the learn-then-act principle if and only if relative risk aversion is *smaller* than unity. Most economists believe that relative risk aversion is larger than unity.

To illustrate, consider the case of global warming. To keep the story simple, we consider two periods, 'today' and 'the future'. Each litre of petrol consumed at any time yields a discounted damage to the environment in the future that is evaluated to be either 1 cent or 30 cents. This is compatible with the degree of uncertainty expressed in the second report of IPCC (1995). The probability that the damage is high is itself uncertain. It can take the values $0.5(1-k)$ or $0.5(1+k)$ with equal probabilities, and the true value of this probability will be revealed in the future. The parameter k measures the speed of scientific progress. The immediate benefit of using petrol is 20 cents per litre.[7]

If relative risk aversion is constant and equal to 2, the optimal initial consumption of petrol is $c_1 = 0.310$ when there is expected to be no scientific progress, that is when $k = 0$. Suppose alternatively that some scientific information is expected to be revealed over the next 10 years (k is positive). If we fix $k = 0.2$, we can check that the optimal initial

[7] Technically, in the absence of more scientific progress between today and the future, the socially optimal consumption plan of fuel maximizes $u(vc_1) + \beta Eu(vc_2 - (c_1 + c_2)\tilde{x})$, where $\beta = 0.8$ is the discounting factor of the first period, which lasts for around 10 years. We consider a utility function $u(c) = (1 + c)^{1-\gamma}/(1-\gamma)$.

consumption of petrol goes *up* to $c_1 = 0.322$. This represents an increase in initial consumption of 3.9%. We can also check that the initial optimal emission is almost doubled by potential scientific progress if $k = 0.8$, i.e., if 80% of the uncertainty is eliminated in the second period. So for plausible values of risk aversion, it appears that prudence is outweighed by the benefits of waiting to learn before we act.

4.2. The option value of flexible actions

The central concept of any dynamic risk analysis is flexibility. When our knowledge about the underlying risk is expected to change over time, it is valuable to be able to adapt our actions to our new information. Thus, any immediate action that reduces our ability to be flexible in the future has a cost. Arrow and Fischer (1974) and Henry (1974) represent this cost in terms of the forgone *option value* of a more flexible, or less irreversible, action.

The concept of option value is best illustrated by the problem of genetically modified seeds in agriculture. It is usually considered that the introduction of GM seeds in agriculture is an irreversible decision. This means that the decision to introduce them today eliminates any degree of freedom for future actions (although the presence of terminator genes tends to diminish this effect; issues raised by terminator genes are discussed by Harhoff *et al.* in this issue of *Economic Policy*). In particular, it will not be possible to remove GM seeds if alarming information about the likely risks becomes available in the future. So a careful cost-benefit analysis should include an opportunity cost of current actions that diminish future flexibility. This is compatible with the Precautionary Principle, since it is often the case that early preventative actions offer more flexibility for the future.

Most of the risks discussed in this paper are subject to irreversible decisions. Reducing biodiversity is an irreversible phenomenon. Most of the carbon dioxide that we emit today cannot be removed from the atmosphere in the future, thereby reducing the ability of future generations to select the best level of concentration given their scientific knowledge. Nordhaus (1994), Manne and Richels (1992) and Ulph and Ulph (1997) have calibrated models of global warming with scientific progress. They found that learning has little or no effect on decisions, mostly because the likelihood is small that future generations will ever want not only to eliminate their own emission, but also to reduce the existing concentration of pollutants. Calculating option values may be a difficult challenge. Techniques exist that are based on dynamic stochastic programming or on the theory of real option values (Dixit and Pindyck, 1994). They have been used to help decision makers for a large variety of investment decisions, from preventative actions to investment in mining.

An interesting question is whether competitive firms take proper account of option values in their own investment decisions. Competing firms may be tempted to introduce innovations immediately in order to pre-empt the market, as shown for example by Weeds (1999). The role of competitive pressure in inducing such pre-emptive behaviour suggests one role for the PP. This argument provides a rationale for public institutions with the power of authorizing the launch of new products on the market.

4.3. Learning by observing trends

In Section 4.1 we assumed that the improvement of our knowledge arose entirely through theoretical progress, unrelated to the observation of the risks themselves. We could alternatively consider situations where our knowledge improves just because we observe early evidence of the risk. A large part of the recent alarm about global warming is due to observing a positive trend in the average temperature on the earth, more than from any improvement in our understanding of the physical determinants of the climate.[8] Similarly, fear of BSE has been substantially triggered by the sharp increase of nvCJD cases in the UK.

In the presence of such a learning process, the effect of a large initial loss is worsened by the prospect of larger losses in the future. In contrast, the absence of damage today is good news about the likelihood of future losses. Thus there is a sense in which this learning process increases the degree of risk, implying an increase in the desirable level of preventative effort today. Thus learning by observing trends is compatible with the Precautionary Principle.

This intuition has been studied formally in the literature on learning in dynamic portfolio management. Merton (1973) was the first to characterize rules for the optimal dynamic portfolio management when the set of future risk-taking opportunities is stochastic. Detemple (1986), Genotte (1986), Brennan (1998) and Brennan and Xia (1999) have examined the specific where the opportunity set is stochastic due to the initial parameter uncertainty of the dynamic stochastic process. McCardle and Winkler (1992) have examined a similar problem in the context of gambling. In a casino, there is an urn of indistinguishable coins, half of which are 'good' and half 'bad'. The good coins land heads, with different probabilities that are known in advance. A single coin is picked at random from the urn that is used for n plays of the game. At each play of the game, you choose how much you want to bet. What is the optimal dynamic strategy in this game against nature? McCardle and Winkler (1992) put this question to over 200 students and found that most people prefer not to bet at first, in order to gather information about the coin. This could be termed precautionary behaviour.

This literature about the effect on investment behaviour of learning the distribution of asset returns shows that, when relative risk aversion is constant, the optimality of a precautionary strategy depends upon whether relative risk aversion is larger or smaller than unity. It is indeed optimal when relative risk aversion is larger than unity (which is the empirically plausible case). Gollier (2000) provides an explanation for why this argument is misleading when relative risk aversion is less than unity.

[8] Globally, the 1990s was the warmest decade on record and 1998 the hottest year.

5. LOSS AVERSION AND RISK SHARING

In the previous section, we implicitly assumed that everyone in society faced the same perfectly correlated risks. In reality, only a fraction of the population will be hit by nvCJD, and not all countries will lose from global warming. When individual risks are idiosyncratic, we need to address the problem of how these risks are shared. This is particularly true when there are asymmetries in individuals' attitudes to gains and losses. One way of understanding such asymmetries is through the concept of loss aversion, a concept introduced by Tversky and Kahneman (1992) in their Prospect Theory. In a static model, loss aversion is a special case of risk aversion where the utility function has a kink at the current level of wealth. From this reference point, a unit loss generates a disutility that is claimed to be roughly equal to twice the utility of a unit gain.

The efficient sharing of risks would completely wash out idiosyncratic risk by diversification, and neither risk aversion nor loss aversion would matter. Full diversification requires that the victims are fully indemnified for their realized losses in such a way as to share all the risk with others. In most cases, this assumption is unrealistic. There are obvious reasons why these risks are difficult to insure on competitive insurance markets. If, for any reason, insurance companies are more pessimistic than consumers, the latter could decline to take out insurance because of the high premium. Moreover, after the asbestosis crisis of the 1980s, insurance companies discovered the legal risk associated with uncertain future interpretations of tort laws. This legal risk is inherent to hazards that have effects far in the future, and insurance companies are now very reluctant to insure against such hazards. Finally, when risks are borne by future generations, as for global warming or waste hazard, it is technically impossible for earlier generations to participate in their sharing. Similar problems arise for a global risk like climate change, since financial instruments that would allow for an international diversification of country-specific risks do not currently exist. The provision of public insurance via implicit solidarity schemes is also imperfect due to the uncertainty of the benefit that will be paid to the victims.

When risks are not efficiently diversified away, each consumer will bear some individual risk. This implies that a risk premium should be added to the benefit of any action that reduces this individual risk. In order to quantify it, let us consider again the case of nvCJD. Victims of nvCJD have their life expectancy reduced by a factor $k = 50\%$. The value of a healthy life is $V = 100$ times the GDP per capita, whereas the value of life is only $(1-k)V$ for a victim. There is no compensation for the disease. Let $\bar{p} = 10^{-4}$ be the expected cumulated frequency of the disease over the next 20 years. In order to eliminate the risk of developing the disease, rational individuals will be ready to pay the expected loss (equal as before to 0.5% of GDP per capita), plus a risk premium. The size of this risk premium depends upon the individual's aversion to risk. Taking the Prospect Theory framework, we can write this risk premium as $\Pi = 0.5\bar{p}(1-\bar{p})kV$. Under the assumed parameter values, this implies that the risk premium amounts to an additional

0.25% of GDP per capita. An important conclusion of this analysis is that more preventative efforts should be undertaken in countries or sectors of the economy where risks are less efficiently shared. *Prevention is therefore a substitute for insurance*.[9]

We have assumed up to now that risk exposure was homogeneous in the population. Suppose instead that it is very unequally distributed among citizens, with some of them facing much higher potential damage in the event of an accident. A simple illustration of this is provided by the health risk induced by nuclear waste disposals. In that case, the distribution of potential damages is very unequal, with the so-called 'target population' living close to the site bearing a much larger damage. The acceptability of the risk to the public would be much reduced by the absence of a proper compensation scheme for these potential victims.[10] A similar point can be made for global warming, where some countries may actually benefit from it, whereas others, like small island nations, are highly vulnerable to rising sea levels.

6. FUTURE GENERATIONS

Genetically modified products, nuclear wastes and the loss of biodiversity all potentially impose costs on society that extend far into the future. The standard way to take future costs into account is to discount them at some socially efficient discount rate which may or may not be equal to the equilibrium risk free rate in the economy. The idea of using a positive discount rate is based mainly on the argument that future generations will be wealthier, so that the same marginal unit of damage is less damaging to their welfare than for the current generation, if the marginal utility of wealth is decreasing. But the consequence of discounting is that damages occurring long in the future, say in more than 100 years, are valued at trivial sums. This reflects the expectation that, at a growth rate of 2% per year, the level of consumption in 100 years will be over seven times what it is now, and in 200 years will be over 50 times what it is now. It implies that the current generations should not really care about these future risks. This is clearly contrary to the PP and is inconsistent with the view of the European Commission (2000) which recommends that 'the potential long-term effects must be taken into account in evaluating the proportionality of measures in the form of rapid action to limit or eliminate a risk whose effects ... will affect future generations'.

It may in fact be sensible to apply a smaller discount rate for longer time horizons than would be implied by exponential discounting. The argument is based on prudence, in the sense defined above: a prudent society should exercise precaution, and leave aside more resources for a more uncertain future. This insight is developed and discussed in Gollier (2001a). Over long time-scales, the levels of GDP per

---

[9] Anxiety, in addition to risk aversion and loss aversion, should be taken into account in the cost-benefit analysis of preventative actions. Caplin and Leahy (2001) have recently proposed a model to deal with anxiety.

[10] This is well understood by many environmentalists, who strongly oppose the idea that areas accepting nuclear waste disposals should be compensated for the potential damages.

capita we can expect to enjoy are indeed increasingly uncertain. If future economic growth trends and volatility are similar to those observed during the twentieth century, I find in that paper that theory suggests taking a 5% discount rate per year over the first 100 years, and then reducing it towards around 2% per year for longer time horizons.

There is an important interaction between natural risk (due for example to global warming), and uncertainty about the growth rate. Gollier and Pratt (1996) show that the presence of the background economic risk should induce us to raise the risk premium associated with the natural risk, even when the two sources of risk are independent. In other words, in cost-benefit analysis, we cannot treat the selection of the discount rate as independent of the measurement of the risk premium associated with the natural risk. Uncertainty about the growth rate of the economy will affect both of them in a direction which is compatible with the PP. However, for short time horizons, uncertainty about future growth is too small to affect risk premium to a significant degree.

7. ASYMMETRIC SCIENTIFIC INFORMATION

Up to now, we assumed that all members of society had the same imperfect information about natural risks. But the Precautionary Principle is also motivated by the difficult problem of assessing the reliability of information from various sources, some of them contradictory. What difference does it make when there is asymmetric information about risks?

7.1. Citizen's freedom, cost of information and traceability

Existing scientific information has the nature of a public good, but it is costly for consumers to acquire. If the cost is high enough, it may be socially efficient for a subgroup of independent citizens to pay this cost in order to formulate guidelines for other consumers. This may suggest the organization of popular juries to determine public policies. However, there is a risk that this subgroup be captured by lobbies. We examine this problem in Section 7.4. Another problem with a centralized decision rule is that it does not take into account the diversity of tastes in the population of risk-bearers. The heterogeneity of risk aversion, ambiguity aversion, prudence and values calls for offering some freedom of choice to citizens.

When the central authority has no more scientific knowledge than the population, it should refrain from imposing choices on the population. In the absence of any asymmetric information (or important externalities between different people), individuals should be allowed to make their own decisions about how to face risk. Nevertheless, there is a case for the State to guarantee the traceability of products in order to allow consumers to make choices on an informed basis.

When existing scientific information is not too costly for the public to acquire, public authorities should provide an efficient mechanism for the production and diffusion of scientific knowledge, and leave to consumers the responsibility for their own risk-bearing decisions. Insurers should be able to cover the associated risk on competitive insurance markets.

7.2. Product liability

One plausible implication of the Precautionary Principle is that firms should not be freed from their liability for deficiencies of a product because they were not aware of any scientific uncertainty at the time of its production. Imposing a no-fault liability system is an incentive for firms to acquire scientific knowledge before delivering a new product to the market. But this may not be enough to guarantee that firms implement an efficient level of precaution. One reason has already been alluded to in Section 4.2: firms will compete to pre-empt new markets, and will tend to innovate too early, without an adequate allowance for the risks.

In the light of the case of asbestos and other so-called ‘long-tailed’ risks (those having costs far into the future), another obvious problem is that firms may fail to take proper account of costs that occur in a far-distant future when most of them will no longer exist. A related argument is limited liability, which caps the amount that firms can pay at levels that may lie far below the cost of the damage to society. This can be partially alleviated by organizing ‘deep pockets’ for industrial firms, possibly by making banks jointly liable for damages caused by borrowers. Nevertheless, this may generate an adverse selection problem on the credit market, as shown for example by Boyer and Laffont (1997). Another possible solution is to organize a market for product liability insurance, but this in turn generates a moral hazard problem, in which firms fail to exert enough effort to prevent risk.

7.3. Imposing the Precautionary Principle on politicians

In the face of asymmetric scientific information between the State and its citizens, politicians sometimes face conflicting interests. This is particularly true when the resolution of uncertainty is slow. Suppose, as in Maskin and Tirole (2000), that the population does not know the best way to fight a global health hazard. Politicians in possession of private information about the nature of the risks may not have the right incentives to reveal this, particularly if it runs counter to the uninformed opinion of the public, if its technical nature makes it hard to explain, and if it is unlikely that the truth will be revealed prior to the politicians’ re-election. Indeed, politicians may have an interest in listening to experts who happen to share opinions that are more palatable to the public. In a range of plausible circumstances, the Precautionary Principle can be considered as a means of discouraging politicians from ignoring serious risks for this reason.

Capture of the regulatory process by industrial and agricultural lobbies is another reason why self-interested politicians may select inefficient policies. This suggests referenda may sometimes play an important role in bringing pressure to bear on the policy process, though their value is likely to be limited to less technical matters. However, referendums may induce politicians to shirk their responsibilities, for the reasons just outlined. And public opinion may be subject to various kinds of herd behaviour (Banerjee, 1992). Maskin and Tirole (2000) examine the relative advantages of direct democracy, representative democracy and judicial power in the face of these pressures.

7.4. How to use and motivate experts

Rational benevolent experts, sharing the same set of basic knowledge, would all aggregate new scientific information in the same way. In consequence, juries of such hypothetical benevolent experts would systematically be unanimous in favour of the most efficient policy recommendation. But in the real world, there is no reason to believe that experts are benevolent. They pursue private goals. They may derive pleasure or profit from promoting their peers' view of the world. They may also be reluctant to make the necessary effort to produce reliable scientific information. Or they may be sensitive to financial incentives from parties with a direct stake in the matter. Because most good experts are, almost by definition, in contact with at least some such parties, there is no doubt that a risk of collusion exists.

In many cases, experts are not rewarded by the State for providing information. We are therefore in the situation that is described by the theory of cheap talk, due to Crawford and Sobel (1982). Experts bias their message in favour of the industry that rewards them. Laypersons who realize this, will counter-bias their own production of information. Because representatives of industry usually want to prove the innocuousness of their own products, the message of the experts employed by them is biased in favour of lower risk (note that this may be true even without dishonesty on the part of the experts; given that differences of opinion exist, only the experts who believe the risks are low will be employed by the industry lobbies). The consequent counter-bias is for catastrophism. This can explain why lay activists have been promoting the Precautionary Principle as a reaction to a perceived loss of credibility on the part of experts.

If experts are biased towards industry, a possible strategy for the public authorities is to follow a bureaucratic procedure in which all information-gathering is undertaken by the authorities themselves. This may be socially efficient if collusion is easy. Alternatively, if collusion is more difficult, it may become efficient for the decision-maker to use experts, provided they can be given the right incentives to tell the truth. Whether to use outside experts depends upon the value of information they can provide, relative to the cost of motivating them to tell the truth (Tirole, 1992). In most cases involving the Precautionary Principle, it is likely that the first dominates the second. Note that requiring scientific reports to be accepted in scientific journals with standardized refereeing procedures, in the spirit of *Nature* or *Economic Policy*, would make collusion more difficult among experts.

More generally, scientific information presented by experts, before being treated as reliable by policy-makers, should be subject to peer review through presentation in seminars and conferences.

Rewarding experts properly may be crucial for the credibility of the system even when there is no risk of collusion. In particular, recent events show that experts may suffer serious damage, at least to their reputations, when they fail to recommend strong preventative action against risks that happen to turn out catastrophic. Unless this is compensated by a positive reward for those experts that advised against expensive preventative action in the face of risks that happen to turn out innocuous, incentives will be systematically distorted. Under these circumstances, experts will systematically prefer to recommend maximum prevention, and the credibility of the system will be lost. Organizing an efficient incentive system would therefore require keeping track of an individual expert's recommendations and evaluating them in the light of subsequent events.

8. CONCLUDING REMARKS

Assessing risks is a complex exercise, particularly when we are uncertain not just about what to do but about the basic scientific data of the decision problem. In this paper, we have explored the implications of this imperfect knowledge for an optimal risk-prevention strategy. The Precautionary Principle asserts that we should never use uncertainty as a reason for postponing our risk-prevention efforts. Reasonable as this sounds, we showed in Section 4.1 that it conflicts with an important consideration in circumstances where current and future preventative actions are close substitutes for each other. This consideration is the learn-then-act principle ('look before you leap'), and it suggests that in these circumstances the Precautionary Principle may be giving us precisely the wrong advice.

However, many other dimensions of the uncertainty problem need to be taken into account, and all of them argue in favour of a reasonable interpretation of the PP, if not always to a degree sufficient to outweigh the learn-then-act principle. First, postponing effort may make potential damages riskier, which induces a precautionary motive for risk-prevention, similar to the precautionary savings motive. Secondly, stronger preventative actions often yield more flexibility for the future, so that acting early has an option value. Thirdly, when better information comes from a process of learning-by-doing, the risk associated with early events is amplified by the information they yield about the future, and this plays a role analogous to that of an increase in risk aversion. Fourthly, because imperfect knowledge of the risk makes it difficult to insure, the social cost of risk should include a risk premium. Finally, uncertainty about the economic environment enjoyed by future generations should be taken into account. This calls for a smaller discount rate for damages occurring over longer time horizons, together with an increase in the risk premium associated with the natural risk. This raises the benefit of acting to prevent long-term risks.

We have shown that sound economic tools exist that can enable cost-benefit analysis to take account of scientific uncertainty, in ways that balance the Precautionary Principle against the benefits of waiting to learn before we act. Furthermore, it is important that they be used to do so, for instinct is an unreliable guide in such circumstances. Maxims such as the Precautionary Principle, if used outside the framework of careful cost-benefit analysis, may sometimes provide us with dangerously misleading advice. Do we really have the willingness to implement such careful methods? The European Commission (2000) appears quite ambivalent on this question. On pages 3 and 19 of its document on the Precautionary Principle, it states that 'finding the correct balance so that the proportionate, non-discriminatory, transparent and coherent actions can be taken requires a structured decision-making process with detailed scientific and other objective information', which sounds encouraging. However, the document also recurrently stresses the fact that 'judging what is an "acceptable" level of risk for society is an eminently political responsibility'.[11] But society's risk is eventually borne by citizens, who have their own preferences towards risk. This paper has provided some elements for determining the acceptability of society's risks from the point of view of individual citizens. Politicians may have a different viewpoint.

Discussion

Benny Moldovanu

Universität Mannheim

The Precautionary Principle (PP) is a relatively vague concept used in the public debate about risk management. The basic problem is that reliable scientific knowledge about potential risks may be available only in the future, while major policy decisions must be taken now. Good examples are the recent debates about the BSE epidemic and global warming. The principle calls for not postponing cost-effective measures to prevent degradation, even if we lack scientific certainty about issues involved. Of course, the whole point is that scientific knowledge is crucial for assessing cost-effectiveness, and therefore the PP is void of content without much more precise, operational steps that should guide us how to assess now the risks for the future.

The present paper constitutes an ambitious attempt to fill the concept with meaning, and to show that simple economic principles can be used in order to make PP operational. I particularly like the use of a formal model (see the Appendix) in order to structure the problem. This is a model of dynamic choice under uncertainty where relevant information appears over time (the model is a variation on one discussed by Epstein, 1980) Most of the basic points in the paper can be explained by varying the parameters in that model. Thus the paper offers a unified framework within which we can address apparently disparate phenomena. This is, in my view, what good science

[11] It is clear that the European Commission did not ask economists to help them to shape their position with respect to the PP. In European Commission (2000), there a lot of bizarre statements for an economist, like 'the decision-maker may, in some circumstances, be guided by *non-economic* considerations such as the protection of health'.

is about. I hope that professional economists, well trained in cost-benefit analysis and decision-making under risk, will play a greater role in assisting policy-making. But economists must also better understand the political economy of decision-making and the particular sensitivities of politicians and the public.

My comments below address several points where some additional thoughts may be helpful.

The main work-horse of the paper is the subjective expected utility (SEU) model of decision-making. The only discussed alternative is a theory developed by Gilboa and Schmeidler (1989) which is based on a maximin criterion. Of course, such a theory is very pessimistic and therefore very conservative in its prescriptions. There are many alternatives to SEU and it may be of interest to know how they perform in, say, Epstein's model if we deal with very large risks occurring with very small probabilities. Finally, ambiguity about probabilities is dismissed as not very relevant. While that may be fine from the point of view of a 'rational' decision-maker, the public at large probably displays different attitudes. Again the question is whether we have some empirical knowledge about such attitudes, particularly when people face large risks occurring with small probabilities. As an example, consider the recent German decision to close down all existing nuclear reactors over a period of about 30 years. Can we justify such a decision by applying standard SEU for agents with a 'typical' degree of risk aversion (see below)?

In several places the author uses very specific numbers for parameters of importance for decision making. For example, there is a claim about a consensus on degrees of relative risk aversion between 2 and 4. Where does this consensus comes from? The fact that some other researchers have used it for calibrations is not very convincing. In another example, discount rates of 5% for 100 years and 2% afterwards are proposed. Again, how should we get reliable estimates about the 'correct' values?

In the section about flexibility there is some confusion about what is, exactly, irreversible. The author offers the example of GM seeds, whose introduction today may lead to irreversible consequences. If this is indeed the case, then introducing an option value in the cost-benefit analysis implies that we should adopt such seeds only for high net present values and not just if the present net value is positive. This is seen as compatible with the precautionary principle. In the BSE example though, killing hundreds of thousands of animals is certainly irreversible. Hence letting those animals live, instead of killing them, seems more consistent with PP! This bidirectionality of PP (in this, and other parts of the paper) is not emphasized enough, and it may lead to a false impression about what PP really says.

The least developed part of the paper is the one dealing with expert procedures. I believe that, together with the decisions about funding of scientific research, this part of the decision-making process plays an important role in practice. There is by now a large literature on aggregation of information and manipulation by experts (Milgrom and Roberts in *RAND*, 1986 is an early contribution). How do the conclusions of that literature relate to the PP?

Tore Ellingsen

Stockholm School of Economics

Technological change is the chief source of material progress. However, new technologies also entail new risks, such as the greenhouse effect, mad cow disease, and bacteria that are resistant to antibiotics. Almost by definition, new risks are difficult to evaluate. Risk assessments vary considerably across experts, and even the best guess may fluctuate widely within short periods of time, as new information accumulates. In short, new technologies entail scientific uncertainty.

How should society deal with scientific uncertainty? Gollier's discussion of this problem comes in three steps, and I'll briefly discuss two of them. The first step concerns whether people care about scientific uncertainty (i.e., ambiguity) as such, or whether they are only worried about the 'best guess' in view of the available information. Experimental evidence reported by Fox and Tversky (1995) indicates that people are not very sensitive to ambiguity, except when the ambiguity is explicitly pointed out to them. In other words, second-order uncertainty is of second-order importance. Gollier therefore sticks to the conventional expected utility framework for the remainder of his paper. (It is perhaps typical that just as 'ambiguity aversion' is losing favour among psychologists and experimentalists, economic theorists have finally come up with a good way of modelling it; see Epstein and Zhang, 2001.)

While I agree that the expected utility framework is a very useful abstraction, I think more could be said about its limitations. There are several ways in which expected utility fails to capture people's attitudes to new risks. First, the concept is essentially atemporal, and therefore does not take into account the fears that people feel during the time interval before the uncertainty is resolved. (For references and careful discussion, see Caplin and Leahy, 2001.) The anticipation of future dangers is sometimes almost as painful as the worst potential outcome. Moreover, anxiety is often much worse among people who have more time to prepare for a danger. As a result, many people avoid obtaining relevant information, preferring reduction of fear to improved preparation for the future.

Importantly, people's worries appear to depend primarily on the *possibility* of a bad event, that is worries are not proportional to the probability of the bad event. One may therefore speculate that scientific uncertainty has a considerable impact on fears; once *some* theory ascribes a large enough probability of a catastrophe, the image will register and sensitive people start fearing the event.

Another limitation of expected utility theory is that it focuses exclusively on outcomes. However, people's attitudes to new technologies are often related to properties of the technology itself. For example, many express the view that one should not 'mess with nature'. It is as if the risks associated with, for example, advanced genetic manipulation, are less acceptable than 'natural' risks, such as earthquakes. (To me, it seems difficult to reconcile Sweden's recent decision to close several nuclear power plants with expected utility theory. Instead, the decision appears to be driven partly by fears that bear little relation to the probability of an accident and partly by aversion to 'unnatural' production processes.)

Indeed, it seems generally true that conventional cost-benefit analysis does not justify observed public regulation of risk. Some risks are reduced at virtually any cost, whereas others are completely neglected despite the availability of relatively cheap remedies; see for example Tengs and Graham (1996). This observation suggests either that our normative theories fail to capture important elements of what people really worry about, or that political processes to a large extent are driven by other considerations than aggregate costs and benefits.

Gollier discusses one set of reasons why actual policies differ from the optimal ones, namely differences in information between policy-makers, experts, and the general population. And he poses the question: how shall we design institutions that ensure well-informed policies, when politicians and experts have privileged access to information? These incentive problems are of course not confined to policies concerning the regulation of risks, but they are particularly severe in this field. The reason is that probabilities are rarely observable; only the outcomes are. Politicians can blame bad outcomes on bad luck, and may thus be tempted to turn a blind eye to risky behaviour by their key supporters. Likewise, if the general electorate is relatively ignorant about the true dangers, public money may be wasted on reducing risks that affect politically influential groups.

In my view, asymmetric information is nonetheless only a part of the problem of opportunistic politicians. Once we admit that politicians are people, with their own interests, even under symmetric information we may expect policies that benefit the current political majority. If in addition we admit the idea that politicians have the same psychological biases as the rest of us – the same vulnerability to excessive fears, and the same propensity to blind denial – it is quite obvious that public risk management will be imperfect.

Of course, these observations should not keep us from developing normative theory. New insights gradually accumulate and eventually affect public debate and political decisions. Gollier's careful treatment of scientific uncertainty that resolves over time is a case in point. The trade-offs between prudence and option value are non-trivial, but they can nonetheless be systematically studied. And systematic scientific analysis remains the best precautionary measure that we have against erroneous thinking.

Panel discussion

Bruno Jullien pointed out that the paper does not discuss the issue of dynamic inconsistency. However, this may be an important aspect of the problem. He also drew attention to the fact that different parts of the paper appeared to come to different policy conclusions, and suggested that it would remain difficult to know how to reconcile such different conclusions.

Kai Konrad suggested that in the absence of asymmetric information, decision-making is not very difficult. The paper considers only one type of asymmetric information, however, which is that between politicians and experts on the one hand,

and private individuals on the other. He claimed that in many cases it would be private individuals that were better informed, in particular about their preferences. Such an asymmetry would tend to imply a greater role for private rather than government decision-making.

Furthermore, Konrad pointed out that in a dynamic framework there may not be much of a choice between taking or not taking a risky decision. As time progresses, not taking the decision, or not innovating, also means taking a decision – the decision not to innovate, which is also risky. The paper suggests that there is a very strong status quo bias in the sense that if we do not innovate, we do not take a risky decision. But it may be just the opposite, that if we do not innovate we may take an extremely risky decision.

Lars Feld remarked that the paper reported very little evidence about people's willingness to pay for risk reduction. He cited the Exxon Valdez case in the US, where the damages in Alaska to the Prince Williamson bay were assessed by econometric estimates. He asked why such methods were not applied more often in practice.

Patrick Honohan felt confused about the title of the paper, which is 'Should we Beware of the Precautionary Principle?'. The message of the paper appeared to be that we should not, because the Precautionary Principle has internal contradictions, among other drawbacks. While he found this argument to be convincing from an economic point of view, he thought it over-optimistic to think that its contradictions would prevent the Precautionary Principle from being misused. For instance, courts might enforce the principle by applying the most cost-effective of various proposed measures, all of them costly and unnecessary, without considering that waiting to act might in some circumstances be better than any of the proposed measures. This could be potentially very undesirable.

APPENDIX. A GENERAL DECISION PROBLEM UNDER RISK AND SCIENTIFIC PROGRESS

All models that are examined in Section 4 can be seen as a special case of the following general model, which generalizes Epstein's (1980) formulation:

$$\max_{c_1 \in K_1} E_1 h_1(c_1, \tilde{x}_1) + E_s \max_{c_2 \in K_2(c_1)} E_2 h_2(c_1, c_2, \tilde{x}_2 \mid \tilde{s}) \tag{1}$$

The timing of the decision problem is as follows: First, the decision-maker (DM) selects $c_1 \in K_1$. The realization of $\tilde{x}_1$ is observed afterwards, which affects the first period welfare $h_1(c_1, x_1)$. At the beginning of the second period, the DM observes a signal $\tilde{s}$ that may be correlated with $\tilde{x}_1$ and $\tilde{x}_2$. This signal represents the (scientific) information produced during the first period about the risk $\tilde{x}_2$ borne during the second period. After observing this message, the DM selects c_2 under the condition that it belongs to some set $K_2(c_1)$. The realization of $\tilde{x}_2$, together with (c_1, c_2) affect the second period welfare h_2. Usually, we compare the optimal c_1 of this problem to the one that would be optimal when $\tilde{s}$ is uninformative, that is when $\tilde{s}$ is uncorrelated to $\tilde{x}_2$. Some authors examine the more general problem of an improvement of the information structure in the sense of Blackwell.

The accumulation model in Section 4.1 corresponds to: $K_1 = K_2(.) = R$, $h_1(c_1, x) = u_1(vc_1)$ and $h_2(c_1, c_2, x) = u(vc_2 - x(c_1 + c_2))$.

The option value examined in Section 4.2 can be seen as a special case of (1) where $K_1 = \{0, 1\}$, $K_2(0) = \{0, 1\}$, $K_2(1) = \{1\}$, $h_1(0, x) = h_2(0, 0, x) = 0$, $h_1(1, x) = -I + v$, $h_2(0, 1, x) = (1 + r)^{-1}(-I + v - x)$, and $h_2(1, 1, x) = (1 + r)^{-1}(v - x)$. Because K_2 depends upon the first period choice, there is some form of irreversibility.

Finally, the learning problem is such that $c_1 = (\alpha_1, s)$ is bidimensional. There is no constraint: $K_1 = R^2$ and $K_2(.) = R$. The first period welfare is $h_1(\alpha_1, s, x) = u(w_0 - s - x(L - h\alpha_1) - \alpha_1)$, where w_0 is a parameter representing initial wealth, s is savings, and α_1 is some investment to reduce the loss in case of a accident $(x = 1)$. The second period welfare is similar with $h_2(s, \alpha_2, x) = \beta u((1 + r)s - x(L - h\alpha_2) - \alpha_2)$. Finally, we assume that $\tilde{x}_1$ and $\tilde{s}$ are perfectly correlated (improvement in knowledge comes only from observing $\tilde{x}_1$). More specifically, we have $\tilde{x}_1 \sim (1, \bar{p};\ 0, 1 - \bar{p})$, and $\tilde{x}_2 \mid x_1$ is obtained by using Bayes rule.

REFERENCES

Arrow, K.J. and A.C. Fischer (1974). 'Environmental preservation, uncertainty and irreversibility', *Quarterly Journal of Economics*, 88, 312–319.

Arrow, K.J. and R.C. Lind (1970). 'Uncertainty and the evaluation of public investment decision', *American Economic Review*, 60, 364–378.

Banerjee, A.V. (1992). 'A simple model of herd behavior', *Quarterly Journal of Economics*, 112, 797–817.

Boyer, M. and J.-J. Laffont (1997). 'Environmental risks and bank liability', *European Economic Review*, 41, 1427–59.

Brennan, M.J. (1998). 'The role of learning in dynamic portfolio decisions', *European Finance Review*, 1, 295–306.

Brennan, M.J. and Y. Xia (1999). 'Assessing asset pricing anomalies', Discussion paper No. 19–99, The Anderson School at UCLA.

Caplin, A. and J. Leahy (2001). 'Psychological expected utility theory and anticipatory feelings', *Quarterly Journal of Economics*, 116, 55–79.

Crawford, V. and J. Sobel (1982). 'Strategic information transmission', *Econometrica*, 50, 1431–52.

Cropper, M.L., W.N. Evans, S.J. Berardi and P.R. Portney (1992). 'The determinants of pesticide regulation: A statistical analysis of EPA decision making', *Journal of Political Economy*, 100, 175–197.

Detemple, J.B. (1986). 'Asset pricing in an economy with incomplete information', *Journal of Finance*, 61, 383–392.

Dixit, A.K. and R.S. Pindyck (1994). *Investment under Uncertainty*, Princeton University Press, Princeton, NJ.

Dupuy, J.-P. (2001). 'Pour un catastrophisme éclairé', unpublished paper, Ecole Polytechnique, Paris.

Eeckhoudt, L. and C. Gollier (1998). 'Which shape for the cost curve of risk?', mimeo, University of Toulouse.

Ellsberg, D. (1961). 'Risk, ambiguity, and the Savage axioms', *Quarterly Journal of Economics*, 75, 643–69.

Epstein, L.S. (1980). 'Decision-making and the temporal resolution of uncertainty', *International Economic Review*, 21, 269–284.

Epstein, L.G. and J. Zhang (2001). 'Subjective probabilities on subjectively unambiguous events', *Econometrica*, 69, 265–306.

European Commission (2000). 'Communication from the Commission on the Precautionary Principle', February 2000, www.europa.eu.int.

Fox, C.R. and A. Tversky (1995). 'Ambiguity aversion and comparative ignorance', *Quarterly Journal of Economics*, 110, 585–603.

Gennotte, G. (1986). 'Optimal portfolio choice under incomplete information', *Journal of Finance*, 41, 733–749.

Gilboa, I. and D. Schmeidler (1989). 'Maximin expected utility with non-unique prior', *Journal of Mathematical Economics*, 18, 141, 153.
Godard, O. (1997). 'Social decision-making under conditions of scientific controversy, expertise and the Precautionary Principle', in C. Joerges, K.-H. Ladeur and E. Vos (eds.), *Integrating Scientific Expertise into Regulatory Decision Making*, Nomos Verl.-Ges., Baden-Baden.
—— (2000). 'The precautionary principle: Matching economic exiomatics and reasonable heuristics to tackle collective risks', 4th GREEN-CIRANO conference on environmental and resource economics, University of Montréal.
Gollier, C. (2000). 'Optimal prevention of unknown risks: A dynamic approach with learning', mimeo, University of Toulouse.
—— (2001a). 'Discounting an uncertain future', *Journal of Public Economics*, forthcoming.
—— (2001b). *The Economics of Risk and Time*, MIT Press, Cambridge, MA.
Gollier C., B. Jullien and N. Treich (2000). 'Scientific progress and irreversibility: An economic interpretation of the Precautionary Principle', *Journal of Public Economics*, 75, 229–53.
Gollier, C. and J.W. Pratt (1996). 'Risk vulnerability and the tempering effect of background risk', *Econometrica*, 64, 1109–1124.
Gollier, C. and N. Treich (2000). 'Decision-making under scientific progress: The economics of the Precautionary Principle', mimeo, University of Toulouse.
Henry, C. (1974). 'Investment decisions under uncertainty: the irreversibility effect', *American Economic Review*, 64, 1006–1012.
IPCC (1995). 'Second assessment report: climate change', www.ipec.ch.
Kimball, M.S. (1990). 'Precautionary savings in the small and in the large', *Econometrica*, 58, 53–73.
Landes, D.S. (1999). *The Wealth and Poverty of Nations*, Norton, New York.
Leland, H. (1968). 'Savings and uncertainty: The precautionary demand for savings', *Quarterly Journal of Economics*, 45, 621–36.
Lucas, R. (1987). *Models of Business Cycle*, Basil Blackwell, London.
Manne, A.S. and R. Richels (1992). *Buying Greenhouse Insurance: The Economic Cost of CO_2 Emissions Limits*, The MIT Press, Cambridge.
Maskin, E. and J. Tirole (2000). 'The politician and the judge: Accountability in government', mimeo, University of Toulouse.
McCardle, K.F. and R.L. Winkler (1992). 'Repeated gambles, learning, and risk aversion', *Management Science*, 38, 807–818.
Merton, R.C. (1973). 'An intertemporal capital asset pricing model', *Econometrica*, 41, 867–887.
Nordhaus, W.D. (1994). *Managing the Global Commons*, MIT Press, Cambridge, MA.
Savage, L.J. (1954). *The Foundations of Statistics*, Wiley. New York. Revised and enlarged edition: Dover, New York, 1972.
Tengs, T.O. and J.D. Graham (1996). 'The opportunity cost of haphazard social investments in life-saving', in R.W. Hahn (ed.) *Risks, Costs, and Lives Saved*, Oxford University Press, New York.
Tirole, J. (1992). 'Collusion and the theory of organizations', in J.-J. Laffont (ed.), *Advances in Economic Theory*, Cambridge University Press.
Tversky, A. and D. Kahneman (1992). 'Advances in prospect theory: Cumulative representation of uncertainty', *Journal of Risk and Uncertainty*, 5, 297–323.
Ulph, A. and D. Ulph (1997). 'Global warming, irreversibility and learning', *The Economic Journal*, 107.
von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behavior*, Princeton University Press, Princeton, 2nd edn, 1947.
Viscusi, K.W. and H. Chesson (1999). 'Hopes and fears: the conflicting effects of risk ambiguity', *Decision Theory*, 47, 153–78.
Weeds, H. (1999). 'Strategic delay in a real options model of R&D competition', mimeo, Cambridge, UK.