The economics of global environmental risk

Chichilnisky, Graciela

1998

7. The economics of global environmental risks

Graciela Chichilnisky*

1 INTRODUCTION

Uncertainty is part and parcel of the human condition. The need to manage climate risk has shaped human institutions for many centuries, giving rise to insurance in agricultural societies and to patterns of land holdings across medieval Europe.[1] Today's concern about global climate change breaks new ground. It challenges conventional wisdom in two ways. One is the worldwide scope of the changes considered. We are concerned with potential catastrophes, such as changes in the sea level and global shifts in the availability of water and fertile soil. No society can ignore such risks. The second challenge is that these changes appear to be driven by human activity, which has now reached levels at which it can affect the earth's fundamental processes, such as its atmosphere and its climate. These two new elements, the global and endogenous nature of these risks, have extended human uncertainty both qualitatively and quantitatively.

Traditional formulations of uncertainty no longer provide an adequate basis for analysis. As shown below, catastrophic risks are not treated adequately by conventional risk analysis, nor are the type of risks that humans induce through their economic activity. Climate risks can be the byproduct of human activity. This adds a new dimension to the uncertainty about earth processes. To the extent that we cannot predict human actions, we cannot predict their potential environmental impact. I call this phenomenon *endogenous uncertainty*. It is a sign of the times, and characterizes our age. For the first time in history, humans are playing a predominant role in determining the world's ecosystems, on land and water, and can influence the planet's atmosphere.[2]

Our impact on the earth's processes is unknown. This impact encompasses much more than the traditional concerns about where humans

* Research support from UNDP, Fondazione Mattei, UNESCO, Sloan Foundation and the US National Science Foundation is gratefully acknowledged.

decide to settle, whether they settle in vulnerable areas such as coasts or near volcanoes. The issue is how we are affecting the overall biodiversity of the planet, how we may be changing the planet's atmosphere and its soil fertility across the different continents and bodies of water. These changes are often caused by economic consumption from far-away human settlements.[3] The economic use of the earth's resources today could induce the main risks that humankind faces in the future.

Environmental stress may be a symptom rather than a disease. Historical circumstances have led to an uneven spread of the industrial revolution across the world. The world contains regions with very different institutions, such as different property rights and markets. Developing countries have common property on resources such as forests and minerals; in industrial countries these tend to be privately owned. This difference plays a major role in the extraction of resources and the production of resource-intensive products – fossil fuels, forest-based products. Differences in property rights lead developing countries, many of which are still agricultural societies, to overextract natural resources and to export them to industrial nations at prices that are below real costs (Chichilnisky 1994a, 1996b). Low prices lead to overconsumption worldwide. The global environmental problem we face today reflects distorted prices and an uneven use of the earth's resources. The matter is further complicated by the emphasis on resource-intensive growth advocated by the Bretton Woods institutions and by many traditional economists in the Western hemisphere since the end of the Second World War (Chichilnisky 1996b; Institute for Policy Studies 1996). This is believed to be at the root of the carbon emissions problem, the destruction of biodiversity, and the attendant climate risks (Chichilnisky 1994a, 1996b). Before we can solve the environmental problem we need to develop economic institutions that value properly the earth's resources and the welfare of the humans that are their stewards.

This chapter focusses on global environmental risks such as climate change, an issue that must be confronted as we move into the future. It proposes sound principles of risk management that make sense in today's society generally, going beyond their role of averting and hedging climate risks. This chapter is about these and related questions. In attempting to answer them, it deals with different aspects of the theory of risk-bearing. I explain current responses to global change, focusing on the new challenges: human-induced or endogenous risks, including potentially catastrophic risks, which are not adequately treated by traditional economic analysis.

There are five key aspects of risk which recur in the analysis. As already mentioned, two of them are essentially new: these risks are driven by human activity and could have catastrophic consequences. The global risks that we face are influenced by our actions and are thus endogenous.[4] Our actions have become an important determinant of the risks we face. The chapter analyses markets with *endogenous uncertainty*, as well as growing economies where today's economic activity determines tomorrow's land productivity. In both cases practical policies are recommended. The other new aspect of climate risks is that they could have catastrophic consequences: they involve small probabilities of events with major, negative consequences. Classical decision theory has neglected decisionmaking involving *catastrophic risks*. Based on Chichilnisky (1996f) I present a rigorous treatment of decisionmaking under catastrophic risks and propose practical solutions.

There are three, more conventional, aspects of climate risks that have been neglected in the economic literature and are analysed here. The first concerns the difficulty in assessing risks. Most climate-related risks are difficult to quantify. Indeed, in a statistical sense the probabilities describing them appear to be unknowable. We may never be able to determine experimentally the probability of global climate change in the relative frequency sense. Such events are inherently unique. It is possible to evaluate the frequency of occurrence of a health risk from morbidity or mortality data, as the outcomes of repeated experiments are available. However we cannot evaluate the risks from CO_2 emissions in this way.[5] The chapter proposes new financial instruments for hedging optimally against this type of uncertainty, which I call 'scientific uncertainty'. These instruments are called 'catastrophe bundles', and were introduced in Chichilnisky (1995, 1996a, 1997a and 1998). The second conventional aspect is the correlation of risks. Events such as climate change could affect large numbers of people in the same way. A rise in sea level could affect low-level coastal communities in most countries. Insurance in the traditional sense of risk-pooling works best for large numbers of small statistically independent or 'individual' risks. We have to ask what types of markets work best with partly individual and partly collective risks. This chapter proposes a way to solve this problem.

Irreversibility is the final issue. In this area, many major economic decisions and their consequences are likely to be irreversible, or if reversible they involve very large costs and lengthy time scales. Climate change, the melting of ice caps, desertification and species extinction are all processes not easily reversible within relevant time scales. This leads to option values. The possibility that future generations will have different preferences from us can also lead to option values, but of a different nature (Beltratti, Chichilnisky and Heal 1998). This chapter explores this question.

In summary, we are dealing with risks that have two major new characteristics: they are *endogenous* and *potentially catastrophic*. In addition, climate risks have three more conventional features: they are *poorly understood*, *correlated* and *irreversible*. In all cases, this chapter proposes ways to advance our understanding of the problems. In policy terms, the nature and extent of uncertainty about global climate change imply that society's position will be dominated by questions such as:

- What cost is it worth incurring to reduce the poorly understood risk of climate change, or to improve our understanding of that risk? and
- How may existing social institutions, such as insurance contracts and securities markets, be used to provide the most efficient allocation of the risks associated with global climate change?

This chapter proposes ways to evaluate decisions under endogenous and potentially catastrophic risks, and incorporates often neglected features of correlated, poorly understood and irreversible risks. The analysis proposed here opens new ways of thinking and at the same time poses new challenges. At the end I indicate new areas of research.

2 SCIENTIFIC UNCERTAINTY

The uncertainty about climate has several sources. There is uncertainty about basic scientific relationships, such as the link between gaseous emissions and global mean temperature. There is also uncertainty about the connection between global mean temperature and climate. Clearly it is climate, a variable encompassing wind patterns, humidity and rain patterns, and not just temperature, that matters from an economic perspective. The floods of 1993 in the US and Bangladesh have reminded us of the profound vulnerability of human settlement to climate, as has hurricane Andrew, which led to about $20 billion of losses in the southern US at about the same time. Climatologists link these to El Nino, the ocean current off the coast of Chile, confirming the global linkages within the earth's climate system.

Future emissions of greenhouse gases and future climate are also highly uncertain. Furthermore, as already pointed out, these emissions can be driven by economic activity and by policy measures: hence the risks faced are endogenous.

Societies can respond to the risks associated with such uncertainty. One way is *mitigation*. The other is *insurance*. We can think of them as broadly equivalent to prevention and cure, respectively, in the medical field.

Mitigation means taking measures to reduce the possible damage. One way of doing this is to take steps that minimize the damage if the harmful event occurs. Building levees, canals and flood drainage systems to reduce the impact of flood waters is an example. An alternative approach to mitigation is to reduce the incidence of harmful events. Of course, if steps are taken to reduce the risk of climate change, then the risks become endogenous, determined by our policy measures. This contrasts with most models of resource allocation under uncertainty, in which probabilities are about acts of nature and are therefore exogenous.[6] In a traditional market approach, there is no scope for mitigation in the second sense of improving odds. The probabilities of states in an Arrow–Debreu market may be subjective and a trader's subjective probabilities may be altered by learning; however, the frequency of incidence of harmful events cannot be altered by traders. The same is true in the classical models of insurance, where the incidence of harmful events is again taken to be exogenous. Mitigation acquires a new meaning when risks are endogenous.

Insurance by contrast does nothing to reduce the chances of damage due to climate change. It only arranges for those who are adversely affected to receive compensation after the event, as in the case of federal disaster relief for flood victims in the US. Insurance is a major economic activity, involving both the insurance industry and large parts of the securities industry, about $1.2 trillion worth of economic activity yearly. Can the existing and very extensive private sector organizations provide those at risk from climate change with adequate insurance cover? If not, why not? What changes in market institutions might be appropriate in this case?

The following provides a brief survey of traditional responses to risk.

3 RESPONSES TO RISK IN TRADITIONAL MARKETS

Economists have two standard models of risk allocation in a market economy. The more general is that of Arrow and Debreu, in which agents trade 'contingent commodities'. The alternative is the model of insurance via risk-pooling in large populations. Neither case addresses the issue of mitigation via a reduction in the incidence of harmful events.

In the Arrow–Debreu market there is a set of exogenous 'states of nature' whose values are random and represent the sources of uncertainty. Classically one thinks of events such as earthquakes and meteor strikes. Agents in the economy are allowed to trade commodities contingent on the values of these exogenous variables. These are called 'state-contingent commodities'. With a complete set of markets for state-contingent commodities, the first theorem of welfare economics holds for economies under uncertainty: an *ex-ante* Pareto efficient allocation of resources can be attained by a competitive economy with uncertainty about exogenous variables.

Arrow (1953) showed that efficiency can in fact be attained by using a mixture of securities markets and markets for noncontingent commodities, so that a complete set of contingent commodity markets is not required. This observation provides a natural and important role for securities markets in the allocation of risk-bearing. The securities used are contracts that pay one unit if and only if a particular state occurs. While the contingent contract approach is in principle all-inclusive and covers most conceivable cases of uncertainty, in practical terms there are cases where it can be impossible to implement. It can be very demanding in terms of the number of markets required. For example, if agents face individual risks (that is, risks whose incidence varies from individual to individual), then in a population of 100 similar agents each of whom faces two possible states, the number of markets required would be 2^{100} (Chichilnisky and Heal 1992a). The number of markets required is so large as to make the contingent contract approach unrealistic.

The use of insurance markets for pooling risks is a less general but more practical alternative. This requires that populations be large and that the risks be small, similar and statistically independent. The law of large numbers then operates and the frequency of occurrence of an insured event in a large sample of agents approximates its frequency in the population as a whole. There is thus a role for insurance companies to act as intermediaries and pool large numbers of similar but statistically independent risks. In so doing they are able via aggregation and the use of the law of large numbers to neutralize the risks faced by many similar agents. The main references on this are Arrow and Lind (1970) and Malinvaud (1972, 1973); recently Cass et al. (1996) updated this analysis to incorporate individual risks with mutual insurance.

The insurance approach is at a disadvantage when risks are correlated. When large numbers of individuals are likely to be affected at once, risk-pooling will not work. However, it does have the advantage relative to the contingent market approach of economizing dramatically on the number of markets needed. In the above example, only two mutual insurance contracts and 2809 securities would be needed instead of 2^{100} contingent contracts (see Chichilnisky and Heal 1992a, and Cass et al. 1996).

When risks are allocated by trading state-contingent commodities securities, or by risk-pooling and insurance, *it is very important that agents know, or believe that they know, the relative frequencies of the states of nature, at least approximately*. This is obvious when trading insurance contracts. The actuarial calculations needed to set insurance premia can only be performed if the parties believe that the relative frequencies of the insured events are approximately known.

In the Arrow–Debreu approach, it suffices to think of agents maximizing expected utility to appreciate the need for them to know, or at least behave as if they know, the relative frequencies of exogenous states. These frequencies are the weights placed on their utilities from state-dependent consumption. The point is simple: if agents cannot assign relative frequencies then their preferences are not well defined and they cannot act to maximize expected utility.

In the context of climate change this may be too demanding. Agents do not know the frequencies of different states, and recognize that they do not know them. They recognize that there are several different opinions about what these are, but feel unable to choose definitively between these alternatives. If they were expected utility maximizers, they would be uncertain about their own preferences. In such a case, it is natural to think of the frequency distribution over climate changes as a state of the world: a risk, in the Savage sense. We do not know what value the frequency will assume, and whatever value this is, it affects economic activity. As shown below, ignorance then assumes the role of a collective risk, and can be treated by the use of state-contingent markets. One sometimes thinks of uncertainty about probabilities being resolved by learning. This is an avenue which is not open when scientific knowledge is incomplete and experiments are not possible, for example, in the case of global warming. In this case an alternative approach is the opening of new markets (see Chichilnisky and Heal 1992a).

In sum: the Arrow–Debreu approach to risk allocation via state-contingent markets is universally applicable. However, it is cumbersome and unrealistically complex when risks have individual components. Insurance markets are more manageable, but leave uncovered collective risks such as the risk induced by ignorance of the true frequency distribution of harmful events. It would be natural to allow agents to trade securities contingent on such collective risks, and cover the individual components of risks by mutual insurance contracts. This approach is developed below. Although new to the economics literature, it is argued below that some of the oldest risk-bearing institutions recorded, agricultural cooperatives, have a similar structure.

4 SCIENTIFIC UNCERTAINTY AS COLLECTIVE RISK

Consider an economy in which agents face risks whose relative frequencies they cannot evaluate. Such risks could derive from the impact of global climate change on income levels via floods, storms or droughts, or from the effects on health of ozone depletion, acid rain or air pollution. There are widely differing opinions about their frequency, on which there is inadequate information. What market structure would ensure efficient risk allocation in this situation?

Chichilnisky and Heal (1992a), formalize the situation in a simple general equilibrium model. Each agent faces the risk of being in one of several states (for example, healthy or sick, productive or unproductive). No one knows what is the true frequency distribution of affected agents. A probability is assigned to each possible frequency. A typical probability distribution of this type might state, for example, that there is a 10 per cent chance that 90 per cent of the population will be harmed by global warming, a 25 per cent chance that 50 per cent of the population will be harmed, and so on. The probability distribution over alternative frequency distributions may be different from individual to individual.

In this framework, there are two levels of uncertainty. The first level is collective: what is the distribution of agents who are harmed in the economy? Will 90 per cent be harmed, or only 30 per cent? This is a question about the true incidence of the phenomenon over the population as a whole. The second level of uncertainty is individual: it is whether a given agent will be harmed or not by climate change. It revolves about questions such as: given that 90 per cent of the population will be harmed, will a particular agent be harmed or not?

In our example of the impact of the depletion of the ozone layer on cancer or the impact of climate change on agricultural productivity, the two levels of uncertainty are: first, uncertainty about the true relationship between ozone depletion and the incidence of individual disease in the population as a whole, or about the true relationship between climate change and agricultural productivity; and second, uncertainty about whether any given person or community will be affected.

Our ignorance of scientific processes (for example, the relation between ozone depletion and skin cancer or between CO_2 emissions and climate change) causes the collective risk, by which we mean the uncertainty about the relative frequency of harmed agents in the population. Uncertainty about this frequency is central to the problem. Even if this was resolved we would still not know who will be damaged and who will not, but we would at least know the frequency of the risk. Once frequencies are known, actuarial calculations can be conducted and the problem is insurable.

I propose an institutional structure which uses two types of financial instruments tailored to these two aspects of the problem. These lead to efficient allocation in the face of such risks. I follow a framework established by Cass et al. (1996) and Chichilnisky and Heal (1992a).

One instrument is a *mutual insurance contract* to deal with the risks faced by agents or communities contingent on each possible distribution of harmful effects worldwide. A mutual insurance contract is an agreement between parties subject to similar risks by which those who are harmed will be compensated by the others. Examples are agricultural cooperatives of the type recorded in Europe at least since the fifteenth century and the nineteenth-century UK workers' associations and friendly societies. These involved agreements between a group of workers that if one were sick and unable to work, he or she would be compensated by the others. In the present context, one could think of groups of communities subject to the possible impact of climate change, with those unharmed compensating the others. Making the terms of such a mutual insurance contract contingent on the distribution of harmful effects worldwide means that there is a different compensation agreement between the parties for each possible aggregate distribution of harmful effects. To know what compensation is due in any particular case, the parties have first to assess the distribution of harmful effects globally. On the basis of this they decide which mutual insurance contract to apply.

Having dealt with individual risks by mutual insurance, we still face collective risks. We need *statistical securities* to deal with these collective risks induced by uncertainty about the overall distribution of adverse effects. We propose to use a framework similar to Arrow securities: these are defined as securities that pay one dollar if and only if a particular state of the world occurs. In our case we use statistical securities which pay one dollar if and only if there is a particular frequency of affected parties in the population. As already noted, the incidence of impacts on the population as a whole is being treated as a 'state of the world' in the Arrow–Debreu sense. We treat each possible distribution of adverse affects as a distinct collective state (called a statistical state), and use securities markets to enable parties to transfer wealth between these states. One Arrow security is needed for each possible distribution of adverse effects worldwide, because to attain Pareto efficiency each separate state must be covered by a security.

The following example helps to make this framework concrete. It is illustrated in Figure 7.1. Consider a world of two countries 1 and 2, in which the climate may be in one of two states α or β. There are two possible probability distributions over these two climate states. These distributions are called A and B, with distribution A giving a probability of 0.1 to climate state α and a probability of 0.9 to climate state β. Distribution B gives the reverse probabilities, that is, it gives probability 0.9 to climate state α and probability 0.1 to climate state β. The endowments of the two countries depend on the climate state, and are as follows: $\omega_1(\alpha)$ is country l's endowment vector if the climate is in state α, and $\omega_2(\alpha)$ is the corresponding endowment for country 2. Similarly, endowments in climate state β are given by $\omega_1(\beta)$ and $\omega_2(\beta)$, respectively. Endowments satisfy $\omega_1(\alpha) > \omega_2(\alpha)$ and $\omega_1(\beta) < \omega_2(\beta)$, so that country 1 is relatively better off in state α and country 2 in state β.

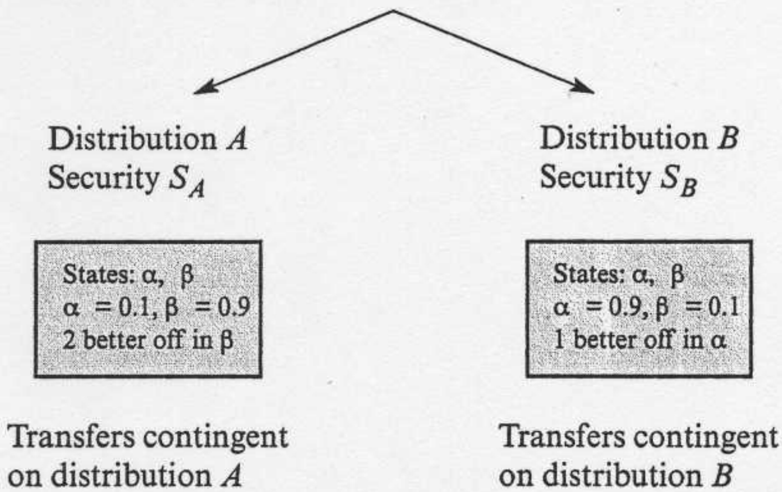

Figure 7.1 Statistical securities pay dependent on the distribution of states; insurance contracts make transfers given a distribution of states

To reach an efficient allocation of risks we need two statistical securities. One, S_A, pays \$1 if and only if the probability distribution over states of the climate is A. The other, S_B, pays \$1 if and only if the probability distribution over states of the climate is B. In practice of course probability distributions are not observable, and we cannot condition contracts on unobservable events. So conditioning on probability distributions means conditioning on frequency distributions consistent with that probability distribution in a sampling sense.

Countries can spread the risk arising from not knowing which is the true distribution over states of the climate by trading these two securities. In addition they make mutual insurance contracts conditional on states of the climate. Such a contract could take the following form. If the distribution over climate states is A (distribution A gives probability 0.1 to climate state α and probability 0.9 to climate state β); then country 1 makes a transfer $\Delta^{\alpha}_{1,2}$ to country 2 if the state of the climate is α, and country 2 makes a transfer $\Delta^{\beta}_{1,2}$ to country 1 if the climate state is β. These transfers satisfy $0.1\,\Delta^{\alpha}_{1,2} + 0.9\,\Delta^{\beta}_{1,2} = 0$ so that the expected transfer is zero and the mutual insurance contract is actuarially fair. There would be a similar contract to cover the case when the distribution over climate states is $\dot{B}$.

To summarize the argument: *our ignorance of the frequency of the impacts of climate change constitutes a collective risk. This collective risk can be allocated through markets for statistical securities, which pay off contingent on that frequency. For the individual risks that remain, it is more practical to use contingent insurance contracts: this is done by having a different individual insurance contract for each possible frequency of impacts.*

There are two features of the results which are of general interest. One is the development of a framework for achieving efficient allocations in the face of uncertain risks. Given rapid changes in technology with potentially far-reaching environmental impacts and health effects, the problem of providing insurance against such risks is particularly important. It is a matter of active concern in the insurance industry. The second interesting feature is the way a combination of securities markets and insurance markets can be used to provide a relatively simple institutional structure for dealing with unknowable risks. These are financial instruments called 'catastrophe bundles', introduced in Chichilnisky (1995, 1996e, 1996d) and Chichilnisky and Heal (1998). Current trends in the securitization of certain risks are consistent with this analysis.

4.1 An Institutional Framework for Hedging Scientific Uncertainty

Our analysis suggests that although the risks associated with global climate change are very difficult to evaluate, there is nevertheless a market framework within which insurance against scientific uncertainty can be provided. It involves, first, identifying the set of possible descriptions of the collective risks. Natural descriptions of risk are frequencies of occurrence of climate-related events such as floods, tropical storms or certain temperature patterns.

Second, this framework involves introducing statistical securities whose payoffs depend on which description of the risk is correct. This amounts to allowing agents to bet on which model of the risk is correct. Betting on which of several alternative descriptions of the way the world works is correct is in effect what one does when choosing one research strategy over another. Corporations, individuals and governments all do this regularly but not efficiently. For example, a market for the securities of high-technology firms pursuing different research strategies towards the same goal is a financial market in which these bets are made.

Finally, our approach involves establishing compensation agreements between harmed and unharmed regions that depend on which description of the risk turns out to be correct. Mutual insurance contracts or mutual compensation agreements are already part of our institutional framework. They date back to the nineteenth century and beyond, and were the foundations of many current insurance companies and trade unions. Consider, for example, agricultural cooperatives, probably the oldest risk-allocation institutions in the world. One of the largest banks in Italy, the Monte dei Paschi dei Siena, was founded to play this role in 1473. They have insured against weather risks since then,[7] and have provided mutual insurance contracts for their members, in that they have arranged transfers from the less to the more fortunate in any given season, the size of the transfer depending on the overall level of prosperity. They have also provided an elementary form of insurance against the overall frequency of poor crop yields in their community by building up reserves to carry over from good to bad years. So they have actually fulfilled both of the insurance functions outlined above – making transfers between agents contingent on the overall incidence of negative events, and allowing a mechanism for transferring wealth between states in the sense of high or low overall incidences of negative events in the population.

4.2 Trading Risks

An interesting aspect of the markets just described is that they can provide a natural mechanism for reconciling differences in assessments, and for testing the conviction behind publicly stated positions.

For many years the US expressed disbelief about the likelihood of climate change. The European Union expressed the opposite belief. Then through a market for securities whose payoffs depend on which description of climate change is correct, the US would naturally sell insurance to the EU. The US would wish to be a seller of securities which pay if climate change is serious, because of its belief that this event will not occur, and a buyer of securities that pay if it is not, because of its belief that this will be the outcome. The EU would be on the opposite sides of these markets.

International markets for the risks of climate change would also provide an objective test of the seriousness with which countries adhere to their publicly professed positions on the risk of climate change. It is possible that a country might publicly profess to a lack of concern about the risks of climate change, in spite of actually being concerned about these risks, in order to 'free ride' on CO_2 abatement policies introduced by others. These issues are discussed in Chichilnisky and Heal (1992b) and the references cited there. The existence of markets for the risks of climate change would place such a country in a dilemma. The country's true beliefs would incline it to sell securities paying off in the event of climate change not being serious, and buy those paying off if it is serious. Consistency with its public positions would require that it be on exactly the opposite sides of these markets. There would therefore be a cash cost to convincing and consistent misrepresentation of true beliefs. These cash costs could offset some of the incentive to free ride on other countries' efforts to reduce greenhouse emissions (Chichilnisky and Heal 1992b).

Note that trading risks is different from the trading of emission permits. The recognition of uncertainty suggests the need to trade state-contingent emission permits, where the state is defined in terms of the frequency of climate-change-related events. Such contingent emission permits could play the role of securities whose payoffs depend on scientific uncertainty.

In the context of emission permits, it is worth noting that climate is a public good. However, it does not fit fully the conventional paradigm because emission abatement, which is the production of the public good 'unchanged climate', is conducted independently in the various countries of the world. It is not produced in a central production facility, as assumed in the usual treatments of public goods. A consequence is that economic efficiency will only imply equalization of the marginal costs of emission abatement across countries if lump-sum transfers between countries are made to equalize the marginal utility of income in all countries. Equalization of the marginal costs of emission abatement across countries is often taken as justification of the superiority of tradable permits as a method for controlling emissions. This point is developed in Chichilnisky (1994b) and Chichilnisky and Heal (1994). More generally, a key issue is that an efficient allocation of a public good such as unchanged climate is achieved through a Lindahl equilibrium and not a competitive equilibrium. In general, competitive markets for tradable emission permits may not decentralize Pareto-efficient allocations of abatement (Chichilnisky, Heal and Starrett, 1993).

5 OPTIMAL ALLOCATION WITH ENDOGENOUS RISKS

What is it worth spending to reduce the probability of harmful climate change? Only if we can answer this question can we judge properly proposals for carbon taxes, alternative energy strategies, and CO_2-reduction protocols. Careful judgement is crucial, as all of these involve very considerable costs, as indicated by Cline (1992) and others. Here I summarize two approaches to this problem, a market approach based on Chichilnisky (1996e) and a growing economy approach based on Heal (1984, 1990).

5.1 Markets with Endogenous Uncertainty: Theory and Policy

The following extends Arrow–Debreu markets to encompass risks induced by the functioning of the economy itself. Examples are the risks of atmospheric and climate change induced by CFC and CO_2 emissions. Economic actions – the consumption and production of goods and services – induce uncertainty, because they are connected with carbon emissions that could alter the atmosphere and potentially change the climate.

Traders maximize expected utility. Their expected utility, however, changes with the aggregate level of output of the economy, because expected utility depends on the probability of the different events, and these change with the activity of the economy.

In this section, endogenous uncertainty is classified into two types: *scientific uncertainty*, which is uncertainty about the impact of production on climate states and their probabilities, and *strict endogenous uncertainty* which is about equilibrium levels of outputs. Below, I formalize a competitive equilibrium with endogenous uncertainty, show that the markets with endogenous uncertainty are typically incomplete, and that the equilibrium allocations are only efficient in a restricted sense. Scientific uncertainty can be fully hedged by financial innovation, for example, CAT Futures, which are new financial instruments recently introduced on the Chicago Board of Trade. However, uncertainty induced by the unknown level of output at an equilibrium, which is strict endogenous uncertainty, cannot be hedged fully. It leads to incomplete markets where the equilibria are not Pareto efficient.

Definitions

A market economy E has $H \geq 2$ traders, and $J \geq 1$ firms which produce $N \geq 1$ commodities over T periods of time. There are S states of exogenous uncertainty. To simplify the exposition, and to isolate the essential features of endogenous uncertainty, the formulation of the market E is identical to the classic Arrow–Debreu formulation in every possible way except in the treatment of uncertainty. All traders in E are competitive, and there is symmetric information.[8]

There is a complete set of assets to hedge exogenous uncertainty. Each asset pays in terms of a numeraire good n, and the span of the economy's asset matrix is S. Each trader h has an initial endowment of goods and assets, $\Omega_h \in R^M$ where $M = N \times S$, an initial endowment $\theta_h = (\theta_h^1 \ldots \theta_h^J)$ of shares in J firms. The economy has a complete set of markets for exogenous uncertainty and is equivalent to a standard Arrow–Debreu economy with commodity space R^M and no uncertainty; to simplify matters and without loss of generality I consider such an economy from now on.[9] The economy's technology is described within every state $s = 1 \dots S$ by a given production possibility set $Y^s \subset R$, and

$$Y = \Pi_{i=1}^{S} Y^s \subset R^M. \tag{7.1}$$

Each price vector p is in $\Delta_M = [(p_1 \dots p_M) \in R^M : p_i \geq 0 \text{ and } \Sigma p_i = 1]$.

Endogenous uncertainty

The economy E discussed until now is identical to an Arrow–Debreu market with production and with exogenous uncertainty about acts of nature. This subsection introduces a new aspect: endogenous uncertainty, which goes beyond the Arrow–Debreu structure and yields different properties.

To motivate the treatment of endogenous uncertainty in E, consider a version of the environmental problem discussed above: each vector of aggregate output of goods in the world economy induces a level of emissions of CO_2 or of CFCs. Corresponding to each level of emissions new states of nature may develop, for example a state where the ozone layer is 50 per cent damaged, or where there is a disruption of the planet's climate pattern known as global climate change. To formalize this, the endogenous uncertainty in the economy E is described as follows: each vector of aggregate production induces a set of states of finite cardinality,[10] $\{1 \dots D\}$, which includes all the states of exogenous uncertainty S and possibly more,[11] describing the risks which traders face.

In markets with exogenous uncertainty there exists either a *fixed* probability distribution over the set S representing the relative frequencies of the events in S, or alternatively different subjective distributions for different traders, each of which is also fixed. Here, instead, the probabilities over the states in D are *variable*: they vary in principle over all possible probabilities over the set D, which is an infinite domain, and do so according to the aggregate vector of output in the economy.

Definition 1 For each aggregate production vector $y \in R^M$ there exists a finite set of states $e(y)$ of cardinality not exceeding $D > S$, $e(y) \subset \{1 \dots D\}$, and a probability density over these states, $\pi(y) = \{\pi_i\}_{i \in e(y)}$, such that $\forall i\ \pi_i > 0$, and $\Sigma_{i=1}^{e(y)} \pi_i(y) = 1$. The set of states $e(y)$ and the density function $\pi(y)$ describe the *endogenous uncertainty* of the economy E at the aggregate production vector $y \in R^M$.

- *Assumption 1* There exists a C^2 function Ψ assigning to each vector $y \in Y$ of aggregate output in the economy a vector $\Psi(y)$ in the unit simplex Δ_D, the positive components of which represent possible states of uncertainty and their respective probabilities:

$$\begin{gathered}\Psi = (\Psi^1, \ldots, \Psi^D) : R^M \rightarrow \Delta_D \\ \Delta_D = [(\pi_1 \ldots \pi_D) : \pi_i \geq 0 \text{ and } \Sigma^D_{i=1} \pi_i = 1] \\ \text{and if } ||y'|| > ||y|| \text{ then} \\ \min_i(\pi_i) > \min_i(\pi'_i), \text{ where} \\ \pi(y) = \{\pi_i\} \text{ and} \\ \pi(y') = \{\pi'_i\}.\end{gathered} \tag{7.2}$$

Market equilibrium with endogenous uncertainty

How do rational traders behave when facing endogenous uncertainty? This depends on the structure of information of the economy. Assume the simplest possible structure in order to isolate the essential features of the problem.

The structure of information What do traders know about endogenous uncertainty? In our model they know that it exists and no more. We are concerned with market equilibria, not with the process by which the economy arrives at one. The discovery process leading to an equilibrium with endogenous uncertainty parallels the treatment of price discovery in Arrow–Debreu theory.[13] There is no assumption about perfect foresight, nor any other form of expectations.

- *Assumption 2* Each competitive trader considers the world's endogenous uncertainty states e(y) and their probabilities Ψ (y) as independent of her or his individual actions.

This is a realistic assumption in economies where uncertainty has some of the characteristics of a 'public good': the intuition is that traders are 'small' so that while the aggregate output of the economy does affect endogenous uncertainty, each trader takes the world's endogenous uncertainty as a parameter. This is similar to the situation with respect to prices in the standard competitive markets: each trader takes prices as given even though everyone's actions determine the prices at an equilibrium.

The trader's choice under endogenous uncertainty Having established the structure of information, the trader's problem of choice under uncertainty is straightforward. Under the standard von Neumann–Morgenstern axioms for choice under uncertainty:

- *Assumption 3* For any given price vector $p \in \Delta_M$ and any given set of states of endogenous uncertainty[14] $e \subset \{1 \ldots D\}$ with probabilities $\{\pi_i\}_{i \in D}$, trader chooses a consumption vector $D_h(p) = (d_1 \ldots d_D) \in R^{M \times D}$ which maximizes the expected utility of consumption

$$\sum_{i=1}^{D} \pi_i u_h^i (d_h^i(p,\pi)) = \text{MAX} \left\{ \sum_{i=1}^{D} \pi_i u_h^i (z_i(p)) \right\} \tag{7.3}$$

where the maximization is restricted to the set of consumption vectors z_i (p) which have a value equal to that of the trader's endowments plus the trader's share of profits:

$$\sum_{i=1}^{D} \langle \text{p}, z_i(\text{p}) \rangle = \sum_{i=1}^{D} \langle \text{p}\Omega_h^i \rangle + \theta_h(\text{p}). \tag{7.4}$$

Since the traders treat the probabilities π_i as given, for each trader h the maximization problem (7.3)–(7.4) has a unique solution as a function of the price vector p when u_h is strictly concave.

Definition 2 An economy E with endogenous uncertainty has H traders, J profit-maximizing firms which produce M goods, a production technology Y as described above, and a structure of uncertainty described by a function $\Psi : R^M \rightarrow \Delta_D$, where for every $y \in Y$ the set of states of endogenous uncertainty $\Psi(y) - S$ is not a singleton,[15] with preferences as described in (7.3) above, and satisfying Assumptions 1–3. Other technical assumptions are in Chichilnisky (1996e: 105–9).

Existence of a competitive equilibrium with endogenous uncertainty

The following definition of a competitive equilibrium with endogenous uncertainty formalizes the notion that the states of endogenous uncertainty and their probabilities are determined as part of the equilibrium concept:

Definition 3 A competitive equilibrium with endogenous uncertainty is: a price vector $\text{p}^* \in \Delta_M$, an aggregate production level of the economy y^* $\in Y$, a set of states of endogenous uncertainty e^*, $e^* \subset \{1 \ldots \text{D}\}$, each state with a corresponding probability $\pi_i^* > 0$, $i \in e^*$, $\Sigma_{i \in e^*} \pi_i^* = 1$, and for each trader h a consumption vector $\text{d}_h(\text{p}^*) \in R^{M \times D}$, $\forall i \in e^* \text{d}_h^i(\text{p}^*)$ $\in R^M$, such that:

1. For each trader h the consumption vector $\text{d}_h(\text{p}^*)$ is optimal for problem (7.3) with constraint (7.4), for the set of endogenous states e^* with associated probabilities π_i^*, $i \in e^*$.
2. The aggregate production vector y^* is profit maximing within Y at the equilibrium prices p^*: $\text{y}^* = y(\text{p}^*)$.
3. All markets clear at each state i of endogenous uncertainty:

$$\forall i \in e^*, \sum_{h=1}^{H} \text{d}_h^i(\text{p}^*, \pi) - \Omega_h^i = \text{y}^i(\text{p}^*),$$

and the states of endogenous uncertainty with probabilities $\{\pi_i^*\}_{i\in e^*}$ are precisely those states and those probabilities which are induced by the aggregate production of the economy at the equilibrium: $\{\pi_1^*\}_{i\in e^*} = \Psi(p^*)$.

The existence of a market equilibrium with endogenous uncertainty has been established under general conditions and generically on technologies:

Theorem 4 There exists a competitive equilibrium for a market with endogenous uncertainty generically on technologies.

Proof See Chichilnisky (1996e, p.123). ■

Risk allocation with endogenous uncertainty

In markets with endogenous uncertainty the concept of Pareto efficiency can be ambiguous. This is because the traders' von Neumann–Morgenstern preferences are defined with reference to the probabilities of the events and these change with the overall level of economic activity. Within Arrow–Debreu markets these are either subjective or objective probabilities, but in any case *fixed*. Here matters are quite different. The probabilities are now endogenously defined as part of an equilibrium. Therefore the traders' preferences themselves vary with the equilibrium, and Pareto efficiency of an allocation becomes a self-referential concept, in the sense that the allocation itself helps determine whether it is more or less valuable than other allocations.

It is, however, possible to define a restricted concept of efficiency in markets with endogenous uncertainty:

Definition 5 An allocation of the economy E is a vector $\mathbf{x} \in R^{M\times H}$; where $\mathbf{x} = (x_h)_{h=1\,\dots\,H}$, $x_h \in R^{M\times D}$; it is called *feasible* if $\Sigma_{h=1}^{H}(x_h - \Omega_h) \in Y$, that is, when the sum of what is allocated in excess of the economy's endowments can be produced.

Definition 6 A feasible allocation $x = (x_h)_{h=1\,\dots\,H} \in R^{M\times D}$ in E is called *exogenously Pareto efficient* when it is Pareto efficient relative to allocations according to the same preferences prevailing at x. That is, when there exists no other feasible allocation $y = (y_h)_{h=1\,\dots\,H} \in R^{M\times D}$ in E such that for all h,

$$\sum_{i=1}^{e(y)} \pi_i(x)u_h^i(y_h) \geq \sum_{i=1}^{e(y)} \pi_i(x)u_h^i(x_h)$$

with strict inequality for some h.

Theorem 7 A competitive equilibrium of a market with endogenous uncertainty is exogenously Pareto efficient.

Proof This follows immediately from the first welfare theorem. ■

Financial innovation and endogenous uncertainty

In the previous sections we saw that markets with endogenous uncertainty have competitive equilibria generally, and these define *exogenously Pareto-efficient* allocations. However, markets with endogenous uncertainty have generally Pareto-inefficient equilibria. One reason is that markets with endogenous uncertainty have no financial assets to hedge endogenous uncertainty. Traders know that at the equilibria different states and probabilities exist, and act appropriately, but they have no means to transfer wealth across states of endogenous uncertainty:

Proposition 8 The market with endogenous uncertainty E is incomplete, in the sense that it has no assets to hedge endogenous uncertainty, that is, no assets which pay contingent on the realizations of endogenously induced risks.

Financial innovation

It seems natural to introduce new assets which pay contingent on the realization of endogenous uncertainty. The introduction of new assets is called *financial innovation*.

In standard markets with exogenous uncertainty and incomplete asset structures (Chichilnisky and Heal 1996) it is always possible to complete the market by introducing new assets. Through the introduction of Arrow securities which allow the transfer of wealth across states between which this was not possible before, markets can be completed. When all possible such assets have been introduced, one says that the markets have been completed. By definition, a completed market economy has the structure of an Arrow–Debreu (complete) market:

Definition 9 A standard market economy with exogenous uncertainty which has S states of nature and where it is not possible to shift income across $S - T$ of its states (the span of its asset matrix is $T < S$) is called *incomplete*. The act of introducing $S - T$ Arrow securities each of which pays a unit of a numeraire in each of the $S - T$ states and zero in all others is called completing the market. A completed market is by definition one which is identical to a standard Arrow–Debreu model. In particular, a competitive equilibrium of such a completed market is always Pareto efficient.

In parallel with the results just quoted on incomplete markets with exogenous uncertainty, I explore the possibility of completing the markets for endogenous uncertainty. Each endogenous risk is represented by an (endogenously determined) probability function π over events in the set D; therefore one should ideally aim at introducing financial instruments which pay contingent on such distributions. The assets we have in mind mimic Arrow securities, but their payoffs are contingent on probabilities: they pay a unit of the numeraire if one probability arises and zero otherwise. Since we have assumed that there is no problem of information, it should be possible to introduce and trade such instruments. The matter may appear at first sight to be somewhat theoretical; for this reason it is desirable to discuss a practical example where similar instruments were introduced and are currently traded.

CAT Futures and catastrophe bundles

Assets which pay contingent on the observed frequencies of occurrence of natural events, were first introduced and analysed in Chichilnisky and Heal (1992a, 1992b and 1998). Assets which pay contingent on the realization of frequencies of natural risks have been recently introduced in the Chicago Board of Trade, called CAT (Catastrophe) Futures. These instruments' payoffs depend *inter alia* on the incidence of tropical storms in the United States, as measured, for example, by the Insurance Service Organization index. The catastrophes contemplated in CAT futures include earthquakes on the West coast, tornadoes on the East coast and floods in the Midwest. The frequencies of these events are unknown, and therefore these frequencies are treated as risks.

Chichilnisky and Heal (1992a) and Chichilnisky (1997b) showed under general conditions that such instruments improve welfare in markets with unknown risks, that is, where the probabilities or frequencies of the events are unknown. Furthermore, Chichilnisky (1995, 1996h, 1997a, 1998) showed that a specific combination of mutual insurance and securities called 'catastrophe bundles' is the most simple and efficient instrument to hedge such risks.

Floods, earthquakes and tornadoes are all exogenous physical events; they are not risks induced by economic actions. Our markets, instead, face endogenous risks. This difference is an important one, as the following result shows:

Theorem 10 It is not possible to complete a market for endogenous uncertainty.

Proof This theorem was established formally in Chichilnisky (1996e); an intuitive explanation follows. Consider an economy with endogenous uncertainty. Assume that all possible contracts contingent on all probabilities $\{\pi_i\}_{i \in D}$ over the sets of states $\{1 \dots D\}$ have been introduced. Equivalently, traders now can trade contingent on the realization of each possible probability distribution $\{\pi_i\}_{i \in D}$. The assumption that the market has been completed leads to a contradiction. If the market were now complete, then by definition it could reach an allocation corresponding to that of an Arrow–Debreu economy with complete asset markets. This implies that the traders must be fully ensured across states of endogenous uncertainty, namely that at such an equilibrium allocation x^*, for each trader h, $x_h^{i*} = x_h^{j*}$ at any two states $i \neq j$ (Chichilnisky and Heal 1994 and Chichilnisky 1996e). But this implies that for each of the two states of endogenous uncertainty i, j each trader has the same consumption vector; therefore the aggregate production of the economy must be the same, $y^{i*} = y^{j*}$. Since the map $\Psi = (\Psi^1 \dots \Psi^D) : R^M \rightarrow \Delta_D$ is a function for each $d = 1 \dots D$, this implies that any two Pareto-efficient allocations lead to the same state of endogenous uncertainty and the same probabilities prevail over these two states. Since this is true for any two states of endogenous uncertainty, this implies that the economy does not have endogenous uncertainty, a contradiction. Therefore a market with endogenous uncertainty cannot be completed. ■

Policy conclusions: the cost of climate risks

Chichilnisky (1997a) formulated and proved the existence of a competitive equilibrium in markets with endogenous uncertainty, where the traders' actions induce changes in the state spaces which represent uncertainty, and in the probabilities of the states. The equilibria exist very generally. Markets with endogenous uncertainty are incomplete. The incompleteness is not assumed; it is proved. It derives from the nature of endogenous uncertainty and it cannot be circumvented. There is no way to complete the market with endogenous uncertainty no matter how many securities are added to it. This is different from the standard literature on incomplete markets in which the incompleteness is assumed and can easily be removed by adding more securities (see Chichilnisky and Heal 1993).

I formulate here a specific instrument to improve the allocation of risk in these incomplete markets: assets which pay contingent on probability distributions over states. Such assets exist in practice although they have been introduced recently. As mentioned above, they were anticipated in Chichilnisky and Heal (1992a, 1992b), and are currently trading in the Chicago Board of Trade under the name CAT Futures. More recently, Chichilnisky (1995, 1996c, 1996e) and Chichilnisky and Heal (1998)) introduced a more sophisticated instrument, catastrophe bundles, and proved that it is the most efficient instrument to hedge catastrophic risks for which there are several possible priors. I showed above that they lead to Pareto-efficient risk allocation with 'scientific risks', but do not fully hedge 'strict endogenous uncertainty'.

A different policy aspect arises in the case of hedging environmental risks which are the byproduct of industrial activity. Here the situation is more favourable: if one seeks constrained Pareto efficiency, it is possible to compute precisely the cost which is worth incurring to decrease the probability of an environmental risk. Assuming a known (or even an approximately known) scientific relation between industrial output y and the probability distribution of the different events (that is, the map Ψ defined above) one computes the manifold of equilibria of the economy with endogenous uncertainty, and finds within it a new equilibrium with the desired value. At this new equilibrium one computes the utility levels achieved by the traders at their new consumption; the difference in the welfare of the traders at the first and at the second equilibrium provides an upper bound (or ceiling) to the willingness to pay for a decreased risk.

5.2 A Growing Economy with Endogenous Uncertainty

Endogenous uncertainty is also present in risk allocation through time. The following provides a framework for a growing economy in which the consumption of fossil fuels should be curtailed because it increases the probability of a change in climate. The economy has three main characteristics. First, the atmosphere may be in one of two states, one favourable to economic activity and one unfavourable (there is a possibility of a future climate change). The favourable and unfavourable states are denoted A_f and A_u, respectively. Second, the atmosphere transits stochastically from the favourable state to the unfavourable, and once there remains there for ever, so that atmospheric change is irreversible: A_u is an absorbing state. The probability of transiting to the unfavourable state is endogenous and increases with the level of cumulative emissions from the use of fossil fuels.

Fossil fuels (use rate R_t), capital equipment (stock K_t) and the atmosphere ($A = A_f$ or A_u) are used to produce output Q_t, which may be consumed C_t or reinvested $\dot{K}_t$ to augment the capital stock. Production generates emissions, which affect the probability of a change in the state of the atmosphere. The atmosphere is a resource that enters into the economy's production function, which may be in a favourable or an unfavourable state. Initially the atmosphere is in the favourable state but may change stochastically to the unfavourable state, and once in this state will remain there for ever. The source of emissions forever is the use of an exhaustible resource in production. The remaining input to production is the capital stock. An obvious example of this structure is the emission of CO_2 generated by the use of fossil fuels.

$$\left.\begin{array}{c} Q_t = Q(K_t, R_t, A) = C_t + \dot{K}_t \\ A = A_f \text{ or } A_u \\ Q(K_t, R_t, A_f) > (K_t, R_t, A_u) \text{ for all } K_t, R_t \end{array}\right\}$$

The probability of a change of atmospheric state depends on cumulative emissions, and emissions are assumed proportional to current use of the fossil fuel. For simplicity we therefore identify emissions and fossil fuel consumption R_t. Let

$$Z_t = \int_0^t R_\tau d\tau, \frac{dZ_t}{dt} = R_t.$$

The evolution of the climate is as follows. There is a date $T > 0$ such that $A = A_f$, $t < T$, and $A = A_u$, $t > T$. Here T is a random variable whose marginal density function f has as its argument cumulative emissions Z_t, $f = f(Z_t)$. The probability that the climate changes, that is, the date T occurs, in an interval (t_1, t_2), is

$$\Pr T \in (t_1, t_2) = \int_{Z_{t1}}^{Z_{t2}} f(Z_t)\, dt.$$

It follows that if $Z_{t1} = Z_{t2}$, so that there is no depletion or emission in the interval (t_1, t_2), then the probability of climate change in that interval is zero. When there is emission in an interval (t_1, t_2), the chance of climate change depends on emissions in that interval and also on cumulative emissions up to that interval. All of this makes good sense.

Output may be consumed or invested. Consumption yields utility and the objective is to maximize the expected present discounted utility of consumption. There is a constraint on the total amount of the resource that can be used, as this is exhaustible. The problem involves maximizing expected utility subject to the resource and national income constraints, where the expectation is over the process governing climate change. Formally:

$$\left.\begin{array}{c} \max E\int_0^\infty U(C_t)e^{-\delta t}\, dt \\ \text{s.t.} \int_0^\infty R_t dt \le S_0 \\ \dot{K}_t = Q(K_t, R_t, A) - C_t \end{array}\right\}$$

The expectation here is over the distribution of the date of change of the climate, T.

For this problem, Heal (1984, 1990) characterizes optimal paths of consumption, capital accumulation and use of fossil fuel. He compares these with those that are optimal in the absence of an atmospheric impact, and also studies the impact of changes in parameters such as the discount rate and degree of risk aversion. He isolates the key parameters in determining the optimal rate of use of fossil fuels.

The introduction of atmospheric impact makes a difference. The time profile of resource use which emerges is flatter than that which emerges from an optimal depletion problem with no atmospheric impact. Initial levels of resource use are lower, and they fall more slowly, than with no atmospheric impact. The difference depends on the degree of risk aversion and on the parameters of the probability distribution relating cumulative emission to climate change.

The behaviour of the shadow price of the resource is also of interest. In the pure depletion case this price rises at the rate of discount; in Heal (ibid.) it may fall and even become negative. We can interpret the difference between the shadow price of the resource in the no-atmospheric-impact case and the current case as an optimal carbon tax. This tax depends on the country's degree of risk aversion and on the parameters of the probability distribution describing the risk of climate change as a function of carbon emissions, as well as on the damage resulting from climate change. The model thus leads to a distinctive approach to characterizing an optimal carbon tax and its evolution over time. Hartwick (1992) gives an analysis of carbon taxes using this framework.

The likelihood of climate change as a function of economic activity is a key relationship in evaluating the choices posed in this model. This is a functional relationship rather than a parameter. Global change R&D leads us to a better understanding of this relationship. It is worth stressing that proper economic analysis requires not just the likelihood of climate change as a result of one particular emission scenario, which is what most scientific analyses provide, but rather *a systematic evaluation of how the nature and likelihood of climate change varies with the pattern of economic activity*. The study and characterization of this likelihood function is an important topic for interdisciplinary research.

It is not surprising that what it is worth paying to reduce the risk of climate change depends *inter alia* on a society's degree of risk aversion and discount rate. However, this has an interesting implication. Even if there were complete agreement about all of the scientific aspects of the global change problem, there could still be disagreement about policy responses. Because of the international externalities associated with climate, so that all countries consume the same climate, CO_2 abatement policies only make sense if coordinated internationally (see Heal 1992).

Different countries' positions with respect to measures to restrict greenhouse gas emissions depend on their discount rates and degrees of risk aversion. The US, for example, has been against global abatement agreements, while Germany has been in favour. This fits with the conventional wisdom that the financial and industrial community in the US have both a higher discount rate and a lower degree of risk aversion (greater willingness to take risks) than those in Germany. The differences in policy positions could then be attributed to differences in preferences rather than, or in addition to, different interpretations of the current scientific evidence. Figure 7.2 portrays such an interpretation of differences in the attitudes underlying policy choices towards global warming.

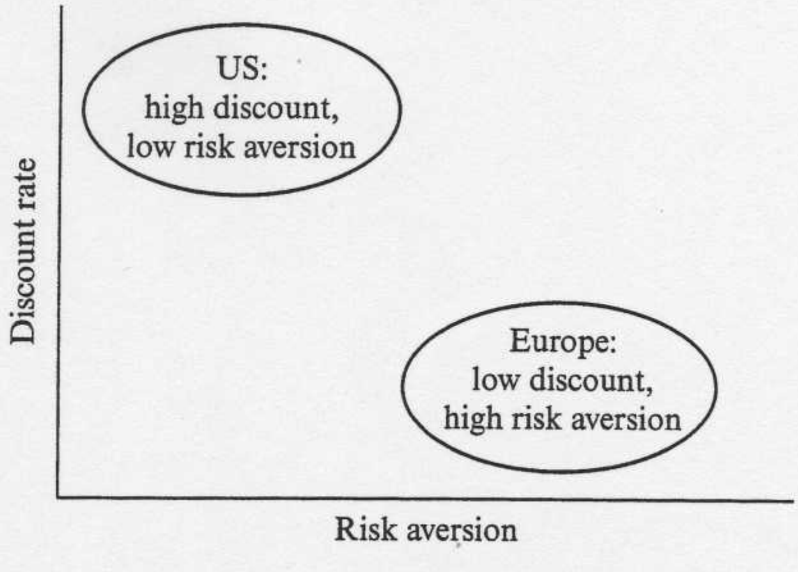

Figure 7.2 Possible systematic differences in parameters determining attitudes towards the risk of climate change

Different perceptions of the risk involved do not, however, preclude efficient solutions. Differences in preferences can lead to gains from trade. In this case differences in attitudes towards risk could be grounds for the introduction of markets in which different risk positions are traded, with efficiency gains, see Chichilnisky and Heal (1992b).

6 OPTION VALUES AND IRREVERSIBILITY

A central issue in valuing environmental resources such as current climate conditions, biodiversity, or complex ecological systems, is the irreversibility of decisions and events. An aspect of these resources is that once altered they cannot easily be restored to their current conditions, at least on a relevant timescale. The decision not to preserve a rich reservoir of biodiversity such as the 60 million-year-old Korup forest in Nigeria is irreversible. The alteration or destruction of a unique asset of this type has an awesome finality, and analysts have sought to capture this in a framework for cost–benefit analysis. This has led to the concept of 'option value': preserving a unique asset in its present state allows us the possibility of changing our minds later. Altering it irreversibly does not. Preserving it has thus to be credited with an 'option value' because it keeps open to us the option of reconsidering our decision. Altering it leaves us no such option in the future.

A concept related to option value is that of 'nonuse' or existence value. We may value environmental goods for which we have no immediate economic use. The existence of certain species is in this category: the Californian condor, the spotted owl, and various snails and fish come to mind. There is no sense in which we can currently use these species: possibly one could argue that the condor and the owl have consumption value for those willing to make the effort needed to see them, but few people come into this category. One doubts that this is a significant issue with the snail.

The two concepts, option value and nonuse value, seem to overlap. Many goods which exemplify one also exemplify the other. At the same time, there are no doubt differences. Nonuse values stem in some degree from ethical considerations, from a recognition that a species has a right to exist even if humanity places no direct value on it. But one suspects that behind many nonuse valuations there lurks an option value: many nonuse valuations stem from an unstated belief that a use value may emerge.

This section reviews two distinct formulations of this issue, one in which the returns to a preservation project are uncertain at present but will be revealed in the future, and one in which the preferences of future generations for environmental facilities are uncertain. The first framework is the one in which the issue of option values has traditionally been studied. I provide an outline of the argument and show that one needs three conditions for an option value to exist. These are irreversibility, the acquisition of information with the passage of time, and an asymmetry of the underlying probability distribution. Similar results apply to the case of uncertainty about the preferences of future generations (Beltratti, Chichilnisky and Heal, 1998).

6.1 Waiting for Information

The option value of preserving an environmental or ecological asset has been explored in the context of uncertainty about the future benefits associated with its existence. A review of the literature is in Fisher and Krutilla (1985).[16] The main issue is that there are benefits that will accrue in the future from the preservation of a resource, but these are currently unknown. If the resource is preserved into the future, then in the future the decision about whether to preserve it can be reconsidered in the light of better information then available about the benefits from its existence. If it is not preserved, then there is no chance of reconsideration when we have better information. In this case conventional decision rules will underestimate the value of preserving the asset. The following example (from Dasgupta and Heal 1979) illustrates the key point in a simple framework. It is illustrated in Figure 7.3.

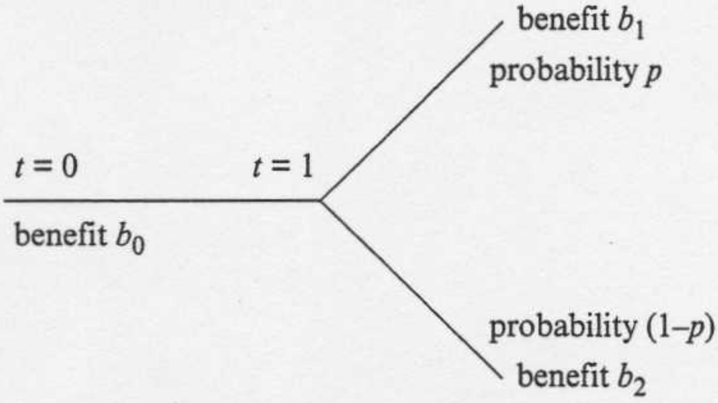

Figure 7.3 The benefits from conservation in different states

Below I shall show that with irreversible decisions there is an option value to conservation in the initial period if and only if there is a positive expected payoff from conservation in that period given that we follow an optimal policy. I contrast this with the reversible case, in which one never conserves in the first period and there is no option value.

Consider two dates, $t = 0$ and $t = 1$. We have one unit of an environmental asset. The benefit from preserving this at time $t = 0$ is b_0. At time $t = 1$ there are two possible states of nature s_1 and s_2. The state of nature is revealed at time $t = 1$. If the state is s_1, the benefit of preserving the asset is b_1: if s_2 is the state, the benefit is b_2. The probabilities of s_1 and s_2 are p and $(1 - p)$, respectively. Decisions about preservation are made at times $t = 0$ and $t = 1$. At $t = 0$ a decision is made on how much of the asset to preserve until $t = 1$: at that date we may either conserve everything conserved at $t= 0$, or conserve less. Given that destruction is irreversible, we cannot at $t = 1$ conserve more than was conserved at $t = 0$. Our options at $t = 1$ are therefore constrained by the decision made at $t = 0$. These data are summarized in Figure 7.3.

Compare the case already described, where the decision made at time $t = 0$ is irreversible, with an alternative case in which this decision can be reversed. In this case the decision made at time $t = 0$ no longer constrains the options available at time $t = 1$. Let us look at this alternative case first, as it is simpler and provides a benchmark. Let c_0 be the amount of the resource conserved at time $t = 0$, and c_1 and c_2 be the amounts conserved at time $t = 1$ in states 1 and 2, respectively. The expected benefit from development (assuming a zero discount rate) is

$$b_0c_0 + pb_1c_1 + (1-p)b_2c_2. \tag{7.5}$$

One has to choose conservation levels c_0, c_1 and c_2 to maximize (7.5). Assume that there is currently no benefit to preservation,[17] $b_0 < 0$, nor is there any benefit in state 1 in the future, $b_1 < 0$. However, there is the possibility of state 2 in which there are positive benefits from preservation, that is, $b_2 > 0$. If decisions are reversible, we conserve nothing at time $t = 0$, that is, we set $c_0 = 0$. Then at time $t = 1$, we conserve nothing in state 1 and everything in state 2, that is, we set $c_1 = 0$ and $c_2 = 1$. In the reversible case we *can* set $c_2 = 1$ because by assumption decisions made at $t = 0$ *are* reversible.

Now consider the real case in which the decision at time $t = 0$ cannot be reversed later. In this case the choice made at $t = 0$ does constrain the choices open at $t = 1$. We have to satisfy the constraint that what is conserved at time $t = 1$ cannot exceed that which was conserved initially, that is, $0 \le c_1, c_2 \le c_0 \le 1$. In particular, if everything is destroyed in the first period, then we have no options in the second. What policies now maximize (7.5)? Is there a value to carrying the option to conserve into the second period? Clearly if in the second period the state of the world is one in which there are positive benefits to conservation, then we will conserve everything left to us by our earlier decisions, that is, we will always set $c_2 = c_0$. If, however, the state is unfavourable to conservation, then we shall conserve nothing and set $c_1 = 0$. Hence the maximand (7.5) reduces to

$$[b_0 + (1-p)b_2]c_0 + pb_1c_1 = [b_0 + (1-p)b_2]c_0$$

and the initial conservation level is positive if and only if

$$[b_0 + (1-p)b_2] > 0. \tag{7.6}$$

The inequality (7.6) has a simple interpretation: the left-hand side is the expected payoff from conservation in the first period. It is the certain payoff in the first period plus the expected payoff from conservation in the second, given that if the state unfavourable to conservation occurs there will be no conservation in the second period. It is the expected payoff to conservation in period one given that an optimal policy is followed subsequently.

It is optimal to conserve in the first period if and only if there is a positive expected payoff from conservation given that we follow an optimal policy. Contrast this with the decision in the reversible case, in which we never conserve and always choose $c_0 = 0$. These two decisions are different if the expected payoff to conservation in the first period is positive.[18] In this case there is an option value to conservation as a means of carrying the resource into the second period and taking advantage of future information.

6.2 Option Values and the Value of Information

The existence of an 'option value'[19] does not depend on risk aversion.[20] The key issues are: first, the irreversibility of the decision; second, the fact that delaying a decision can let one take advantage of better information; and third, the asymmetry represented by (7.6). This latter condition implies that on average there will be benefits from conservation in the first period, provided that we choose optimally later.

Important practical implications flow from the analysis that we have just completed. Climate change is likely to be irreversible if it occurs. So in a cost–benefit analysis of preventing climate change (that is, preserving the atmospheric environment), it may be appropriate to assign conservation (preventing climate change) an option value. This will be the case if the passage of time is likely to bring significant new information about the likelihood of climate change or about its consequences and the expected payoffs satisfy (7.6) above.

A thorough study of the costs and benefits of reducing climate change is in Cline (1992). It seems worth noting that although this study refers many times to the scientific uncertainties associated with predicting climate change, it at no point attributes an option value to preservation, that is, to preventing climate change. This means that it may systematically underestimate the benefit–cost ratio of preservation of the atmosphere in its status quo. There is also an analysis in Manne and Richels (1992) of the value of waiting for scientific information about the greenhouse effect. They consider two possibilities: acting strongly now to reduce the emission of greenhouse gases, or taking very limited action now and waiting until there is further scientific evidence. Taking major steps towards emission abatement now amounts to conserving the atmospheric environment in its present state, and should again be credited with an option value. Manne and Richels fail to do this, and so again underestimate the value of buying insurance against the greenhouse effect by acting strongly now. As the value of an option generally increases with increasing uncertainty about the future, and as uncertainty looms large in any projections regarding global warming, the extent of the underestimate could be important.

6.3 Uncertainty about Future Generations

The concept of option value can be generalized or refined in several ways. A key consideration seems to be the possibility that future generations will value environmental resources more than we do. If this is simply a statement that these resources will be scarcer, and so more valuable on the margin, then this effect is captured in the usual approach to cost–benefit analysis (see Heal 1992).

It may, however, be a statement that future generations could have different preferences from us, and might value environmental assets differently. Because they might value them more, one should, it is argued, attribute a value to leaving them the option of high consumption levels. Solow (1992) argues that an important element in the definition of sustainability is recognizing the possibility that the preferences of future generations about environmental assets may be different from ours. This seems close to the concept of option value set out above, and indeed it is, though there are some differences that are revealing. Recent results in Beltratti et al. (1998)[21] establish that uncertainty about future preferences alone is not sufficient to produce an 'option value' for increasing the resource left to the next generation. In addition to pure uncertainty, there must be asymmetry in the distribution of possible changes in preferences. Neutral uncertainty, in the sense that increases and decreases in intensity of preferences are equally likely, does not generate a case for leaving more to the future in case their preferences for the resource are more intense than ours. Uncertainty makes a case for conservation only when the expected return to postponement of consumption is positive.

7 RESPONSES TO CATASTROPHIC RISKS

As already mentioned, traditional formulations of uncertainty do not deal adequately with catastrophic risks, namely with low probability events with major adverse consequences. Traditional decision theory, is based on von Neumann–Morgenstern axioms. I show below why those axioms are not adequate for ranking catastrophic risks. Then I introduce and develop new axioms and derive the optimization criteria which they imply, following Chichilnisky (1996f). Finally I discuss practical responses to climate change that emerge from this analysis.

7.1 Von Neumann–Morgenstern Axioms

Mathematical axioms introduced half a century ago by John von Neumann and Oscar Morgenstern gave rise to a now classical tool for decisionmaking under uncertainty.[22] The axioms give rise to procedures to rank or evaluate risky outcomes. The von Neumann–Morgenstern (VNM) axioms provide a mathematical formalization of how to rank lotteries. Optimization according to such a ranking defines decisionmaking under uncertainty.

A system with uncertain characteristics can be in one of several possible states; each state is represented by the value of a random variable. For example, the average temperature of the planet's surface is a state.

For simplicity the system's states are described by real numbers. To each state $s \in R$ there is an associated outcome, for example to each temperature level there is an associated vector describing soil fertility, so that one has $f(s) \in R^N$, $N \geq 1$. A description of outcomes across all states is called a 'lottery'. A lottery is a function $f : R \rightarrow R^N$, and the space of all lotteries is therefore a function space L.

A main result obtained from the VNM axioms is a representation theorem which characterizes *all* possible rankings satisfying their axioms. These rankings are given by a specific type of functions $W : L \rightarrow R$, known as 'von Neumann–Morgenstern (VNM) utilities'. The decision procedure obtained by optimizing such utilities is called 'expected utility maximization' and has the form,

$$W(x) = \int_{s \in R} u[x(s)] d\mu(s) \tag{7.7}$$

where the line R is the state space, the variable $x : R \rightarrow R^N$ is a lottery, $u : R^N \rightarrow R$ is a (bounded) utility function describing the utility provided by the outcome of the lottery in each state s, $u(s)$, and where the measure $d\mu(x)$ is a probability distribution over measurable subsets of states in R.

According to the VNM representation theorem, rational choice under uncertainty must take the following form: a lottery x is ranked above another y if and only if W assigns to x a larger real number, that is,

$$x \geq y \Leftrightarrow W(x) > W(y),$$

where W satisfies (7.7).

The optimization of expected utility is a widely used procedure for :valuating choices under uncertainty. Functions such as W are amenable o a large body of knowledge which goes back several centuries: the cal- :ulus of variations. The Euler–Lagrange equations are typically used to :haracterize optimal solutions. Such mathematical tools are widely used o find and describe choices under uncertainty.

'.2 Catastrophic Risks

)espite their frequent use, the classic methods defined above are not ade- uate for lotteries involving catastrophic risks. The reasons are both ractical and theoretical. From the practical point of view, it has been hown that using such criteria undervalues catastrophic risks and hence onflicts with the observed evidence of how humans evaluate such risks Chichilnisky 1996f; Lowenstein and Thaler 1989; Lowenstein and Elster 992). For example, using VNM utilities, the most damaging scenarios of lobal climate change induce little if any economic loss. The ntergovernmental Panel on Climate Change (IPCC), the main interna- onal scientific organization in this area, recently predicted a highly ontested figure of about 2 per cent loss of economic value from a dou- ling of CO_2 concentration in the atmosphere. This is a symptom of a iore general phenomenon: a simple computation shows that the hypo- ietical disappearance of all irrigation water in the US and all the ountry's agricultural produce would have at most a 2.5 per cent impact n its gross domestic product (Cline 1992). This finding underscores the nportance of using appropriate criteria for evaluating catastrophic risks.

Mathematically the problem is that the measure μ which emerges from ie VNM representation theorem is countably additive. Since the utility inction $u : R^N \rightarrow R$ is bounded $(\sup_{x \in R} u(x)) < \infty$, the countable additiv- y of μ implies that any two lotteries $x, y \in L$ are ranked by W quite dependently of the utility of the outcome in states whose probabilities e lower than some threshold level $\varepsilon > 0$, where ε depends on x and y.[23] ich a function is called 'insensitive to small probability events'. ormally:

Definition 11 The function W is called insensitive to small probability events if

$$W(x) > W(y) \Leftrightarrow \exists \varepsilon > 0 \ \varepsilon = \varepsilon(x, y) : W(x') > \mathrm{W}(y') \qquad (7.8)$$

r every x' and y' such that
$= x$ and $y' = y$ a.e. on a set $A \subset R$, $\mu(A^c) < \varepsilon$.

This means that W ranks x above y if and only if it ranks x' above y' for any pair of lotteries x' and y' which are obtained by modifying arbitrarily x and y in sets of states A with probability lower than ε.[24] The interpretation of this property is that the ranking defined by W is 'insensitive' to the outcomes of the lottery in small probability events. The following result shows why VNM utilities are not adequate for evaluating catastrophic risks.

Lemma 12 Von Neumann–Morgenstern utilities are insensitive to low probability events. Therefore they are not adequate for ranking catastrophic risks.

Proof Consider two lotteries $x, y \in L$, where x is superior to y according to a VNM utility W. Formally:

$$W(x) = \int_{s\in R} u[x(s)]\, d\mu(s) > W(y) = \int_{s\in R} u[y(s)] d\mu(s).$$

Obviously

$$W(x) > W(y) \Leftrightarrow \exists\alpha(x,y) > 0 : W(x) > W(y) + \alpha.$$

To show that a VNM utility is insensitive to small probability events we must show that the ranking between x and y is insensitive to the outcomes in sets of small enough measure, say ε, as defined in (7.8). I will show that W satisfies this property. Let $\varepsilon = \alpha/3.M$, where $M > \mathrm{Sup}_{x\in R} | u(x) |$. Two new lotteries x' and y' are now obtained by altering arbitrarily the lotteries x and y, respectively, on two arbitrary sets V_1 and V_2, each of which has a measure smaller than ε. Formally: $x(s) = x'(s)$ a.e. in V_1^c, and $y(s) = y'(s)$ a.e. in V_2^c. By construction

$$|W(x) - W(x')\ | < \int_{s\in v_1} u[x(s)] d\mu(s) < M.\varepsilon \le \alpha/3,$$

and equally

$$|W(y) - W(y')\,| < \int_{s\in v_2} u[y(s)] d\mu(s) < M.\varepsilon < \alpha/3,$$

Since $W(x) > W(y) + \alpha$, it follows that

$$W(x) > W(y) \Rightarrow W(x') > W(y');$$

reciprocally

$$W(x') > W(y') \Rightarrow W(x) > W(y)$$

so that for the chosen $\varepsilon = \varepsilon(x, y)$

$$W(x) > W(y) \Leftrightarrow W(x') > W(y'),$$

as we wished to prove. Since this is true for any two lotteries x and y, W is insensitive to small probability events. ■

A consequence of this lemma is that VNM utilities are not well suited for evaluating catastrophic risks. The problem is general. It can be shown formally that cost–benefit analysis under uncertainty based on expected utility maximization (which follows from VNM axioms) underestimates the outcomes of small probability events. It is therefore biased against environmental projects which are designed to avert catastrophic events. Experimental evidence shows that humans treat choices under uncertainty somewhat differently from what the VNM axioms would predict, suggesting the need for alternative axioms which describe more accurately humans' valuations (Lowenstein and Thaler 1989; Lowenstein and Elster 1992).

7.3 Updating Von Neumann–Morgenstern Axioms

Recently a new set of axioms was proposed to update VNM axioms for catastrophic risks (Chichilnisky 1996f). These axioms take a more balanced approach towards small probability events. They contrast with VNM axioms in the treatment of small probability events (ibid.). On the basis of these axioms a new representation theorem has been obtained that fully characterizes the functions to be maximized under uncertainty.[25] This defines a new decisionmaking tool, one that appears to conform the evidence of how humans evaluate catastrophic risks (Lowenstein and Thaler 1989; Lowenstein and Elster 1992).

7.4 New Axioms of Choice for Catastrophic Risks

The three axioms introduced in Chichilnisky (1996f) are simple. The first axiom is standard, and is certainly satisfied by VNM utilities:

1. continuity and linearity of the ranking of lotteries x with respect to the utility $u(x)$.

The following two axioms are new; and the second axiom (2) is not satisfied by VNM utilities:

2. sensitivity to low probability events. This axiom rules out (7.8), as defined below.
3. sensitivity to large probability events, as defined in (7.9) below.

Definition 13 A ranking is sensitive to low probability events when it does not satisfy Definition 11.

Definition 14 A ranking is said to be insensitive to large probability events[26] when

$$W(x) > W(y) \Leftrightarrow W(x') > W(y'), \tag{7.9}$$

for any two lotteries x' and y' that are obtained by modifying arbitrarily x and y on a bounded set of states $S \subset R$ of arbitrarily large probability.

Definition 15 A ranking is sensitive to large probability events when it does not satisfy (7.9).

Example 16 The following is a ranking that concentrates on events of vanishing probability and neglects large probability events. Define a measure $\mu(s)$ on measurable sets of the line R as follows: every bounded set has measure zero, and every complement of a bounded set has measure one. This is a finitely additive measure since the measure of a union of finitely many disjoint sets is the sum of the measure of the sets. However, it is not a countably additive measure, since the measure of a union of countably many bounded sets which equals R is one, and clearly this is different from the countable sum of the measure of the bounded sets, which is zero. Define now the ranking $W(x) = \int_{s \in R} u[x(s)]d\mu(s)$. Such a function is insensitive to bounded sets of events which have positive probability according to any standard countably additive measure of the line. Let $L_\infty(R)$ be the set of measurable and essentially bounded real valued functions on the line. The 'dual' of L_∞, denoted L_∞^*, is the set of all real valued, continuous linear functions on L_∞. It has been shown (see, for example, Chichilnisky 1996f) that this dual contains two types of elements, both types being defined by measures on the line R: standard or 'countably additive' measures, and 'purely finitely additive' measures. The latter assigns measure zero to any bounded set in R, as is the case with the measure constructed in the beginning of this example. Such measures define continuous, real valued linear functions on lotteries in L_∞, $W : L_\infty \to R$, called 'integral operators', given by $W(x) = \int_{s \in R} u[x(s)]d\mu(s)$; the measure μ (whether countably or finitely additive) is called a 'kernel' because of its role inside the integrand in the definition of the operator W. Such functions on lotteries satisfy (7.9), and by definition they are insensitive to large probability events. They are therefore ruled out by axiom (3) which requires sensitivity to large probability events. Indeed, such a function W puts all the 'weight' on infinity, that is, on events which have arbitrarily small probabilities according to any standard countably additive measure μ on R.

7.5 A New Representation Theorem

Like the VNM axioms, the new axioms defined here lead to a section theorem. It has been shown in Chichilnisky (1996f) that there exist functions $\Psi : L_\infty \to R$ which rank all lotteries and satisfy all three axioms in 7.4. As in the VNM case, these are given by integral operators. However, rather than having countably additive kernels as in the VNM representation, these functions are a convex combination of integral operators with countably additive measures and integral operators with purely finitely additive measures. Both measures (countably and finitely additive) are nonzero.

Theorem 17 There exist rankings of lotteries that satisfy the three axioms (1), (2), (3). VNM expected utilities do not. Every ranking of lotteries that satisfies the three axioms admits a representation by a function $W : L_\infty \to R$, of the form

$$W(x) = \int_{s \in R} u[x(s)]d\mu(s) + \Phi\ \{u[x(s)]\}.$$

where μ is a standard countably additive measure on the reals R, $\int_{s \in R} d\mu(s) < \infty$, and where $\Phi \ \varepsilon L_\infty^*$ is a purely finitely additive measure on R.

Proof See Chichilnisky (1996a and f) ■

An example will fix ideas and illustrate the result.

Example 18 For simplicity, consider here discrete states indexed by the integers Z. Now a lottery is an element of l_∞, the space of bounded sequences of real numbers. Define a continuous linear functional on lotteries $\Psi : l_\infty$ as follows:

$$\Psi(x) = \gamma \sum_{s=1}^{\infty} \lambda^{-s} u[x(s)] + (1 - \gamma) \lim_{s \to \infty} u[x(s)], \qquad (7.10)$$

where $0 < \gamma < 1$, and where $\lim_{s \to \infty} u[x(s)]$ is defined below. This function satisfies all the axioms. The interpretation of the two parts of the function Ψ in (7.10) is as follows. The first part is an integral operator with an integrable kernel $\{\lambda^{-s}\}_{S \in z}$ which defines a countably additive measure on Z, and therefore emphasizes the weight of large probability events in the ranking of a lottery $x \in l_\infty$. The second part defines a purely finitely additive measure on Z which assigns positive weight to possible catastrophic and small probability events. The second part of (7.10) ensures that no matter how small is the probability that a limiting value will be achieved by the lottery, the weight that this fact has in the criterion ensures that if a lottery x is preferred to another y, changing the lottery x and y on a set of events of arbitrarily small probability can reverse the ranking. Therefore the criterion is not insensitive to small probability events. Formally, $\mathrm{Lim}_{s \to \infty}$ is defined as the Hahn Banach extension of the standard limit function, extended using the Hahn Banach extension theorem from the subset of sequences in l_∞ that do have a limit, to all of l_∞. On the sequences that have a limit, $\mathrm{Lim}_{s \to \infty} u\,(x(s))$ is the utility level corresponding to the limiting value of the sequence.

Intuitively the second term in (7.10) defines a measure with 'heavy tails'. Since both parts are present in (7.10), the corresponding function Ψ is sensitive to small and to large probability events. Catastrophic risks are therefore evaluated more realistically by such functions.

Remark 1 Observe that the second terms in (7.10) is a continuous function on l_∞, with respect to the standard norm of l_∞, but does not admit a representation as an expected value.[27] The optimization of functions such as Ψ is not amenable to standard tools of calculus of variations. This must be developed in new directions. Some results already exist (Chichilnisky 1996f), but much work is still needed.The study of optimal solutions of these types of functions has led to asymptotically autonomous dynamical systems, which occur naturally when one extends the Euler–Lagrange analysis of optimal solutions to encompass the type of operators defined here. Statistical analysis of such systems also requires new tools.

7.6 Responses to Catastrophic Climate Risks

How to employ the new criterion in hedging catastrophic risks? In practical terms, what should one optimize under this criterion? Certainly we should not maximize expected utility, since we have shown that it underestimates catastrophes. The representation theorem provided above gives us the clue: it tells us what to optimize when making decisions involving catastrophic risks.

Example 19 The criteria derived from the axioms involve maximizing a convex combination of expected value *plus* maximizing the infimum utility value that the lotteries can achieve, that is, averting the small probability events that are the 'catastrophes'. The behaviour of a rational agent is therefore more conservative under these new criteria than it is under the von Neumann–Morgenstern axioms.

8 CONCLUSIONS AND OPEN QUESTIONS

The foundations are in place for understanding the economics of global environmental risks, but certain aspects require more attention.

From a policy perspective, scientific uncertainty is one of the most prominent issues. The risks we face are largely unknown: we do not know the probability that the climate will warm up, or of its consequences. Decisionmaking under these circumstances is a challenging endeavour. I have proposed financial instruments to aid in this task, including a combination of insurance and securities, called 'catastrophe bundles' (Chichilnisky 1995, 1996c, 1996h, 1997b and Chichilnisky and Heal 1997, 1998). These instruments work well under scientific uncertainty, but not with strict endogenous uncertainty. An open question is how to extend these results to the case where, in addition to the frequencies being unknown, these frequencies are altered by human activity.

I have reviewed the use of options for handling irreversibilities, and proposed a new type of options to hedge the risk that future generations will have preferences different from our own. Option values can make our behaviour more conservative. However, we do not know whether this is true when future preferences change in a symmetric fashion: this is still an open question.

Often the purpose of policy is to change the risks that we face. Examples are policies that aim at decreasing the risk of climate change. There is an explicit acknowledgement that we are facing endogenous risks, namely risks that are responsive to our actions. Endogeneity is a fundamental feature of global environmental risks, one that is little understood and only recently explored in economics. Traditional economic models for risk management, which take risks as exogenous, do not provide an adequate framework for analysing and hedging global environmental risks. The global carbon tax investigated by the OECD and reviewed in Chichilnisky (1994b) is a policy designed to change the odds; another is the creation in the Kyoto Protocol of global trading of permits for the emission of CO_2, which was originally proposed in May 1994 to the FCCC by Chichilnisky and Heal (1995) and Chichilnisky (1997). The explicit objective of these proposals is to decrease the risk of a climate change. A systematic study of markets with endogenous risks started with the existence theorems in Chichilnisky and Wu (1992), Chichilnisky (1992) and Chichilnisky et al. (1996e). Endogenous risks in connection with global warming have been analysed within growing economies (Heal 1984).[28] Two main results were given in this chapter, one analysing the existence and welfare properties of markets with endogenous uncertainty, and the other of growing economies where productivity in the future changes with today's carbon emissions. More work needs to be done in this area.

Another feature of global environmental risks is that they are correlated, as they affect many indivduals at once. Therefore they cannot be hedged solely by insurance, which is specifically designed for individual risks and relies on the law of large numbers. Global environmental risks can be catastrophic. Therefore traditional decision theory, which is based on expected utility criteria and underestimates small probability events, is not appropriate. I have proposed a new set of axioms to deal with catastrophic risks, derived the decision criteria that they imply and suggested how to apply this to management of catastrophic risks. These criteria have been applied to characterizing optimal behaviour in optimal growth models with exhaustible resources and with renewable resources, and a new 'turnpike' theorem has been obtained (Chichilnisky 1997c and Heal 1997). More work remains to be done on the optimal management of catastrophic risks.

NOTES

1. In medieval England, a peasant farmer's land was broken into many widely-dispersed parcels. Economic historians interpret this as a way of hedging climate risk (see references in Bromley 1992). Land in different locations would be affected differently by droughts, floods and frosts. By spreading land holdings over different locations, as well as by organizing agricultural cooperatives, these societies hedged against climate risks.
2. The August 1997 issue of *Science* is dedicated to the topic of human-dominated ecosystems.
3. Human location is not the issue here. This is not about the choice of humans to settle near natural hazards. This latter choice is important for human welfare, but it is not the issue that drives today's global environmental concerns. Indeed, the human activity that affects the atmosphere and the soil most is caused by humans who live elsewhere. Most ecosystem destruction is for agricultural production and resource extraction for export to far away lands. Many scientists believe that the separation of the location of the consumers and the location of the producers, which may diminish sensitivity for the negative 'externalities' or environmental damage caused by resource extraction, leads to faulty cost–benefit analysis and thus economic activity that is harmful to the environment.
4. Endogeneity of risks leads to moral hazard when risks depend on actions which cannot be observed by the insurers and will be influenced by the insurance offered. In the present context such problems are not central to the analysis: asymmetric information is not a characteristic of climate risks. Endogenous uncertainty is more general than moral hazard.
5. In this respect there may be a difference between the various aspects of climate risk. There are historical data on the relation between atmospheric CO_2 and climate from tree-ring and ice-core studies. With ozone depletion the phenomenon is so new that such data are not available.
6. See Chichilnisky and Wu (1992), Chichilnisky (1992, 1996e, 1997a and b), Kurz (1974).
7. They also supported part of this research, when Chichilnisky held the Salinbemi Chair at the University of Siena, 1995.
8. Therefore no 'moral hazard' exists.
9. A possible difference is that short sales on assets is generally allowed, while short sales on real goods may be restricted. We can allow for this by enlarging the trading space to include negative quantities of some of the commodities traded. In this case we must require an additional condition for existence of a competitive equilibrium, see Chichilnisky and Heal (1993).
10. The cardinality of the set of states could be extended to be infinite without changing the results in any way, but at the cost of more notation. In the case that the cardinality is infinite one needs to work in economies with infinite dimensional commodity spaces, ideally in Sobolev spaces; see, for example, Chichilnisky and Heal (1993) for a general

theorem of existence and characterization of a competitive equilibrium in Sobolev spaces with or without short sales.

11. It would suffice to consider the case where $D = S$ because the case $D > S$ can be reduced to the former if we consider that the probabilities in the original Arrow–Debreu model with exogenous uncertainty are zero outside the set s, that is, $\forall i \notin s$, $\pi_i = 0$. However, it seems preferable to keep a distinction between exogenous states and endogenous states.
12. See Chichilnisky (1996e: 105).
13. For example, the endogenous states and their probabilities could be announced by an auctioneer, as prices are announced by an auctioneer in a Walrasian market; the role of the Walrasian auctioneer is to ensure that no trading takes place until an equilibrium is reached. The same here: the (expanded) role of the auctioneer is now to ensure that no trading takes place until an equilibrium with endogenous uncertainty has been reached. The auctioneer's announcements about sets of states and their probabilities are neither correct nor false, in the same way that a Walrasian auctioneer announces any prices, and not just the equilibrium prices. In the same vein, here the announced states and probabilities may or not be the ones which will eventually emerge in a market equilibrium.
14. To simplify notation, and without loss of generality, we set from now on $e = \{1, ..., D\}$, by allowing some probabilities $\pi_i = 0$.
15. That is $\exists i, j \in \Psi(y) - S$, $i \neq j$, $\pi_i > 0$ and $\pi_j > 0$.
16. Among the studies of this issue are Weisbrod (1964), Krutilla (1967), Cichetti and Freeman (1971), Schmalensee (1972), Arrow and Fisher (1974), Henry (1974) and Bohm (1975).
17. If $b_0 > 0$, there are benefits to conservation in the first period, so that $c_0 = 1$, that is, we conserve in the first period. We concentrate on the interesting case of $b_0 < 0$, when the only incentive to conserve in period 1 is the possibility of a positive return in period 2.
18. An important simplifying assumption in this example is the linearity of payoffs in the level of preservation. Fisher and Krutilla (1985) discuss the role of linearity.
19. Pindyck (1991) considers a similar example in the case of irreversible investment decisions, and shows that the option value of delaying an investment decision to take advantage of information that will become available in the future can be computed using the formula used in finance for valuing an option to buy a stock. See also Dixit (1992).
20. This is because, as seen in the previous subsection, it is obtained by maximizing the expected value of benefits.
21. The model is different from most other models in which option values have been studied. It is an infinite-horizon stochastic dynamic optimization model in which the maximand is the expected present value of utility and future preferences evolve stochastically.
22. Several other mathematicians and economists, such as Hernstein and Milnor (1953) developed related axioms.
23. The space of lotteries L is the space of all measurable and essentially bounded real valued functions on the line.
24. A^c denotes the complement of the set A.
25. Chichilnisky (1996f). See Machina (1982) for an alternative analysis to the von Neumann–Morgenstern treatment of decisionmaking under uncertainty. Machina does not provide an axiomatic treatment.
26. Large probability events are events with probability close to one. Probabilities of events are bounded above by the number one.
27. The space of bounded sequences l_∞ is a Banach space with the norm $\| z \| = \sup_{s \in z} | z_s |$.
28. The impact of traders' expectations on risks was studied by Green, Grandmont, Kurz, and others.

REFERENCES

Arrow, K.J. (1953), 'The role of securities in the optimal allocation of risk-bearing', in French: *Econometrie: Proceedings of the Colloque sur les Fondements et Applications de la Théorie du Risque en Econometrie*, Centre National de la Recherche Scientifique, Paris. English Translation: *Review of Economic Studies* (1964), **31**, 91–6.

Arrow, K.J. and Anthony C. Fisher (1974), 'Environmental preservation, uncertainty and irreversibility', *Quarterly Journal of Economics*, **88**, 312–19.

Arrow, Kenneth J. and Robert C. Lind (1970), 'Uncertainty and the evaluation of public investments', *American Economic Review*, **53**, 941–73.

Barrett, S. (1990), *Environmental Directorate*, Paris: OECD.

Beltratti, A., G. Chichilnisky and G. Heal (1998), 'Uncertain future preferences and Conservation', in G. Chichilnisky, G.M. Heal, and S. Vercelli (eds) *Sustainability: Dynamics and Uncertainty*. Boston: Kluwer, 1998.

Bohm, Peter (1975), 'Option demand and consumer's surplus: comment', *American Economic Review*, **65**, 733–6.

Bromley, D. (1992), *Making the Commons Work*, San Francisco: ICS Press.

Cass, D., G. Chichilnisky and H. Wu (1996), 'Individual risks and mutual insurance', *Econometrica*, **64** (2), 333–41.

Chichilnisky, G. (1992), 'Existence of a general equilibrium with price uncertainty', Discussion Paper No. 651, Columbia University, December in G. Chichilnisky (ed.), *Markets, Information and Uncertainty: Essays in Honor of Kenneth Arrow*, New York: Cambridge University Press, forthcoming.

Chichilnisky, G. (1994a), 'North–South trade and the global environment', *American Economic Review*, **84** (4), 851–74.

Chichilnisky, G. (1994b), 'The taxation of carbon emissions in industrial and developing countries', proceedings of a Conference on the Economics of Climate Change, OECD, Paris, June 1993, in Tom Jones (ed.), *The Economics of Climate Change*, Paris: OECD.

Chichilnisky, G. (1995), 'Catastrophe bundles can hedge unknown risks', *Bests' Review*.

Chichilnisky, G. (1996a), 'An axiomatic approach to sustainable development', *Social Choice and Welfare*, **13**, 321–257.

Chichilnisky, G. (1996b), 'The economic value of the Earth's resources', Invited Perspectives article, *Trends in Ecology and Evolution* (TREE), 135–40.

Chichilnisky, G. (1996c), 'Financial innovation in property catastrophe reinsurance: the convergence of insurance and capital markets', *Risk Financing Newsletter*, **13** (2), 1–7.

Chichilnisky, G. (1996d), 'The future of global reinsurance', *Global Reinsurance*.

Chichilnisky, G. (1996e), 'Markets with endogenous uncertainty: theory and policy', Invited Plenary Session of FUR VII, International Conference on the Foundations and Applications of Utility, Risk and Decision Theory, Norwegian School of Management, Sandvika, 1 July 1994, *Theory and Decision*, **41**, 99–131.

Chichilnisky, G. (1996f), 'Updating von Neumann axioms for catastrophic risks', Invited Presentation, Fields Institute for Mathematical Sciences, Toronto, Canada, 9–11 June.

Chichilnisky, G. (1996g) 'Development and Global Finance: the Case for an International Bank on Environmental Settlements', ODS Discussion Paper No 10, United Nations Development Program, New York.

Chichilnisky, G. (1996h), 'An axiomatic approach to sustainable development', *Social Choice and Welfare*, **13**, 231–57.

Chichilnisky, G. (1997a), 'Markets with endogenous uncertainty', Keynote address, Australasian Meetings of the Econometric Society, University of Melbourne, Australia, July.

Chichilnisky, G. (1997b), 'A radical shift in risk management', *Contingencies*, Fall.

Chichilnisky, G. (1997c), 'What is Sustainable Development?' *Land Economics*, **73** (4), 467–91.

Chichilnisky, G., J. Dutta and G.M. Heal (1991), 'Price uncertainty and derivative securities in general equilibrium', Discussion Paper No. 574, Columbia University, revised 1993 and 1996.

Chichilnisky, G., F. Hahn and G.M. Heal (1992a), 'Price uncertainty and incomplete markets', Working paper, Columbia Business School in G. Chichilnisky (ed.), *Markets, Information and Uncertainty: Essays in Honor of Kenneth Arrow*, New York: Cambridge University Press, forthcoming.

Chichilnisky, G. and G.M. Heal (1992a), 'Financial markets with unknown risks', in G. Chichilnisky, G. Heal and S.Vercelli (eds), *Sustainability: Dynamics and Uncertainty*, Boston: Kluwer, 1998.

Chichilnisky, G. and G.M. Heal (1992b), 'Global environmental risks', *Journal of Economic Perspectives*, **7**, (4), Fall, 65–86.

Chichilnisky, G. and G.M. Heal (1993), 'Competitive equilibrium in Sobolev spaces without bounds on short sales', *Journal of Economic Theory*, **59** (2), 369–84.

Chichilnisky, G. and G.M. Heal (1994), 'Who should abate carbon emission? An international perspective', *Economics Letters*, Spring, 443–9.

Chichilnisky, G. and G.M. Heal (1995), 'Markets wih tradeable CO_2 emission quotas: principles and practice', Economics Department Paper No. 153, OECD, Paris, OECD/GD (95) 9.

Chichilnisky, G. and G.M. Heal (1996), 'On the existence and topological structure of the pseudoequilibrium manifold', *Journal of Mathematical Economics*, **26**, 171–86.

Chichilnisky, G. and G.M. Heal (1998), 'The future of global reinsurance', *Journal of Portfolio Management*, Summer.

Chichilnisky, G., G. Heal and D. Starrett (1993) 'International Markets with Emissions Rights of Greenhouse Gases: Equity and Efficiency' Center for Economic Policy Research, Publication No. 81, Stanford University, Autumn.

Chichilnisky, G., G.M. Heal, P. Streufert and J. Swinkels (1992), 'Believing in multiple equilibrium', Working Paper, Stanford Institute for Theoretical Economics.

Chichilnisky, G., G.M. Heal and D. Tsomocos (1996), 'Option values with endogenous uncertainty: ESOP's, MBO's and asset backed securities', *Economic Letters*, **48** (3–4), 379–88.

Chichilnisky, G. and H.-M. Wu (1992), 'Financial innovation and endogenous default in incomplete asset markets', Stanford Institute for Theoretical Economics Technical Report No. 50, Stanford, California, USA.

Cichetti, Charles J. and A. Myric Freeman III (1971), 'Option demand and consumer surplus: further comment', *Quarterly Journal of Economics*, **85**, 528–39.

Cline, W. (1992), *The Economics of Global Warming*, Washington, DC: Institute for International Economics.
Dasgupta, P. and G.M. Heal (1979), *Economic Theory and Exhaustible Resources*, Cambridge: Cambridge University Press.
Dixit, A.K. (1992), 'Investment and hysteresis', *Journal of Economic Perspectives*, **6** (1),107–32.
Fisher, A. and J. Krutilla (1985), 'Economics of nature preservation' in Alan V. Kneese and James L. Sweeney (eds), *Handbook of Natural Resource and Energy Economics*, Amsterdam, New York and Oxford: North-Holland.
Fudenberg, D. B. Holstrom and P. Milgrom (1990), 'Short-term contracts and long-term agency relationships', *Journal of Economic Theory*, **51** (1), June.
Hahn, F. (1973), *The Concept of Equilibrium in Economics*, Inaugural Lecture, Churchill College, Cambridge University, Cambridge: Cambridge University Press.
Hahn, F. (1991), 'A remark on incomplete market equilibrium', in G. Chichilnisky (ed.), *Markets, Information and Uncertainty: Essays in Honor of K.J. Arrow*, New York: Cambridge University Press, forthcoming.
Hahn, F. (1992), 'Making general equilibrium theory more plausible', *Second Annual Arrow Lecture Series*, Department of Economics, Stanford University, 5 and 6 May.
Hartwick, J.M. 'Declining biodiversity and risk-adjusted NNP', Working Paper, Department of Economics, Queens University.
Heal, G. (1984), 'Interaction of the economy and climate: a framework for policy design under uncertainty', in Kerry Smith and Ann Dryden White (eds), *Advances in Applied Macroeconomics*, Greenwich, CT: JAI Press, pp. 151–8.
Heal, G.M. (1990), 'Economy and climate: a preliminary framework for microeconomic analysis' in Richard Just and Nancy Bockstael (eds), *Commodity and Natural Resource Policies in Agricultural Systems*, New York and Oxford: North-Holland, pp. 855–80.
Heal, G.M. (1992), 'Optimal resource use', in Alan V. Kneese and James L. Sweeney (eds), *Handbook of Natural Resource and Energy Economics*, Amsterdam, New York and Oxford: North-Holland.
Heal, G.M. (1993), 'International negotiations on emission controls', *Economic Development and Structural Change*, **1** (1), 1–19.
Heal, G.M. (1997), 'Valuing the Future', Book manuscript, Columbia University Program on Information and Resources.
Henry, Claude (1974), 'Option values in the economics of irreplaceable assets', *Review of Economic Studies, Symposium on the Economics of Exhaustible Resources*, 89–104.
Hernstein, N. and J. Milnor (1953), 'An axiomatic approach to measurable utility', *Econometrica*, **21**, 291–7.
Institute for Policy Studies (1996), 'Changing the earth's climate for business',. Report by the International Trade Information Service (US) and Halifax Initiative (Canada), Institute for Policy Studies, 733,15th Street NW, Suite 1020, Washington DC 20005.
Krutilla, John V. (1967), 'Conservation reconsidered', *American Economic Review*, **57**, 777–86.
Kurz, M. (1974), 'The Kesten–Stigum model and the treatment of uncertainty in equilibrium theory', in M.C. Balch, D. McFadden and S. Wu (eds), *Essays in Economic Behavior Under Uncertainty*, Amsterdam: North-Holland.

Lowenstein, G. and J. Elster (eds) (1992), *Choice Over Time*, New York: Russell Sage Foundation.

Lowenstein, G. and R. Thaler (1989), 'Intertemporal choice', *Journal of Economic Perspectives*, **3** (4), 181–93.

Machina, M. (1982), 'Expected utility analysis without the independent axiom', *Econometrica,* **50**, 277–323.

Malinvaud, Edmond (1972), 'The allocation of individual risks in large markets', *Journal of Economic Theory*, **4**, 312–28.

Malinvaud, Edmond (1973), 'Markets for an exchange economy with individual risks', *Econometrica*, **41**, 383–410.

Manne, A. and R. Richels (1992), *Buying Greenhouse Insurance*, Cambridge, MA: MIT Press.

Pindyck, Robert S. (1991), 'Irreversibility, uncertainty and investment', *Journal of Economic Literature*, **XXIX** (3), 1110–48.

Schmalensee, Richard (1972), 'Option demand and consumers' surplus: valuing price changes under uncertainty', *American Economic Review*, **62**, 813–24.

Solow, Robert M. (1992), 'Sustainability: an economist's perspective', eighteenth J. Seward Johnson Lecture, Woods Hole Oceanographic Institution, Woods Hole, MA.

Stein, J. (1992a), 'Cobwebs, rational expectations and futures markets', *Review of Economics and Statistics*, **LXXIV** (1), 127–34.

Stein, J. (1992b), 'Price discovery processes', *Economic Record*, 34–45.

Weisbrod, Burton A. (1964), 'Collective consumption services of individual consumption goods', *Quarterly Journal of Economics*, **77**, 71–7.

Wilson, Robert (1992), 'Strategic analysis of auctions', in R. Auman and S. Hart (eds), *Handbook of Game Theory*, Vol. 1, Amsterdam: Elsevier Science Publishers, BV.