

Global Environmental Risks

Graciela Chichilnisky and Geoffrey Heal

Managing climate risk is not a new activity. 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 and buying insurance, farmers have managed climate risk for many centuries.

Today's concerns about global climate change break new ground in two ways. One is the global scope of the potential changes considered, such as changes in the atmosphere of the planet. The second is that these changes appear to be driven by human activity, which has now reached levels at which it can affect the global climate. Climate has always been unpredictable, but the inclusion of these two new elements has extended this uncertainty both qualitatively and quantitatively. Classical formulations of uncertainty in economics no longer suffice. 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 a socio-economic perspective. The floods of 1993 in the United States and Bangladesh have reminded us of the profound vulnerability of human settlement to climate. Climatologists link these to El Niño, 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. In addition, these emissions can be driven by economic activity and by policy measures: hence the risks faced are endogenous.

■ *Graciela Chichilnisky is Professor of Economics, Columbia University, and Geoffrey Heal is Professor and Senior Vice Dean, Columbia Business School, both in New York City, New York.*

There are two standard ways in which 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.¹ The probabilities of states in an Arrow–Debreu framework may be subjective and an agent’s subjective probabilities may be altered by learning. However the actual frequency of incidence of harmful events cannot be altered by agents: there is no scope for mitigation in the second sense of improving odds. The same is true in the classical models of insurance (see for example Malinvaud, 1972, 1973), 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 U.S. flood victims. Insurance is a major economic activity, involving both the insurance industry and large parts of the securities industry. 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?

This paper is about these and related questions. In attempting to answer them, we deal with many different aspects of the theory of risk-bearing. Four key issues recur in our analysis.

The first issue concerns the difficulty in assessing risks. Most climate-related risks are difficult to quantify. Indeed, in a classical statistical sense the probabilities describing them are unknowable. We may never be able to observe enough experiments to approximate 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 CO₂ emissions in this way.²

¹Kurz (1974) set out a framework for endogenous uncertainty: recently this has been an active field, e.g., Kurz (1990, forthcoming 1993), Chichilnisky and Wu (1991), Chichilnisky, Dutta and Heal (1992), Chichilnisky (1992), Chichilnisky, Hahn and Heal (1992).

²In this respect there may be a difference between the various aspects of climate risk. There is historical data on the relation between atmospheric CO₂ and climate from tree ring and ice core studies. With ozone depletion the phenomenon is so new that such data is not available.

The second issue is endogeneity of risks. The risks that we face are affected by our actions.³

The third issue is correlation of risks. Climate changes will affect large numbers of people in the same way. A rise in sea level, for example, will 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 risks. We thus have to ask what types of markets work best when risks are connected and collective.

Irreversibility is the final issue. In this area, many major economic decisions and their consequences are likely to be irreversible. Climate changes, the melting of ice caps, desertification, species extinction, are all processes not reversible, or at least not on relevant time scales.

In summary, we are dealing with risks that are *poorly understood, endogenous, collective and irreversible*. In policy terms, the nature and extent of uncertainty about global climate change implies that society's position will be dominated by two questions: What cost is it worth incurring to reduce the poorly-understood risk of climate change, or to improve our understanding of the risk? 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?

Risk-Allocation in a General Equilibrium Framework

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 framework 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." As examples, think of "water in a drought" and "water in a flood." With a complete set of markets for state-contingent commodities, the first theorem of welfare economics holds for economies under uncertainty: a Pareto efficient allocation of resources can be attained by a competitive economy with uncertainty about exogenous variables.

³Endogeneity of risks leads to moral hazard when risks to agents depend on their actions, which cannot be observed by the insurers and will be influenced by the insurance offered. This leads to arguments for coinsurance (see for example Stiglitz, 1983). In the present context such moral hazard problems are not central to the analysis: asymmetric information is not a characteristic of climate risks. Endogenous uncertainty is more general than moral hazard.

Arrow (1953) showed that efficiency can in fact be attained by using a mixture of securities markets and markets for non-contingent 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 (Chichilnisky and Heal, 1992). For example, if agents face individual risk (i.e., risks whose probabilities vary 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} . 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 incidence 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).

The insurance approach is at a disadvantage when risks are correlated. When large numbers 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 (Chichilnisky and Heal, 1992; Cass, Chichilnisky and Wu, 1991).

When risks are allocated by trading state-contingent 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 it 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. In this case an alternative approach is the opening of new markets (Chichilnisky and Heal, 1992, and the discussion below).

In sum: the Arrow–Debreu approach to risk allocation via state-contingent markets is in principle universally applicable. However, it is cumbersome and perhaps unrealistic in the case of risks with individual components. Insurance markets are more manageable, but leave uncovered collective or correlated risks such as the risk induced by ignorance of the true frequency distribution of harmful events. So 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 is precisely the approach that we develop below. Although new to the economics literature, it is by no means new in practice; we argue below that some of the oldest risk-bearing institutions recorded, agricultural cooperatives, have exactly this structure.

Ignorance as Collective Risk

Consider an economy in which agents face risks whose relative frequencies they know that 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. What market structure would suffice to assure efficient allocations in this situation? There are widely differing opinions about these impacts, on which there is inadequate information.

Chichilnisky and Heal (1992) formalize this type of situation in a general equilibrium model. Each agent faces the risk of being in one of several states (e.g. healthy or sick, productive or unproductive). No one knows what will be 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 percent chance that 90 percent of the population will be harmed by global warming, a 25 percent chance that 50 percent 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, we have two levels of uncertainty. The first level of uncertainty is collective: what is the distribution of agents who are harmed in the economy? Will 90 percent be harmed, or only 30 percent? This is a question about the aggregate incidence of the phenomenon over the population as a whole. The second level of uncertainty is individual: it is uncertainty about whether a given agent is harmed or not by climate change. It devolves about questions such as: given that 90 percent of the population will be harmed, will a particular agent be harmed or not?

In our examples 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 secondly, uncertainty about whether any given person or community will be affected.

Our ignorance of scientific processes (like the relation between ozone depletion and skin cancer or between CO₂ emission 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. When this is resolved we will still not know who is damaged and who is not, but we will at least know the frequency describing this. Once frequencies are known, actuarial calculations can be conducted and the problem is insurable.

We propose an institutional structure which uses two types of financial instruments which are tailored to these two aspects of the problem. These can lead to efficient allocation in the face of such risks. We follow a framework established by Cass, Chichilnisky and Wu (1991) and Chichilnisky and Heal (1992).

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 that 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 U.K. workers' associations and friendly societies. These involved agreements between a group of workers that if one were sick and unable to work, that worker would be compensated by the others who were not harmed. 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, and on the basis of this decide which mutual insurance contract to apply.

Having dealt with individual risks by mutual insurance, we still face collective risks. We need *Arrow securities* to deal with these collective risks induced by uncertainty about the overall distribution of adverse effects. Arrow securities are usually defined as securities that pay one dollar if and only if a particular state of the world occurs. Here they 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, 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 will help to make this framework concrete. 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, i.e., 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 1’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 .

To reach an efficient allocation of risks we need two Arrow securities. One, S_A , pays \$1 million if and only if the probability distribution over states of the climate is A . The other, S_B , pays \$1 million 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 a probability distribution 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 $\Delta_{1,2}^\alpha$ to country 2 if the state of the climate is α , and country 2 makes a transfer $\Delta_{2,1}^\beta$ to country 1 if the climate state is β . These transfers satisfy $0.1\Delta_{1,2}^\alpha + 0.9\Delta_{2,1}^\beta = 0$ 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 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 securities that pay off contingent on that frequency. For the individual risks that remain, it is more practical to use mutual 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 individual 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 very 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 and correlated risks.

An Institutional Framework

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 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.

Next this framework involves introducing 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. 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. An additional example is provided by the Chicago Board of Trade, which is in the process of introducing securities called "catastrophe futures" whose payoffs depend inter alia on the incidence of tropical storms in the United States, as measured by the Insurance Service Organization's index. This is a particular example of the instruments initially proposed in Chichilnisky and Heal (1992).

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. In fact, 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. In fact, one of the largest banks in Italy, the Monte dei Paschi di Siena, was founded to play this role in 1473. They have insured against weather risks

since then,⁴ 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 also hedge 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.

Trading Risks

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

Suppose for example the United States believes it most likely that there will be little climate change, and the European Community believes otherwise. Then through the market for securities whose payoffs depend on which description of climate change is correct, the United States will naturally sell insurance to the E.C. The United States 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 E.C. 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 CO₂ abatement policies introduced by others. These issues are discussed in detail in Heal (1993) 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.

Note that trading risks is different from the trading of emission permits, as discussed in the articles by Poterba and Weyant in this volume. The recognition of uncertainty suggests the need to consider state-contingent emission permits,

⁴They also supported part of this research and provided access to their records.

where the state is defined in terms of the frequency of climate-change related events. Permits to emit CO₂ would for example be contingent on the incidence of indicators of climate change. Such contingent emission permits could play the role of securities whose payoffs depend on the collective risk.

In the context of emission permits, it is worth noting the public good nature of the climate. 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: it is a privately produced public good. 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, a somewhat unlikely contingency. Equalization of the marginal costs of emission abatement across countries is often taken as justification of the superiority of tradeable permits as a method for controlling emissions. This point is developed in Chichilnisky (1993) and Chichilnisky and Heal (1993). More generally, a key issue is that an efficient allocation of a public good such as unchanged climate must be supported by a Lindahl equilibrium and not a competitive equilibrium. In general competitive markets for tradeable emission permits may not decentralize Pareto efficient allocations of abatement.

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 CO₂-reduction protocols. Careful judgment is crucial, as all of these involve very considerable costs, as indicated by Weyant's paper in this symposium, Cline (1992) and others. Here we shall summarize one approach to this problem, based on Heal (1984, 1990). This is a model that examines the extent to which the consumption of fossil fuels should be curtailed because it increases the probability of a change in climate. The model in Heal (1984) is based on three assumptions. First, the atmosphere may be in one of two states, one favorable to economic activity and one unfavorable (there is a possibility of a future climate catastrophe). Secondly, the atmosphere moves stochastically from the favorable state to the unfavorable, and once there remains there forever, so that atmospheric change is irreversible. Finally, the probability of a transition from the favorable to the unfavorable state is endogenous and increases with the level of cumulative emissions from the use of fossil fuels. The probability that this change will definitely occur at some date is less than one, so there is a positive probability of the change never occurring.

Fossil fuels, capital equipment and the atmosphere are used to produce output, which may be consumed or reinvested 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 favorable or an unfavorable state. Initially the atmosphere is in the favorable state but may change stochastically to the unfavorable state, and once in this state will remain there forever. The source of emissions which affects the probability of climate change 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 CO₂ by the use of fossil fuels.

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; there is also a national income identity. The overall problem solved involves maximizing expected utility subject to the resource and national income constraints, where the expectation is over the process governing climate change.

Heal (1984) characterizes optimal paths of consumption, capital accumulation and use of fossil fuel for this problem. 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 and uses this to isolate the key parameters in determining the optimal rate of use of fossil fuels. The introduction of atmospheric impact makes a fundamental 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 in the no-atmospheric-impact case. The difference depends on the degree of risk aversion and on the parameters of the probability distribution relating cumulative emission to climate change.

The behavior 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 (1984) 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. Because it is an intertemporal model, it provides this characterization in terms of the internalization of probabilistic intergenerational externalities. 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 are providing, 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 and important 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, CO₂ abatement policies only make sense if coordinated internationally (Poterba's paper in this symposium; Barrett, 1990; Carraro and Siniscalco, 1991; Heal, 1993b).

The fact that different countries need not agree on policy choices even if they agree on the scientific evaluation of the problem could clearly make such agreement difficult to obtain. Different countries' positions with respect to measures to restrict greenhouse gas emissions depend on their discount rates and degrees of risk aversion. The United States, for example, has been against global abatement agreements, while Germany has been in favor. This fits with the conventional wisdom that the financial and industrial community in the United States has both a higher discount rate and a lower degree of risk aversion (greater willingness to take risks) than that 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.

Different perceptions of the risk involved do not however preclude efficient solutions. Economics is about differences in preferences leading to 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, as discussed earlier.

Option Values and Irreversibility

In valuing environmental resources such as current climate conditions, biodiversity, or complex ecological systems, the irreversibility of decisions and events can be central. A key 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 the option of reconsidering the decision.

A concept related to option value is that of non-use value 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 non-use value, seem to overlap. Many goods which exemplify one also exemplify the other. At the same time, there are no doubt differences. Non-use 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 non-use valuations there lurks an option value: many non-use valuations stem from an unstated belief that a use value may emerge.


In this section we review 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. We provide an outline of the argument in this case and illustrate the fact 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.

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).⁵ The central issue is that benefits 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 follow-

⁵Amongst the studies of this issue are Weisbrod (1964), Krutilla (1967), Cichetti and Freeman (1971), Schmalensee (1972), Arrow and Fisher (1974), Henry (1974a, b) and Bohm (1975).

Figure 1

Uncertain Benefits from Preservation in Period 2

ing example (from Dasgupta and Heal, 1979) illustrates the key point in a simple framework.

We 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. We contrast this with the reversible case, in which we never conserve in the first period and there is no option value.

Consider two dates, the present $t = 0$ and the future $t = 1$. We have one unit of an environmental asset. The benefit from preserving this now at time $t = 0$ is b_0 . At time $t = 1$ in the future 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 both currently at time $t = 0$ and in the future at $t = 1$. 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 of course at $t = 1$ conserve more than was conserved at $t = 0$. Our options at $t = 1$ are therefore constrained by the decision made at $t = 0$. This situation is summarized in Figure 1.

We shall compare the case already described where the decision made at time $t = 0$ is irreversible with a hypothetical alternative case in which this decision can in fact be reversed. In this case the decision made at time $t = 0$ no longer constrains the options available at time $t = 1$. We look at this alternative case first, as it is simpler and provides a benchmark. Let c_0 be the amount of the resource conserved now at time $t = 0$, and c_1 and c_2 be the amounts conserved in the future at time $t = 1$ in states 1 and 2 respectively. The

expected benefit from development (assuming a zero discount rate) is $b_0c_0 + pb_1c_1 + (1-p)b_2c_2$. We have to choose conservation levels c_0 , c_1 and c_2 to maximize (the expected benefit). We focus on the interesting case in which there is currently no benefit to preservation,⁶ $b_0 < 0$, nor is there any benefit in state 1 in the future, $b_1 < 0$. However, there is the future possibility of state 2 in which there are positive benefits from preservation, $b_2 > 0$. If decisions are reversible, we preserve nothing at time $t = 0$, i.e., we set $c_0 = 0$. Then at time $t = 1$, we conserve nothing in state 1 and everything in state 2, i.e., 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, $0 \leq c_1$, $c_2 \leq c_0 \leq 1$. In particular, if everything is destroyed in the first period, then we have no options in the second. What policies now maximize (the expected benefit)? 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 in that period we will conserve everything left to us by our earlier decision, that is we will always set $c_2 = c_0$. If however the state in the future is unfavorable to conservation, then we will conserve nothing and set $c_1 = 0$.⁷

It is optimal to conserve in the first period if and only if there is a positive expected payoff to future conservation, given that we choose optimally later, i.e., we set $c_1 = 0$ and $c_2 = c_0$. Contrast this with the decision in the reversible case, in which we never conserve and always chose $c_0 = 0$. These two decisions are different if the expected payoff to conservation in the first period is positive.⁸ 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.

Option Values

Note that the existence of an "option value" does not depend on risk aversion, as we assumed throughout the previous subsection that the maxi-

⁶If $b_0 > 0$, there are benefits to conservation in the first period, so that $c_0 = 1$, i.e., 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.

⁷Noting that $C_1 = 0$ and $C_2 = C_0$ in the irreversible case, the expression for expected benefit at the top of this page reduces to $\{b_0 + (1-p)b_2\}c_0$. As the expression is linear in C_0 , the initial conservation level is positive if and only if the derivative with respect to C_0 is positive, i.e. $\{b_0 + (1-p)b_2\} > 0$. The inequality 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 unfavorable 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.

⁸An 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.

mand is the expected value of benefits.⁹ The key issues here are: first, the irreversibility of the decision; second, the fact that delaying a decision can let one take advantage of better information, and third, that on average there will be benefits from conservation in the first period, provided that we choose optimally later (as shown in note 7).

There are important practical implications of 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 (i.e., preserving the atmospheric environment), it may be appropriate to credit preservation (preventing climate change) with an option value. This could 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.

The most thorough study of the costs and benefits of reducing climate change is 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 preventing climate change. This means that Cline's study 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. It could for example be decisive in the endorsement of a global carbon tax.

Uncertainty about Future Generations

There are several ways of generalizing or refining the concept of option value. 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 (Heal, 1993a).

It may, however, be a statement that future generations could have different preferences from us, and might value environmental assets differently.

⁹Pindyck (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).

Because they might value them more, we 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 very different from ours (see also Chichilnisky, 1993b; Beltratti, Chichilnisky and Heal, forthcoming 1993). This seems close to the concept of option value set out above, and indeed it is, though there are some differences that are revealing. We next study this problem, drawing heavily on results in Beltratti, Chichilnisky and Heal (1992), and using a highly simplified version of that model.¹⁰


We shall illustrate the conclusion that uncertainty about future preferences alone is not sufficient to produce an "option value" case 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 (increases and decreases in the intensity of preferences are equally likely) does not generate a case for leaving more to the future in case they value the resource more highly than us. Uncertainty makes a case for conservation only when the expected return to postponement of consumption is positive.

Consider a two period world where there is a fixed total stock of a natural resource to be consumed in the two periods. The initial stock is known to be s_0 . The amounts consumed in the first and second periods are c_1 and c_2 respectively: these must obviously satisfy $c_1 + c_2 = s_0$. As the stock is irreplaceable, anything that is consumed in the first period is not available in the second, so that consumption here is an irreversible depletion of the stock. The utility from period one consumption is $u(c_1)$, which is an increasing strictly concave function. The utility from second period consumption is unknown: it may be either $(1 + \alpha)u(c_2)$ with probability p or $(1 - \beta)u(c_2)$ with probability $(1 - p)$. Here $0 < \alpha, \beta < 1$. So there is a probability p that the utility derived from future consumption will be "scaled up" by a factor α , and a probability $(1 - p)$ that it will be "scaled down" by a factor β .

Consider first as a benchmark the case in which there will be no change in preferences, so that we just have to pick c_1 and c_2 to maximize $u(c_1) + u(c_2)$. Figure 2 shows this situation: the length of the horizontal axis is s_0 , the initial stock of the resource. Consumption in the first period c_1 is measured to the right from the left-hand origin and consumption in the second period to the left from the right-hand origin. Marginal utility in each period is plotted, and the optimal levels of c_1 and c_2 are those at which the marginal utility curves cross. In the absence of discounting, and with utility functions the same in each period, these curves will of course be symmetric, as shown in Figure 2.

¹⁰The full model is radically 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.

Figure 2

Uncertainty about Future Preferences

Now suppose that there is uncertainty about preferences in the second period—think of this as uncertainty about the preferences of a future generation. Also simplify matters by assuming that $\beta = 0$, so that the only possible change in preferences is a “scaling up” of the utility of consumption. This corresponds to the case that we mentioned at the start of this section, namely the possibility of an increase in the appreciation that people have for the resource.¹¹ The expected marginal utility curve for second period consumption is now above the certain second period marginal utility curve, as shown in Figure 2.

The optimal first period consumption level is now lower than before, as a result of the possibility of a shift in future preferences towards the natural resource. Generally one can show that the amount of this reduction depends on the probability distribution governing the change in preferences, the discount

¹¹Now we have to choose c_1 and c_2 to maximize the expectation of utility, which as $\beta = 0$ is $u(c_1) + p(1 + \alpha)u(c_2) + (1 - p)u(c_2)$. The solution to this requires that the marginal utility of consumption in the first period equals the expected marginal utility in the second, i.e.,

$$\frac{\partial u}{\partial c_1} = \{p\alpha + 1\} \frac{\partial u}{\partial c_2}$$

Clearly $\{p\alpha + 1\} > 1$, so that the period two marginal utility curve is moved up by uncertainty, and period one consumption falls, as in Figure 2.

rate and the degree of risk aversion (Beltratti, Chichilnisky and Heal, 1992). Does this reduction in period one consumption reflect an "option value" in the sense of the previous section?

To understand this, we have to consider the more general case that we posed initially.¹² Suppose first that there is uncertainty about future preferences but on average we expect no net change, in the sense that increases and decreases balance out in probability terms. Then the second period marginal utility curve is unaltered by the uncertainty about future preferences. In this case uncertainty about future preferences will not lead to a reduction in present consumption. Indeed it will lead to no changes in any consumption levels, even if agents are strictly risk averse in the sense that their utility functions are strictly concave.¹³

If on the other hand there is an expectation of an increase in the utility of consumption in the second period, then the period two expected marginal utility curve will lie above that under certainty and consequently the optimal period one consumption level will be lower than under certainty. Conversely, if there is an expectation of a decrease in the utility of consumption in the second period, then there will be a decrease in the period one consumption relative to its level under certainty.

In conclusion, uncertainty about future preferences alone is not sufficient to produce an "option value" case 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 with increases and decreases equally likely does not generate a case for leaving more to the future in case their preferences for the resource are stronger than ours. Uncertainty makes a case for conservation only when the expected return to postponement of consumption is positive.

One final observation and an indication of possible future research: the Beltratti, Chichilnisky and Heal model of option values summarized here is one in which utility is derived only from the flow of consumption of the environmental resource. In practice the stock may also enter as an argument of the utility function. For example, we value the current climate as an asset: we value the current stock of species or of rainforests. In this case there are likely to be two qualitatively different types of optimal consumption path, depending on

¹²Equalization of expected marginal utilities in the two periods requires in this case that

$$\frac{\partial u}{\partial c_1} = \{p(1 + \alpha) + (1 - p)(1 - \beta)\} \frac{\partial u}{\partial c_2}.$$

Now the period two expected marginal utility curve may in general lie above or below the first period curve: it will be exactly the same as the first period curve, i.e., the curve in the absence of uncertainty, if and only if the coefficient of $\partial u / \partial c_2$ is one, i.e., $\{p(1 + \alpha) + (1 - p)(1 - \beta)\} = 1$. This condition means that the expected shift in period two utility is zero.

¹³Technically, this is because uncertainty here is about the utility functions, and the maximand is linear in these. Society is not risk-averse about utility levels, even though it is about consumption levels.

the size of the initial stock of the resource. If this is large, the optimal path will involve the maintenance of positive stocks of the resource indefinitely: if it is small, then the entire stock will eventually be consumed. The size of the critical initial stock at which this qualitative change in solution occurs, will depend on preferences. In this case, it is possible that a change in uncertainty about future preferences will tip the economy from one optimal consumption regime to another. Such a phenomenon would make a dramatic difference to the computation of the option value.

Conclusions and Open Questions

While some of the foundations are in place for an understanding of the economics of global environmental risks, there are certain aspects that require more attention. From a policy perspective, the endogeneity of the risk faced is important. The purpose of many recommended policies is precisely to change the risks that we face. An example is the global carbon tax proposed by the OECD and reviewed in Chichilnisky (1993a). Only recently has there been systematic study of the welfare economics of markets with endogenous risks, and many questions remain open.

The policy implications of many of the issues that we have reviewed need further attention. How important, for example, are the option values associated with global changes? It would be interesting to see some of the large models reviewed by Weyant in this symposium address this question. We need to study further the institutional implications of trading environmental risks on financial markets such as contingent claims markets and mutual insurance markets. The public good aspect of the global climate has not been analyzed adequately. The structure of emissions permit markets needed for efficiency is still not understood. Our analyses of option values are still based on very limited assumptions, in spite of a long and distinguished literature in this area. Finally, economists clearly need more exchange of information with physical scientists: it is fair to say that the two areas fail to make any significant contact in a field that needs the skills of both.

The prospect of climate change induced by human activity faces societies with demanding issues in risk management and risk assessment: at the same time, it faces economics with challenges and opportunities. The challenge is to develop intellectual tools, communicate them to society at large and prove that they can add value to the analysis of a complex and possibly fundamental problem.

■ *Both Chichilnisky and Heal acknowledge financial support from the NSF (grants 92-16028 and 91-0460 respectively), the Fondazione Mattei and the Monte dei Paschi di Siena. Both authors appreciate valuable comments from Ken Arrow, Andrea Beltratti, Mordecai Kurz, Richard Sandor, and Bob Solow. We are particularly grateful to the editors of this issue, Dick Schmalensee, Carl Shapiro and Timothy Taylor, for valuable comments and suggestions which go way beyond the call of duty.*

References

- Arrow, Kenneth J.**, "The Role of Securities in an Optimal Allocation of Risk-bearing." *Econometrie, Proceedings of the Colloque sur les Fondements et Applications de la Theorie du Risque en Econometrie*, Centre National de la Recherche Scientifique, Paris, 1953. English translation in *Review of Economic Studies*, April 1964, 31, 91-96.
- Arrow, Kenneth J.**, and **Anthony C. Fisher**, "Environmental Preservation, Uncertainty and Irreversibility," *Quarterly Journal of Economics*, May 1974, 88:2, 312-19.
- Arrow, Kenneth J.**, and **Robert C. Lind**, "Uncertainty and the Evaluation of Public Investments Decisions." *American Economic Review*, June 1970, 60:3, 364-78.
- Barrett, Scott**, "The Paradox of International Environmental Agreements." Paris: Environment Directorate, OECD, 1990.
- Beltratti, Andrea, Graciela Chichilnisky, and Geoffrey M. Heal**, "Irreversibility, Uncertain Future Preferences and Option Values," Stanford Institute for Theoretical Economics, Technical Report No. 53, Stanford University, 1992.
- Beltratti, Andrea, Graciela Chichilnisky and Geoffrey M. Heal**, "Sustainable Growth and the Green Golden Rule," working paper, Columbia Business School, 1993. In Goldin, Ian and L. Alan Winters, *Economic Approaches to Sustainable Development*, OECD, Paris, forthcoming 1993.
- Bohm, Peter**, "Option Demand and Consumer's Surplus: Comment," *American Economic Review*, September 1975, 65:4, 733-36.
- Bromley, Daniel W.**, *Making the Commons Work*. San Francisco: ICS Press, 1992.
- Carraro, Carlo, and Domenico Siniscalco**, "Strategies for International Protection of the Environment," working paper, Fondazione Eni Enrico Mattei, Milano, 1991.
- Cass, David, Graciela Chichilnisky, and Wu Ho-Mou**, "Individual Risks and Mutual Insurance." CARESS Working Paper #91-27, Department of Economics, University of Pennsylvania, 1991.
- Chichilnisky, Graciela, and Wu Ho-Mou**, "Financial Innovation and Default in Incomplete Markets," working paper, Columbia University, Department of Economics, 1991.
- Chichilnisky, Graciela**, "Existence and Optimality of Equilibrium with Endogenous Uncertainty," working paper, Columbia University, Department of Economics, 1992.
- Chichilnisky, Graciela**, "The Abatement of Carbon Emissions in Industrial and Developing Countries." Paper presented at the International Conference on the Economics of Climate Change, OECD/IEA, Paris, June 1993a.
- Chichilnisky, Graciela**, "What is Sustainable Development?" Stanford Institute for Theoretical Economics Technical Report, 1993b.
- Chichilnisky, Graciela, Jayasri Dutta, and Geoffrey M. Heal**, "Price Uncertainty and Derivative Securities in a General Equilibrium Model," First Boston Working Paper, Columbia Business School, 1992.
- Chichilnisky, Graciela, Frank H. Hahn, and Geoffrey M. Heal**, "Price Uncertainty and Incomplete Markets," First Boston Working Paper, Columbia Business School, 1992. In Chichilnisky, Graciela, *The Formulation of Economic Theory: Essays in Honor of Kenneth Arrow*. Cambridge, U.K.: Cambridge University Press, forthcoming.
- Chichilnisky, Graciela, and Geoffrey M. Heal**, "Financial markets for unknown risks." First Boston Working Paper Series, Columbia Business School, 1992.
- Chichilnisky, Graciela, and Geoffrey M. Heal**, "Who Should Abate Carbon Emissions? An International Viewpoint." First Boston Working Paper, Columbia Business School, 1993.
- Cicchetti, Charles J., and A. Myric Freeman, III**, "Option Demand and Consumer Surplus: Further Comment," *Quarterly Journal of Economics*, August 1971, 85:3, 528-39.
- Cline, William R.**, *The Economics of Global Warming*. Washington, D.C.: Institute for International Economics, 1992.
- Dasgupta, Partha S., and Geoffrey M. Heal**, *Economic Theory and Exhaustible Resources*. Cambridge: Cambridge Economic Handbooks, Cambridge University Press, 1979.
- Dixit, Avinash K.**, "Investment and hysteresis," *Journal of Economic Perspectives*, Winter 1992, 6:1, 107-32.
- Fisher, Anthony C., and John V. Krutilla**, "Economics of Nature Preservation." In Kneese, Alan V., and James L. Sweeney, ed., *Handbook of Natural Resource and Energy Economics*. Amsterdam, New York and Oxford: North Holland, 1985, 165-89.
- Hartwick, John M.**, "Decline in Biodiversity and Risk-adjusted NNP," working paper, Queens University, Department of Economics, 1992.
- Heal, Geoffrey M.**, "Interaction Between

Economy and Climate: A Framework for Policy Design under Uncertainty." In Smith, V. Kerry, and Ann Dryden White, eds., *Advances in Applied Microeconomics*. Greenwich: J.A.I. Press, 1984, 151–58.

Heal, Geoffrey M., "Economy and Climate: A Preliminary Framework for Microeconomic Analysis." In Just, Richard E., and Nancy Bockstael, eds., *Commodity and Resource Policies in Agricultural Systems*, Berlin, Heidelberg and New York: Springer, 1990, 196–212.

Heal, Geoffrey M., "Optimal Resource Use" In Kneese, Alan V., and James L. Sweeney, ed., *Handbook of Natural Resource and Energy Economics*. Amsterdam, New York and Oxford: North Holland, 1993a, 855–80.

Heal, Geoffrey M., "International Negotiations on Emission Control," *Economic Development and Structural Change*, 1993b, 1:1, 1–19.

Henry, Claude, "Option Values in the Economics of Irreplaceable Assets," *Review of Economic Studies*, *Symposium on the Economics of Exhaustible Resources*, 1974a, 89–104.

Henry, Claude, "Investment Decisions Under Uncertainty: 'The Irreversibility Effect'," *American Economic Review*, December 1974b, 64:6, 1005–12.

Krutilla, John V., "Conservation Reconsidered," *American Economic Review*, September 1967, 57, 777–86.

Kurz, Mordecai, "The Kesten-Stigum Model and the Treatment of Uncertainty in Equilibrium Theory." In Balch, M. S., D. L. McFadden, and S. Y. Wu, *Essays on Economic Behavior Under Uncertainty*. Amsterdam: North Holland, 1974, 389–99.

Kurz, Mordecai, "On the Structure and Diversity of Rational Beliefs," Technical Report No. 39, Stanford Institute for Theoretical Eco-

nomics, Stanford University, 1990.

Kurz, Mordecai, "General Equilibrium with Endogenous Uncertainty," working paper, Department of Economics, Stanford University, forthcoming 1993. In Chichilnisky, Graciela, ed., *The Formulation of Economic Theory: Essays in Honor of Kenneth Arrow*. Cambridge, U.K.: Cambridge University Press, forthcoming.

Malinvaud, Edmond, "The Allocation of Individual Risk in Large Markets," *Journal of Economic Theory*, April 1972, 4:2, 312–28.

Malinvaud, Edmond, "Markets for an Exchange Economy with Individual Risk," *Econometrica*, May 1973, 41:3, 383–410.

Manne, Alan S., and Richard G. Richels, *Buying Greenhouse Insurance*. Cambridge: MIT Press, 1992.

Pindyck, Robert S., "Irreversibility, Uncertainty and Investment," *Journal of Economic Literature*, September 1991, XXIX: 3, 1110–48, 1991.

Schmalensee, Richard, "Option Demand and Consumer's Surplus: Valuing Price Changes under Uncertainty," *American Economic Review*, December 1972, 62:5, 813–24.

Solow, Robert M., *Sustainability: an Economist's Perspective*, the Eighteenth J. Seward Johnson Lecture, Woods Hole Oceanographic Institution, Woods Hole, Massachusetts, 1992.

Stiglitz, Joseph E., "Risk Incentives and Insurance: The Pure Theory of Moral Hazard," *The Geneva Papers*, January 1983, 8:26, 4–32.

Weisbrod, Burton A., "Collective-Consumption Services of Individual-Consumption Goods," *Quarterly Journal of Economics*, August 1964, 78, 471–77.