1 A DECISION-MAKING FRAMEWORK

Key points

The central policy issue facing the Review can be simply stated: what extent of global mitigation, with Australia playing its proportionate part, provides the greatest excess of gains from reduced risks of climate change over costs of mitigation?

The mitigation costs are experienced through conventional economic processes and can be measured through formal economic modelling.

Only some of the benefits of mitigation are experienced through conventional market processes (types 1 and 2) and only one is amenable to modelling (Type 1). Others take the form of insurance against severe and potentially catastrophic outcomes (Type 3), and still others the avoidance of environmental and social costs, which are not amenable to conventional measurement (Type 4).

The challenge is to make sure that important, immeasurable effects are brought to account.

The long time frames involved create a special challenge, requiring us to measure how we value the welfare of future generations relative to our own.

This chapter puts forward a framework for looking at these issues.

How do we assess whether Australian mitigation action is justified? Would the substantial costs of mitigation be exceeded by avoided costs of climate change? What degree of mitigation would lead to the largest net benefits?

These turn out to be immensely complex questions. The answers depend on our judgments about the prospects for effective international mitigation. They depend on the efficiency of measures to achieve reductions in greenhouse gas emissions, including supporting measures that affect the market response to the mitigation regime, and therefore the costs of achieving various levels of abatement. They depend on the efficiency of supporting measures to share the costs of mitigation across the Australian community, and on the international distribution of the mitigation burden. They depend on the options for and costs of adaptation. These decisions need to be taken under conditions of uncertainty and risk.

The answers also depend on our ability to measure accurately the conventional economic effects of climate change, and the likely reduction in those effects due to mitigation. Not all of the effects on output and consumption through market processes are amenable to precise quantification. Our conclusions depend on our ability to form sound judgments about the magnitude of any changes that are excluded from attempts at formal measurement because adequate information is not available at this time. The answers depend fundamentally on the approach taken to decision making under conditions of risk and uncertainty, and in particular, on the insurance value that is placed on avoiding the possibility of large negative outcomes.

The answers depend also on the value we place on outcomes not related to consumption of goods and services, but on Australians' valuation of environmental amenity in many dimensions. These assessments are affected by how we view the inter-relationship between these and other non-material values with conventional consumption in determining welfare.

The answers are affected by the relative value that is placed on the welfare of people living in the future relative to the welfare of those living at present.

This chapter introduces an approach to decision making to openly deal with these immensely complex and difficult issues. This allows people who are uneasy or unhappy about the conclusions to understand or take issue with the underlying premises and logic.

We are seeking to assist community choice on the extent of mitigation that provides the greatest excess of gains from reduced climate change over costs of mitigation. The complexity of the influences on that choice makes simplicity especially challenging and particularly important. Here, even more than in other areas of public policy choice, focus on the central underlying issues is essential if we are to reach conclusions through a transparent process, open to challenge, as a basis for long-term community support, policy continuity and stability.

Climate change mitigation decisions in 2008, and for the foreseeable future, are made under conditions of great uncertainty. There is great uncertainty about the climatic outcomes of varying concentrations of greenhouse gases; about the impact of various climate outcomes; and about the costs and effectiveness of adapting to climate change. There is uncertainty about the costs of various degrees of mitigation in Australia; about the extent to which the international community will make effective commitments to mitigation; and about the relationship of global to Australian mitigation efforts.

Under conditions of such uncertainty, it is sensible to ask whether it would be better to delay decisions while information is gathered and analysed. However, it is as much a decision to do nothing, or to delay action, as it is to decide to take early action. The issue is whether delay would be a good decision.

When global warming first became a major international public policy issue nearly two decades ago, it may have been good policy to take modest and low-cost steps on mitigation, while investing heavily in improving the information base for later decisions.

In 2008, the costs of delay—in the probabilistic terms that frame a good decision under conditions of uncertainty—are high. The work of the Review has contributed to changing international perceptions on the rate at which emissions will grow over the next several decades under business as usual. Australia and the world are running towards high risks of dangerous climate change at a more rapid rate than was previously understood. The opportunity costs of delaying decisions are high.

Australia and its partners in the international community will, for good reasons, make historic and fateful decisions about their approaches to climate change mitigation in the three years ahead. They will do this on the basis of currently available information and analysis, however sound or weak that may be.

The sceptical economist—and the Review counts itself within this tradition—insists on equally rigorous evaluation for a decision to delay as for a decision to take action now.

The Review's approach to the important questions about mitigation policy starts with scientific assessment of the costs of climate change to Australia and Australians. We have to be able to compare the costs of climate change without mitigation, and with varying degrees of effective mitigation and adaptation effort. These costs include indirect costs through effects on other countries, to the extent that these feed back into impacts on Australia, or in themselves are valued by Australians. The scientific assessments are highly uncertain, and their impacts on human activity and welfare even more so. We have no alternative to making decisions on complex issues of valuation under conditions of great uncertainty.

1.1 The costs of mitigation

The increase in greenhouse gas emissions is a product of the advances in science, technology and economic organisation that have transformed humanity as well as its natural context over the last two centuries. In the history of life on earth, and even of human life, we are talking about an almost infinitesimally short period of extraordinary dynamism.

A modern acceleration in rates of human-induced greenhouse gas emissions is the source of contemporary concerns about climate change.

Economic development over the past two centuries has taken most of humanity—but certainly not all—from lives that were insecure, ignorant and short, to personal health and security, material comfort and knowledge unknown to the elites of the wealthiest and most powerful societies in earlier times.

In the first millennium after the life of Jesus Christ, global economic output increased hardly at all—by only one sixth. All of the small increase was contributed by population growth, and none by increased production per person. By contrast, output increased 300-fold in the second millennium, with population increasing 22 times and per capita production 13 times. Most of the extraordinary expansion took place towards the end of the period. From 1820 until the end of the 20th century, per capita output increased more than eight times and population more than five times (Maddison 2001).

In most of its first two centuries, the cornucopia of modern economic growth was located in a small number of countries, in Western Europe and its overseas offshoots in North America and Oceania, and in Japan. In the third quarter of the 20th century it extended into a number of relatively small economies in East Asia.

A new era began in the fourth quarter of the last century, with the rapid extension of the beneficent processes of modern economic development into the heartland of the populous countries of Asia, including China, India and Indonesia. From this has emerged what can be described as the Platinum Age of global economic growth in the early 21st century (Garnaut & Huang 2007).[1] Incomes are growing rapidly in a large proportion of the developing world. In the absence of a major dislocation of established trends, this is likely to continue for a considerable period. There will be a greater absolute increase in annual human output and consumption in the first two decades of the 21st century than was generated in the whole previous history of our species. Similarly strong growth in output can be expected in the next following decade to 2030.

Increasingly through the 21st century, the expansion of production will be associated with rising output per person, rather than increase in population. In all of the economically successful countries, higher incomes, the increased survival rates of children and the expansion of education and choice for women are leading to declining rates of population increase. Before the end of the 21st century, a continuation of these processes should lead to stabilisation (by about 2080), and then, at least for a while, a gradual decline in global human population. By that time, nearly three billion will have been added to the global population.

The era of modern economic growth has been intimately linked to rapid expansion in the use of fossil fuels. This is returning to the atmosphere a part of the carbon that was sequestered naturally over billions of years, through a process that created the conditions necessary for the emergence of human life on earth. While the share of carbon returned to the atmosphere is small relative to the stock, it is large enough to throw the equilibrium of heat trapping in the atmosphere out of balance.

The amount of fossil fuel in the earth's crust, in the forms of petroleum, natural gas, coal, tar sands and shale, is finite. However, the amount is so large that its limits are of no practical importance for climate change policies.

However, there is a much tighter engineering limit to the availability for human use of fossil fuels: the point at which the energy used to extract the resources would be greater than their energy content.

Tighter still is the economic limit: the availability of fossil fuels in forms and locations that can be extracted for human use at costs below the prices of oil, gas and coal in global markets. There is debate on whether the economic limits will constrain global economic growth in the period immediately ahead or in the foreseeable future. The limit will be reached much earlier for liquid petroleum than for natural gas, and for gas much earlier than for coal.

It was once common for economists to see constraints on the availability of natural resources and in particular fossil fuels as placing limits on modern economic growth (Malthus 1798; Jevons 1865). The success of technological improvement and economic processes in easing supposed constraints in the first centuries of modern economic growth established confidence that these constraints could be overcome in ways that allowed global economic growth to continue.

Rapid growth from the early 1950s to the early 1970s, and extraordinary Japanese growth at the end of that period, rekindled old concerns about resource constraints on growth.

Fossil fuel resource availability was one element in the cautions of the Club of Rome, and their prophecy about limits to growth in the early 1970s (Club of Rome 1972). The extraordinary growth in demand for fossil fuels in the early years of the Platinum Age—and the immense and unexpected increases in prices that accompanied it—have rekindled interest in the issue. Will the supply of fossil fuels slow down the growth in greenhouse gas emissions enough to do the mitigation task?

It is clear from the present state of knowledge—as it was not to earlier generations—that it would be possible for humanity to break the link between economic growth and combustion of fossil fuels. This would make it possible for the world economy to adjust to the approach of economically relevant limits to fossil fuel availability, without bringing the increase in human consumption of goods and services to an end.

For the time being, the pervasive and rapidly growing use of fossil hydrocarbons in economic activity is a matter of economic optimisation and not of technological necessity. If the human species avoids some catastrophic truncation of the triumphs of modern economic development, it will need to pursue a transition out of reliance on fossil fuels—and it will succeed in doing so.

The constraints on the economic availability of fossil fuels will aid the climate change mitigation process. But the Review's analysis suggests that in the time available, the reduction in use of fossil fuels, associated with scarcity and high prices, will be nowhere near enough to avoid high risks of dangerous climate change.

To the extent that mitigation is effective, reduced demand for petroleum and other fossil fuels associated with effective mitigation would reduce the global price of these resources, improve the terms of trade of importing countries, and probably have favourable effects on global economic growth. This would be an offset for some countries against the cost of mitigation.

The beneficiaries of lower fossil fuel prices would not include Australia, whose terms of trade rise with high global energy prices. Lower export prices for resources hurt producers in resource-based industries, and the beneficiaries of government revenue generated from these industries. But they also tend to lower interest rates and the exchange rate, and increase incomes, in some rural manufacturing and service industries, and for many households.

Adjusting to limits on the use of fossil fuels required to mitigate climate change would be less costly than adjusting to naturally imposed economic constraints on the availability of fossil fuels. This is because sequestration through physical processes (geosequestration) or biological processes (biosequestration) can ease the mitigation task but cannot ease natural constraints on fossil fuel supply.

However, mitigation needs to be imposed through political processes. Such decisions in single countries are hard enough. Achieving mitigation outcomes through cooperation of many sovereign entities, each with an incentive to shift the cost of adjustment to other countries, is more challenging.

A dramatic transformation in humanity's use of fossil-fuel-based energy would be necessary sooner or later to sustain and to extend modern standards of living. It will be required sooner if the world is to hold the risks of climate change to acceptable levels. The costs incurred in making an early adjustment will bring forward, and reduce for future times, the costs of the inevitable adjustment away from fossil fuels. How much sooner and at what extra cost are central questions before the Review.

The costs of mitigation depend on the extent to which, and the time over which, reductions in emissions are achieved. Costs depend on the efficiency of the chosen policy instruments. There are cost advantages in having a single price on emissions as the main instrument of policy, supported by measures to correct market failures in utilisation of the commercial opportunities created by the price on emissions.

If mitigation is approached through an efficient set of policies, its costs are determined by the extent and the rate of emissions reductions to be achieved. These, in turn, are determined by the ambitions of a global effort to which Australia has subscribed, and by what Australia is prepared to do in the context of global action.

The costs of mitigation can be calculated for various levels and rates of reductions in emissions. Each level and rate of Australian mitigation can be related to a global mitigation outcome. The global mitigation outcome will define a benefit to Australia in terms of reduced risks of climate change. The benefits of reduced risks of climate change to Australia can be identified. The costs and benefits of mitigation can then be compared. The policy task in setting Australian mitigation objectives, therefore, begins with identification of the costs and benefits (in reduced risks of loss from climate change) for various mitigation ambitions.

The higher the market prices of petroleum, coal and natural gas, the lower the costs of mitigation will be. The costs of business as usual, compared with the costs of using alternative, low-emissions technologies, will be higher. The historically high fossil fuel prices make this of current interest.

The more ambitious the extent and speed of reductions in emissions, the higher the costs of mitigation will be. The costs of mitigation will be lower the more efficient the instruments chosen to give effect to policy.

An economically efficient approach to mitigation would generate a rising carbon price over time, imposing increasingly strong pressure for adjustment out of high-emissions technologies, and increasingly strong incentives for sequestration. For a given abatement task, emissions costs will be lowest if the emissions price rises at the interest rate, which will lead to optimal timing in investment in the mitigation effort.

The challenge is to allocate efficiently over time access to a limited global capacity to absorb additional greenhouse gases without unacceptably high risks of dangerous climate change. The allocation problem is familiar as one of optimal depletion of a finite resource. This frames the economics of the timing of the mitigation effort, and suggests the relevance of the 'Hotelling curve' to the price curve for the right to emit (Hotelling 1931).

The annual costs of mitigation are likely to rise for some time, as a rising carbon price forces deeper abatement. While the price would be expected to continue to rise over time at the interest rate, the cost to the economy would not rise at that rate. At some point, the tendency for costs to rise would be moderated and eventually reversed by improvements in low-carbon technologies.

At some time in the future—no later and perhaps much earlier than the time when economic constraints on the use of fossil fuels would be forcing structural change comparable with what had been achieved for mitigation purposes—the incremental costs of mitigation will become negative. The sunk costs of technological improvement and structural change associated with mitigation will avoid the need for investments to accommodate the constraints on availability of fossil fuels.

Above all else, the cost of mitigation in Australia, and not only the benefits in avoided climate change, will be shaped by the nature of the global mitigation effort. An effective global effort would open a wide range of opportunities for trade in mitigation responsibilities, assigning greater reductions in emissions to countries in which it can be achieved at lowest cost. A global effort would increase and distribute more efficiently and equitably the world's investment in new technologies to develop lower-emissions paths to consumption and production. It would obviate the need for special policy measures to avoid carbon leakage—the shift of emissions-intensive industries from high-mitigation to low-mitigation countries—a policy requirement that is likely to distort both domestic economic efficiency and political integrity.

1.2 Risk and uncertainty

Climate change policy requires us to come to grips with both risk and uncertainty. Keynes (1921) and Knight (1921) drew a distinction between the two that is still useful today.

Risk relates to an event that can be placed on a known probability distribution. When we toss a coin, we do not know whether or not we will see a head. If we toss the coin enough times, it will fall as a head about half of the time.

In many spheres of human life, an activity has similarities with others that have been repeated many times, so that participants have a reasonable idea of the odds. A piece of surgery with some risk of death and short-term investments in financial markets have some similar properties to the toss of a coin. No new piece of surgery, and no new investment, is exactly the same as any other. But there have been enough similar events for players to feel that they can form judgments with some confidence about the probabilities.

There is uncertainty when an event is of a kind that has no close precedents, or too few for a probability distribution of outcomes to be defined, or where an event is too far from understood events for related experience to be helpful in foreseeing possible outcomes. Humans are often required to form judgments about events that are unique, or so unusual that analysis based on secure knowledge and experience is an absent or weak guide. Columbus sailing west in search of China, or Oxley heading west along the rivers of Australia in search of an inland sea, are historically important examples (Figure 1.1).

Figure 1.1 The risk–uncertainty spectrum

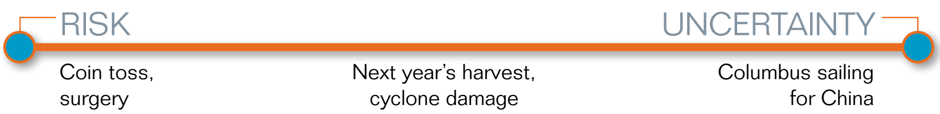

The 18th century British philosopher Bayes has given his name to a well-developed approach to decisions under uncertainty. Bayesian decision theory encourages us to treat decisions under uncertainty as if we were taking a risk (Raiffa 1968; Raiffa & Schlaifer 1961). We will make the best possible decisions under uncertainty if we force those who are best placed to know to define subjective probabilities that they would place on various outcomes, and work through the implications of those assessments as if they were probability distributions based on experience (Figure 1.2). These subjective probability distributions can then be updated on the basis of experience.

Figure 1.2 A probability distribution

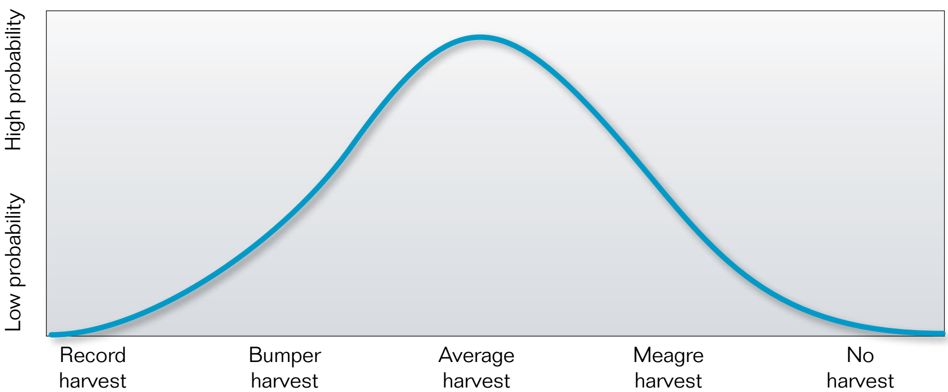

While the distinction between risk and uncertainty is analytically helpful, it does not distinguish discrete and separate phenomena. Rather, risk and uncertainty are the extreme ends of a single spectrum. Next year's harvest can be assessed as a risk on the basis of past experience but carries an element of uncertainty, because it is affected by various climatic parameters that are not at all predictable from experience or with current knowledge. The risk of a cyclone hitting a tropical city can be assessed using data on past occurrence of cyclones, although aspects of the potential damage are uncertain.

If it is correct to treat a subjectively formed assessment of a probability distribution as if it were drawn from a distribution based on repeated experience, what is the difference between risk and uncertainty? Perceptions of the probability distribution formed under conditions of uncertainty are more likely to change materially with a small number of new observations or amount of experience or further analysis.

The Review's work on climate change has made some contact with risk, more with uncertainty, and most of all with the wide territory between them. The mainstream science, embodied in the work of the Intergovernmental Panel on Climate Change (IPCC), sometimes discusses possible outcomes in terms of fairly precise probability distributions, yet describes its assessments in terms of 'uncertainties'. This suggests that they are applying Bayesian approaches to decisions under uncertainty. The decision framework is rarely made explicit, and sometimes is not clear.

The climate models on which the assessments are based are themselves diverse. They provide numerous observations on possibilities out of their diversity; in addition, each generates numerous results from repeated experiments. These are the senses in which the IPCC science draws from probability distributions. There are many points at which judgment rather than experience informs the model relationships. The resulting conclusions are therefore located somewhere on the uncertainty side of the middle of the risk–uncertainty spectrum.

1.3 Four types of benefits from mitigation

Three types of benefit from avoided climate change—that is, mitigation—can be measured in monetary values, as a change in the value of output or consumption. The fourth type of benefit of mitigation requires a different measurement unit.

The four types of climate change impacts, which in part can be mitigated, are illustrated in Figure 1.3.

1.3.1 Type 1: currently measurable market impacts

The first type of benefit from mitigation comprises currently measurable market impacts of climate change, which are avoided by mitigation. The measurement can be brought together through a computable general equilibrium economic model. The starting point for assessment is the estimation of climate impacts based on the means of the relevant probability distributions for these outcomes. These effects are typically measured as an impact on GDP or consumption, with monetary values as the unit of measurement.

Figure 1.3 The four types of climate change impacts

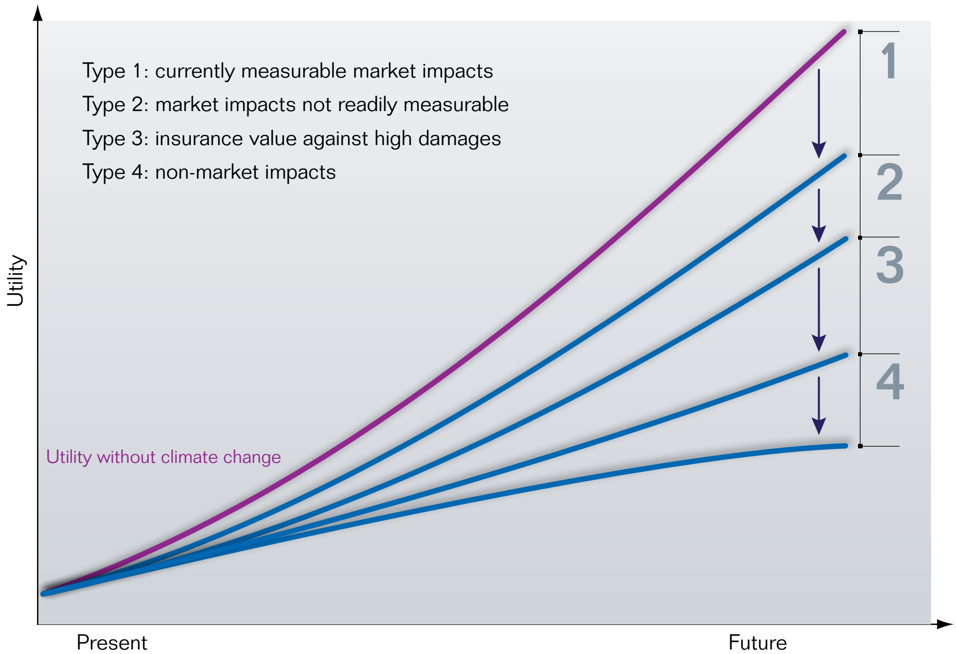

1.3.2 Type 2: market impacts not readily measurable

The second type of benefit of mitigation comprises market impacts similar in nature to the first, but not amenable to measurement in the current state of knowledge. For the Review, these impacts were not defined precisely enough in time for the modelling, but are, in principle, amenable to quantitative analysis. We seek to use what we know of these effects roughly to compare their possible size with the impacts that have been subject to formal modelling. As with the effects that are subject to modelling, we focus on the medians of the probability distributions of possible outcomes. We are drawing these judgments from views of the impacts that are closer to the uncertainty than the risk end of the risk–uncertainty spectrum. There is no reason to expect our estimates of these impacts to be too low rather than too high, but they are more likely than the estimates of the first type of benefit to be subject to large adjustments, in one direction or another, with the advance of knowledge. Examples from the Review include the impact of climate change on the tourism industry. As with the first type of benefit, the estimation of these effects would be in monetary values of GDP or consumption.

1.3.3 Type 3: insurance value against high damages

The third type of benefit of mitigation is the insurance value that it provides. On many impacts, there is large asymmetry between human evaluation of outcomes that are much more benign or much more damaging than the median. Humans tend to be risk averse when the outcomes include the possibility of large loss. Some of the possible outcomes at the bad end of the probability distribution would be thought by many people to be catastrophic. In such cases, mitigation has additional insurance value. What would we be prepared to pay to avoid a small probability of a highly damaging or possibly catastrophic outcome? It is probably more than we would be prepared to pay to avoid the certain prospect of the catastrophic event's contribution to the median outcome. Uncertainty strongly plays into this category of benefits, as the probability of extreme or catastrophic climate impacts is not known from experience, and must instead be based on expert judgment.

It is not a new idea for governments to make large financial commitments for insurance against low-probability, high-impact events. Defence absorbs several percentage points of GDP per annum, most of it on insurance against genuinely low-probability developments.

The possibility of outcomes that most people would consider to be catastrophic makes this a particularly important element of the assessment. Weitzman (2008a) sees it as the main element:

> [T]he burden of proof in the economics of climate change is presumptively upon whoever wants to model or conceptualize the expected present discounted utility of feasible trajectories under greenhouse warming without considering that structural uncertainty might matter more than discounting or pure risk per se. Such a middle-of-the-distribution modeler should be prepared to explain why the bad fat tail of the … distribution is not empirically relevant and does not play a very significant role in the analysis.

1.3.4 Type 4: non-market impacts

The fourth type of benefit is more difficult to conceptualise, and quantify. Let us evoke an old tradition in economics, and talk about units of utility. (We could just as well call it welfare, if we removed ourselves from a modern interpretation in terms of social security for disadvantaged people.) That unit, 'utility', is used in economics to represent welfare in a demonetised fashion.

The focus of Australian policy making, as in other countries, is on maximising the welfare—or utility—of Australians. We can think of a utility function as rising with Australian consumption of goods and services, and also with a number of non-monetary services, such as environmental amenity (which itself may have a number of components), longevity, health, and welfare of people in other countries. If the comparisons of costs and benefits of the first three types of benefits from mitigation suggest a particular outcome, and it is clear from inspection that inclusion of the fourth might lead elsewhere, it is necessary explicitly to compare the monetary with the non-monetary effects on welfare of a particular position. This could in principle be done by forcing a monetary value onto particular non-monetary outcomes. An alternative is to leave the comparison of the monetary and non-monetary outcomes until after the possible conflict between them is known.

Examples of such non-market impacts include Australians' valuation of environmental amenity. They include the value that Australians place on the integrity of the Great Barrier Reef, Ningaloo, Kakadu and other features of the

Australian and international landscapes, on known shorelines, on genetic diversity and on the survival of species. They include the value that Australians place on long-established communities and social structures built around particular patterns of climate, or the use of green urban gardens and playing fields for recreation. To include such elements in an Australian utility function is not to place intrinsic value on environmental conservation, as some argue that we should. It is only necessary to accept that many Australians value such things, including as options for their offspring and future generations, and would be prepared to sacrifice some consumption of goods and services to retain them. Australians also value the avoidance of poverty and trauma in other countries, as demonstrated in their continued support of public and private international development assistance and disaster relief.

Non-market elements in a utility function, and the level of utility when the function includes market and non-market elements, are in their nature difficult to measure. Any politically derived mitigation policy decision will implicitly value them alongside changes that can be measured more easily.

Traditional welfare economics contains a few important insights into the roles that non-market factors, such as environmental amenity and concern for the welfare of others, might play in determining utility in a world of climate change.

The non-market values are likely to be 'superior goods', in that the relative value that people assign to them rises with incomes. In the late 21st century, when the average purchasing power of incomes over material goods and services is likely to be several times the present level, much higher relative value will be placed on any truncation of the natural estate that has occurred in the intervening years.

It is likely that at higher incomes, the price elasticity of substitution between conventional consumption and access to such non-market values as environmental amenity and concern for others' welfare will be low, and much lower than today. Near subsistence levels of consumption, few people would willingly sacrifice goods and services for greater environmental amenity, or for improved development prospects of others at home or abroad. In the likely material affluence of the late 21st century, many more people are likely to trade substantial amounts of access to material consumption for small amounts of improved values of services not available through market processes.

An extremely low rate of substitution between non-market services and conventional consumption of goods and services at high incomes, in the presence of large impacts from climate change, would challenge the proposition that continuing economic growth would necessarily lead to higher average utility in the distant future.

One implication is that the utility of Australians under policies that allocate priority to such non-market values as the services provided by the natural environment, and development in poor countries, is likely to be much higher than the application of today's preference systems at today's material consumption levels would suggest.

1.4 How effective adaptation reduces the costs of climate change

Some of the costs of climate change can be diminished by the adaptive behaviour of individuals and firms, and by policies that support productive adaptation.

Effective adaptation requires a strong applied science base; good markets for reallocation of resources, goods and services; and capital for investment in defensive structures and new productive capacity that is more suitable to the new environment.

All of these capacities are more abundant in developed than low-income developing countries. For the latter, the impact of climate change is likely to be undiluted and more severe. Australia's location in an immediate region of vulnerable developing countries will make some of its neighbours' challenges our own. Investment in adaptive responses in the arc of island countries and regions from Timor-Leste through eastern Indonesia, Papua New Guinea and the South Pacific is likely to become an important component of Australia's own cost-reducing adaptation to climate change.

The modifying impact of adaptation is illustrated by Australian agriculture. Better and earlier knowledge will allow farmers to make timely decisions on whether new money should continue to be invested in locations that seem to be severely damaged by climate change, or whether it is better to find new livelihoods in less challenging locations. Investment in plant and animal genetics may be able to diminish the loss of productivity associated with higher temperatures and changing rainfall patterns. Investment in water retention or storage will sometimes be an economically sensible response to more variable rainfall.

Hardest of all, the most effective adaptive responses in agriculture to climate change will sometimes require fundamental changes in attitudes, policies and institutions. For example, the loss in irrigated agricultural value under moderate warming and drying scenarios could be greatly reduced by shifting from established to free market allocation patterns of water allocation, so that limited water resources are directed without qualification to their most productive uses. Livestock industries in these same circumstances would suffer less if established patterns of quarantine on feed imports were to be relaxed. We can presume that change of such a fundamental kind would not be achieved without rancour and disputation over policy, and would require public policy management of exceptional dexterity and quality.

In assessing the costs and benefits of mitigation, the costs of adaptation need to be subtracted from business as usual and mitigated output and consumption. The benefits of adaptation through reduced climate change damage need to be subtracted from the gains from mitigation. The Review takes this partly into account by presuming a substantial adaptive response in assessing the costs of climate change at various levels of mitigation. For example, the presumed wheat yields are based on the expectation that planting times and new seed varieties will be developed rapidly for changing conditions.

The costs of adaptive responses will generally come early, and the benefits from reduced costs of climate change later. On the whole, the Review has only been partially able to take account of the costs of adaptation; and the assessment of reduced costs of climate change on output and consumption is incomplete—a task for future analysts.

Some of the most important adaptive responses to climate change, and the most difficult to bring to account in an analysis of optimal levels of mitigation, involve changes in attitudes and values. The city dwellers of densely populated regions of Northeast Asia have long been accustomed to life that is almost entirely separated from the natural environment. If climate change separates more and more humans from natural environmental conditions, will they simply change in their values and preferences, and learn to accept without a feeling of loss what they have never known, or knew only in distant memory? Will it matter to Australians of the future if their children do not enjoy the grassed playing fields that were once formative in what we imagined as our community culture? Could Australians learn to love living in a country and a world without many of the natural features that are now enjoyed with comfortable familiarity?

Humans adapt to changed and difficult environments when they must. There will still be joys of learning and life even in the most unhappy scenarios of future climate change. But it must be doubted that humans will have changed so much that they fail to regret what they still had in the Australian natural environment in the early 21st century. For want of reason to do otherwise, the Review will assess the value that Australians place on environmental amenity in terms of today's perspectives and preferences, manifested in a future world of greatly increased consumption of conventional goods and services.

1.5 Measuring the benefits of mitigation against the costs

To a sceptical economist, the case for action is not made simply by comparing the cost of unmitigated climate change with the cost of mitigation.

The relevant comparator is the reduction in the cost of climate change that is achieved as a result of the mitigation action. If we are evaluating Australian mitigation action, the reduction in costs of climate change that is relevant is that associated with the total global mitigation action that it enables—either by Australia, or by the set of countries that are undertaking joint action.

The benefit from mitigation is the costs of climate change avoided, after the costs and ameliorating effects of adaptation had been taken into account. Do the benefits of mitigation exceed the costs for Australians?

The costs of mitigation come earlier and are more certain. The benefits come later and are less certain. How do we compare later with earlier benefits? How do we compare more with less certain outcomes?

The costs and benefits of mitigation, in Australia and in other countries, fall on and accrue to different groups in the community. They are also felt and valued in various ways by different people. How do we weigh the relative effects on welfare of different people? In particular, what relative weight do we give to costs and benefits to the rich and to the poor? It may be that an overall assessment of whether mitigation is worthwhile will depend on the distribution of costs and benefits across the community.

The Stern Review (2007) addressed the question of whether mitigation action was justified for the world as a whole. This turns out to be an easier question than whether mitigation action is justified from the point of view of an individual country. An assessment of whether mitigation action is justified for an individual country must deal with all of the complexities that Stern addressed for the world as a whole, plus one. And that additional source of complexity is perhaps the most difficult of all.

The relevant mitigation is global. A single country's action is relevant only in its direct and indirect contribution to global mitigation. The costs of various levels of mitigation for a single country depend mainly on the extent of its own mitigation—although these costs are substantially reduced in a global agreement within which at least major economies apply similar emissions pricing regimes. The benefits depend overwhelmingly on what other countries are doing. Each country's evaluation of whether some mitigation action of its own is justified depends on its assessment of the interaction between its own decisions and those of others. Thus its own decision framework must depend on its assessment of the dynamics of complex games, among many countries. The games are framed within an awful reality, that each country has a narrow national interest in doing as little as it can, whatever others do, so long as its own action does not diminish the mitigation action that others actually take.

The global mitigation effort is the sum of the separate but inter-related mitigation decisions of individual sovereign countries. It is the sum of implicit or explicit decision processes in all countries, of the kind that we are attempting for Australia. The sum of the decision processes in many countries—democratic and authoritarian, soft and hard states, rich and poor—will determine the global mitigation effort.

1.6 A graphical representation of the benefits and costs

Let us plot our expectations of the level of national utility or welfare over time in the absence of mitigation, national or global (Figure 1.4). National utility will generally rise over time, in line with the Australian experience through its history. On the same graph, now plot expectations of welfare over time at a given level of national mitigation, which is associated with a defined degree of global mitigation.

Figure 1.4 Utility with and without mitigation

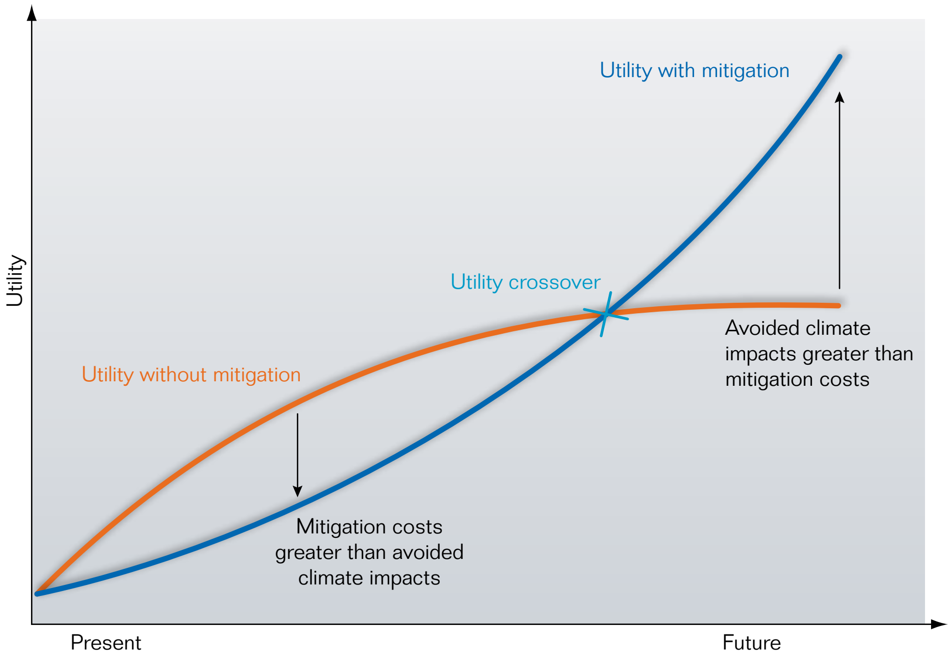

Figure 1.5 Utility under a more ambitious level of mitigation

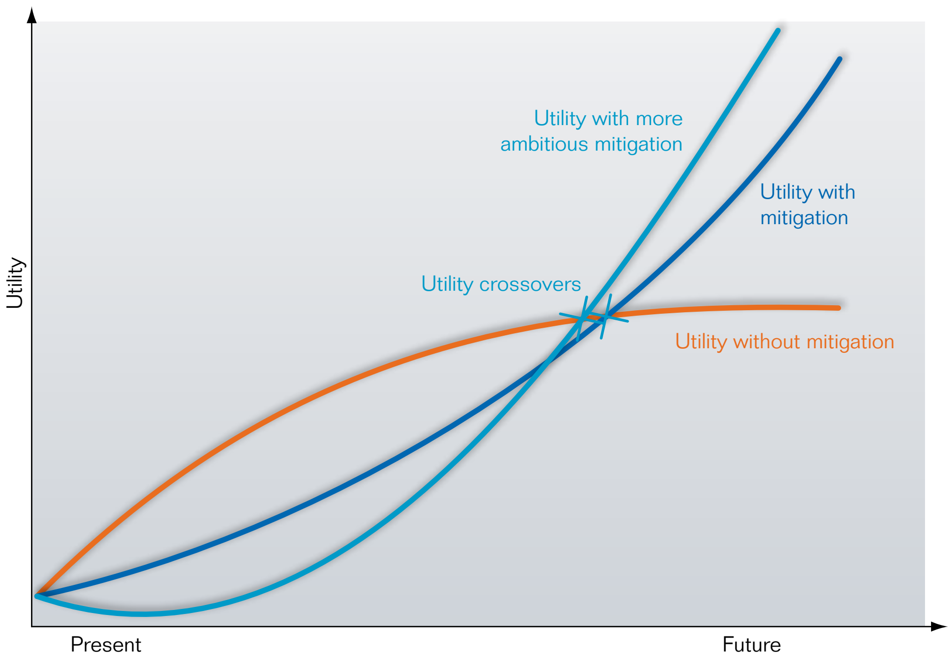

Figure 1.6 Utility with more climate change impacts taken into account

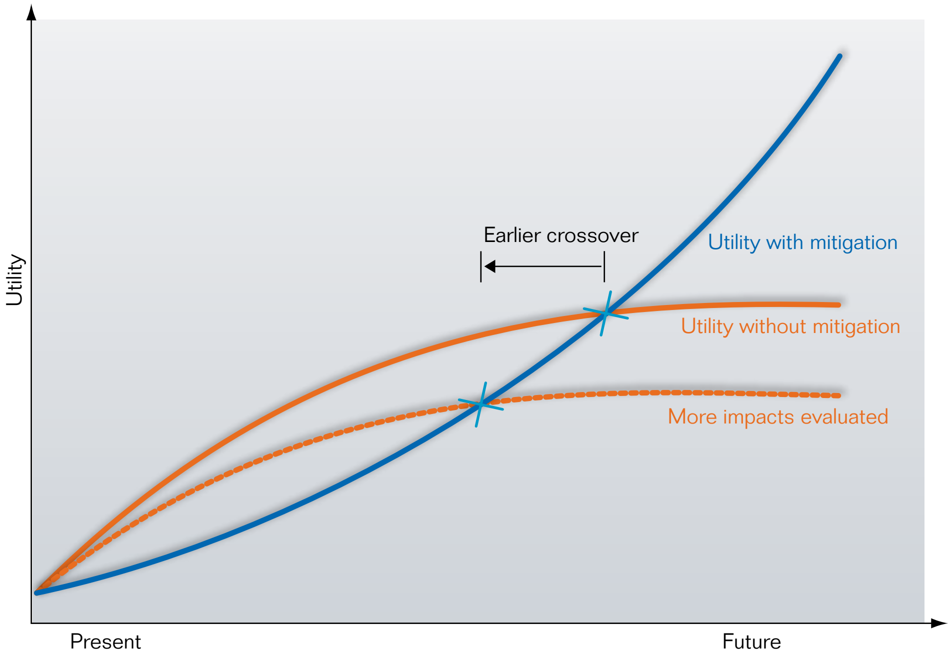

As shown in Figure 1.4, the utility curve without mitigation is above the utility curve with mitigation in the early years. Mitigation has a cost. However, in cases in which national mitigation is associated with substantial global mitigation, at some point the mitigation utility curve may rise above the utility curve in the absence of mitigation. We can call the point at which utility associated with mitigation exceeds utility in the absence of mitigation for the first time, the utility crossover point.

The two curves together describe the shape of a fish. The body of the fish covers years in which the net benefits of mitigation are negative. The area of the body of the fish represents the excess costs of mitigation in the years to the crossover point. The tail of the fish covers years in which the net benefits of mitigation are positive. The tail of the fish grows in depth and total area with time.

Figure 1.5 shows the utility curves with and without mitigation for a more ambitious level of mitigation. The fish will tend to have a fatter body and tail, as both the costs of mitigation and the benefits are increased. The relative sizes of body and tail, and locus of the crossover point, are matters for empirical analysis and judgment. Figure 1.6 shows the utility curves with a greater degree of climate change impacts being taken into account. As more costs of climate change are quantified, the utility curve for the unmitigated case shifts downward.[2] The crossover point comes earlier and the tail becomes larger relative to the body of the fish—implying higher net benefits from mitigation. These are also matters for empirical analysis. The policy question is whether the area of the body of the fish exceeds that of the tail of the fish. Future utility has to be valued at present values, so the difference in the annual levels of utility between the two curves defining the fish have to be discounted to the present at an appropriate rate. The choice of discount rate will have a major influence on the size of the body relative to the size of the tail. We turn to the discount rate after looking more closely at other influences on the sizes of the body and tail of the fish.

The cost of mitigation (augmented somewhat, and more in the low-mitigation cases, by the greater cost of investment in adaptation in the low-mitigation scenarios) will determine the depth of the fish's body.

For any chosen mitigation outcome, there will be an optimum distribution of the mitigation effort over time. For purposes of presentation, we assume that ideal mitigation policies, including ideal allocation of emissions reduction over time, are chosen for implementation.

The length of the body is the time it takes to get to the crossover point.

Beyond the crossover point, the length of the tail is determined only by any limit that society may place on the future time over which it remains concerned for the utility of Australians.

1.7 Valuing the future relative to the present

Should society place any limit on the future time over which it remains concerned for the utility of Australians? It is not obvious why it should do so.

The value of avoided irreversible effects of climate change extends forward to the point at which the life of the human species may have been extinguished by some separate influence. There is some chance of extinction at any time, at least in the contemporary human state of knowledge about weapons of mass destruction and capacity to control their use, and the low level of knowledge and capacity relevant to avoidance of the risks of the earth colliding with extra-terrestial bodies. The probability of extinction is not high in any year and perhaps in any century, but it is above zero.

So if we are to include the welfare of all future generations in our assessment of utility, how should we value the future relative to the present?

How much value should be attached to climate change impacts that occur beyond the lifetimes of most of those alive today? In comparing utility across generations, we need to determine the discount rate. In the Review's quantitative assessments, only the market costs and benefits of median impacts (types 1 and 2) of climate change mitigation are assessed, and only to the year 2100. Nevertheless, any view formed on discounting is important also in an assessment of qualitative climate change impacts in the longer term, and their implication for mitigation policy today.

There are two reasons why society may place less value on income and consumption in the future than on the same income and consumption today. The first element in the discount rate is the rate of pure time preference, that is, the rate at which future utility is discounted simply because it is in the future. Many of the philosopher kings of economics, from Ramsey (1928) to Sen (1961), have argued for a pure rate of time preference that is close to zero, thus placing no discount on the utility of people in the future just because today these people are young or not yet born. For example, as DeLong (2006) pointed out, a rate of pure time preference of 3 per cent means that somebody born in 1995 'counts for twice as much as someone born in 2020'. This flies in the face of the utilitarian principle underlying most economic analysis: that equal weight should be placed on each person. Dasgupta (2006), in his critique of the Stern Review's approach to discounting, also supports a rate of pure time preference close to zero. By contrast, some economists, including Nordhaus (1994, 2008), a pioneer of economic modelling of climate change, use a pure rate of time preference of 1.5 per cent or higher, calibrated to yield an overall discount rate that matches the overall cost of capital to the economy (see below).

The Review judges that a near-zero pure rate of time preference is appropriate. The only reason for a positive rate of pure time preference is the risk of human extinction in any one year. This should be a low number, and the Review uses a rate of pure time preference of 0.05 per cent. This is similar to the parameter value used in the Stern Review (2007).

The second element in the discount rate is the marginal elasticity of utility with respect to consumption. This is a measure of society's concern for equity in income distribution. We accept that a dollar of incremental income means less to the utility of the rich than of the poor. The people of tomorrow will have higher material incomes and wealth than people today, although this is likely to be offset by reduced environmental amenity in assessment of utility. It is reasonable to value future income at a lower rate than current income, insofar as future income is higher. How much less? Higher and lower values have been suggested, but no one contests that income has diminishing marginal utility with increased income. There are compelling theoretical reasons for using an elasticity of 1, and Quiggin (2008) argues that this is the most common choice in the literature (see also Jensen & Webster 2007). By contrast, Dasgupta (2007) argued that an elasticity of 1 implies that 'distribution of well-being among people doesn't matter much', and that higher values more adequately reflect distributional concerns and observed savings rates. These savings rates, however, vary greatly between countries and through time. Stern used a parameter value of 1, Nordhaus of 2.

The Review uses two alternative parameter values for the marginal elasticity of utility, 1 and 2, a range that accommodates strongly diverging views on how much should be spent now to benefit future, presumably richer, generations. Under an elasticity of 1, future income is discounted at the same rate as the increase in per capita income (plus the rate of pure time preference), while at an elasticity of 2 it is discounted at twice that rate. The average annual growth rate in Australian per capita income from 2013 to 2100 in the base case is 1.3 per cent. Thus the two real discount rates used by the Review for assessment of discounted net costs of mitigation of climate change in Australia are 1.35 per cent and 2.65 per cent.

The argument for being careful about the sacrifice of current utility through expenditure on mitigation in pursuit of future income is a powerful one. But there is one important qualification of this case for caution about strong mitigation on intergenerational income distribution grounds. The rate of substitution between conventional consumption and non-market services is likely to be low when incomes and material consumption are much higher than they are today. Climate change may greatly diminish the availability of non-market services for future generations. As a result, one cannot be sure that, despite much higher material consumption, the average utility of people in future will be greater than the average utility today. Hence, linking the marginal elasticity of utility to the growth in per capita income may lead to higher than intended discount rates. Furthermore, if considerable weight is given to the bad end of the probability distribution of outcomes from climate change, there is a possibility that utility may be lower for many people in future than at present.

There is another view, that market discount rates reflect the time preferences that are revealed in actual decisions on savings and investment, which are the vehicles for arbitrage between future and current economic activity. This raises two questions. What is the appropriate market discount rate? The other and more fundamental question is whether discounting is a normative or a positive issue, and so whether it is appropriate to use a market rate (Baker et al. 2007).

On the first question, a case can be made for using the market rate for sovereign debt in countries like Australia. The mid-point of the range of normative discount rates discussed above roughly coincides with the inflation-adjusted long-term market rate of return of government bonds in Australia and the United States, which stands at 2.2 and 2.1 per cent respectively. These would seem to be more appropriate than equity market rates, which are much higher, reflecting perceptions of firm-specific and other risks that are not relevant to the current analysis.

Is there a contradiction between using a discount rate in the range 1.25 to 2.65 per cent in summing utility of income over the generations, and applying a higher rate, 4 per cent, in pricing emissions permits and allocating capital over the 21st century in the Garnaut–Treasury modelling that is discussed later in this report? No. The Garnaut–Treasury rate is that at which investors choose to allocate capital between permits and other financial investments over time. It is our assessment of the 'Hotelling rate' that the market would come to apply in decisions that determine the optimal rate of depletion of the atmosphere's limited capacity to absorb greenhouse gases. It is advisedly a rate embodying all of the commercial risks involved in holding permits and investing in emissions-related activities. Its nearest analogue in currently operating markets is the gold contango, where the implicit real discount rate is typically lower than the 4 per cent applied in the Garnaut–Treasury modelling for the permit market contango. This reflects the view agreed between the Review and Treasury, that there would be additional risk and therefore a higher risk premium in the permit market. The Review uses a discount rate of 4 per cent in the modelling analysis of permit price trajectories, embodying a risk-free real interest rate of 2 per cent and a risk premium in markets for permits of 2 per cent (see Chapter 11).

On the second question, strong arguments can be made against any approach using observed market rates in the case of climate change. They include that a market portfolio approach averaging over high-yielding and underperforming projects is not applicable to climate change where winners and losers live in different generations (Beckerman & Hepburn 2007), and that the genuine uncertainties over which interest rate should be used in the long term favour lower discount rates (Weitzman 2007).

The Review judges that a normative approach is warranted on an issue that affects society as a whole over long time frames and on fundamental issues. Yet the justification of the discount rates used does not rely on using a normative approach. Rates that the Review derived from analysis, presented above, straddle the market rate that is judged to be most appropriate. In this case at this time, there is no conflict between normative and positive approaches.

The analysis for the Review has been calibrated with percentage points of GNP or consumption. The use of a discount rate that is higher than the rate of growth of GNP will cause the present value of a percentage point of current GNP to be greater than that of a percentage point of future GNP. The use of a discount rate that is lower than the rate of growth of GNP causes the present value of a percentage point of future GNP to exceed that of a percentage point of current GNP.

In Australia's case, unlike for most developed countries, the modelling points to the expected rate of GNP growth (2.1 per cent over the remainder of the 21st century) falling within the middle or higher than the middle of the range of discount rates thought to be relevant (1.35 to 2.65 per cent). It follows that at the lower discount rate, the present value of a percentage point of GDP in the early 22nd century will exceed that of a percentage point of GDP now.

Notes

1 The 'Platinum Age' is so named because global economic growth in this period has been and is expected to continue to be more expensive and stronger than in the 'Golden Age' of the 1950s and 1960s.

2 For ease of exposition, Figure 1.6 assumes that utility with mitigation already accounts for the lesser economic consequences of the additional impacts. If this were not the case, the utility-with-mitigation curve would shift downwards but the downward shift would be less than the downward shift of the utility-without-mitigation curve. The crossover point would still be earlier.

References

Baker, R., Barker, A., Johnston, A. & Kohlhaas, M. 2007, *The Stern Review: An Assessment of its Methodology, Productivity Commission Staff Working Paper*, Australian Government, January.

Beckerman, W. & Hepburn, C. 2007, 'Ethics of the Discount Rate in the Stern Review on the Economics of Climate Change', *World Economics*, vol. 8, no. 1, pp. 187–210.

Club of Rome 1972, *The Limits to Growth: A report for the Club of Rome's project on the predicament of mankind*, D.H. Meadows, D.L. Meadows, J. Randers & W.W. Behrens (eds), Universe Books, New York.

Dasgupta, P. 2007, 'Commentary: the Stern Review's economics of climate change', *National Institute Economic Review* 199: 4–7.

DeLong, B. 2006, *Partha Dasgupta Makes a Mistake in His Critique of the Stern Review*, November 30, <http://delong.typepad.com/sdj/2006/11/partha_dasgupta.html>.

Garnaut, R. & Huang, Y. 2007, 'Mature Chinese growth leads the global Platinum Age', in R. Garnaut & Y. Huang (eds), *China: Linking Markets for Growth*, Asia Pacific Press, Australian National University, Canberra.

Hotelling, H. 1931, 'The economics of exhaustible resources', *Journal of Political Economy* 39(2): 137–75.

Jensen, P. & Webster, E. 2007, *Inter-generational Justice, Discount Rates and Climate Change*, report commissioned for the Garnaut Climate Change Review, Melbourne Institute of Applied Economic and Social Research, Melbourne.

Jevons, W.S. 1865, *The Coal Question: An inquiry concerning the progress of the nation, and the probable exhaustion of our coal-mines*, Macmillan, London.

Keynes, J.M. 1921, *A Treatise on Probability*, Macmillan, London.

Knight, F. 1921, *Risk, Uncertainty and Profit,* Houghton Mifflin, Boston.

Little, I.M.D. & Mirrlees, J.A. 1968, *Manual of Industrial Project Analysis in Developing Countries*, OECD, Paris.

Maddison, A. 2001, *The World Economy: A millennial report,* OECD, Paris.

Malthus, T. 1798, *An Essay on the Principle of Population*, 1st edn, J. Johnson, London.

Nordhaus, W. 1994, *Managing the Global Commons: The economics of climate change*, MIT Press, Boston.

Nordhaus, W. 2008, *The Challenge of Global Warming: Economic models and environmental policy,* Yale University, New Haven, Connecticut.

Quiggin, J. 2008, 'Stern and his critics on discounting and climate change: an editorial essay', *Climatic Change,* doi:10.1007/s10584-008-9434-9.

Raiffa, H. 1968, *Decision Analysis: Introductory lectures on choices under uncertainty*, Addison-Wesley, Reading, Massachusetts.

Raiffa, H. & Schlaifer, R.O. 1961, *Applied Statistical Decision Theory*, Harvard Business School, Boston.

Ramsey, F.P. 1928, 'A mathematical theory of saving', *Economic Journal* 38: 543–59.

Sen, A.R. 1961, 'On optimising the rate of saving', *Economic Journal* 71: 479–96.

Stern, N. 2007, *The Economics of Climate Change: The Stern Review*, Cambridge University Press, Cambridge.

Weitzman, M.L. 2008a, 'On modeling and interpreting the economics of catastrophic climate change', Harvard University, Cambridge, Massachusetts, available at <www.economics.harvard.edu/faculty/weitzman/files/modeling.pdf>.

Weitzman, M.L., 2008b. 'A review of the Stern Review on the economics of climate change', *Journal of Economic Literature* 45(3): 703–24.