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

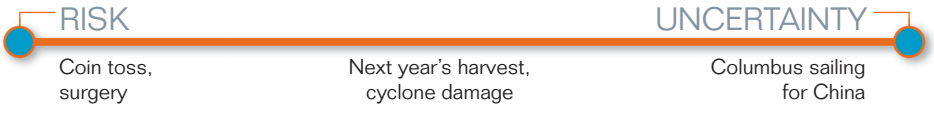

The 18th century British philosopher Bayes has given his name to a well-developed approach to decisions under uncertainty. Bayesian decision theory encourages us to treat decisions under uncertainty as if we were taking a risk (Raiffa 1968; Raiffa & Schlaifer 1961). We will make the best possible decisions 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

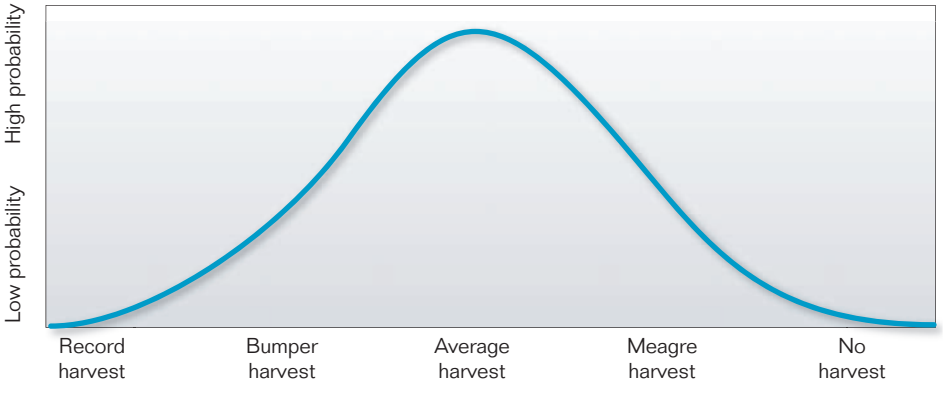

While the distinction between risk and uncertainty is analytically helpful, it does not distinguish discrete and separate phenomena. Rather, risk and uncertainty are the extreme ends of a single spectrum. Next year's harvest can be assessed as a risk on the basis of past experience but carries an element 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

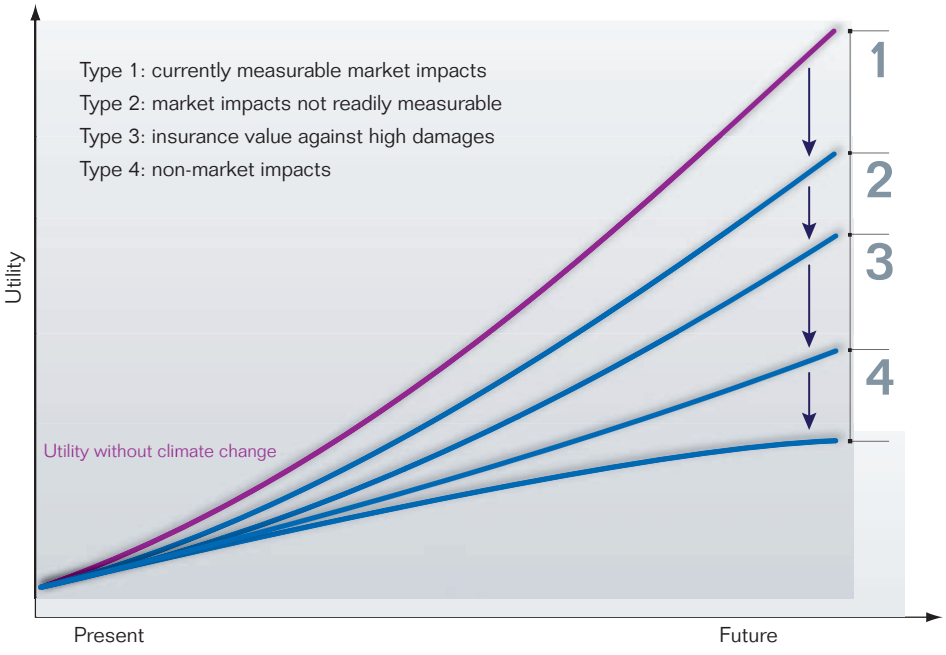

1.3.2 Type 2: market impacts not readily measurable

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