

**Measuring the Immeasurable:
The Costs and Benefits of Climate Change Mitigation**

Professor Ross Garnaut

The Sixth H.W. Arndt Memorial Lecture

The Australian National University

5th June 2008

Heinz Arndt's intellectual interests were eclectic and passionate. He grounded his academic life in economics because that was where the solutions to problems of poverty were likely to be found, but he remained interested in the scientific, political and philosophical context. By mid-career he was forming the view that to work on the economic problems of rich countries was akin to membership of the groups that sprang up in industrialising Britain in the nineteenth century, to assist genteel ladies in distress. That was when he moved to the Coombs Building, here, to specialise in development, and began the work that leaves this place today the main world centre outside Indonesia for the study of the Indonesian economy.

Heinz's eclectic and wide interests and his early concentration on development made him unfashionable in the economics profession, as it narrowed its focus to increasingly technical issues, and to the reliable data series that could only be found then in advanced economies.

Climate change was a Heinz Arndt sort of issue. It draws to itself all of the threads of human intellectual endeavour. And it arouses political passion. Heinz would have become deeply interested and argumentative about the science, the international relations, and the philosophy. His passion would have been the father to his intellectual positions, rather differently as he moved famously across the Australian political spectrum over time.

He and I didn't once talk about climate change. It hadn't caught either one of us as a central issue before his death in 2002. I had given the opening address to the Australian Mining Industry Council annual conference in 1991, introducing to the Australian policy discussion the idea that now goes by the name of "carbon leakage". I had said in a policy paper for the Australian Business Council in 1997 that international policy responses to climate change could damage the future of coal exports. But that was at the margins of each of our interests.

Over nearly 40 years, I managed to discuss issues from widely different perspectives with Heinz without ever having had anything like an argument. On central issues in economic policy, I suppose our views did not differ much, after he followed my

generation of ANU economists away from the old Australian view in sympathy with protection, to unequivocal commitment to free trade. We managed never to descend to what could be described as argument through many discussions from widely different positions on the big international political issues of our times in our region: the Vietnam War, recognition of China, and East Timor. It would have been like that with climate change, as it came to occupy more of my mind last year and this.

But Heinz would have been having arguments with lots of others, and sending letters off to the Canberra Times and articles to Quadrant. Never as a sceptic. Always as a believer or an atheist.

I thought that I would use this occasion to talk about the conceptual framework of my Climate Change Review for the Federal, State and Territory Governments.

How do we assess whether Australian mitigation action is justified? Would the substantial costs of mitigation be exceeded by the 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 Garnaut Climate Change Review seeks to form and to present judgements relating to the key mitigation choices for Australia in a transparent way.

In the Draft Report that will be released to the Australian public on July 4 will provide a sense of the early tendencies of analysis. The Supplementary Draft Report at the end of August will present a comprehensive set of results from the quantitative analysis completed for the Review: the work on the costs of mitigation jointly with the Australian Treasury, and separate analysis of the benefits of mitigation through diminished climate change. The Review's conclusions will be drawn together in the Final Report, to be presented to Governments at the end of September and released soon after.

Climate change mitigation decisions in 2008, and for the foreseeable future, are made under conditions of great uncertainty. There is large 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. There is large uncertainty 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 such uncertainty, it is always sensible to ask whether it would be better to delay decisions, while information relevant to the decision is gathered and analysed. However, it is as much a decision to delay action, as it is to decide to take early action. The issue of course is whether delay would be a good decision.

In 2008, the costs of delay – on a balance of probabilities, in the probabilistic terms that frame a good decision under conditions of uncertainty – are high. The mainstream science, the tendencies in global economic development and the state of the international decision-making process suggest that “business as usual” is running Australia and the world towards high risks of dangerous climate change at a rapid rate. 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.

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 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 are valued by Australians in themselves.

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 great uncertainty.

RISK AND UNCERTAINTY

Risk relates to an event that can be placed on a known probability distribution. When we toss a coin twice, we do not know whether or not we will see two heads. If we double toss the two coins enough times, they will both fall as heads around a quarter of the time.

In many spheres of human life, an activity has many 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, both have these same properties. 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 judgements 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 judgements 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 is an historically important example.

Bayesian decision theory advises us that 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.

In truth, 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.

If it is correct method 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 or amount of new observations or experience or further analysis.

After a year living as an outsider in close proximity to the climate science, I have 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 United Nation's Inter-governmental 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. The climate models provide numerous observations on possibilities out of their diversity, as well as from each generating numerous results from repeated experiments. These are the senses in which the IPCC science draws from probability distributions. There are many points at which judgement 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.

Every climate scientist has his or her own views on some issues that differ from the mainstream in detail. But the broad findings of the IPCC have general support amongst scientists with relevant specialist expertise.

The broad wisdom of the IPCC is strongly contested by a small number, and a small minority, of reputed climate scientists. It is not contested by the large majority of specialists, and by the leaders of the relevant learned academies in the countries of great scientific accomplishment.

It is sometimes observed by Dissenters that Galileo turned out to be right as a minority of one against the intellectual establishment of his time. Does not this establish that the intelligent Dissenter can be right?

Yes, it does. But the establishment of Seventeenth century Catholic Europe was not learned in scientific method. Would not Galileo be with the majority of established science today?

Probably.

Mainstream science is right on a balance of probabilities.

The Dissenters are sometimes called sceptics. This is a misnomer in general. Many hold to their views with profound belief that is independent of external information or analysis.

The Dissenters are possibly right, and probably wrong.

To conclude this discussion of uncertainty and belief, I recall the perspective offered by the former Australian Science Minister, Barry Jones. In his World Meteorological Day Address in 1992, he applied the famous wager of the seventeenth century French scientist, Blaise Pascal, to the climate change problem. If there were no God and one believed, pondered Pascal, what is the loss? But if there were a God, and He rewards belief or denial in Heaven and Hell, the absence of belief is catastrophic. It is rational to act as if there were a God.

I should say here that Heinz would always be prepared to back the passions of his mind even at these awful odds!

Pascal's wager would seem to make the case against the Dissenters.

But, as we will see, it is not quite so easy with climate change. Belief, acted upon, could be costly, and wasted if it is all a warp in the modern history of science. There is no alternative to seeking to measure the costs and benefits of efforts to mitigate climate change, while being mindful of uncertainty. And, regrettably, there is no alternative to acting on the results of that analysis now, actively or passively, as the passage of time is rapidly reducing the scope for choice amongst policies affecting climate outcomes.

THE COSTS OF MITIGATION

A modern acceleration in rates of human-induced greenhouse gas emissions is the source of contemporary concerns about anthropogenic global warming.

Economic development over the past two centuries has taken most of humanity from lives that were brutal, ignorant and short, to personal health and security, material comfort and knowledge that were 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 twentieth century, per capita output increased more than eight times and population more than fivefold (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 twentieth 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 I have described as the Platinum Age of global economic growth in the early twenty first century (Garnaut, 2007). 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. Analysis presented in the Draft Report points to the Platinum Age contributing a greater absolute increase in annual human output and consumption in the first two decades of the twenty first century than was generated in the whole previous history of our species, and then adding almost that much again in the next following decade to 2030.

The era of modern economic growth has been intimately linked to rapid expansion in use of fossil fuels. This is returning to the atmosphere a small part of the carbon that was sequestered naturally over millions of years, through a process which created the conditions that were necessary for the emergence of human life on earth.

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

There are, however, much tighter limits to what might be thought of as an 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 about 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.

The success of technological improvement and market processes in expanding supply and easing demand for scarce natural resources in the first centuries of modern economic growth established confidence that global economic growth was unlikely to be constrained by the availability of fossil fuels in any time frame that was relevant to current decisions.

The extraordinary growth in demand for fossil fuels demand in the early years of the Platinum Age – and the immense and unexpected increases in prices which have accompanied it – have rekindled interest in resource limits to growth. Will the supply conditions of fossil fuels slow down the growth in greenhouse gas emissions enough to do the mitigation task for humanity?

It is clear from the present state of knowledge – as it was not to earlier generations – that it would be 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. One day, humanity would have to make this transition. But it would be easier and cheaper, and less disruptive to the continued growth in incomes, if it were done gradually, with strong focus on efficient policies to promote the emergence of commercially viable technological alternatives to the use of fossil fuels.

For the time being, the pervasive and rapidly growing use of fossil hydro-carbons 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 make 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 this will be nowhere near the extent necessary to avoid high risks of dangerous climate change.

Adjusting to limits on the use of fossil fuels required to mitigate climate change would be less costly than to economic constraints on the availability of fossil fuels. This is because sequestration through physical (geo-sequestration) or biological (photosynthesis) processes can ease the mitigation task but cannot ease economic constraints on fossil fuel supply. But mitigation would be more difficult politically than reduction in emissions from higher market prices, because any constraints to force the adjustment would need to be imposed through human political processes. Decisions of this kind in single countries are hard enough. The necessity to achieve mitigation outcomes through cooperation of many sovereign entities, each with an incentive to shift as much of the costs of adjustment as possible onto other countries, increases the challenge.

A revolution 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 we are to hold the risks of climate change to acceptable levels. The costs that we bear in making an early adjustment will bring forward, and reduce for future times, the costs of the inevitable eventual adjustment away from fossil fuels. How much sooner and at what extra cost is the central question before the Review.

Costs of mitigation depend on the extent to which, and the time over which, reductions in emissions are achieved. They depend on the efficiency of the instruments chosen to implement policy. As I have argued in the Interim Report and the ETS Discussion Paper, and as I will confirm in the Draft Report, 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.

The cost 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, and the costs and benefits of mitigation 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 costs of mitigation will be lower the higher are the market prices of petroleum, coal and natural gas. This is because the costs of “business as usual”, to be compared with the costs of using the alternative, low-emissions technologies, will be higher. This is a matter of high current interest, at this time of historically high fossil fuel prices.

The costs of mitigation will be higher the more ambitious the extent and speed of reductions in emissions. It 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, and therefore impose 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 annual costs of mitigation are likely to rise for some time, as a rising emissions price forces deeper abatement. At some time, this tendency would be moderated and eventually reversed by improvements in the technologies that emerge to replace fossil fuels and other sources of emissions.

At some time in the future – when economic constraints on the use of fossil fuels would in any case 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 some 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 affected by the nature of the global mitigation effort. An effective global effort would make available a wider range of opportunities for trade in mitigation responsibilities, assigning higher effort to countries in which it can be achieved at lowest cost. A global effort would increase and distribute more equitably the world's investment in new technologies to develop lower emissions paths to consumption and production. And it would obviate the need for special policy measures to avoid carbon leakage, or the shift of emissions-intensive industries from high-mitigation to low-mitigation countries—a policy requirement that is likely to be profoundly distorting of domestic economic efficiency and political integrity.

It is important to see any period in which an Australian mitigation effort is in place prior to an effective global effort as short, transitional, and directed at achievement of a sound global agreement.

SOME CONTEMPORARY FACTORS IN THE COST OF MITIGATION

It may be useful to share some developments in my thinking about public discussion and real developments surrounding the costs of mitigation since the Interim Report in February and the Emissions Trading Scheme Discussion Paper in March.

The big external development has been the continued lift in oil, gas and coal prices, to levels that are several times forward market expectations only a few years ago. Community concern so far has focussed mainly on the petrol and diesel prices, where the pass-through to consumers is rapid.

Rising gas prices—to go much further with the internationalisation of the Eastern Australian gas market that is now in process—and coal prices have begun to affect

households' utility bills. They will have much larger effects over the next few years, and would do so even if there were no Emissions Trading Scheme in contemplation.

As I have just noted, higher fossil fuel prices lower the costs of mitigation. They reduce demand for fossil fuels. If an ETS were already in operation, the higher energy prices would reduce demand for permits and lower permit prices. The ETS would to some extent cushion the shock of exogenous increases in energy prices.

In this period before the introduction of the Australian ETS in 2010, higher fossil fuel prices will cause Australian emissions to shift downwards. If we had been more or less in line with the Kyoto requirements, we will now be tending below.

The Draft Report will examine in detail the implications of and opportunities created by the likely overperformance of Australia during the Kyoto period for the optimal design of the ETS up to the end of 2012.

I have considered carefully the proposal of Warwick McKibbin, my long-time friend and ANU colleague, to base a long-term permanent mitigation system on a "hybrid" of an ETS and a price cap. Warwick was a pioneer of serious work on the economics of climate change in Australia. If his proposals for gradual and constrained action had been applied in the late 1990s when first suggested, we would be in a better position today.

The world is now some way down the track on an international system based on emissions reduction targets, starting with developed countries. As discussed in the Interim Report, there are many imperfections in the Kyoto agreement that must be corrected in its successors if there is to be worthwhile progress towards reducing risks of dangerous climate change to acceptable levels. But the focus needs to be on the improvement of the system that has been emerging within the UNFCCC. There is no time to start again. A price cap is not consistent with the emerging international approach.

The idea that a cap could be put on the carbon price, and Australia and the world accept the amount of mitigation that happens to come from that is inconsistent with the urgency of the emissions reduction task. The work of the Review on the reality of

the rate of growth in emissions in the Platinum Age has led to realisation that the time available for effective action is considerably shorter than previously assessed. Warwick's own work, adopting different methods, has arrived at a position on "business as usual" emissions that is as similar to ours as we both are different from the earlier conventional wisdom.

In the context of over-performance on Australia's Kyoto targets, there is a case for the 2010 ETS to begin with a fixed, low price for permits as a purely transitional measure in the start-up period. For this to work, the period to the end of 2012 would need to be separated from the more demanding mitigation effort from the beginning of 2013. A fixed, low price would be better than a price cap because the separation from the post-2012 mitigation effort may lead to a zero price under an ETS, that would not allow the building of the ETS institutions and processes. The merits of such an arrangement will be discussed in detail in the Draft Report.

There is indeed need and room for flexibility in the system, to take account of any current misjudgements in the future costs of mitigation, fluctuations over time in economic variables affecting the permit price, and the shocks that from time to time enter any market. These were discussed in the Interim Report, and that discussion will be taken further in the Draft Report. This comes from opportunities for international trade in permits, from opportunities for hoarding and lending permits, and from occasional discretionary interventions by the regulator, the Independent Carbon Bank. The closest analogue of the independent regulator is the Reserve Bank of Australia and its interventions in the foreign exchange market at times of high and misjudged speculative activity. Unlike a price cap, each of these instruments of flexibility is consistent with Australia meeting firm emissions trajectories as will be required under international agreements.

Nor does the logic of analysis of decisions under uncertainty in the presence of risk aversion argue for delay and slow starts on mitigation while we gather more information. Uncertainty, the possibility of extreme outcomes and risk aversion point towards firm early action to keep options alive, while the advance of science narrows the range of uncertainty.

An ETS with a price cap does not allow, and Warwick does not support, large-scale trade in permits. Such trade can substantially reduce the costs of specified levels of mitigation in Australia, and in the world as a whole. It is an essential condition for substantial mitigation effort in many developing countries, certainly including our neighbours Papua New Guinea and Indonesia.

I have learned much from Warwick's work on climate change over the past decade and more. Some parts of his contribution have continuing relevance. His model is one of three being applied in the joint Garnaut Review-Treasury work that is now in progress. But, regrettably, large parts of Warwick's "blueprint" are not the best fit in the world of international climate change policy in 2008.

The public discussion of Australian mitigation has become pretty ragged as the important times for decisions approach. In this world of Chinese whispers, it may be useful simply to reiterate the perspective on emissions reduction trajectories that I presented for discussion in the Interim Report last February. The Interim Report was introducing ideas that we would be examining and analysing, including though modelling. I talked about the target being the Kyoto obligation up to the end of 2012. That would be followed by a period in which the target or trajectory would be in line with the average effort of developed countries. The positions of McCain and Obama would suggest that this criterion is likely to lead to confirmation of trajectories along the lines of the Australian Government's announced policies.

There has been a lot of public interest in the timing of the modelling work that the Garnaut Review is undertaking with the Treasury. I said in a speech to a Sydney coal conference a month ago that the main modelling results would come out in a Supplementary Draft Report in late August. The modelling has taken longer than we had originally hoped, entirely as the result of the complexity of the task of integrating the use of three general equilibrium models, and their application to an exercise of unprecedented ambition, involving amongst others the projection of developments in the Australian and global economy over a century. Mitigation costs under a range of emissions reduction scenarios are being analysed. None of the scenarios has led to the overwhelming of the models.

FOUR KINDS OF BENEFITS FROM MITIGATION

Three kinds of benefit from climate change mitigation can be measured in monetary values, as a change in the value of output or consumption. Assessing welfare that is affected by the fourth in addition to the first three kinds of benefit of mitigation requires a different measurement unit. 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.)

The first kind of benefit comprises currently measureable market impacts. 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. We will typically measure these effects as an impact on GDP or consumption, with monetary values as the unit of measurement.

The second kind of benefit comprises market impacts that are similar in nature to the first, but which are not amenable to precise measurement in the current state of knowledge. For the review, these comprise impacts that were not defined precisely enough in time for the Review's modelling, but which 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 means of the probability distributions of possible outcomes. We are drawing these judgements 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 kind of benefits to be subject to large adjustments, in one direction or another, with the advance of knowledge. Examples from the Review include increased demands for defence expenditure in response to geo–strategic instability associated with climate change; or the impact of climate change on the tourism or wine industries. As with the first type of benefits, the estimation of these effects will be in monetary values of GDP or consumption.

The third kind 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 and much more damaging than the mean. Some of the possible outcomes near the more damaging 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 may be more than we would be prepared to pay to avoid the certain prospect of the mean outcome. Like the first and second kinds of benefit from mitigation, insurance value of mitigation can also be measured in monetary value of GDP or consumption.

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 sees it as the main element:

“...the burden of proof in the economics of climate change is presumptively upon whomever wants to model optimal-expected-utility growth under endogenous greenhouse warming without having structural uncertainty tending to matter much more than risk. Such a middle-of-the-distribution modeller needs to explain why the inescapably thick tails of the ...distribution...is not the primary focus of attention and does not play the decisive role in the analysis.”

We are focused on maximising the 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 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 two outcomes and the possible conflict between them is known.

Examples of such non-market impacts include Australians' valuation of environmental amenity (the Great Barrier Reef and other features of the Australian and international landscapes, inherited shorelines, genetic diversity and the survival of species, 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 people argue that we should. It is only necessary to accept that many Australians value such things and would be prepared to sacrifice some consumption of goods and services to have more of them. Another example of a non-market impact that Australians value would be the avoidance of poverty and trauma in other countries (again, as valued by the sum of individual Australians). The proof of the importance of such matters lies in the continued support by Australians 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 natures difficult to measure. Any 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 and of possibilities of mitigation.

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 twenty first 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 will be low, and much lower than today. Near subsistence levels of consumption, few people would willingly sacrifice much access to material goods and services for greater environmental amenity, or for improved development prospects of others at home or abroad. But in the likely material affluence of the late twenty first century, many more people are likely to want to trade substantial amounts of access to material consumption for improved values of services that are 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 would challenge the proposition that continuing economic growth would necessarily lead to higher average utility in the distant future.

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

HOW EFFECTIVE ADAPTATION REDUCES THE COST OF CLIMATE CHANGE AND THE BENEFITS OF MITIGATION

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.

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, and the benefits of climate change avoided through mitigation diminished correspondingly. The benefits of adaptation through reduced climate change damage need to be subtracted from the gains from mitigation.

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 able

partially to take account of the costs and benefits of adaptation in its quantitative analysis.

MEASURING THE BENEFITS OF MITIGATION AGAINST THE COSTS

To a sceptical economist, like me, 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 which are undertaking joint action.

The relevant benefit of mitigation in reduced costs of climate change is the costs of climate change, below what would have been the case after the costs and ameliorating effects of adaptation had been taken into account.

The benefit from mitigation is the costs of climate change avoided. 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? Here we must come to grips with challenging issues about discounting for time, and about conversion of probability distributions into expected values, and of comparing the value of a certain outcome with an uncertain distribution of outcomes with the same mean.

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 Report 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 the 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 for any given level of mitigation by its embodiment 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, amongst many countries, 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.

The Review's terms of reference require it to analyse the degree of Australian mitigation effort that would be necessary to support a global agreement to hold greenhouse gas concentrations to 550ppm, and separately to 450ppm. The Final Report, with the support of the modelling reported in detail in the Draft Supplementary Report, will examine the choice of Australian mitigation ambitions comprehensively.

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 we should do so.

The value of avoided, and of irreversible, effects of climate change extends forward to the point at which the life of the human species would have been extinguished by some separate influence. 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?

In comparing utility across generations, we need to determine the discount rate. There are two key variables: the pure rate of time preference; and the elasticity of the marginal utility of consumption.

The rate of pure time preference is the rate at which future utility is discounted simply because it is in the future.

Many of the philosopher kings of economics, from Ramsay to Sen, have argued for a pure rate of time preference that is close to zero. This approach was followed by Stern, who placed it at 0.1 percent, corresponding to a view on the probability in any year of human extinction. Some have commented that if the probability of extinction were as high as that, it is unlikely that the human species would still be here. Stern might respond that the human capacity to construct truly fateful weapons of mass destruction has not been with us long.

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 was the main criticism of Stern's approach.

The issue is whether the pure rate of time preference is a normative or a positive issue.

The second determinant of 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 poor than of the rich. How much less? Higher and lower values have been suggested, but no-one contests that income has diminishing marginal utility with increased income.

In the expected circumstances of continually rising incomes, this argues for placing higher value on current than future income.

This argument for being careful about the sacrifice of current utility through expenditure on mitigation in pursuit of future income is a powerful one. It would be the more powerful if a substantial part of the burden of current mitigation were to be placed on people in low-income countries to the extent that their prospects for economic development were to be significantly diminished.

There is one qualification of this case for caution about strong mitigation on inter-generational income distribution grounds. To the extent that there is a low rate of substitution between conventional consumption and non-market services when incomes are high, and to the extent that climate change introduces a possibility that the availability of non-market services may be greatly diminished for future generations, 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.

If anything like a market discount rate were used, developments over the time periods that are necessary for mitigation policy to have substantial effects do not matter much at all. At a real discount rate of 4 percent – the rate that the Garnaut Review has judged to be appropriate for modelling the future price curve for emissions permits (the “Hotelling curve”) – a dollar in 50 years time is worth just 14 cents today. At market rates for equities and other assets that are judged to be risky, the value of incremental income or consumption in half a century is negligible.

Are we at all comfortable about living for the moment to the extent suggested by the use of market rates?

Is there tension between our normative (what should be done) and positive (what seems to happen in markets) view of discounting for the future? Should we treat the interest rates generated in financial markets as market failures?

The application of a social discount rate, lower than market rates, is recommended as best practice in one highly practical area of applied, policy-oriented economics: cost-benefit analysis. Little and Mirrlees recommend the application of lower rate of

discount for future income than is generated in financial markets (Little and Mirrlees, 1968, OECD, Paris).

If there is market failure in debt markets, it has two elements. These work in opposite directions in the conventional case where average incomes and utility rise over time. The market's myopia might be expected to price in a significant rate of pure time preference. On the other hand, the preferences of the poor are under-represented in financial markets, where a dollar of savings and investment is the basis of the franchise. The enfranchisement of the poor in financial markets would argue for higher weight being given to current incomes, and therefore for the application of a higher discount rate.

Of course, 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. The future poor get no votes anywhere, and least of all in Wall Street, the City of London, and Puxi.

My own inclination is towards the use of a low pure rate of time preference, alongside recognition that in dealing with the means of the probability distributions, future incomes should be valued at substantially less per dollar on inter-generational equity grounds. The net result may justify the application of something like a market rate of interest for good sovereign debt to the discounting of outcomes near the middle of the distributions from the mainstream science. This outcome reflects coincidence of conflicting empirical influences, rather than the logic of debt markets. The Final Report will seek to show sensitivity of the policy conclusions to variations in the discount rate.

A different calculus becomes necessary for consideration of the future values of the truly awful possibilities.

THE REVIEWS RECOMMENDATIONS IN A WORLD OF UNCERTAINTY AND IMPORTANT IMMEASUREABLE IMPACTS

The Review successively in the Draft Report, the Supplementary Draft Report and the Final Report will present quantitative measures where it can, and estimate the

potentially measureable effects when the data are not available for elaborate modelling of the potentially measureable.

The Draft Report on July 4 and the Final Report will discuss the implications of taking into account the possibility of outcomes being much worse than is suggested by the means of the probability distributions. They will seek to bring to account the value of various non-market services that are valued by Australians and which would be substantially affected by realisation of outcomes predicted by mainstream science.

Doing all of these things in a transparent way will, I hope, reveal to the Governments to which I will be reporting, and to the Australian community, the implications of the climate change policy choices that will be made over the period ahead.

An observation of daily debate and media discussion in Australia could lead one to the view that this issue is too hard for rational policy-making in Australia. The issues are too complex, the vested interests surrounding it too numerous and intense, the relevant time-frames too long.

Following the Lee Lecture last year, Climate Change policy remains a diabolical problem. There is a chance – just a chance – that Australia and the world will manage to develop a position that strikes a good balance between the costs of dangerous climate change and the costs of mitigation. The consequences of the choice are large enough for it to be worth a large effort to take that chance, in the short period that remains before our options diminish fatefully.