Risk & Sustainable Management Group

Climate Change Working Paper: C07#1

Research supported by an Australian Research Council Federation Fellowship
http://www.arc.gov.au/grant_programs/discovery_federation.htm

Stern and his critics on discounting and climate change

by

John Quiggin

Australian Research Council Federation Fellow, University of Queensland

Schools of Economics and Political Science
University of Queensland
Brisbane, 4072
rsmg@uq.edu.au
http://www.uq.edu.au/economics/rsmg

THE UNIVERSITY OF QUEENSLAND
AUSTRALIA

This version: 19 June, 2007

Stern and his critics on discounting and climate change

John Quiggin

Australian Research Council Federation Fellow
School of Economics and School of Political Science and International Studies
University of Queensland

EMAIL j.quiggin@uq.edu.au
PHONE + 61 7 3346 9646
FAX +61 7 3365 7299
http://www.uq.edu.au/economics/johnquiggin

I thank Terry Barker, Roger Jones and Nancy Wallace for helpful comments and criticism.

This research was supported by an Australian Research Council Federation Fellowship.

Stern and his critics on discounting and climate change

The Stern Review of the economics of climate change (Stern 2007*)* has had a substantial impact on public debate over policy responses to climate change. This is in part a matter of timing. Although economists have analyzed issues related to climate change for years, this discussion has been overshadowed by debate over scientific issues, initially among climate scientists and then, for much of the last decade, between scientists and political critics of climate science, primarily associated with US-based think tanks. Although disputes continue, serious debate about the reality of global warming has ended. The Fourth Assessment Report of the Intergovernmental Panel on Climate Change (IPCC 2007a) has indicated, with more than 90 per cent probability, that human-caused emissions of carbon dioxide (CO2) and other greenhouse gases were primarily responsible for observed and projected global warming.

The Stern Review was undertaken by Sir Nicholas Stern at the request of the Treasury of the United Kingdom. Stern's report was released in November 2006, and subsequently published as Stern (2007). Stern changed the terms of the debate by presenting the issues in economic rather than scientific terms. The effects of global warming, previously discussed in qualitative terms, were argued by Stern to correspond to large losses in economic welfare.

Some aspects of the Stern Review have been accepted with little controversy. Most significantly, the estimate (p. xvi) that stabilising global atmospheric CO2-equivalent concentrations in the range 500–550 parts per million (ppm) would reduce world consumption by around 1 per cent has been generally accepted as reasonable, if somewhat optimistic. For comparison, the IPCC (2007b, p. 27) gives cost estimates ranging from slightly negative to 4 per cent of global income to achieve stabilisation in a range of 535–590 CO2 equivalent ppm. These relative modest estimates contrast with political discussion of mitigation policies, which has commonly presupposed much larger, though vaguely specified, economic costs.

There has also been little discussion of the Stern Review's projections of climate change, which generally followed those of the Intergovernmental Panel on Climate Change (2007a) or of estimates of the economic value of damage to natural ecosystems, which broadly followed those of Nordhaus and Boyer (2000). Nordhaus and Boyer have been criticised for undervaluing damage to natural environments (Quiggin 2006), and a similar undervaluation is evident in Stern (Neumayer 2007). There has been some criticism of Stern's estimates of impacts on human population, mainly on the grounds that Stern underestimates the potential for adaptation.

The main focus of discussion of the Stern Review has been the way in which future costs and benefits have been converted to present values through discounting. Stern's approach to discounting has been criticised by a number of economists including Dasgupta (2006), Nordhaus (2006), Tol and Yohe (2006) and Yohe (2006) as well as by popular writers such as Lomborg (2006). More favourable assessments include DeLong (2006) and Weitzman (2007a).

This review will focus on the question of discounting. The paper begins with an outline of the expected utility model used in the Stern Review. Next, the question of 'inherent discounting', that is, the idea that future outcomes should be discounted simply because they are in the future, is examined, along with the closely related concept of the pure rate of time preference. This discussion forms the basis for an assessment of the approaches to discounting and climate change adopted by the Stern Review and by its critics.

Expected utility

Stern, like most contributors to the debate over discounting and climate change, uses a model based on expected utility theory. There are good reasons for this. First, expected utility has the property of dynamic consistency, which means that, if you make a plan, anticipating all possible contingencies, you will want to continue with that plan over time, whichever contingency arises. No other choice model has this property except under special conditions.

Second, expected utility theory allows a single utility function that simultaneously determines attitudes to intertemporal wealth transfers, interpersonal redistribution and risk reduction (transfers of income and consumption between states of the world). Under a plausible technical assumption, broadly supported by empirical evidence (Pålsson 1996), all of these attitudes are determined by a single parameter, denoted η (eta) in the Stern Review, which measures the reciprocal of the elasticity of substitution for consumption.

The technical requirement, referred to as constant relative risk aversion, is that attitudes to risk, expressed in terms of proportions of income, should be independent of income levels. So, if a typical individual with a low income by current standards is unwilling to accept a bet offering equal chances of a 100 per cent increase in income and a 50 per cent reduction, the same should be true for typical individuals with high incomes and for people in the future when average incomes are likely to be higher. Given this assumption, it makes sense to undertake analysis as Stern does, in terms of percentages of income.

The central idea of expected utility theory is that people are not concerned ultimately with money income but with the utility derived from consumption. The assumption of diminishing marginal utility means that a dollar of extra income or consumption is worth less for rich people than for poor people. So:

(i) assuming rising incomes, a dollar of extra income is worth less in the future than it is today;

(ii) under uncertainty, a dollar of extra income in a bad (low consumption) state of the world is worth more than a dollar in a good (high consumption) state of the world; and

(iii) transferring income from rich people to poor people improves aggregate welfare.

The first assumption means that, in the presence of technological progress that allows rising incomes and consumption, we would expect a positive discount rate; that is, a given increment to money income should be worth less in the future than in the present. The second means that a risky asset (more precisely, a risky asset with returns that are correlated with aggregate consumption) should be worth less than a riskless asset, since it yields high returns when the marginal utility of income is low and vice versa. The third means that even when redistributive taxes and international aid are costly (for practical purposes, nearly always) they can improve welfare.

Assuming that the combination of the expected utility model and inherent discounting (discussed below) captures all the issues under consideration, the riskless social discount rate is determined by a simple formula

$$r = \delta + \eta^* g, \qquad (1)$$

where g is the rate of growth of consumption per person, and δ is the inherent discount rate, discussed below. A similar, slightly more complex formula can be used to derive the rate of return for a risky asset, based on its correlation with aggregate consumption.

The parameter η, the elasticity of the marginal utility of consumption, represents the proportional rate at which the marginal utility of consumption is reduced as consumption increases. The choice of η is central to the debate over discounting. The larger is η, the stronger all of effects (i)–(iii) become. So high values of η imply a high preference for current consumption, high aversion to risk and large benefits from redistribution.

Even economists familiar with the mathematical derivation of η often have problems understanding the implications of different choices of η, particularly when time, uncertainty and interpersonal redistribution interact. So it may be useful to consider the most common single choice, $\eta = 1$, which is used in the Stern Review.

The special case $\eta = 1$

Expected utility theory was first developed to analyse choice under uncertainty. In discussions focused on risk, the most common single parametric choice is $\eta = 1$, corresponding to a logarithmic utility function of the form $u = \log(c)$ where u is utility, c is consumption and log is the natural logarithm. This is a particularly tractable function, which also gains popularity from tradition, having been proposed as a utility function for money by Bernoulli (1738/1954). Arrow (1971) also suggested, on theoretical grounds, that η should be approximately 1.

Subsequent empirical evidence has been mixed. The assumption of logarithmic utility ($\eta = 1$) seems to give a good match for calibrated macroeconomic models designed to match growth and business cycle facts (Gomme and Rupert 2005). By contrast, data on the equity premium, discussed below, is often taken to imply that the value of η must be greater than 1. Estimates based on observations of labour market responses to risk suggest values of η around 0.5 (Kaplow 2005).

The range of values observed in different fields suggests that Stern's choice of $\eta = 1$ is a reasonable compromise. On the other hand, it implies that an expected utility model with a single value for $\eta = 1$ is unlikely to match all the available data. Thus for any choice of parameters, it will be possible to point to contrary evidence.

There is a natural way of interpreting logarithmic utility in the intertemporal context. With this specification (and ignoring inherent discounting as discussed below) one per cent of consumption now has the same value as one per cent of consumption at any time in the future. So, for example, a policy that reduced consumption (not the rate of growth of consumption!) by one percentage point from 2000 to 2050, relative to some baseline, then increased consumption by one per cent relative to the same baseline until 2100, would come out exactly neutral. Logarithmic utility is implicit in much discussion of intergenerational equity, notably the intergenerational accounting analysis popularised by Kotlikoff and Burns (2004), which focuses on the proportion of income paid in tax by each generation.

At this point a numerical example might be useful. The world's mean income per person is currently around \$US7000, though the majority of people get much less and the billion or so in rich countries much more. Under the projections used in the Stern Review, average world income in 2100 is estimated at about \$US 100 000. Consumption is approximately equal to income, and in what follows, these numbers will be treated as consumption levels.

With logarithmic utility, a sacrifice of \$70 per person (1 per cent of consumption) today would be justified if (and only if) it increased the consumption of our great-grandchildren in 2100 by at least \$1000. If this trade-off appears reasonable, then a value of η =1 is appropriate. If the future payoff required is higher (or lower) then so is the preferred value of η.

Intuition about the far future tends to be cloudy, so it is worth observing that, under expected utility theory, exactly the same arguments apply to redistribution within the current generation. To illustrate, it is useful to turn around the direction of redistribution. Consider a redistributive program that takes $1000 from the well-off (in this example households with income of $100 000 per person) and uses the proceeds to benefit the poor (those with $7000 per person). (Alternatively to keep the focus on redistribution from the poor to the rich, suppose that such a program already exists, and consider scrapping it.)

Such redistribution always involves a range of costs including administration, compliance, efforts at avoidance and evasion, and incentive costs. Suppose that, in a particular case, these costs amount to 93 cents per dollar initially taxed, so that for a net loss of $1000 to the well-off, the net benefit to the poor is $70. If such a program is exactly marginal, so that any program with a larger net benefit is acceptable, and any program with a smaller net benefit is unacceptable, then the implied social preferences have η = 1. If the minimum acceptable net benefit is larger (smaller), then we require η to be less than (greater than one).

Problems with expected utility

The expected utility model is neat, logically compelling and tractable, but it suffers from two big problems. First, at least some of the time, most people do not behave in a way that is consistent with the expected utility model. For example, people simultaneously gamble at unfavorable odds and take out costly insurance, which violates the predictions of expected utility theory with respect to uncertainty.

In decisions under uncertainty, individuals often seem to put more weight on low-probability extreme outcomes than would be implied by expected utility (Kahneman and Tversky 1979). Although various responses to this observation have been proposed, the most popular has been to use a rank-dependent weighting scheme, as proposed by Quiggin (1981, 1982, 1993) and incorporated in the cumulative prospect theory model of Tversky and Kahneman (1992).

Weitzman (2007a) observes that placing a high weight on extreme outcomes may be sensible in the presence of structural uncertainty about parameters such as the sensitivity of climate to increased CO2 emissions. He shows that, with an underlying lognormal probability distribution, such uncertainty gives rise to ‘fat-tailed’ subjective probability distributions with unbounded means, so that evaluations of uncertain choices are dominated entirely by events in the tail of the distribution. Bounding the worst-case outcome using economic estimates of the value of statistical lives allows the derivation of finite evaluations of alternative policies, but these still require a very high weight on extreme outcomes. In the case of climate change, only the right-hand (high-warming) tail of the distribution is relevant, since events in the left tail correspond to little or no change in global average temperature.

Similar kinds of problems arise in discounting over time. People often apply a high discount rate to trades between the present and the near future, but a low discount rate for trades between the near and far future. This is called hyperbolic discounting. Substantial bodies of literature on generalised expected utility theory and behavioural economics attempt to address this problem, as discussed below.

The second problem is that observed market outcomes are not consistent with expected utility theory as it is commonly applied. This problem is partly because people do not act in accordance with expected utility theory and partly because markets do not work in the smooth and frictionless way assumed in standard finance theory models.

The most important problem in this respect is the 'equity premium puzzle', and the closely related 'risk-free rate puzzle'. The equity premium puzzle, first observed by Mehra and Prescott (1985), is that, for plausible choices of η, the real bond rate should be somewhat higher than it is, and the rate of return to equity substantially lower.

Historically, real returns to investors from the purchases of US government bonds have been estimated at one percent per year, while real returns from stock ('equity') in US companies have been estimated at seven percent per year, a difference of six percentage points. By contrast, for reasonable choices of η, the difference should be no more than half a percentage point. The equity premium puzzle can be resolved by assuming very high values of η since risk aversion increases the premium. But high values of η imply a high discount rate, so the risk-free rate puzzle is made worse.

As is noted by Grant and Quiggin (2005), failure to take account of the equity premium puzzle can lead to mistaken and inconsistent policy judgements. The inconsistencies between the expected utility model and observed choices and market outcomes mean that, for any possible choice of parameters, it is possible to present hypothetical choices for which most people will reject the implications of the model, or to point to market outcomes inconsistent with the proposed parameters. In these circumstances, using expected utility theory to derive policy implications from particular market outcomes (such as average returns to equity) can lead to paradoxical or untenable conclusions.

Inherent discounting and weighting

Although expected utility provides a complete theory of allocation of consumption across individuals, time periods and states of the world, it is often supplemented by some sort of weighting scheme. This is true whether expected utility is used positively, to model actual behaviour, or normatively as a guide to rational individual decisionmaking and ethical social decisionmaking.

In the case of allocating consumption over time we need to consider whether we should discount future consumption simply because it is in the future, even with the same marginal utility. Such discounting is represented by the term δ in equation (1).

The parameter δ is commonly referred to as the 'pure rate of time preference', but usage of this term is not consistent, and some writers use it to refer to the rate of discount for monetary flows (r in equation (1)). Hence, in this paper the term 'inherent discounting' will be used.

Inherent discounting

One of the longest running controversies in welfare economics has concerned the appropriateness of applying different weights to people in different generations, and, more generally of discounting future utility whoever receives it. Ramsey (1928), whose work is the starting point for formal analysis of intertemporal choices, rejected inherent discounting as ethically unjustified, and this viewpoint is shared by most philosophical advocates of utilitarianism. On the other hand, a good deal of evidence suggests that individuals tend to discount their own future consumption.

Before discussing inherent discounting, it is worth observing that standard expected utility theory suggests one reason for discounting future consumption; namely the possibility that we will not be around to enjoy it. As individuals, we face a typical annual mortality risk of around 1 per cent, and it makes sense to discount future utility by this amount. But at least some of the time people (most notably teenagers) discount the future much more than this.

For society as a whole, there is a comparable risk arising from the possibility of nuclear annihilation, a killer meteor and so on. The risk need not involve a total extinction of the species; it is sufficient that the disaster be great enough that 'all bets are off' in terms of calculations about the future.

With this point addressed, there remains the question of whether we do and should, discount future utility. The evidence on individual behaviour is far from clear. On the one hand, there is a lot of evidence to support the idea of 'hyperbolic discounting'. However, this is offset by a notion of 'mental accounts' (Thaler 1990). Individuals may allocate resources between activities and follow inconsistent rules in different activities. For example, the same person may allocate money to an automatic saving scheme offering low or even negative real returns, while displaying hyperbolic discounting with respect to the remaining cash flow.

Leaving such phenomena to one side, the evidence for high inherent rates of discount is not strong. The most obvious market measure to use in assessing intertemporal tradeoffs is the real rate of interest on low-risk bonds (government or AAA corporate). This rate has generally been between one and two per cent and is currently around two per cent. Given that the rate of growth of average consumption per person is between one and two per cent, this is consistent with zero discounting and $\eta = 1$.

Even if individuals do display inherent discounting, that does not necessarily mean that this is appropriate as a basis for social decisions. Future individuals presumably will not share the view that utility in our time is inherently more valuable than utility in theirs. In fact, as individuals, introspection and casual observation suggest that we generally regret decisions made in the past on the basis of inherent discounting. Such decisions represent selfishness on the part of our past selves at the expense of our current selves, analogous to individual selfishness with respect to others.

The case against inherent discount is summarised by DeLong (2006)[1]:

> A δ of 3% per year is unconscionable--it means that somebody born in 1970 "counts" for twice as much as somebody born in 1995, who in turn "counts" for twice as much as somebody born in 2020.

A crucial point, often overlooked in discussions of intergenerational equity is that members of different generations are alive at the same time. Any policy that discounts future utility must discriminate not merely against generations yet unborn but against the current younger generation,

Stern and the critics

The analysis in the Stern Review follows the general approach set out above. The value of η is set to 1, which is, as noted above, the most common single choice for

1 A numerical error in the original statement of this point has been corrected

this parameter. The value of δ is set to 0.001 (0.1 per cent), reflecting a rejection of inherent discounting, except insofar as it reflects the possibility of extinction. Similarly, there are no interpersonal weights, but changes in the consumption of low-income individuals and countries are weighted more highly at the margin because of the assumed diminishing marginal utility of consumption.

In Stern's analysis, the rate of growth of consumption, g, is derived from the economic scenarios. Typical values of g are between 1.5 per cent and 2 per cent, so the corresponding values of r are between 1.6 per cent and 2.1 per cent.

The effect of choosing $\eta = 1$ and δ near 0 is that concern for future generations extends more or less indefinitely into the future, when changes in welfare are expressed in terms of percentages of income or consumption. On the other hand, the discounted value of payments expressed in monetary terms declines quite fast. At a rate of 2.1 per cent, a dollar of (constant price) income received in 2100 is worth approximately 12 cents today. A income stream of a dollar a year, received for a million years into the future is worth a little under $50.

Criticisms of δ

One part of the debate over δ can be dismissed quickly. Many of the critics on this point have confused δ and r, apparently assuming that δ is a discount rate, rather than a measure of inherent discounting, which is only part of the discount rate. Examples of this confusion include Lomborg (2006) and Leonhardt (2007).

A similarly weak criticism is the observation that benefit–cost analysis of public investment proposals, including that undertaken by the UK Treasury, commonly uses discount rates higher than that considered by Stern (Tol 2006). This is a neat debating point, but has little practical relevance. The political structure of project appraisal is that estimates of costs and cash flows rely on inputs from project proponents that are almost always over-optimistic (Flyvbjerg, Bruzeliu and Rothengatter 2003) Treasury controls the choice of discount rate and uses it to adjust for downside risk as well as for discounting. Hence, the official discount rate is substantially higher than the true social rate of discount. The use of discount rates to adjust for risk in this way has long been recognised as inappropriate (Little and Mirrlees) but persists because of its political convenience.

Among the more serious critics, both Nordhaus (2006) and Yohe (2006) focus on the sensitivity of the results to changes in the value of δ, but do not give any specific argument for inherent discounting. Yohe does not present any argument for a high value of δ, simply observing that others have used high rates.

Nordhaus (2006) cites earlier work (Nordhaus 1994; Nordhaus and Boyer 2000) in which the value $\delta = 0.03$ (3 per cent) was derived as a residual. Nordhaus assumes that the discount rate r should match the observed rate of return to (debt and equity) capital, and that this rate should be used to explain observed levels of aggregate savings. The historically observed average rate of return to capital is around 5.5 per cent. For reasons of technical tractability, Nordhaus chose to set $\eta = 1$. With an observed growth rate of consumption, g, of about 1.5 per cent, equation (1) implies that δ must be equal to 3 per cent.

Such an analysis implies a very low level of concern for our descendants (certainly for anyone more than two generations removed from us). In the discounting procedure proposed by Nordhaus, the welfare of our great-grandchildren (whether or not they have yet been born) has about a tenth the weight we accord ourselves. Not surprisingly, this translates into a 'do nothing now' approach to global warming.

Leaving aside the ethical difficulties here, Nordhaus' analysis fails to take account of the equity premium, which he mentions briefly as a factor that can be neglected. As was first pointed out by Cline (1991), if market rates of return are to be taken as correct measures of social welfare, the natural choice for the riskless discount rate is the rate of interest on government bonds, which has historically averaged between 1 and 2 per cent. Using this rate, equation (1) implies that δ must be close to zero.

In the absence of any convincing justification for inherent discounting, the case for a low rate such as that chosen by Stern seems overwhelming. Hence, if there is a problem with the ultimate outcome it is necessary to look elsewhere in the analysis. From here on, the value of $\delta = 0.001$ will be assumed, and discussion of the implications of other choices is conditional on this.

The intertemporal elasticity of substitution

A more plausible criticism concerns Stern's choice of η and suggestions that higher values should have been considered. The most direct criticism is that, in a growing economy, a low value of η underweights the welfare of the current generation, at the expense of succeeding generations who will be much richer. This point is made by Dasgupta (2006), who considers the case when society has available an unlimited supply of projects yielding a riskless rate of return of 4 per cent. As Dasgupta shows, with $\eta = 1$, the implied policy recommendation is that the vast majority of current income (around

97.5 per cent) should be saved in order to allow for greater consumption in the future.

The underlying problem is observed by DeLong (2006). Looking at current rates of savings and economic growth, Dasgupta's estimated rate of return to marginal investment of 4 per cent seems conservative for a classical growth model based on factor accumulation. To achieve 1 per cent growth in consumption per person in such a model, it would be necessary to generate net additions to the capital stock equal to 25 per cent of total income each year (since 4 per cent of 25 per cent is 1 per cent). By contrast, China is achieving annual growth rates of around 8 per cent, with savings equal to 40 per cent of economy. So the fact that we see more rapid growth with lower rates of net saving seems to imply that there must exist many projects with rates of return greater than or equal to 4 per cent.

However, once technical progress, generated either exogenously or through the existence of increasing returns to scale in knowledge, is taken into account, the picture changes radically. In the real economy, where most growth in consumption arises from technical progress, the optimal rate of saving is far lower than that derived by Dasgupta (2006).

A more direct way of refuting Dasgupta's argument is to observe that the major premise must be false. If there existed an infinite supply of projects with riskless returns of 4 per cent, the rate of return on riskless bonds would have to equal 4 per cent, rather than the 1 to 2 per cent observed in practice. Although this difference may appear small, it is critical in practice.

Criticism of the time horizon used in discounting

The most plausible criticism of low rates of discount is that they require us to take account of developments more than 100 years into the future about which we can in practice, know very little. This is a reasonable criticism, but its main effect is to point up the limitations of utilitarian benefit–cost analysis for a problem like global warming.

We know that the effects of global warming will be felt far into the future. We can either mitigate these effects, at very modest cost to ourselves, or leave the problem to future generations. However, our understanding of the extent to which our current actions will cause damage that is irreparable, or very costly to correct, is limited once we extend the time horizon far into the future. There is no easy way of getting useful probability and cost numbers here.

One partial solution might be to end the analysis at, say 2050 or 2100, with future effects being measured as a diminution in the capital stock (including natural capital). Although logically equivalent to the discounting procedure employed by Stern and his critics, this might turn out to be more tractable and intuitively comprehensible.

This issue has been addressed by Weitzman (2007b), who argues that Stern's choice of a low discount rate and a long time horizon may be understood as a roundabout way of addressing the problem of irreducible uncertainty.

Concluding comments

Critics of the Stern Review, such as Dasgupta (2006), Nordhaus (2006), and Yohe (2006) have relied mainly on the claim that the parameters used in discounting are extremely low, yielding implausible results. Stern's specification of the elasticity of the marginal utility of consumption is standard, and both lower and higher values are commonly considered in sensitivity analysis. As a result, most controversy has centred on Stern's choice of a very low rate of pure time preference.

Stern's modelling approach is primarily the result of applying the standard utilitarian view that all people count equally and that there is no justification for treating people more favorably simply because they were born earlier. If this view is accepted, the pure rate of time discount is simply a device to take account of possibilities that would render all calculations irrelevant, such as a nuclear catastrophe. Since the probability of such an event in any given year is close to zero, the pure rate of time discount must be similarly close to zero. Given Stern's premises, there is nothing remarkable about the parametric value of 0.1 per cent chosen to represent the pure rate of time discount.

As Weitzman (2007b) indicates, the real difficulty here is that the analytical methods of expected utility theory are ill-equipped to handle issues involving uncertainty about crucial parameters evolving over long periods of time. Yet there is no well-developed alternative.

In analysing such problems we are pushing economic analysis to its limits. Economists lack an analytical procedure to deal with problems involving unforeseen outcomes, or even to explain the large price premium associated with risky investments such as corporate equity. While such basic problems remain unresolved, any choice of discounting procedure to assess long-term risks such as climate change will yield some implications that are intuitively unappealing or inconsistent with observed market outcomes. Economists can help to define the issues, but it is unlikely that economics can provide a final answer.

Ultimately, the response to climate change is a social and political choice facing the global community, to be determined through processes such as the United Nations Framework Convention on Climate Change or, if such processes fail, through a default decision to do nothing. Stern and others have shown that the price of stabilising the global climate through economically efficient policies will be modest. The decision on whether to pay that price will determine, implicitly at least, whether the discount rates proposed by Stern accurately represent the collective judgements of the global community.

References

Arrow, K. (1971) *Essays in the Theory of Risk-Bearing,* Markham, Chicago.

Bernoulli, D. (1954), 'Exposition of a new theory of the measurement of risk', *Econometrica*, 22(1), 23–36 (originally published 1738).

Cline, W. (1991), 'Comments on Schelling', pp. 222–28 in *Global Warming: Economic Policy Responses*, (Eds, Dornbusch, R. & Poterba, J.) MIT Press, Cambridge, Mass.

Dasgupta, P. (2006), 'Comments on the Stern Review's Economics of Climate Change', mimeo, University of Cambridge

DeLong, B. (2006), 'Partha Dasgupta makes a mistake in his critique of the Stern Review', http://delong.typepad.com/sdj/2006/11/partha_dasgaptu.html,

Flyvbjerg, B., Bruzeliu, N., and Rothengatter, W. (2003). *Megaprojects and Risk: An Anatomy of Ambition.* Cambridge: Cambridge University Press.

Gomme, P. and Rupert, P. (2005), *Theory, Measurement, and Calibration of Macroeconomic Models*, Federal Reserve Bank of Cleveland Working Paper 05-05.

Grant, S. and Quiggin, J. (2005), 'What does the equity premium mean?', *The Economists' Voice*, 2(4), Article 2.

Intergovernmental Panel on Climate Change (2007a), *Working Group I Report (WGI): Climate Change 2007: The Physical Science Basis*, IPCC, Geneva.

Intergovernmental Panel on Climate Change (2007b), *Fourth Assessment Report*, IPCC, Geneva.

Kahneman, D. and Tversky, A. (1979), 'Prospect theory: An analysis of decision under risk', *Econometrica*, 47(2), 263–91.

Kaplow, L. (2005), 'The value of a statistical life and the coefficient of relative risk aversion', *Journal of Risk and Uncertainty*, 31(3), 23–34.

Kotlikoff, L. and Burns, S. (2004) *The Coming Generational Storm: What You Need to Know About America's Economic Future,* MIT Press, Cambridge, Mass.

Leonhardt, D. (2007, Feb 21). A battle over the costs of global warming. *New York Times*.

Lomborg, B. (2006), The dodgy numbers behind the latest warming scare*, WSJ Opinion Journal,* Nov 2, http://www.opinionjournal.com/extra/?id=110009182

Mehra, R. and Prescott, E. (1985), 'The equity premium: a puzzle', *Journal of Monetary Economics*, 15(2), 145–61.

Neumayer, E. (2007), 'A missed opportunity: the Stern Review on climate change fails to tackle the issue of non-substitutable loss of natural capital', unpublished paper, London School of Economics.

Nordhaus, W. (1994) *Managing the Global Commons: The Economics of Climate Change,* MIT Press, Cambridge, Mass.

Nordhaus, W. (2006), 'The Stern Review on the economics of climate change',unpublished paper, Yale University.

Nordhaus, W. and Boyer, J. (2000) *Warming the World: Economic Models of Global Warming,* MIT Press, Cambridge Mass.

Pålsson, A.-M. (1996), 'Does the degree of relative risk aversion vary with household characteristics', *Journal of Economic Psychology*, 17, 771–87.

Quiggin, J. (1981), 'Risk perception and risk aversion among Australian farmers', *Australian Journal of Agricultural Economics*, 25(2), 160–69.

Quiggin, J. (1982), 'A theory of anticipated utility', *Journal of Economic Behavior and Organization*, 3(4), 323–43.

Quiggin, J. (1993) *Generalized Expected Utility Theory: The Rank-Dependent Model,* Kluwer, Amsterdam.

Quiggin, J. (2006), 'Assessing the costs and benefits of reducing emissions of greenhouse gases', submission to the Stern Review on the economics of climate change, University of Queensland, Brisbane.

Ramsey, F. (1928), 'A mathematical theory of savings', *Economic Journal*, 38, 543–59.

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

Tol, R. (2006), 'The Stern Review of the economics of climate change: a comment', unpublished paper, Economic and Social Research Institute, Dublin.

Tol, R. and Yohe, G. (2006), 'A Review of the Stern Review', *World Economics*, 7(4), 233–50.

Thaler, R. (1990), 'Anomalies: saving, fungibility, and mental accounts', *Journal of Economic Perspectives*, 4(1), 193–205.

Tversky, A. and Kahneman, D. (1992), 'Cumulative prospect theory: An analysis of attitudes towards uncertainty and value', Journal of Risk and Uncertainty, 5(3), 297–323.

Weitzman, M. (2007a) 'Structural uncertainty and the value of statistical life in the economics of catastrophic climate change', unpublished paper, Princeton University.

Weitzman, M. (2007b), 'The Stern Review of the economics of climate change', Journal of Economic Literature, forthcoming.

Yohe, G. (2006), 'Some thoughts on the damage estimates presented in the Stern Review', *Integrated Assessment*, 6(3), 65–72.

PREVIOUS WORKING PAPERS IN THE SERIES
CLIMATE CHANGE

C05_1 Counting the cost of climate change at an agricultural level, John Quiggin (October, 2005).

C06_1 Fiddling while carbon burns: why climate policy needs pervasive emission pricing as well as technology promotion, John C.V. Pezzey, Frank Jotzo and John Quiggin (December, 2006).

C07_1 Stern and his critics on discounting and climate change, John Quiggin (June, 2007).

C07_2 Declining inflows and more frequent droughts in the Murray-Darling Basin: climate change, impacts and adaption, John Quiggin (2007).

C07_3 Complexity, climate change and the precautionary principle, John Quiggin (2007).

C07_4 The practicalities of emissions trading, Joshua Gans and John Quiggin (August 2007).

C08_1 Equity between overlapping generations, John Quiggin (June, 2008).

C08_2 Uncertainty and climate change policy, John Quiggin (2008).

C08_3 The impact of climate change on agriculture, John Quiggin (August, 2008).

C08_4 Grandfathering and greenhouse: the roles of compensation and adjustment assistance in the introduction of a carbon emissions trading scheme for Australia, Flavio Menezes, John Quiggin and Liam Wagner (December, 2008).

C09_1 Agriculture and global climate stabilization: a public good analysis, John Quiggin (August, 2009).

C10_1 Climate change, uncertainty and adaption: the case of irrigated agriculture in the Murray-Darling Basin in Australia, John Quiggin, David Adamson, Sarah Chambers and Peggy Schrobback (2010).

C11_1 Stabilizing the global climate: a simple and robust benefit cost analysis, John Quiggin (2011).

C13_1 Carbon pricing and the precautionary principle, John Quiggin (March 2013).

C13_2 How I learned to stop worrying and love the RET, John Quiggin (May 2013).