The 47th Annual Conference of the Australian Agricultural and Resource Economics Society
Fremantle, Western Australia, 12-14 February 2003

Deciding with long-term environmental impacts: what role for discounting?

Steven G.M. Schilizzi[1]

'Modelling inter-temporal concern is [yet] to be transformed from an art to a science'.

Jack Pezzey, 1997

In politics we have so firm a faith in the manifestly unknowable future that we are prepared to sacrifice millions of lives to an opium smoker's dream of Utopia or world dominion or perpetual security. But where natural resources are concerned, we sacrifice a pretty accurately predictable future to present greed. We know, for example, that if we abuse the soil it will lose its fertility; that if we massacre the forests our children will lack timber and see their uplands eroded, their valleys swept by floods. Nevertheless, we continue to abuse the soil and massacre the forests. In a word, we immolate the present to the future in those complex human affairs where foresight is impossible; but in the relatively simple affairs of nature, where we know quite well what is likely to happen, we immolate the future to the present.

Aldous HUXLEY, 1945
"Time Must Have A Stop", 1945. Triad/Granada ed. 1982, page 286.

ABSTRACT

The problem of how to discount values in the far future is reviewed, and shown to lead down a blind alley. An alternative is proposed that allows long term consequences to be addressed by decisions using a relatively short term time horizon. A simple model investigating the optimal containment of radioactive waste in a deterministic world is used to show that current generations can indeed cater for the interests of the far future while optimising over the short term; however, this is not always possible. The proposed method seems to address most of the critiques of long term discounting found in the literature.

[1] School of Agricultural and resource Economics, Faculty of Natural & Agricultural Sciences, The University of Western Australia, 35 Stirling Highway, Crawley (Perth) WA 6009.
Steven.Schilizzi@uwa.edu.au

Foreword

This paper has its origins in a presentation given at the 2001 AARES[2] Pre-Conference Workshop in Adelaide (South Australia), organised by David Pannell and myself, on 'Discounting in Theory and Practice'[3], the title of which was 'Long Term Discounting: Why So Many Voices and So Little Consensus?'[4]

1. Introduction

The nature of the problem is well known. Over the very long term, spanning centuries and many generations, the use of a constant positive discount rate reduces future values, however large, to next to nothing. For instance, with a standard discount rate of 7%, any rational individual should only be willing to pay $1.70 for avoiding the destruction of all the world's wealth in 200 years. The source of the problem is also well known: the power of exponential compounding, which reflects the common assumption of a constant rate of unlimited economic growth. However, the problem has received many different and often conflicting 'solutions', with some authors claiming it has no solution - at least , no rational solution!

Three paradigmatic examples that epitomise the issue are habitat destruction leading to species extinctions, global warming, and radioactive nuclear waste disposal. They all involve long time frames, respectively decades, centuries, and millennia - the half-life of plutonium-based radioactive waste spanning over 24,000 years. In each case, when protective measures are envisaged, short term benefits are balanced against (very) long term costs, most of which, though difficult to quantify, are thought to be 'potentially huge'.

Any discussion involving the (very) long term begs the question: how long is the (very) long term? Three criteria at least have been offered, all of which result, perhaps not by coincidence, in the same approximate number. These are time frames

1) greater than the average distance between two generations,

[2] AARES: Australian Agricultural and Resource Economics Society
[3] The Proceedings will be edited and published as a book by Edward Elgar Publishing Ltd.
[4] I wish to acknowledge the following people for the very fruitful discussions they have provided me with over the last three years, often by email: David Pannell, Greg Hertzler, John Quiggin, Jack Pezzey, Mike Young, David W. Pearce, Ari Rabl, and Martin Weitzman. All errors, omissions and inaccuracies remain of course mine.

2) greater than the longest term for government treasury bonds,

3) those for which there exist no market mechanisms for intertemporal trading.

All three criteria result in a threshold time frame of around 30 years. Anything beyond that time horizon appears as the 'very long term'. Let us refer to it as T.

The purpose of this contribution is to try to find a way around the lack of consensus over the issue of 'long term discounting', with the aim to help clear the way for 'modelling inter-temporal concern if is to be transformed from an art to a science' (Pezzey, 1997: p. 464). The book edited by Portney and Weyant in 1999, 'Discounting and Intergenerational Equity', exemplifies the lack of consensus among leading economists, and has been taken as a landmark reference for this contribution. The aim here is first to clarify the setting within which the problem is posed, and then suggest a framework for further progress towards resolving the problem.

The analysis is practically oriented, in that it addresses the concerns of a (public) decision-maker. We also take a step back from the question "how to value (very) long term consequences", which presumes the need for such a valuation, and ask, rather, "how can we make decisions that take into account (very) long term consequences?" At this stage, although they are by no means excluded, we avoid theoretical discussions of inter-temporal social welfare and criteria of sustainability. In other words, we start from the practical user-oriented question of 'how do we make decisions when (very) long term consequences are at stake' and work backwards towards justifications of one's choice.

The rest of the paper is organised as follows. Section 2 reviews the responses offered to the long term discounting problem. Section 3 proposes an alternative approach to discounting the long term. Section 4 presents a simple model of radioactive waste containment and discusses the main results. Section 5 concludes.

2. The long term discounting problem

Conventionally, the long term discounting (LTD) problem is defined as the rate at which values *beyond* the long term time threshold, T, however determined, are to be compared to values today. In practice, this means values at least 30 years into the future, and typically very much further.

2.1 The three long term discounting approaches

The literature shows there are three possible responses to the LTD problem. These are, in their pure form:

- No discounting or zero discounting of (very) long term impacts.
- Use of a single constant positive discount rate.
- Use of a declining sequence of positive discount rates.

The last category subdivides as follows [5]:

- Ø Use of two discount rates, a short term and a long term one.
- Ø Use of more than two rates, in declining order
- Ø Use of a non-constant declining discount rate as a function of time. (This case corresponds to an infinite number of declining discount rates.)

Hyperbolic discounting is a typical example of the last category.

There are also hybrid responses, such as the Chichilnisky criterion (1996, and Åsheim, 1996), which is a weighted average between a constantly discounted and a non-discounted term for future values.

The no- or zero-discounting response may be seen as two variants of the same option. However, they are not theoretically equivalent, even if they are so in practice. No discounting reflects a refusal of the principle of discounting for the long term, whereas zero-discounting accepts this principle but argues that the LT discount rate should be zero. The refusal variant is generally made on ethical grounds. Environmental activists or ideologists have often adopted this position, but Thomas Schelling (1995, 1999), an economist, also argues for this solution. Zero discounting, if chosen, must however be justified on economic grounds. Both rationales will be examined further below.

The single (positive) discount rate response expresses the fact that both the short term and the long term, that is, values *before* and *after* the threshold T, must be discounted in a similar manner. This is the standard exponential discounting used in finance, but also in optimal growth theory, and more generally by mainstream economists working within the so-called neo-classical paradigm.

[5] This categorisation is simplified on practical grounds. An elaboration is offered in Appendix 1.

The multiple discount rate response is a more recent idea, although Ainslie's challenge to (single rate) exponential discounting dates back to 1975. The number n of discount rates can be directly related to the number n-1 of time thresholds T_i, $i \in [1, n-1]$. Thus the use of two rates, a higher and a lower one, imply a single time threshold. A different use of the idea is found, implicitly, in Fisher and Krutilla (1975), where technological decay and increased scarcity of natural resources can be interpreted as two different discount rates over the *same* time horizon.

If n rates are used, they appear in decreasing order, with lower rates used further into the future. The use of a decreasing, rather than increasing sequence is based on empirical knowledge: people have been systematically observed to discount far future values less than near future values (Loewenstein and Prelec, 1992, 1993)[6]. In the limit, as the number of time thresholds and corresponding discount rates increase to infinity, one obtains a continuously declining discount rate function. Hyperbolic discounting is a typical example. An example of a discrete sequence of (five) declining discount rates is given by Weitzman (1998, 2001). However, this sequence is underlain by a continuous gamma distribution, leading Weitzman (2001) to call it 'gamma discounting'. The reasons underpinning the use of multiple declining discount rates are several, and will be examined later on.

Given the three – explicit or implicit – discounting possibilities, how are we to choose? As it turns out, each holds surprises of its own.

1) Not discounting or zero-discounting the long term leads to problems of short-term ($t<T$) inefficiency. This is because the computed value of investments is distorted towards benefits arising after T, and there results under-investment in projects yielding benefits before T (where typically $T \approx 30$ years).
2) The use of a single constant discount rate seems to lead to problems of intergenerational equity (IG-equity). If too high, future generations may be sacrificed; if too low, the current generation may be sacrificed (see Huxley's quote). This suggests the existence of an (unique) 'optimal' value of the discount rate, which is indeed the solution to an (unique) intertemporal optimisation problem. This however

[6] Frederick (1999) found evidence of increasing discount rates, but their interpretation (Rubinstein, 2000) is unclear, and may be due to the specific experimental setup.

raises questions of what function is to be optimised over generations, and, through the choice of an intertemporal social welfare function, reintroduces ethics and subjective, if not arbitrary, choices.

3) The use of a declining sequence of discount rates seems to lead to problems of time inconsistency: decisions which are optimal from today's perspective are no longer optimal from tomorrow's perspective, even if nothing new happens (no new information). Strotz (1955/56) qualified this phenomenon as a form of myopia.

Two issues are at stake here:

- the relative role of efficiency and equity concerns;
- time inconsistency.

The first issue runs much deeper than the second. At first sight, it seems that the long term discounting problem can be said to stem from a conflict between criteria of (intertemporal) efficiency and (intergenerational) equity. This is indeed a problem, as discounting is a tool for achieving economic efficiency, not equity. If equity is a concern, how then can it be catered for?

2.2 *Efficiency vs. equity: lower the discount rate, somehow!*

The literature offers a number of strategies that have been adopted to address the problem of intergenerational equity over the long term. However, all share a common feature: they involve lowering the discount rate somehow. Unfortunately for a unified understanding of the issue, the rationales for doing so vary widely. Let us briefly review them.

a) Set pure time preference equal to zero

One strategy, in public decision making, is to start with Ramsey's well-known equation,

(1) $r = d + eg$

where r is the social rate of discount, or the social rate of time preference (SRTP), d is the pure rate of time preference (PRTP), or utility discount rate, e is the elasticity of marginal utility of consumption or income (depending on the numeraire used), and g is the annual growth rate of consumption or income, that is, the growth rate of the economy[7]. As Ramsey himself argued in 1928, d can, and should, be set to zero on ethical grounds. Although individuals may show impatience for personal consumption, society as a whole may not. This is a normative, rather than a descriptive, stance. It implicitly gives individuals of future generations the same weight as individuals today.

Besides being normative, this approach focuses on society's consumption stream, rather than on its production capacity. In a perfect market with complete futures, both aspects would equate: the rate at which the value of future consumption declines, and the rate of return on current investments. However, such is not the case. Thus the question remains as to which aspect should be given precedence, consumption or production.

b) Consider the possibility of future negative growth

The elasticity factor (e) in equation (1) above is an empirically observed quantity, varying mostly between 0.5 and 1.5, and often taken to be equal to 1 (which implies a logarithmic utility function, which was Bernouilli's (1738) initial proposal). The only other solution left for lowering r is to have a pessimistic view of g. After all, annual growth rates can be, and have been, negative. However, over long periods of time, it has been positive, and typically ranges from 1 to 3% in real terms. One is left with the question of over what period of time g should be averaged, as well as over what scale g should be measured: sectoral, national, or global? The above figures represent the 200 years following the industrial revolution (1800-2000); who is to say that they will be valid for the next 200 years?

c) Offset discounting by increasing future costs

A related approach is to consider, as Fisher and Krutilla did as early as 1975, that costs associated with an investment will increase over time. In environmental applications, this could represent increased natural resource scarcity, such as the disappearance of wilderness areas. Thus the burden put on future generations by our current discounting of future impacts, both benefits and costs, are offset by increasing

[7] If C is consumption or income at time t, and U(C) is utility of consumption or income, then $g = C'/C$ where $C'=dC/dt$ and $e = -CU''/U'$ where $U'=dU(C)/dC$ and $U''=d^2U(C)/dC^2$.

future scarcity costs. Note that this is absolute scarcity. Discounting is further offset by a technological factor, which reflects the increasing obsolescence of current manmade assets. As and Pearce and Turner (1990: chapter 20) point out, this double whammy on discounting can be reinterpreted as the use of two different discount rates: one, augmented, for development projects, and one, reduced, for nature conservation projects. However, this approach targets relatively small publicly funded investment projects that do not alter the economy as a whole, a point made by Quiggin (1993, 1997).

d) Use a declining discount rate function

Yet another strategy to lower the discount rate for the far future is to use a declining discount rate function. Typically, a hyperbolic function has been used, mainly on grounds of observed behaviour (e.g. Loewenstein & Prelec, 1992). They observed that decision makers account for both the absolute and the relative timing of events, that is, the duration between two events and their timing with respect to the present. The problem here is the normative status of this descriptive human behaviour, as well as that of time inconsistency (Page, 1977) and its effects on project appraisal (Cropper & Laibson, 1999). There is also some debate as to whether behaviour in non-market laboratory conditions is a valid reproduction of the market-sanctioned decisions we make in real-world environments.

Recently, empirically observed hyperbolic discounting has been given an alternative explanation (Read, 2001; Frederick et al., 2002), that of ‘sub-additive discounting’. The fall of the discount rate over time then appears as an illusion produced by the product of two factors: the interval-density effect, and the zero-time interval origin. The first factor describes the extent to which a time interval (T, T’) is partitioned into sub-intervals (T, T_1, T_2, …,T_n, T’). Read (2001) showed that the total amount of discounting over (T, T’) increases with n, even though (T, T’) remained unchanged. The second factor refers to the fact that in practice we discount from the present, that is, over the interval (0, T’). Thus, as T’ recedes into the future, the (implicit partitioning) of (0, T’) becomes increasingly ‘coarser’, leading to an apparently falling discount rate.

Although the end effect is a hyperbolic function, the underlying explanation is different from the usual one that relies on absolute and relative positions of T and T'[8].

e) Account for uncertain future discount rates

Weitzman's (1998) idea, which he later presented under a different angle (Weitzman, 2001) and which has recently been advocated by Resources for the Future (Newell & Pizer, 2001), uses yet another strategy to lower the discount rate for long-term consequences. His idea is best illustrated in reference to the Ramsey formula, although reference to this formula is not necessary. It is based on the following mathematical insight. Suppose several growth rates are possible in the future (g in equation 1), and therefore several discount rates are also possible. Suppose for simplicity that each rate is equally likely (uniform probability distribution). Then it turns out that, the further you compute into the future, the quicker the higher discount rates discount themselves away. Far enough into the future, you may as well use the lowest of the possible discount rates: it will be the only one 'surviving' its self-erosion effect. The effect typically becomes noteworthy 200 time units into the future (years if annual discount rates are used).

The outcome of this approach is a declining sequence of discount rates as a function of time – that is, the further one computes (today) into the future, the lower the discount rate one must use. Weitzman (1998, 2001), by considering only integer-valued rates, defines five time horizons and, correspondingly, five discount rates. Beyond 300 years, it is zero.

This is a very neat idea, and one that does the job pretty well. The problem with it is that it is hard to relate it to any concept of equity, intergenerational or otherwise. It is purely a consequence of our ignorance of the future. Now, it may be that there is some intrinsic link between attitudes to uncertainty and attitudes towards time as well as (intra- or inter-generational) equity, as some recent work seems to suggest[9], but no such link is

[8] It is intriguing also to relate these studies to French philosopher Henri Bergson's treatment of the measurement of psychological time, as opposed to physical time. Bergson, in 'Durée et Simultanéité', wrote after Einstein produced his Special Theory of Relativity in 1905, and speculated that in the human mind time itself was elastic, so that an hour in a year's time should be measured differently to an hour starting now. This intermingling of psychology and physics has taken new twists with quantum theory.

[9] See Kroll & Davidovitz, 1999; Amiel & Cowell, 2000 and 2001, and the paper by Quiggin and Horowitz (1995).

implied by Weitzman's idea[10]. So if we are happy to use this approach to reduce the discount rate for the far future, we can do so, but not on the grounds of intergenerational equity. In other words, we are not addressing our concerns for future generations. The fact that the result happens to favour them is a mathematical fluke. This may be seen as hardly satisfactory.

More importantly, since the choice of the long-term discount rate is extremely sensitive to the lowest possible rate, how are we to identify this lowest rate? If it is by using past data (e.g. on treasury bonds), how are we to choose the period of time over which to define a lowest rate? Note too that Weitzman's model does not accommodate a negative lowest rate – a not inconceivable case with prolonged negative growth[11].

From the review above, the conclusion that seems we should draw is that, however we decide to deal with very long term consequences, it would be wise to clearly distinguish efficiency concerns from (intergenerational) equity concerns. This would allow, as is customary in economics, to use the full-efficiency, no-equity solution as a benchmark, relative to which one can measure the degree of equity one is willing to show. This is what Schelling (1995, 1999) forcefully recommended. For all practical purposes, it is useful to define as 'purely economic' decisions those that only consider costs and benefits, however determined, to the decision maker's own interests or utility function. This does not necessarily exclude altruism, if the decision maker happens to derive utility from others' welfare; however, considerations of altruistic utility functions needlessly muddy the waters, in that it is always possible to reinterpret them as 'enlarged selfishness'[12].

[10] There is also some questioning as to whether, on purely mathematical grounds, such a procedure is totally legitimate, as it seems to imply the computation of a stochastic integral, which is not possible.

[11] Prolonged negative growth rates are especially plausible when one considers some countries with high levels of resource destruction with little capital build-up and a measure of growth in terms of 'green' net national product, or NNP.

[12] For example, catering for one's family, clan, company or nation, rather than for one's self, can be considered as altruistic if the individual is taken as a point of reference, but as selfish if the community of all other families, clans, companies or nations is taken as the reference.

2.3 *Discounting, pure time preference, and time inconsistency*

The second issue is that of time inconsistency. If the discount rate is non-constant over time, decisions become inconsistent in the sense that they will be reversed at some point in the future, and the decision maker will make the decision today knowing she will reverse it later on, even in a perfectly foreseeable world. As Solow (1999) puts it, ‘this is a poor way to run a railroad’. However, this is not the end of the story.

In the above discussion over how the discount rate behaves over time, reference was implicitly made to the ‘social rate of time preference’, as defined for example in the Ramsey formula. It is a quantity that combines both the rate of pure time preference (RPTP), or utility discount rate, and a rate depending on the productivity of capital, linked to the average market interest rate. Although not clearly spelled out, basically all of the empirical literature on hyperbolic discounting deals with the RPTP of individuals concerned with their own private utility function. Experiments typically involve trade-offs over time of two *fixed* quantities, whether of money or otherwise: a smaller quantity sooner versus a larger quantity later, *all else being held constant*: the individual’s wealth, the productivity of her assets, and the state of the economy. In other words, the capital growth or productivity component of the total discount rate is held constant and not accounted for. The conclusion is that hyperbolic discounting, however interpreted, applies only to the RPTP or utility discount rate. It does not apply, at least not on the grounds of this empirical literature, to any social discount rate. And in particular, it is difficult to see in it any justification for its normative use for public decision making.

This should come as no surprise to anyone who has read the original paper that Paul Samuelson wrote in 1937. This is where the concept of a constant utility discount rate was first proposed, leading to the conventional use of the exponential utility discount function. Although, perhaps due to its mathematical elegance, economists took up this proposal and have run with it ever since, Samuelson took great pains in that paper to warn against any normative or descriptive interpretation of his formulation, stressing the fact that it was absolutely arbitrary and that ‘there is no reason why people should behave in this way’. As a result, we should be happy to adopt, from now on, a hyperbolic function for utility discounting. Any constant ‘rate of pure time preference’ can be dismissed as an irrelevant parameter. Note that Samuelson was very clear that he was dealing with utility discounting, not consumption or income discounting.

This helps us clarify the issue of time inconsistency. If utility discounting is hyperbolic, how do we deal with the fact that it leads to time inconsistent decisions? But this is forgetting that utility is utility of something, whether of income, consumption, or otherwise; strictly, consideration of future utility cannot be made independently of some economic state variable. In a remarkable paper, Pezzey (2002) manages to dispel the confusion by showing that when the discount factor (not rate) is made to depend not only on time itself, but also on some economic state variable s_t at time t, time inconsistency disappears. Its existence is due to the fact that the *instantaneous* discount rate is not reoptimised at every t as a function of s_t. Thus, if future utility is discounted hyperbolically, the resulting value of future income, consumption or returns on capital can also be discounted hyperbolically, provided that the discount factor at time t is made to depend on the state of the economy, and therefore on absolute time (the historical date). This is compatible with Read's origin-fixed interval interpretation of sub-additive discounting on (0, T). However, there is no reason why the state-dependent part of the discounting function should behave hyperbolically; indeed, given its dependence on the real rate of growth, it is logical that it behaves exponentially with a fixed growth discount rate. As a result, the combination of hyperbolic time preference and exponential income or consumption discounting could take on a complex form.

The last consideration, of course, is that time preference itself depends on some economic state variable, namely some measure of wealth. In their paper, Becker and Mulligan (1997) explore such a possibility.

The above discussion on time preference mainly refers to individual or private decision making. In matters of public decision making, consideration of pure time preference, which is simply a measure of impatience, is rather controversial. Ramsey, in his path-breaking paper (1928), dismissed it as unethical. The question, though, is whether we are interested in its descriptive or normative use for decision making. My view is that if we are trying to explain, ex post, how public decisions are or have been made, then we must consider a descriptive approach to time preference, and work out whether government agencies or other public decision makers reveal some positive pure time preference or not. However, especially considering how imperfect, if not corrupt, many government decisions are, it is the normative approach that should be used when we are interested in recommending, ex ante, a decision with community-wide implications. Unless we have a model of endogenous time preference, we must follow Ramsey's advice (and many others' since) that the only justifiable choice is zero time preference. In this case, whether time preference is hyperbolic or not becomes irrelevant. If government, on the one hand, represents the interests of the whole community and, on the other hand, can and should supersede 'myopic' private preferences, then it must impose its own (absence of) time preference.

2.4 *Capital markets, discounting, and the long term future*

If we accept for now that public decision making should neglect pure time preference and focus on the values underlying the utility function, the question becomes to what extent can capital markets, which determine the marginal productivity of capital and the return on investments, deal with long term consequences? From the decision maker's point of view, whatever her concerns for the far future, will she be able to afford her concerns given the constraints imposed by the cost of capital? This is especially relevant as most of the capital for long term investments is borrowed on the financial market.

In essence, capital or financial markets are institutions for trading values over time (with an essential component of risk trading as well). In practice, values cannot be traded over very long periods of time, like centuries. Capital markets have de facto a time horizon, of the order of about 30 years. There is little, if any, willingness to trade values over a longer time period, although some forms of investment (superannuation, life insurance, in-family intergenerational lending) can exceed this time; however, there are no proper 'markets' that trade such values. As a result, beyond the market's time horizon, there are no prices of capital. At first sight, markets cannot cater for the very long term.

A) On closer inspection, however, this view rests with how efficient capital markets are deemed to be. If they are perfectly efficient, then each generation will be able, in a perfectly overlapping generations model, to trade values with the next generation, and so on into the future. Individuals behave as if the were infinitely lived and as if they were maximising their benefits over an infinite time horizon. This is essentially the approach adopted by the neo-classical optimal growth literature. Compensation between earlier and later generations happens gradually and automatically, and welfare or utility can remain non-decreasing over time.

Of course, such a view presumes perfect micro-macro coordination at all scales, and that all assets, whatever their nature, can be traded in capital markets. This also implies perfect substitutability between tradable assets. As well we know, such a view is overly optimistic, especially when natural resources and environmental assets are at stake[13].

B) The opposite view states that the existence of non-substitutable or non-tradable assets, such as most environmental values, creates a time barrier for capital markets (the economic time horizon), in which case they are totally incapable of catering for the long term. The only way to do so is to use ethical motives and show concern for future generations. However, this also means a willingness to pay for the far future that will in any event be limited by the current income stream, both private and public. It is easy to understand why proponents of this view advocate strong governmental action to force private interests into catering for the long term. But they run against the possibility that government failure be at least as bad as market failure, if not worse.

In an ingenious paper, Ari Rabl (1996) proposes notwithstanding a two-stage 'inter-generational' discounting method that can handle the absolute time barrier. This amounts to using two discount rates, a higher short term one and a close-to-zero long term one (beyond the 30 year time horizon). He applies his method to nuclear plant construction and radioactive waste disposal (Rabl, 1999). His idea is based on compensating future generations with the extra capital build-up (or economic growth) that such investments will have made possible. Philibert (1998), in a review article, points out that this 'impossible' intergenerational compensation, based on such a mechanism, runs into problems of time inconsistency at T=30 as well as future continuation of growth. Rabl's paper is still a valuation exercise in long term cost benefit analysis.

[13] The efforts put into environmental valuation and into creating tradable property rights will help, jointly, to bring about a state of affairs tending, however imperfectly, towards this ideal situation.

C) There is of course a middle view, whereby capital markets are seen as imperfectly efficient, with the tradability of assets linked to their degree of substitutability and to the irreversibility of their possible future loss. These characteristics are closely related to the assumptions one makes over the possibilities of technical progress, which, in terms of endogenous technological change, must depend both on the incentives to generate R&D and on the economic efficiency of R&D. Thus the key variable becomes the willingness to invest in those types of investments that can offset in some way the losses imposed on future generations. This offsetting must however deal with the severity of the so-called 'commitment problem', where generation 2 can step in and destroy (or consume) what generation 1 invested for generations 3, 4, and beyond. What is needed is an irreversible investment that intermediate generations cannot undo. There is only one form of investment that stands up to the challenge: investment in knowledge and R&D. To the extent that the knowledge produced has any value, and that it is governed by increasing returns to scale, it will, as discussed below, be safeguarded and available for the future.

3. An alternative to the long term discounting approach

3.1 *Introducing the idea: investments with irreversible benefits*

Let's face it: the long term valuation problem is an intractable one. At least, it has been so far. This suggests the idea that if we cannot base long term decisions on long term valuations, we need to base them on something else. In other words, we need to seek a way around the long term discounting problem and avoid it altogether, albeit in a way that allows for the opportunity costs of investments.

Compounding this, Schelling (1995, 1999) makes the point that costs and benefits to future generations are akin to costs and benefits to 'foreigners' in programs of international aid, and therefore the issue is purely one of equity, be it intra- or inter-generational. As discounting is a tool for maximising efficiency, discounting has no place in the valuation of long term consequences. I find Schelling's argument compelling. As stated earlier, it is wise for clarity's sake to keep efficiency and equity concerns clearly separate. If long term decisions cannot be based on long term valuations and yet depend on some type of (perhaps implicit) valuation, the only solution left is perhaps a startling one: they must be based on short term valuations!

This amounts to asking how a decision maker with no ethical or altruistic concerns for future generations can nevertheless cater for their welfare while making his decisions within his own time horizon, and with presumably positive discounting. Can medium term (< 30 years) decisions handle long term costs and benefits without the need to value these costs and benefits? More, can the decision maker optimise within his own limited economic decision making framework as if he were optimising over the very long term? And if not, can he be led in a short enough time to do so in a way that is beneficial to him and today's society?

The only way this can be possible is to use the R&D investment idea suggested earlier (section 2.4,C), and invest in specifically targeted cost mitigating R&D effort. The importance of this control variable lies in the irreversibility of its benefits. If its benefits were to decline in the future, it must be that other knowledge and other technologies have superseded its value, but thereby reducing future costs even further. (Knowledge, know-how and technical progress lead to further knowledge, know-how and technical progress, continually modifying the relative value of existing technologies.)

This notion of investments with irreversible benefits is different from Arrow's discussion of irreversible investments (Arrow, 1999). Arrow, rather quickly it seems, dismisses the idea of irreversible investments on the grounds of the commitment problem. This is when there is no guarantee that the next generation, or some other after it, will continue to invest at the same rate as the current one. But with irreversible benefits, and with otherwise increasing future costs, every successive generation will have an incentive to continue investing in irreversible benefits, simply because, as discussed below, it will itself benefit both directly and indirectly. Does it look like we are about to reduce our R&D efforts towards mitigating the possible effects of climate change? The only reason why such an effort might decline in the future would be because, for one reason or another, the costs of climate change would start decreasing. Arrow's argument does not address this aspect, and his reference to 'elementary theory' might explain why. Holmes might in this case refrain from answering Watson's questions in the usual manner!

It also matters that such R&D investments be specifically targeted at reducing the specific costs that will weigh on future generations. If radioactivity is the future risk, then R&D targeted at reducing the impacts of radioactivity on living organisms is such an investment. (This would entail fundamental research on the biophysical interactions between radioactivity, molecular biology and genetics.)

The key point, however, is that any such investment will not only produce benefits (e.g. in the form of reduced risks) in the long term, but may also do so in the medium or short term, within the economic time horizon. This will happen in two ways: directly, through reduced radioactive risks within an individual's own lifetime, and indirectly, through positive spin-offs creating value in other spheres of health and environment. These two sources of benefits will add up and will constitute the 'internal' benefits of the R&D investment program. The 'external' benefits are those to people beyond the economic time horizon.

Within this framework, the problem becomes the following: can the direct plus indirect inner benefits be sufficient so as to initiate a stream of (R&D) investments that will also offset future damages? And if so, under what conditions? How is this problem to be properly formulated, and can a negative answer suggest ways for overturning it?

3.2 *Analysis of an illustrative case: the containment of highly radioactive waste*

As always, the problem will be simplified and a specific case will be examined. The containment of a given quantity of highly radioactive waste appears as the simplest problem of this kind. The primary decision problem can be taken as the level of structural investment necessary to achieve a socially optimal level of risk from today's point of view, without consideration of the long term future. This means optimising the containment investment in terms of the risks faced by the current generation (all those already born, or at least conceived). In practice, this sets the time limit for the accounting of costs and benefits at about 100 years, which spans 3 to 4 generations (assuming very few people alive today will be alive then). Such an optimum from today's generations' standpoint will be suboptimal for people further into the future if, as is the case, risks increase with time due to the aging of the containment system and to other, external causes (land movements, earthquakes, leaks into groundwater, terrorism, etc.).

The problem is then to find under what conditions the long term ‘external’ costs can be offset in a way that will yield positive net ‘internal’ benefits within this the decision maker’s economic time horizon. This implies a second decision problem and the use of an ‘auxiliary’ control variable, a stream of R&D investments in radioactive damage mitigation. Unlike the primary (structural) investment (S), which is a static control, this one is a dynamic control over time – a level of investment for each year over a certain period of time{I(t)}. The time over which {I(t)} will be made needs to be optimised. Their cost is a function of the opportunity cost of capital, linked to its marginal productivity in the current state of the economy. Obviously, the relative rates of return on capital and of increase in future risks will matter, but this increase in future risks is a function of the level of investments. The benefit side of the equation is measured by the reduction in the level of risk relative to the no-investment case (I(t)=0)[14].

R&D investments may be more or less efficient in reducing risks over time, and will exhibit decreasing marginal efficiency with respect to the level of I(t). This efficiency is of course a crucial parameter, but it depends on the efficiency with which the R&D effort (measured in economic terms) is itself organised. In other words, part of the R&D effort can be directed at increasing the efficiency with which it can mitigate future impacts. As a result, the extent to which an increase in I(t) reduces future risks can itself increase over time, as the use of knowledge becomes increasingly efficient. This is best modelled as a trend in time.

Uncertainty is an intrinsic aspect of this problem, and a number of variables should be considered as random. However, for the time being and to keep things simple, we shall only work with expected quantities, implying risk neutral decision makers throughout. This is done in order to focus on the main concepts, but it must be borne in mind that the consideration of uncertainty will have nontrivial consequences and is likely to affect whatever results a deterministic treatment may achieve.

---

[14] As will be specified below, both S and I(t) reduce future radioactive risk, but in different ways. If risk is a function of p, the probability of an event, and of C, the damage cost of the event, then S reduces p while I(t) reduces C.

A second simplification is that zero pure time preference is assumed throughout. This does not affect the essence of the problem. There are good reasons, stated earlier, that justify such a choice (radioactive risk is a public issue). However, a (hyperbolic) time preference function could easily be superimposed on the values of risks at time t.

The two control variables S and I(t) are not on the same footing. An optimal stream I(t) is sought subject to a given level of investment S. However, S is optimised given the impact of an increased probability of an event on the current and future level of risks. If S is reduced (to save on up-front expenses), today's risks will be higher and future risks, all else equal, will increase more rapidly (the containment system being of poorer quality). This could be offset by a higher stream of I(t), which would have the advantage of spreading the expenses over time; but today's people would have to take higher risks in the meantime (R&D efforts take time to produce a noticeable effect), and it would take some time until risks are brought back to the level they were with the initial S. This said, the main problem considers S as given and seeks to offset given external costs using the I(t) control.

We will now specify a simple model to investigate this problem. There are some technical details that pose a conceptual challenge over which we shall skip. One such problem is the measurement of expected risk when probabilities are very small but the (contingent) costs are potentially huge. In our case, the probabilities are of the order of 10^{-6} and damages of the order of $\$10^{8}$. The expected value appears to be worth \$100, but (see Collard, 1988) the principles of conventional expected utility theory with such extreme events become very controversial. Still, in this exercise we shall happily skip over the issue and consider the value of the risk to be simply worth \$100.

The key variable that will determine the outcome is the opportunity cost of capital, or the rate of return on capital, given the state of the economy. Because the problem is micro-economic in nature, this rate is considered as exogenously given. In further developments it will be interesting to make it endogenous to the investment decision, for example if the externality is related to climate change. Important parameters are the size of the investment relative to the available capital stock (b), the value (C) of damages if an event happens (in this exercise a simple number, but in a more developed model a probability distribution across events of varying severity), the parameter function k(I) governing the cost reduction efficiency of I(t), and that (π) governing the probability of event reduction of an increase in S.

The name of the game is to look, as a function of r given S, π and k(I), for investment streams I(t) that yield positive net benefits within the current decision maker's time horizon and which nevertheless reduce long term risks. The trick here is that the time horizon does have an upper bound, but is otherwise to be optimised along with the investment stream. Thus the optimisation problem is two-dimensional in ({I(t)}, T), being understood that {I(t)} stops at t = T* ($\leq$ T), the endogenous. time horizon.

4. A simple model of radioactive waste containment

4.1 *Model description*

Let the amount of radioactive waste Q be given, for example at the decommissioning of a nuclear plant. This could be made up of plutonium and other highly radioactive material. Half-life of plutonium is about 24,000 years. Because, by economic standards, the decay rate is so small (less than 0.003% p.a.), the amount of radioactive material can be considered as constant over the first few centuries (it will have decreased by only 1.5% in 500 years!), a period well beyond both private and public decision making. This waste can be confined and stored in various ways, each having different degrees of safety now and in the future. The more is invested in the confinement system, the lower are the risks. Whatever the confinement system, two things are true: the probability of a radioactive mishap, p, no matter how small (say 10^{-6}), cannot be exactly zero, and, secondly, this probability increases over time. We assume it does so exponentially at a very slow rate, π (say 1% p.a.), although this rate is 3 orders of magnitude greater than the radioactive decay rate:

(2) $$p_t = p_0 e^{\pi t}$$

and in this particular case

(2') $$p_t = 10^{-6} e^{0.01 t}$$

Let S be the structural investment in the confinement system, possibly deep under ground. And let S_o be the level of investment that is optimal for current risks today. S_o

reflects current willingness to pay for today's safety, regardless of the future. It determines the probability p that a radioactive event (accident, leak, fracture, explosion) could happen in period one (this year or this day). S_o also determines the rate π of increase in the probability of a mishap. An increase dS will both decrease p and π, in a way assumed to be in the simplest way hyperbolic:

(3a) $p(S) = \phi \frac{1}{S}$ with $p_0 = \phi \frac{1}{S_0}$ determining the value of φ,

(3b) $\pi(S) = \psi \frac{1}{S}$ with $\pi_0 = \psi \frac{1}{S_0}$ determining the value of ψ.

φ and ψ are scaling parameters determined by the initial values, and $S = S_0 + dS$.

Investment S_o is drawn from a pool of initial capital, K_0, leaving available for further investments $K_0^S = K_0 - S$. S_o has an opportunity cost determined by the general health of the economic system, which determines the rate of return on, or the marginal productivity of, capital, r. In a perfectly competitive financial market, this also corresponds to the riskless interest rate. At time t, the opportunity cost of dS can be written as

(4) $C_S(t) = dSe^{rt} - dS = dS(e^{rt} - 1)$

Let {I(t)} be the stream of R&D based mitigation investments to reduce the future impacts of radioactive events. I(0) is the initial mitigation investment at time t=0 and I(t) is any future mitigation investment at time t. This investment, specifically targeted at reducing the impacts of future radioactive events, is made possible by diverting a portion of available capital stock $K_s(t)$ at time t, which would otherwise have been reinvested in K(t+1). Let b be this proportion of $K_s(t)$, such that

(5) $I(t) = b.K_S(t)$.

b would, given the orders of magnitude used here ($K_0 \approx \$10^9$), typically be a small proportion of the available $K_S = K_0 - S$, of the order of 0.1% to 0.5%,. In this study, we have kept the value of b constant over time. This means that the investments will increase in proportion to capital build-up, and decrease with any erosion of the capital stock.

The effect of such an investment is to reduce the costs of any future radioactive impact, whether due to damage to people or the environment, or both. Assume that without any such investment, costs remain unchanged; that is, if an event were to happen at time t, the cost of its impact, C, is a given independent of t:
If I(t) = 0, C(t) = C = constant.

A more rigorous model should of course write C as a random variable, and also, insofar as the valuation of damages depends on current wealth, let its value depend on K(t)[15]. Instead, we assume for now no trend in either its mean or variance. In this model, we are only considering the expected values and considering them constant, measured in real terms. (This is not a necessary assumption, but it keeps things clearer.)

We define contingent costs as the product:
(6) CC(t) = p(t).C
Because p(t) is increasing over time at a rate dependent on S, CC(t) also increases at rate $\pi = \pi(S)$.

The damage costs at time t will, relative to no mitigation effort, be reduced by I(t) not only as a function of the current level of effort, but of the cumulative effort up to time t. This is because, if we neglect forgetfulness which today is insignificant, the benefits of R&D outcomes are for all practical purposes irreversible. (This avoids the commitment problem which plagues other long term strategies.) The effect of cumulative knowledge is described by the following equation:

(7) $$C(t) = Ce^{-k(I_0)\int_0^t I(t)d\tau - I_0}$$

[15] We could for example allow C to grow or fall at the same rate as K(t), given K_0 and C_0.

where $k(I_o)$ is the mitigation efficiency parameter of the mitigation effort (or investment). It depends on the level of investment in a way that is increasing at a decreasing rate.
In the current version of the model, we make k depend only on the initial I_o:

(8) $$k_I = k(I_0) = m(1+\ln I_0)^n$$

where m and n are scaling parameters, and ln(x) is the natural logarithm of x.

A more accurate description, to be used in a later version, makes k depend on I_t and on the flow of time itself. This second dependence represents the growth of the efficiency with which knowledge is put to use. It makes the cost reduction efficiency of mitigation effort a function of the efficiency of the use of knowledge. This knowledge efficiency is simply modelled as a trend in time. The complete description is:

(9) $$k_{(I,t)} = k(I(t),t) = k(I_t)(1+hk(I_t))^{qt}$$ where $0 < q \leq 1$ is a scaling parameter.

h is the R&D efficiency parameter which commands the time trend for k(I). With this specification, the damage cost at time t becomes:

(7') $$C(t) = Ce^{-k(I_t,t)\int_0^t I(t)d\tau - I_0}$$

The actual risk at time t is then defined using the above equations:

(10) $$R(t) = p(t)C(t)$$

With respect to the mitigation investment effort, we are interested in how much risk is reduced at time t. Risk reduction dR(t) is given by:

(11) $$dR(t) = CC(t) - R(t) = p(t)C - p(t)C(t) = p(t)[C - C(t)]$$

This can be considered as the benefit side of the mitigation investments. The cost side is given by their opportunity cost C_m, which at time t is:

(12) $$C_m(t) = K_S^0 e^{rt}(1-b)^t$$

This completes the description of the model. We are then interested basically in two things. First, given the initial structural investment S, what is the optimal stream of mitigation investments? Secondly, if S is allowed to vary, that is, if we consider ex-ante the decision of how much to invest in the storage and containment of radioactive waste, we are interested in knowing which combination of S and {I(t)}, given r, will yield the maximum net benefit over a period of time?

This requires the definition of net benefit. In this model, it can be defined as the net capital outstanding at t plus the surplus value of cumulated risk reduction over the cumulated opportunity cost of the mitigation investments. The net capital at t is the amount left after having each period diverted a share b of $K_S(t)$, relative to what would have been available with b=0; that is, if all available capital had been reinvested in the economy at the going interest rate (assumed constant over time). The surplus to be added to this quantity is the difference, up until t, between the cumulated value of risk reduction (assuming an additive utility function, which is linear) and the sum total of opportunity costs of {I(t)}up until t.

(13) $$NB(t) = K_S e^{rt}(1-b)^t + C\int_0^t p(t)(1-e^{Z(t)})d\tau - \int_0^t [K_S e^{rt}\{1-(1-b)^t\}]d\tau$$

where $Z(t) = -k(I_t, t)\int_0^t I(t)d\tau - I_0$ (see equation 7')

This quantity is akin to that of a net present value, with the discount rate being the rate of return on capital, r. This r could be considered random, rather than constant, and this could have non-trivial effects on the results; but we shall leave this issue for later work.

Model parameters have been chosen to represent plausible orders of magnitude:

K_0 = \$1100m
S_0 = \$ 100m
$K_{0\_S} = K_0 - S_0$ = \$1000m
$p_0 = 10^{-6}$
π = 1% per annum

C = $100m (to be interpreted as an expected value, written down in a single period)
Scaling parameters m = n = 1, φ=100, $\psi=10^6$

The values of m and n determine the damage reduction efficiency of mitigation effort, k(I). Set both at 1, a value of I = 1 ($1 million) puts k at 0.69%, while I = 5 puts k at 1.79%. This represents a moderately efficient scenario for mitigation investments, rather on the optimistic side regarding technological prospects. Setting m = 0.5 would represent a much more inefficient outlook. Then I = 1 would put k at 0.35% and I = 5 only at 0.90%. Setting m = 2 would represent a very efficient R&D impact. To get a feel of how the model responds to these values, set r = 4% and compare I(0) = 1 ($1 million) and I(0) = 2. Then the results in terms of radioactive risk at a given time are as follows, with initial risk at 100 (Table 1):

Table 1 With r = 4% and RR = risk reduction ratio (initial R = 100)

m	I(0)	Risk(T=30)	RR
0.5	1	110	
	2	72	1.5
1	1	90	
	2	39	2.3
2	1	60	
	2	11	5.5

The model was implemented using Excel®, but only some preliminary calculations have so far been done. Full results using this model will be presented in a later publication.

4.2 *Indicative results and discussion*

What follows are only some preliminary indications of qualitative significance. Their aim is at this early stage only to highlight the sort of outcomes one can encounter, and to suggest further questions. This is by no means a full treatment of the problem, but rather a specific experiment carried out for heuristic purposes.

A first outcome of the model is that because of the nonlinearities involved, the risk function over time is not necessarily monotonic (see figures 1, 2, 3). For intermediate values of I(t), and depending on other model parameters, it may exhibit either a U shaped or an inverse-U shaped curve, and this may happen at all time scales, in the short as in the long term. Only with the lower and higher values of I(t) can risk, respectively, increase and fall monotonically. This result, especially for intermediate values of I(t), makes the problem of a ‘sustainable’ or ‘non-decreasing’ level of risk a thorny one. With (inverse) U-shaped curves, the peak or the trough changes in maximum height and in distance to the origin on the time axis.

In the falling monotonic cases, resulting from the optimal short term investment schedule, all is well: the short term and long term goals are aligned. In the opposite case, where the optimal I*(t) produces a monotonically increasing risk curve, both current and future generations are in the same sinking boat. In such a case, it will be wise, if there is still time, to revise the structural investment S, so as to reduce, via a reduced rate of probability increase, the rate of increase of future risks. This action is favourable to both current and future generations.

In the U-shaped and inverse U-shaped cases, two situations can happen (Table 2). The U can happen in the short term, before T, or later, in the long term. Both inverse U curves are good for the long term future. This means that the current generation, and in particular the decision maker, can cater for the interests of future generations while happily maximising their net benefits over their short time horizon. This holds even though values before the horizon T are discounted at the rate r.

Figure 1: The case of a short term inverse U-shaped curve

r = 3%, I(0) = 0.85, k(I) = 0.62%

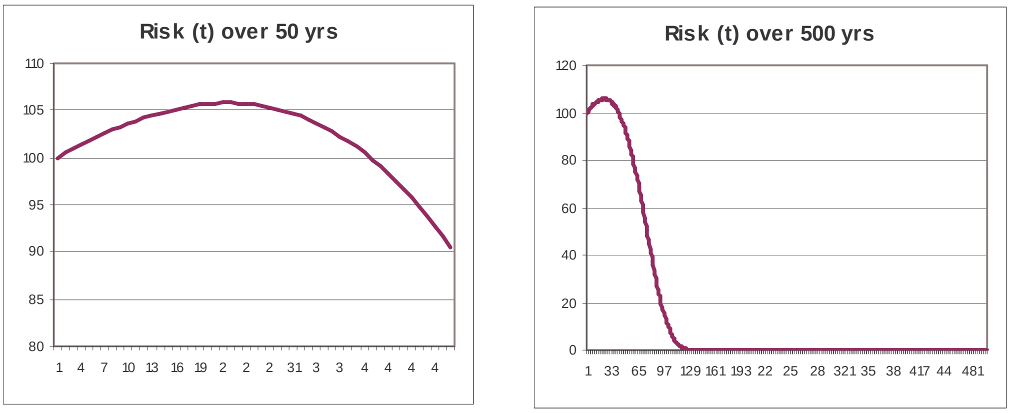

Figure 2: The case of a long term inverse U shaped curve, monotonic in the short run

r = 3%, I(0) = 0.10, k(I) = 0.10%

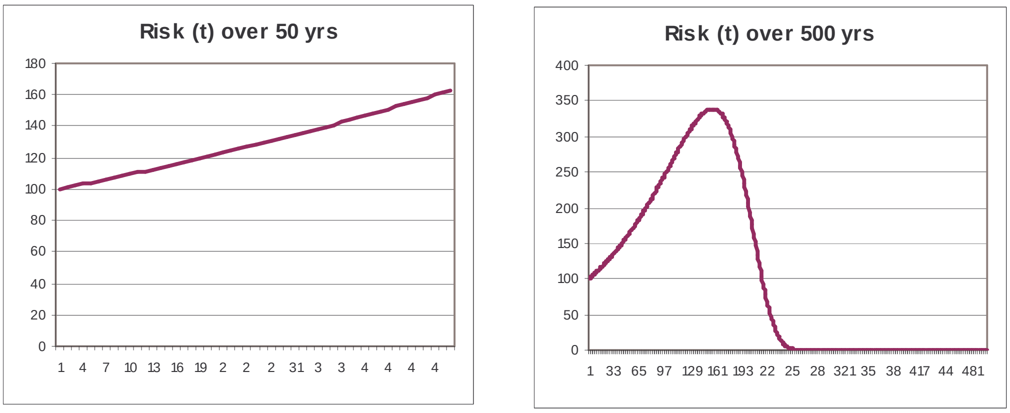

Figure 3: the case of a monotonic falling risk curve

r = 3%, I(0) = 1.5, k(I) = 0.92%

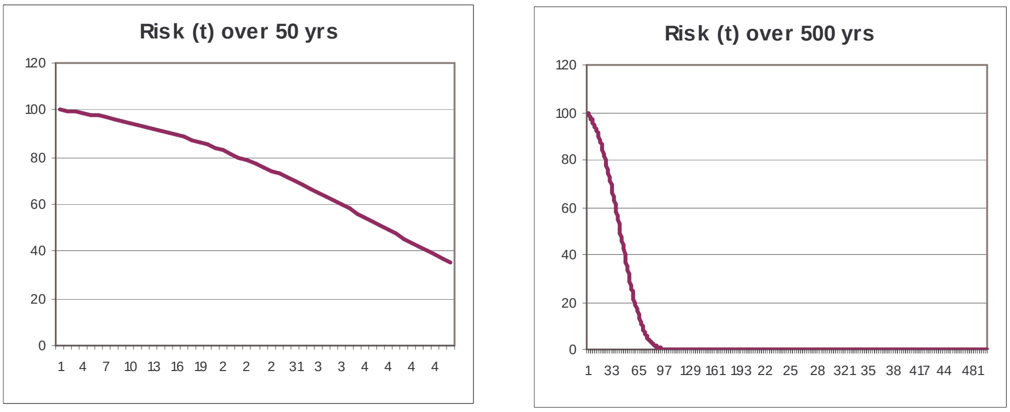

In the U shaped cases, a short term reduction in risk will correspond to a long term increase in risk. Current and future generations have opposite interests. The optimised investment strategy is favourable for current generations but not for future ones. The short term case (see Table 2), where risks first decline then rise before T, is not 100% bad. It may happen that the investment strategy, though optimal, is nevertheless unfavourable to the present generation: the net benefits are negative because the risk curve rises sharply before T, in which case there is an incentive to resort to the second control variable, S, by increasing the quality and security of the containment system. The model clearly shows, through the importance of the relative values of S and K, that we are dealing with a small scale problem: both I(t) and S must remain small relatively to K. Otherwise, the problem has macro-economic implications, with r becoming endogenous and total K becoming variable through, for instance, borrowing in international capital markets.

The long term U case (see Table 2) is hopeless. This is a dynamic ‘trap’ against the far future. It is the only case where, without some ethical concern for future generations, the R&D investment strategy will not solve the problem.

There is however another way, which is to modify the knowledge efficiency parameter h or increase the parameter m governing the k(I) function. The first approach involves the more efficient use of knowledge; the second involves institutional and technological restructuring for increasing the responsiveness of damage cost reduction to increases in R&D investments. Of course, such efforts would need to make economic sense and yield positive net benefits within the current time horizon; the problem here is that there is no easy way to measure such benefits, other than to run a model of this kind!

Table 2: Risk outcomes as a function of short term (t<T) optimisation (all potential cases)

	Rising LT risks	Falling LT risks
Situation LT Outcome	Monotone rising risks ↑ => Action: increase S **GOOD or BAD**, depending on dS	Monotone falling risks ↓ **GOOD**
Situation LT Outcome	LT, U-shaped curve ∪ = falling in the ST = rising in the LT **BAD**	LT, inverse U-shaped curve ∩ = rising in the ST = falling in the LT **GOOD**
Situation LT Outcome	ST, U-shaped curve ∪↑ = falling then rising in the ST = rising in the LT **BAD** most of the time (dS = ?)	ST, inverse U-shaped curve ∩↓ = rising then falling in the ST = falling in the LT **GOOD**

Notes: These results only consider I(t) as a control variable, and S and r as exogenously given.
"Action" means using S as a second control for re-optimising I(t).

The above outcomes are quite sensitive to the long term (real) rate of return on capital, r, which in this exercise is considered as fixed over time (a long term average). This parameter also determines the optimal time horizon, when it is shorter than the maximum relevant for the current generations' cost benefit analysis. T* is defined as the time horizon corresponding to the maximum net benefits from I(t) investments. T* and I*(t) form a pair, given r. As Table 3 shows, when the given r increases, T* declines, maximum benefits increase, but, surprisingly, the necessary (minimum) investment I*(t) increases only very slightly. This is due to the built-in assumption of total reinvestment of non-I(t) capital in further capital formation, and to the fact that r represents real productivity increase. As r increases, the rate at which capital builds up also increases.

Table 3: How T* decreases with r

r	T*	Max NB(T*)	I*(0)
1.1%	30	276	2.6
1.5%	26	356	2.7
2%	22	447	2.8
4%	15	730	3
6%	12	965	3.1
8%	10	1145	3.1

It is noteworthy that a short T* (e.g. smaller than 30 years) does not necessarily exclude the possibility of catering for long term consequences; several of the above six cases are favourable, and the monotonically risk increasing case can lead to an upward revision of S, favourable to both current and future generations. However, the optimal change in S, dS*, can be such that short term risks are reduced and fall with time while long term risks continue to increase, albeit at a slower rate than before. Changing S can change the scene from a monotonically risk increasing scenario to a U shaped scenario, "trapping" the situation against the interests of the far future.

Although the model can deliver much more than this, I shall stop here, as these general results are sufficient to answer the title question: what role for discounting in long term decision making?

5. Conclusion: what role for discounting?

We are now in a position to provide an answer to the title question, provided we specify the type of discounting. If the question is, what role for long term discounting, the answer is: none. If the question is, what role for short term conventional discounting, the answer is: a role limited to the relatively short term, as defined by the (often endogenous) economic time horizon. Not only is a long term discount rate questionable in principle, but there is no need for one. The objection will arise that we cannot then value the costs and benefits in the far future. The answer to this is again, not only that we should not do so as a matter of principle, but there is no need for such a valuation.

It must be remembered that discounting is a means to an end, the end being the valuation and comparison of different values over time, and that this valuation is itself a means to a higher end, that of trying to decide with long term consequences. In the case where these consequences are negative, there needs to be the possibility of a counter-action to mitigate their impact. Provided some basic conditions are met, such as the irreversibility of mitigation efforts and the possibility of using a dynamic control variable, conditions met by investments in R&D, only the valuation of consequences up to the economic time horizon is necessary. We have no need for valuing far future events. The optimal mitigation effort will reflect the extent to which we, current living generations, value long term consequences if, and only if, we do so for ethical reasons. It is remarkable that it is possible, even without ethical concerns for the far future[16], to cater for the long term while deciding only within one's current time horizon.

Possible, but not always. The model built to examine a simple case of radioactive waste containment has shown that, depending on different parameter configurations, current and future generations may have convergent or divergent interests. In the first case mitigation efforts by the current generations will reduce both the short and the long term costs; in the second case, they will reduce the short term but increase the long term costs. There are also cases where the situation is initially unfavourable to both (optimal short term decision leading to increasing risks in both the short and ling term), but leads to a trap where future generations are immolated, using Huxley's wording, on the altar of short term benefits. It is important to understand that this result allows for the role of increased knowledge and technological improvements, so that, in this special case, the outcome does indeed look like a trap. The best response in this case is to know what must be done beforehand in order to avoid falling into it.

This study has investigated a fundamental question in a very simple way, artificially neglecting an important aspect of the problem: uncertainty. The average long term productivity of capital, the efficiency of R&D mitigation efforts, and the extent of damages, should an event happen, are all uncertain and should be considered as random variables. The application of the basic idea, because it considers the interest rate as given, is limited to relatively small scale projects and excludes issues like climate change, which are macro-economic in scope. As for the example treated, it is but a first step towards a more general treatment which would include a productive component. The appropriate example would be the construction of a nuclear plant producing highly radioactive waste during its lifetime. A further step would be to include the possibility of irreversible losses, like the extinction of species or ecological communities, that science could not bring back to life. Even with this last complication, provided that ex-ante risks of extinction are considered, the method proposed here should be of some help.

[16] We need to remember that such ethical concerns come at the cost of a reduction in net benefits before T. Also, though not mentioned in this paper, the model used in this study shows it is sensible to consider *two* time horizons, T_S and T_L, where T_L represents the time beyond which the decision maker no longer considers the effects of the current decision. In our case, T_L could be for example 500 years. Once this number is big enough, it does not matter much for model results what the precise value is.

REFERENCES

Ainslie, G. (1975) "Specious reward: a behavioural theory of impulsiveness and impulse control", *Psychological Bulletin* 82(4): 463-96.

Amiel, Y. and Cowell, F. (2000) "Attitudes towards risk and inequality: a questionnaire-experimental approach" Discussion Paper No. DARP 56 (June), Distributional Analysis Research Programme, London School of Economics, ISSN 1352-2469

Amiel, Y. and Cowell, F. (2001) "Rrisk and inequality and perceptions" Discussion Paper No. DARP 55 (June), Distributional Analysis Research Programme, London School of Economics, ISSN 1352-2469

Arrow, K.J. (1999) "Discounting, morality and gaming" in in P.Portney and J.Weyant (eds): *Discounting and intergenerational equity*: 13-21.

Åsheim,G. (1996) "Ethical preferences in the presence of resource constraints" *Nordic J. of Political Economics* 23: 55-67.

Becker, G.S. and Mulligan, C.B. (1997) "The endogenous determination of time preference" *Quarterly Journal of Economics* 112(3): 729-58.

Chichilnisky, G. (1996) "An axiomatic approach to sustainable development" *Social Choice and Welfare* 13: 231-57.

Collard, D. (1988) "Catastrophe risk: or the economics of being scared" in D. Collard Pearce D.W. and Ulph D. (eds.) *Economics, Growth and Sustainable Environments*: 67-83. MacMillan Press, London.

Cropper, M and Laibson, D. (1999) "The implications of hyperbolic discounting for project evaluation" in P.Portney and J.Weyant (eds): *Discounting and intergenerational equity*: 163-72.

Frederick, S. (1999) "Discounting, time preference, and identity" PhD Thesis, Dept. of Social & Decision Sciences, Carnegie Mellon University.

Frederick, S., Loewenstein, G. and O'Donoghue, T. (2002) "Time discounting and time preference: a critical review" *J. of Economic Literature* 40 (June): 351-401.

Hartwick, J.M. (1977) "Intergenerational equity and the investing of rents from exhaustible resources" *American Economic Review* 67(5): 972-4.

Kroll, Y. and Davidovitz, L. (1999) "Choices in egalitarian distribution: inequity aversion versus risk aversion" Discussion Paper No. DARP 43 (March), Distributional Analysis Research Programme, London School of Economics

Fisher, A.C. and Krutilla, J.V. (1975) "Resource conservation, environmentral preservation, and the rate of discount" *Quarterly J. of Economics*: 359-70.

Loewenstein and Prelec, (1992) "Anomalies in intertemporal choice: evidence and an interpretation", *Quarterly J. of Economics* 107(2): 573-97.

Loewenstein and Prelec, (1993) "Preferences for sequences of outcomes" *Psychological Review* 100(1) 91-108.

Newell R. and Pizer, W. (2001) Discounting the benefits of climate change mitigation. How much do uncertain rates increase valuations? Pew Center on Global Climate Change, Economics Technical Series, (Dec. 2001) 37 p.

Page, T. (1977) *Conservation and Economic Efficiency*, published for Resources for the Future, The Johns Hopkins University Press: Baltimore and London.

Pearce, D.W. and Turner, R.K.(1990) *Economics of Natural Resources and the Environment*, Harvester Wheatsheaf.

Pezzey, J. (1997) “Sustainability constraints versus “optimality” versus intertemporal concern, and axioms versus data” *Land Economics* 73(4): 448-66.

Pezzey, J. (2002) “Exact measures of income in a hyperbolic economy” Centre for Resource and Environmental Studies, Australian National University, Canberra ACT 0200.

Philibert, C. (1998) “The economics of climate change and the theory of discounting” *Energy Policy* 27: 913-27.

Portney P.R. and Weyant, J.P., eds. (1999) *Discounting and intergenerational equity,* Resources For the Future: Washington D.C.

Quiggin, J. (1993) “Discounting and sustainability” *Review of Marketing and Agricultural Economics* 6(1): 161-4.

Quiggin, J. (1997) “Discount rates and sustainability” *International J. of Social Economics* 24(1): 65-90.

Quiggin, J. and Horowitz, J. (1995) “Time and Risk” *J. of Risk and Uncertainty* 10(1): 37-55.

Rabl, A. (1996) “Discounting of long term costs: what would future generations prefer us to do?” *Ecological Economics* 17: 137-45.

Rabl, A. (1999) “Who pays what as a function of discount rate in intergenerational cost-benefit analysis: the example of nuclear power” Mimeo, Ecole des Mines, 60 Boul. St. Michel, 75272 Paris, France.

Ramsey, (1928) “A mathematical theory of saving” *The Economic Journal* 38(152): 543-559.

Read, D. (2001) “Is time discounting hyperbolic or subadditive?” *J. of Risk and Uncertainty* 23: 5-32.

Rubinstein, A. (2000) “Is it ‘economics and psychology’? The case of hyperbolic discounting” Mimeo, Tel Aviv University and Princeton University.

Samuelson, P. (1937) “A note on the measurement of utility” *Review of Economic Studies* 4 : 151-61. .

Schelling, T.C. (1995) “Intergenerational discounting” *Energy Policy* 23(4/5): 395-401.

Schelling, T.C. (1999) “Intergenerational discounting” in P.Portney and J.Weyant (eds): *Discounting and intergenerational equity*: 99-102.

Solow, R. (1999) Foreward in Portney & Weyant (eds): *Discounting and intergenerational equity* : vii-ix.

Strotz, (1955-56) “Myopia and inconsistency in dynamic utility maximisation” *Review of Economic Studies* 23(3) : 165-80.

Weitzman, M. (1998) “Why the far distant future should be discounted at its lowest possible rate” *J. of Environmental Economics and Management* 36(3) : 201-8.

Weitzman, M (2001) “Gamma discounting” *American Economic Review* 91(1): 260-71.

Appendix 1 : Categorising the different long term discounting possibilities

Note that not discounting the long term (beyond T), or using a zero discount rate for it, is a particular case of the multiple rate case. One uses some positive rate before T and an (explicitly or implicitly) zero rate after T. A more rigorous categorisation of the multiple rate case should therefore be,

- in the discrete form:

Ø two or more decreasing discount rates, with the last one being zero;

Ø two or more discount rates, with the last one being the lowest but remaining positive;

- in the continuous form:

Ø a continuously declining discount function, with the discount rate tending towards zero as time goes to infinity;

Ø a continuously declining discount function, with the discount rate tending towards some strictly positive value as time goes to infinity.

From this perspective, the overall categorisation should comprise only two categories, namely:

- Use of a single positive discount rate (we ignore the case of a single zero discount rate for all time, as this runs counter to economic fundamentals);
- Use of a discretely or continuously declining sequence of rates, with the final rate being either positive or zero.

However, because of the underlying theoretical justifications, discussed in the text, the initial categorisation into three cases remains the most practical, and we stick to it.