

Valuing environmental costs and benefits in an uncertain future: risk aversion and discounting

FABIEN MEDVECKY

University of Sydney

Abstract: A central point of debate over environmental policies concerns how future costs and benefits should be assessed. The most commonly used method for assessing the value of future costs and benefits is economic discounting. One often-cited justification for discounting is uncertainty. More specifically, it is risk aversion coupled with the expectation that future prospects are more risky. In this paper I argue that there are at least two reasons for disputing the use of risk aversion as a justification for discounting when dealing with long-term decisions, one technical and one ethical. Firstly, I argue that technically, it implies an inconsistency between theory and practice. And secondly, I argue that discounting for uncertainty relies on a form of individualism which, while reasonable in standard microeconomic theory where an agent chooses how to spread her own consumption over her own lifetime, is inappropriate in the context of inter-generational social decisions.

Keywords: discount rate, risk aversion, uncertainty, inter-generational justice, economics, environmental decisions

JEL Classification: B41, D61, D63, D81, Q51

Long-term environmental policy decisions, such as those concerning climate change and nuclear energy, raise many important scientific, social, and economic issues. A central point of debate over such decisions concerns how future economic costs and benefits should be assessed. The methods chosen to assess future costs and benefits, and the justifications we have for choosing such methods have deep social,

AUTHOR'S NOTE: I am grateful to the Centre of Excellence for Environmental Decisions and in particular Hugh Possingham. This paper was presented at the Harvard Climate: Science + Humanities conference, and the Eindhoven University of Technology workshop Cost-Benefit Analysis: Uncertainty, Discounting and the Sustainable Future where it benefitted greatly from insightful comments. I am also very thankful to Simon Caney, Katie Steele, Kim Sterelny, Mark Colyvan, and two anonymous referees for further helpful comments.

political and ethical implications (Gardiner 2004, 572; Caney 2009, 163). Arguably the most commonly used method for assessing the value of future costs and benefits in both scientific and economic models is discounting, a cost-benefit analysis tool that decreases future costs and benefits by a yearly rate, the discount rate. The discount rate is an aggregate of various parameters, each representing a different motivation and justification for discounting, and these motivations and justifications are each a point of debate amongst experts in the field (Stern 2007, 41; Nordhaus 2007, 689; Quiggin 2008, 200).

One such justification for discounting is uncertainty. The rationale for this view is that all things being equal, we are better able to assess what the state of affairs will be like tomorrow than next year, and we are better able to assess what the state of affairs will be like next year than the following year, and so on for any subsequent year. Given this increasing uncertainty over time, we are justified in placing greater value on present consumption than on future consumption, as present consumption is more certain. However, it is not uncertainty in any and all of its forms that is used to justify discounting; it is uncertainty over the likelihood of hitting the expected value of a project, which is used to justify discounting. It is a preference for a guaranteed outcome over a gamble that has the same expected value as the guaranteed return, or simply risk aversion, which is usually presented as a justification for discounting (Smith 2011, 4; Howarth 2009, 24; Brent 1996, 168; Department of Infrastructure and Planning 2007, 9).

I will argue that there are at least two reasons for disputing the inclusion of risk aversion in our justifications for discounting when dealing with long-term decisions. Firstly, appealing to uncertainty and risk aversion as a justification for discounting is technically problematic because it implies an inconsistency between theory and practice. This has implications for cost-benefit analysis which have not been noted before. Secondly, the premise that we are justified in preferring the more certain present to the less certain future is ethically problematic as this justification relies on a form of individualism. This form of individualism might be reasonable in standard microeconomic theory where an agent chooses how to spread her own consumption over her own lifetime. However, when faced with long-term environmental decisions, one moves from the individual decisions to the social decisions and from intra-generational decision to inter-generational decisions.

Unlike discounting in an intra-generational setting, discounting in an inter-generational setting presents us with a discontinuity between the beneficiaries of the decision (the present generation) and the subsidisers of the decision (the future generations). This discontinuity—which has played an important role in the ethical discourse concerning climate change—challenges the microeconomic rationale for using discounting when dealing with uncertainty. Before arguing my case, I will briefly discuss discounting, and how uncertainty in the form of risk aversion is used to justify discounting.

1. DISCOUNTING AND UNCERTAINTY

Most, if not all of our economic decisions involve comparing commodities. We might compare one commodity at a price against another commodity at a price, such as an apple at \$1 against a pear at \$1.05. Or we might compare one commodity against the same commodity, but at different points in time; such as an apple at \$1 now against an apple at 90c next year. In cost-benefit analysis (CBA), discounting is a way of assessing the value of consumption—either as costs or as benefits—at different points in time. It is a way of assessing how much next year's apple is worth to me relative to today's apple. As part of this paper is concerned with the technical application of discounting, let me briefly explain the mechanics of discounting. Consider the following pair-wise comparison:

- 1) \$200,000 now and further \$100,000 in 10 years time or
- 2) \$280,000 now and nothing further.

Assuming both options cost the same, in order to work out which option offers the highest return, we need to assess whether \$100,000 in 10 years time is greater than, less than, or equal in value to \$80,000 now. Working out the Net Present Value (*NPV*) of \$100,000 in 10 years is just what discounting does. We can think of discounting as the reverse of interest; how much money would I have to invest today at a given interest rate r such that in 10 years time I have \$100,000? Put in a general form, the equation for the net present value of a good $NPV(G)$, with future value $fV(G)$ in t years, at discounting rate r is:

$$NPV(G) = fV(G) / (1+r)^t$$

At a discount rate of 2.8% p.a., the rate suggested by the Office of Management and Budget, \$100,000 in 10 years time has a *NPV* of \$75,870 (OMB 2011). Thus option 2) is better, paying \$280,000, compared to option 1) paying \$275,870. But how do we work out the discount rate? A common way to arrive at a discount rate is to consider the various motivations we have for discount independently, and to work out how much discounting each motivation justifies. The amount of discounting justified by the various motivations is then aggregate into a single rate. These motivations are usually given as impatience or pure time preference, economic growth, opportunity cost, and risk and uncertainty.

In this paper, I will only be concerned with how much discounting can be justified by uncertainty. More specifically, I will be concerned with uncertainty spelled out in terms of risk aversion, as this is a common model. For example, Robert Brent, in his textbook on applied CBA, explains, “the risk premium is added to the discount rate to correct for the uncertainty characteristic of the benefit being in the future” (Brent 1996, 168). In regard to long-term environmental decisions, such as the climate change debate, Kathryn Smith states that one of the roles of η , a parameters of the social discount rate, “is to measure relative risk aversion. In a stochastic model, higher η implies more disutility from exposure to risk” (Smith 2011, 4). One point to note before proceeding is that there has been substantial discussion on whether the discount rate should be constant for through time or not. In this paper, I will be assuming constant rates as this is the practice advocated by the OMB amongst others (OMB 2011).

Risk aversion and risk premiums

The economic definition of a decision under uncertainty is a decision in which the probability of each outcome is known, but in which no single outcome is certain (Perloff 2004, 574). According to proponents of discounting for uncertainty, uncertainty leads us to discount because we prefer more certain returns to less certain returns, and the present is more certain than the future. As codified in the expression “a bird in the hand is worth two in the bush”, we place greater value on what we know we have over what we think we might get.¹

¹ In fact this preference for the certain option over the uncertain (even if the certain option is of lesser value) is the economic definition of risk aversion; see Perloff 2004, 582.

Consider the returns of a forestry project and deciding whether to harvest now, or in five years, assuming both options have the same expected return. We might be quite certain of the current market price of the wood and the current harvesting rate, but both of these are liable to change over time. Those changes in price and harvesting rate may be either to our benefit, or to our detriment. The prices might go up or down, and the harvested quantities might increase or decrease. But as far as discounting is concerned, it does not matter that uncertainty cuts both ways; that events might equally turn out to be better or worse than expected. What concerns us is that it might turn out worse. In effect, we want to give more weight to the possibility that the returns will be lower than the expected returns. To express our concern over the negative possibility, we turn to expected utility theory (EU) and include some risk aversion.

Risk aversion is the preference for investments which offer returns with smaller deviation from the mean over investments which offer returns with larger deviation from the mean, where the mean is the expected returns. For example, consider choosing between two coin-toss games with pay-offs as described in table 1 (assuming all games cost the same to play and the coin is fair). Both games are uncertain and both have an expected return of \$50, but game *B* is less risky in that the payoffs do not diverge as far from the expected return as they do in game *A*.

Table 1

	Heads ($Pr=0.5$)	Tails ($Pr=0.5$)	Expected returns
Game <i>A</i>	$\$100 \times 0.5 = \50	$0 \times 0.5 = 0$	$\$50 + \$0 = \$50$
Game <i>B</i>	$\$60 \times 0.5 = \30	$\$40 \times 0.5 = 20$	$\$30 + \$20 = \$50$

The worst case scenario of game *B* is that the player wins \$40, which is \$10 less than the expected return, whereas the worst case scenario of game *A* is that the player wins nothing, which is \$50 less than the expected return. A risk-averse individual would prefer to play game *B* to game *A*. Such an individual would value game *B* more highly than game *A*, and would be willing to pay a premium to play game *B*. This premium, the risk premium, varies according to the situation and the decision maker's risk aversion utility of wealth function. An individual's utility of wealth function tells us how much good (broadly construed) an

individual derives from a given return. Using a text-book risk-averse utility of wealth function such as $U(W)=\sqrt{W}$ for the above coin bet, we get the following risk-averse table:

Table 2

	Head ($Pr=0.5$)	Tail ($Pr=0.5$)	EU
Game A	$U=\sqrt{\$100} \times 0.5 = 5$	$U=\sqrt{\$0} \times 0.5 = 0$	5
Game B	$U=\sqrt{\$60} \times 0.5 = 3.872$	$U=\sqrt{\$40} \times 0.5 = 3.162$	7.034

While both games have the same expected return, for risk-averse individuals, the riskiness of game A reduces its utility compared to game B. In order to perform a CBA on this game, we turn the expected utility back in monetary value by calculating the *risk-free* (or *certainty equivalent*) of the expected utility. This is done by converting the utility back into dollar values using the reverse of the utility of wealth function, in our case *risk-free equivalent* (W)= U^2 . In the coin-toss games above, the risk-free equivalent of game B's payoff is \$49.47, and game A's risk-free equivalent payoff is \$25. In other words, an individual exhibiting the risk-averse utility of wealth function $U(W)=\sqrt{W}$ would derive the same amount of utility out game B as she would out of a guaranteed \$49.47, and should be indifferent between game B and a guaranteed \$49.47. Given that game B has an expected monetary value of \$50 and the agent is indifferent between game B and \$49.47, the risk premium she is willing to pay in this instance is \$0.53 (\$50 - \$49.47).

While different degrees and structures of risk aversion can be represented by altering the utility of wealth function, all of them result in a certainty equivalent amount smaller than the expected monetary value. Moreover, for the same expected value the greater the uncertainty of the returns in terms of variance, the lower the expected utility of the investment. This decrease in utility leads to a decrease in the risk-free equivalent value of the return, and the more the risk-free equivalent decreases, the greater the risk premium becomes. If uncertainty increases over time, then the risk premium a risk-averse individual is willing to pay also increases over time. This increase in the risk premium over time leads to progressively decreasing the value of future returns over time. It is this progressive decrease in the value of future returns that discounting for uncertainty is supposed to represent.

2. DISCOUNTING FOR UNCERTAINTY: THE TECHNICAL PROBLEM

Recall the forestry example above. I know what I get if I harvest the wood now, but much can happen in five years (for better or for worse). The longer I delay my harvesting, the greater the possibility of deviation from my expected return. Being risk-averse, I would be willing to pay a premium to guarantee my returns, and discounting the expected return of the harvest in five years just reflects my willingness to pay a risk premium. However, in this forestry example, I was solely concerned with the benefits or returns on my investment. But when we discount, we discount both costs and benefits.

If we discount future costs and benefits because we are risk-averse, then the future costs demand a different treatment to the future benefits. This is contrary to the standard practice in CBA, where the discount rate is insensitive as to whether it is applied to costs or to benefits (Department of Finance and Administration 2006, 41). For example, assume that we are assessing a forestry project with expected returns of \$20M a year and expected costs of \$15M a year for the next ten years. The standard practice in CBA is to discount both the expected returns and the expected costs at the same rate before balancing them, or (equivalently) to first subtract the costs from the returns, and then to discount the result. The problem is that discounting costs with the same positive rate used to discount benefits in fact involves a risk-loving preference, not a risk-averse preference. To accurately reflect risk aversion, the part of the discount rate which reflects risk aversion in regards to future costs should be *negative*.

For risk-averse individuals, discounting reflects a willingness to pay a premium for certainty. According to risk aversion, the more uncertain we are, the more we are willing to pay a premium to compensate for the uncertainty. And to pay a premium is to incur a cost. In the case of benefits, placing a cost on the returns reduces the magnitude of the benefits, so a positive discount rate accurately reflects our willingness to pay a premium. However, in the case of costs, placing a cost (for the premium) on existing costs *increases* the magnitude of the existing costs. This is contrary to discounting costs, which leads to a *decrease* in the magnitude of the costs.

Consider a case where a risk-averse individual is unsure about the amount of a cost, say a forthcoming tax bill. Assume the tax bill could be \$1000, \$1500, or \$2000, each with equal probability 1/3.

The expected monetary value of the forthcoming tax bill is clearly \$1500, but a risk-averse individual who is willing to pay a premium for certainty would be indifferent between a tax bill T which is greater than \$1500 but certain, and the present scenario (which has lower expected monetary value than T but greater variance). A risk-averse individual would be willing to pay $T - 1500$ premium for certainty, a premium which *increases* the magnitude of the costs.

As is explained in the Office of Management and Budgets (OMB) Circular No A-94, “a risk-averse individual may have a certainty-equivalent for an uncertain set of costs that is larger in magnitude than the mathematical expectation of costs” (OMB 1992). Yet, the standard use of a positive discount rate leads to a decrease in the value of the costs. If we were to discount the expected return of -\$1500 for one year, at say 7%, the *NPV* would be -\$1400. If we want to reflect an agent’s willingness to pay a premium for risk aversion in regards to costs, we should include a negative rate in the agent’s discount rate.

The fact that uncertainty leads to a different treatment of costs and benefits opens a veritable can of worms. There are four issues I will consider in the following subsections, namely:

- 1) Is risk aversion the right description of why we discount for uncertainty?
- 2) If costs and benefits require different treatment, should we start by discounting them separately, and then subtracting the costs from the benefits, or should we first subtract the costs from the benefits, and then discount the result accordingly?
- 3) Do all motivations for discounting lead to costs and benefits requiring different discounting rates, or is this unique to discounting for uncertainty?
- 4) Does uncertainty always lead to different discount rates for costs and benefits?

My answers will be that 1) risk aversion may not be descriptively accurate of empirical evidence regarding individual choices under uncertainty, but trying to be descriptively accurate does not help, 2) we should start by discounting the costs and benefits, and then balancing them out, 3) the need for different discount rates for costs and benefits is unique to uncertainty, and 4) uncertainty does not always lead to different rates for costs and benefits, but even in the cases where we

could use the same rate for both costs and benefits, we ought to reject uncertainty as a motivation for discounting on pragmatic grounds.

Is risk aversion the right description of why uncertainty leads to discounting?

Having a uniform risk-averse value function for both costs and benefits requires that we use a different discount rate for each, which is contrary to standard practice. In order to rectify this problem, two options are open to us. On the one hand we could accept the implication of the above discussion and endorse a ‘two rates’ approach to discounting, with a lower rate for costs to reflect the negative risk-averse component. Alternatively, we could hold on to the one discount rate for both costs and benefits and argue that risk aversion fails to accurately represent how agents reason under uncertainty. Indeed, empirical evidence suggests that agents making decisions under uncertainty do not behave according to an expected utility model as is normally assumed in economics (Loewenstein and Prelec 1992, 573; Rabin and Thaler 2001, 220). If we want discounting to accurately reflect how people behave when making decisions under uncertainty, we should consider the empirical evidence about how people make such decisions.²

The best known evidence for how people behave when making decisions under uncertainty comes from Kahneman and Tversky (Kahneman and Tversky 1979, 263). Participants in Kahneman and Tversky’s experiments were offered hypothetical choices between pairs of options, with different values and probabilities in each case. The results of the experiments showed that individuals making decisions under uncertainty were not only sensitive to payoffs and probabilities, which is what would result from adhering to expected utility theory, but individuals are also sensitive to other factors such as the direction of the payoffs. For example, one of the choices faced by participants was a choice between either A: \$6,000, with $P=0.45$, or B: \$3,000, with $P=0.90$. In that case, the overwhelming majority of participants (86%) exhibited a risk-averse attitude and preferred option B (Kahneman and Tversky 1979, 267). When offered a negative version of the above pairwise choice, namely either A: -\$6,000, with $P=0.45$, or B: -\$3,000, with $P=0.90$, the overwhelming majority of participants, 92%,

² Whether we do in fact want this is an interesting question, but sadly it lies outside the scope of this paper.

preferred the riskier option A (Kahneman and Tversky 1979, 267; Loomes and Sugden 1982, 805).

Experimental results show that participants' utility of wealth function was asymmetric across gains and losses; individuals value uncertainty over costs differently to uncertainty over benefits, which could be good news for those wanting to hold on to discounting unaltered. If individuals making decisions under uncertainty do not have a uniform risk aversion for gains and losses, then we may not need to resort to positive discount rates for benefits and negative discount rates for costs. We could hold on to a uniform rate for costs and benefits so long as individuals making decisions under uncertainty displayed 1) a risk-averse attitude in regard to gains, and 2) a risk-loving attitude in regard to losses, and 3) an equal degree of risk aversion in regard to gains as the degree of risk love displayed in regard to losses.

As it turns out, individuals making economic decisions under uncertainty did usually display a risk-averse attitude in regard to gains (all things being equal), so a positive discount rate for benefits seems appropriate. Moreover, the empirical evidence also shows that in regard to economic losses, individuals do display the risk-loving attitude which is required to justify a positive discount rate for costs. But while individuals might be risk-averse when considering gains and risk-loving when considering losses, the empirical evidence shows "that agents are more sensitive to losses than to gains, resulting in a utility function that is steeper for losses than for gains" (Köbberling and Wakker 2005, 120). This has come to be called loss-aversion (Rabin and Thaler 2001, 226). As Kahneman and Tversky explain:

A salient characteristic of attitudes to changes in welfare is that losses loom larger than gains. The aggravation that one experiences in losing a sum of money appears to be greater than the pleasure associated with gaining the same amount (Kahneman and Tversky 1979, 279).

If the exhibited rate of risk aversion in regard to gains does not match the exhibited rate of loss aversion in regard to losses, then we again require a different discount rate for costs and benefits. To accurately reflect the empirical evidence about individuals' risk preferences, we should use a positive discount rate for both gains and losses, but the discount rate for losses should be of a greater magnitude. Thus, while it is true that the standard risk-averse expected

utility model is an inaccurate (or at least incomplete) description of how agents behave under uncertainty, appealing to descriptive accuracy would not help justify the present standard practice of using the same positive discount rate for both costs and benefits. Moreover, even if using the same positive discount rate for both costs and benefits was descriptively accurate, it does not follow that it is ethically defensible.³

Should we start by discounting costs and benefits separately, and then balancing them, or should we first balance the costs and benefits, and then discount the result?

One advantage of using the same discount rate for both costs and benefits is that it simplifies the process of CBA by making it insensitive to the order in which we proceed. It makes no difference whether we balance the costs and benefit first and then discount, or whether we discount first, and then balance the costs and benefits. However the introduction of different discount rates for costs and benefits complicates this. Now, the order in which we proceed does make a difference. Recall the above forestry example with expected returns of \$20M, and expected costs of \$15M. If we first balance the costs and benefits, and then discount the result at 5% for one year, we have a *NPV* of \$4.75M. If, on the other hand, we first discount the benefits and the costs at 5% and -5% respectively, and then balance them, we get a *NPV* of \$3.25M. Clearly, it cannot be the case that both these methods are right. Either we first balance the costs and benefits, and then discount, or we discount the costs and benefits first, and then balance them.

I argue that if we want to discount for uncertainty we must first discount all costs and benefits separately, and then balance them out. Ordering the process the other way around by first balancing and then discounting is untenable; here is why. Let us assume for a moment that we do in fact balance the costs and benefits first and discount second. The balancing of the costs and benefits merges the two figures into single number: the balance. Presumably, if the balance is positive it means we anticipate a profit or benefit, so we discount the result with a positive discount rate. And if the balance is negative, we discount it with a negative discount rate because we anticipate a loss or cost from our project. Yet this can only be reasonable if the costs and benefits occur at approximately the same time. If costs and benefits occur at different times, the 'balance first' option becomes problematic.

³ I will consider this ethical problem in section 3, below.

Again consider the forestry example above, but with the following variations: a) all costs occur in 6 months from now, and all benefits occur in 1 year, b) all benefits occur in 6 months from now, and all costs occur in 1 year. If we first balance out the costs and benefits and then discount them, these work out the same. If, on the other hand, we begin by discounting the costs and benefits, and then balance them (at -5% and 5% respectively), option a) has a *NPV* of \$3.65M, while option b) has a *NPV* of \$3.75M. This may not be a great difference, but nor is it negligible, and it is a difference which is lost if we balance out the costs and benefits before we discount them.

While the question of how often we should discount (every year, every six month, every quarter) is a vexing issue for discounting *even* if we use the same discount rate for both costs and benefits (OMB 1992), the inclusion of different discount rates for costs and benefits make this issue even more vexing. Discounting the costs and benefits first (with their respective discount rate) before balancing them out has the great advantage of not being susceptible to this problem. Since costs and benefits are discounted independently, we can accurately reflect the time at which the costs or benefits occur.

Do all reasons for discounting lead to different discounting rates for costs and benefits, or is this unique to uncertainty?

While uncertainty requires that we treat costs and benefits differently, other motivations for discounting do not raise the same demands. The reason why uncertainty is unique in requiring a positive discount rate for benefits and a negative rate for costs is that uncertainty is the only reason for discounting which leads to a preference for *both* costs and benefits to be as close to the present as possible. Whether we appeal to opportunity costs, economic growth or pure time preference (impatience) as justifications for discounting, all of these lead us to prefer benefits as early as possible and costs as delayed possible. Put simply, in all cases other than uncertainty, we have a preference for early returns over delayed returns. Benefits are returns, so we prefer them as early as possible; costs on the other hand, are the opposite of returns, and hence are preferred as late as possible.

As an example consider opportunity costs as a motivation for discounting. Opportunity costs lead us to discount costs and benefits because of the financial returns we could have had if we had invested the resources under consideration (Torgerson and Raftery 1999, 914).

Recall the forestry example above with \$20M returns and \$15M costs. Assuming I could get 5% p.a. return on an investment, if I had the \$20M benefits now, by next year that would be worth \$21M. Conversely, all I need to have today to generate \$20M next year at 5% is \$19M—namely the discounted or *NPV* of the \$20M, so the sooner the returns, the better. As for the costs, the process is reversed; the later the costs the better. If I had to pay the costs today, I would have to pay \$15M. However, if I do not have to pay the \$15M till next year, then I can invest the \$15M such that by next year, I have \$15.75M. And if I do not have to pay the \$15M till next year, but I can get 5% return on my investment, all I need to invest today to generate \$15M for next year is \$14.3M—again the discounted or *NPV* of the \$15M.

With opportunity costs, the nearer (temporally) the benefits are, the better off I am, hence the positive discount rate. And with the costs, the further in the future the costs are, the better off I am, which again requires a positive discount rate. As stated above, all motivations for discounting other than uncertainty behave like opportunity costs. They all lead us to prefer benefits to be in the near future and costs to be as far into the future as possible. This set of opposing preferences justifies a positive discount rate for both costs and benefits. Uncertainty over returns, on the other hand, leads to parallel preference across costs and benefits. The more certain a cost or benefit is, the more desirable it is; the earlier a cost or benefit is, the more certain it is; therefore both costs and benefits are preferred as early as possible.

Does discounting for uncertainty always lead to different rates for costs and benefits?

So far I have proceeded under the (unstated) assumption that the reasons for the uncertainty over the costs and benefits were independent of each other. Now let us consider what happens if costs and benefits are correlated. Recall one more time the forestry project with expected returns of \$20M and expected costs of \$15M. Let us assume that there are only two uncertain variables: 1) how much timber will be sold—these are the benefits—and 2) how many new seedlings will be required—which are the costs. Let us further assume that trees are only felled to fill the timber orders and seedlings are only bought to replace felled trees. In this example, the costs and benefits are correlated; if we sell more timber than expected (and hence have a greater turnover than expected), we will need to fell more trees than

expected. If we fell more trees than expected, we will need to buy more seedlings to replace those trees (and hence have a greater cost than expected). Conversely, if we sell less timber than expected (and hence fell fewer trees than expected) we will need to buy fewer seedlings (and hence have a lower cost than expected).

In such a case, using different discount rates for costs and benefits seems unreasonable because any increase or decrease in benefits relative to the expected benefits will be mirrored by a proportionally equivalent change in the costs. Therefore, it cannot be that we will have lower than expected benefits and higher than expected costs. In cases where the uncertainty over the costs and benefits exactly correlates (when the uncertainty in regard to cost equals the uncertainty in regard to benefits), it seems reasonable to use the same discount rate for both. When the uncertainty in regard to cost does not exactly correlate with the uncertainty in regard to benefit, we are no longer justified in using the same rate for both (for the reasons given at the beginning of this section). However, I will argue that even in such cases where costs and benefits exactly correlate, we have good pragmatic reasons to reject discounting for uncertainty.

Correlated costs and benefits

In cases where costs and benefits are exactly correlated we are permitted to use the same discount rate for both costs and benefits. But in such cases we are faced with a new decision in regards to the discounting for uncertainty, namely: what form of risk aversion should our discount rate reflect? The discount rate could reflect an aversion towards uncertainty in regards to benefits or an aversion towards uncertainty in regards to costs. Put differently, we could use a positive or a negative discount rate to account for uncertainty. While the standard practice in CBA is to use a positive discount rate, this does not need to be the case.

I will refer to a preference for certainty in regards to benefits over a preference for certainty in regards to costs simply as *risk aversion*. *Loss aversion* will be used to denote a preference for certainty in regards to costs over a preference for certainty in regards to benefits. Recall that all we require when costs and benefits are correlated is that we use the same rate across both costs and benefits. We can use a positive discount rate, a negative discount rate, or a zero discount rate depending on

whether we want to reflect risk aversion, loss aversion, or no preference either way.

Consider a 'break-even' environmental programme with correlated expected costs and benefits of \$4 billions each. Using a positive rate to discount for uncertainty would reflect risk aversion; the thought would be something like "we might not get the benefits we expect, so we better under value our returns and let the costs follow suit". We could use a negative rate to discount for uncertainty, which would reflect loss aversion; the reasoning would be "we may have higher costs than we expect, so we better over value our costs and let the benefits follow suit". Alternatively, we may choose to use a zero discount rate, and not discount at all: "costs and benefits might not be what we expect them to be, but they will move together, so we may as well be neutral towards risk". All we required to be economically consistent is that the discount rate be the same across costs and benefits, whether it be positive, negative, or neutral. I will argue that while we have no economic reasons to sway one way or the other, we have pragmatic reasons to use a zero discount rate for uncertainty.

Consider what the consequences might be of applying each of these attitudes to risk in CBA to see whether this can help guide our decision. In many and possibly most cases—when costs and benefits are not correlated—discounting for uncertainty is a technical minefield. Until we either revise our theory of discounting (allow for two rates—one for costs and one for benefits—and work through the effects of this revision) or revise our theory of risk aversion (and revise it in such a way that would warrant a uniform discount rate), we ought not to include a rate for uncertainty in our discount rate. I will argue that not including a rate for uncertainty in the discount rate is the best option in all cases, as it is the only rate that reconciles cases when costs and benefits are not correlated with cases when they are correlated.

Most projects involve many economic measures, all of which are to be discounted. In order to remain consistent throughout a project, one needs to use the same discount rate for all measures. But this can only be justified if all of the measures in the project are correlated. Furthermore, projects often interact with other projects. Often, one project relies on the data and outcomes of previous project, making uniformity a desirable attribute of the discount rate. Consider a project, *P*, in which all the internal costs and benefits are correlated. The inclusion of some non risk-neutral discount rate in the assessment

of P rules out including the outcomes of P into any project which has costs and benefits that are not correlated with P . This gives us a reason to reject both risk aversion and loss aversion, at least on some occasions. Given that we have no reason to endorse either risk aversion or loss aversion over risk neutrality, but that we sometimes have reasons to reject both risk aversion and loss aversion, risk neutrality stands as the consistently least troublesome contender, including the cases where costs and benefits are correlated.

In this section, I have shown that in order to correctly reflect risk aversion and the willingness of risk-averse agents to pay a risk premium in the discount rate, the part of the discount rate which represents risk aversion should be positive when discounting benefits, but negative when discounting costs. The current practice, however, is to use the same, usually positive rate to account for risk aversion in the discount rate for both costs and benefits. I have argued that we should therefore be risk-neutral and not include a rate for risk aversion in the discount rate. Proponents of discounting for uncertainty might counter that this is an unnecessarily rash move and that we should just use different rates for costs and benefits to correctly reflect risk premiums. But, aside from the technical difficulties using different rates would create, I believe there is also an ethical problem with discounting for uncertainty in inter-generational social decisions.

3. DISCOUNTING FOR UNCERTAINTY: THE ETHICAL PROBLEM

One important difference between individual and social choices is that with individual decisions, decision theorists are only concerned with the demands of rationality. In social decisions, however, decision theorists also have to consider “the relationship between the demands of rationality and those of justice or fairness” (Resnik 1987, 177).⁴ While I will not consider the various definitions of “justice” and “fairness”, I will argue that discounting for uncertainty runs contrary to our intuitions of what is just or fair under almost any definition of these terms (Marglin’s view aside).⁵

⁴ This is not to say that there are no ethical considerations to individual decisions; rather, the suggestion is that individual decision theory does not take into account notions of justice or other ethical considerations. Social decision making does.

⁵ Marglin has famously (or infamously) argued that “I want the government’s social welfare function to reflect only the preferences of present individuals. Whatever else democratic theory may or may not imply, I consider it axiomatic that a democratic government reflects only the preferences of the individuals who are presently

Standard CBA discounts for uncertainty under the assumption that individuals prefer more certain returns to less certain returns and that this preference justifies decreasing the value of less certain returns relative to more certain returns. While I will accept these assumptions, I believe this is a poor justification for discounting in long-term social decisions. A central factor of the ethical dimension of inter-generational decisions is the discontinuity between decision makers and those that will be affected by the decision. It is this discontinuity which leads me to reject discounting for uncertainty in long-term decisions.

Discounting for uncertainty is simply a form of risk aversion and can be compared to single-time period decisions with different probability distribution, as I will do below. Indeed, comparing risk aversion in inter-temporal decisions as is the case in discounting to single-time period decisions is a fruitful approach as it shows explicitly what we are doing (Parfit 1983, 33).

One problem with uncertainty as a justification for discounting in long term social decisions is that this justification is founded on individual choice theory, and this translates poorly into a justification for inter-generational social decisions. With regard to individual choices, the argument is that individuals have certain attitudes to risk, and this attitude justifies their preferences. The standard microeconomic explanation of discounting is a representation of an individual's degree of risk aversion in regard to *her* returns; whether I prefer *A* with probability distribution *D* or *A* with probability distribution *D'*, where *D'* has the same expected value as *D*, but greater deviation from the mean than *D*. This justification relies on a form of individualism based on an agent's right to chooses how to spread her own consumption over her own lifetime.

However, an inter-generational social decision involves a preference over uncertainty *and* across persons; whether I prefer *A* with probability distribution *D* for me or *A* with probability distribution *D'* for some future individual, again where *D'* has the same expected value as *D*, but greater deviation from the mean than *D*. The discontinuity between decision makers and those that will be affected by the decision is reflected in the fact that discounting in inter-generational choices also involves preferences across persons. It is this dimension of discounting

members of the body politic" (Marglin 1963, 97). Almost all other economists and philosophers believe that social distribution decisions should extend further than simply reflecting the views of the electorate. For some other views, see Scanlon 1982; Sibley 1953; Rawls 1993; Dobson 1996; Ekeli 2005.

for uncertainty in inter-generational decisions which I believe is of concern.

Since the temporal dimension, *when A* occurs, is only relevant in so far as it determines the probability distribution, we can replace the temporal dimension with a change in probability distribution (Parfit 1983, 32). A preference for \$1000 now over \$1000 at some future time *t* can be represented as a preference for \$1000 guaranteed over some gamble with an expected monetary value of \$1000, for example, a 50/50 chance at \$1500 or \$500. Now if we translate this into a social decision and make our preference over uncertainty *and* across persons, we get a preference for \$1000 guaranteed for some individual *A* over a 50/50 chance at \$1500 or \$500 for some other individual *B*. This already seems problematic; why should the fact that there is uncertainty over the value of the outcome affect *who* gets the outcome.

It is worse still if we translate our decision into an inter-generational decision, where we make our preference over uncertainty and across persons, and in which we (the decision makers) are often some of the persons involved. We get a preference for \$1000 guaranteed for some individuals, often ourselves, over a 50/50 chance at \$1500 or \$500 for someone else, namely future generations. In the cases where we are some of the persons involved, we will prefer the guaranteed money for us to the possible money for someone else. But because we are considering choices which distribute goods across persons including ourselves, we cannot differentiate between the preference for certainty (which leads to discounting) and the preference for self-interest (which does not necessarily lead to discounting). Indeed, we will prefer money for us with any probability distribution (all things being equal) over money for someone else. This can lead us to run against our preferences for the more certain over the less certain. We would likely prefer a 50/50 chance at \$1500 or \$500 for us over \$1000 guaranteed for someone else, even though there is greater variance in our preferred option. This is comparable to preferring future goods for us over present goods for someone else, which runs contrary to the theory of discounting.

Our incapacity to differentiate between our preference for certainty and our preference for self-interest poses an ethical challenge. A preference for self-interest has nothing to do with uncertainty. It can, however, be a reason to prefer one outcome over another. We might prefer the present because *we*, the current generation, want the

consumption. But such an approach to social decision making seems irreconcilable with most conceptions of fairness and justice. Indeed, most philosophers and economists explicitly object to the use of such self-interest in social decision making (Ramsey 1928, 543; Smart 1973, 63; Smith 1976 [1776-1789], III-2; Stern 2007, 31).

For the majority not willing to endorse such self-interest, discounting for risk aversion must be distinguishable from discounting for self-interest. But it is not possible to distinguish between the two. This creates a dilemma for those wanting to discount for uncertainty but not for self-interest: either they must accept that self-interest might be “smuggled in” with uncertainty; or they must reject discounting for uncertainty. As I am not willing to accept self-interest as a motivation for discounting in social decisions, I side with those who reject discounting for uncertainty.

Discounting as compensation

Proponents of discounting for uncertainty might reply to my ethical objection by arguing that future individuals might accept the imposition of risk if they were to be compensated for this imposition. According to this line of reasoning, to impose a risk on another individual is morally permissible if that other individual is willing to accept the risk for an agreed amount of compensation. Now the issue turns not on whether we can discount for risk, but on what the discount rate should be, such that it incorporates the agreed compensation.

While this argument might have some weight in intra-generational decisions, it has little traction when applied to inter-generational issues. The challenge this argument faces, when applied to inter-generational decisions, stems from our incapacity to assess what counts as compensation. As the individuals to be compensated (future generations) are not part of the decision making framework, they cannot state what they would accept as a just compensation, thus, there is no agreed upon compensation. At best, we can allow for what we believe would be a just compensation for the imposition of risk. But, as was stated previously, in many inter-generational distribution decisions, we—the current generation—are not only decision makers, but are also potential beneficiaries of the decision. This self-interest skews our capacity for detached and fair judgement, much as it did in regard to uncertainty. The ‘compensation argument’ does not solve the problem, it simply shifts the problem. Instead of imposing our view of acceptable

risks, we impose our view of acceptable compensation for acceptable risks. This approach, in my view, creates more problems than it solves.

Consider uncertainty over the cost of decommissioning a nuclear power station. Currently the average cost of decommissioning a nuclear plant is around US\$400M (Nuclear Energy Institute 2011). If we were to discount this value for risk aversion, we would reduce its magnitude and have a possible shortfall. According to the argument above, we could compensate for the risk by altering the discount rate. But I simply would not know how much to compensate people in 80 or 100 years from now for the possible shortfall they might face when dealing with the decommissioning of nuclear power stations. In fact, it has been argued that the uncertainty over issues such as nuclear power should lead us to be cautious about our assessment of future costs and benefits and, if anything, we should err on the side of over-allocation of funds, not under-allocation of funds (Caney 2009, 176; Gardiner 2006, 3; Steele 2006, 19).⁶

Whichever way we go, we impose views of acceptable compensation for acceptable risks. While some imposition of our view of what sits as acceptable compensation is an inevitable part of inter-generational decisions, the imposition of compensation for our risk aversion is not. We do not need to impose our view of acceptable risks onto future generations, thus we do not need to impose our view of acceptable compensation for acceptable risks. All we need to do is remain risk-neutral when choosing a discount rate.

CONCLUSIONS

In the economic debate surrounding environmental decision making, discounting stands as one of the most controversial issues. In this article, I have focussed on one of the justifications for discounting, uncertainty. I argued that uncertainty spelled out in terms of risk aversion cannot justify discounting in the case of inter-generational social decisions. I argued that there are two reasons to reject discounting for uncertainty in such cases.

Firstly, I argued that on technical grounds, discounting costs and benefits by the same rate inaccurately reflects risk aversion. I noted an exception to this case, namely when costs and benefits are correlated and have argued that while in such cases we have no economic reasons to outright reject discounting for uncertainty, we do have a pragmatic

⁶ This issue is closely related to the precautionary principle.

reason: that only a risk-neutral discount rate (or no discounting) can lead to a uniform discount rate across projects.

Secondly, I argued that on ethical grounds, discounting for uncertainty relies on a form of individualism. This individualism might be reasonable in standard microeconomic theory, where an agent chooses how to spread her own consumption over her own lifetime. However, when applied to long-term environmental decisions, the discontinuity between the decision makers (the present generation) and those who will bear the consequences of the decision (the future generations) undermines the rationale for discounting for uncertainty.

Since discounting for uncertainty is both technically and ethically problematic in its current form, I believe we ought to reject uncertainty as a justification for discounting in cases of long-term environmental decisions and not include a rate for uncertainty in our discount rate.

REFERENCES

- Brent, Robert. 1996. *Applied cost-benefit analysis*. Cheltenham: Edward Elgar Publishing.
- Caney, Simon. 2009. Climate change and the future: discounting for time, wealth, and risk. *Journal of Social Philosophy*, 40 (2): 163-186.
- Department of Finance and Administration. 2006. *Introduction to cost-benefit analysis and alternative evaluation methodologies*. Canberra (ACT): Commonwealth of Australia.
- Department of Infrastructure and Planning. 2007. *Project assurance framework: cost-benefit analysis*. Brisbane (QLD): Queensland Government.
- Dobson, Andrew. 1996. Representative democracy and the environment. In *Democracy and the environment: problems and prospects*, eds. William M. Lafferty, and James Meadowcroft. Cheltenham: Edward Elgar, 124-139.
- Ekeli, Kristian Skagen. 2005. Giving a voice to posterity: deliberative democracy and representation of future people. *Journal of Agricultural and Environmental Ethics*, 18 (5): 429-450.
- Gardiner, Stephen. 2004. Ethics and global climate change. *Ethics*, 114 (3): 555-600.
- Gardiner, Stephen. 2006. A core precautionary principle. *Journal of Political Philosophy*, 14 (1): 33-60.
- Howarth, Richard B. 2009. Discounting, uncertainty, and revealed time preference. *Land Economics*, 85 (1): 24-40.
- Kahneman, Daniel, and Amos Tversky. 1979. Prospect theory: an analysis of decision under risk. *Econometrica*, 47 (2): 263-291.
- Köbberling, Veronika, and Wakker, Peter P. 2005. An index of loss aversion. *Journal of Economic Theory*, 122 (1): 119-131.
- Loewenstein, George, and Drazen Prelec. 1992. Anomalies in intertemporal choice: evidence and an interpretation. *The Quarterly Journal of Economics*, 107 (2): 573-597.

- Loomes, Graham, and Robert Sugden. 1982. Regret theory: an alternative theory of rational choice under uncertainty. *The Economic Journal*, 92 (368): 805-824.
- Marglin, Stephen. 1963. The Social rate of discount and the optimal rate of investment. *The Quarterly Journal of Economics*, 77 (1): 95-111.
- Nordhaus, William D. 2007. A review of the Stern review on the economics of climate change. *Journal of Economic Literature*, 45 (3): 686-702.
- Nuclear Energy Institute. 2011. Costs: fuel, operation and waste disposal. NEI Website. http://www.nei.org/resourcesandstats/nuclear_statistics/costs/ (accessed April 2012).
- Office of Management and Budget (OMB). 1992. Guidelines and discount rates for benefit-cost analysis of federal programs, Circular No. A-94. U.S.A. Executive Office of the President Website. http://www.whitehouse.gov/omb/circulars_a094 (accessed April 2012).
- Office of Management and Budget (OMB). 2011. Discount rates for cost-effectiveness, lease purchase, and related analyses, Circular A-94, Appendix C (Revised December 2011). U.S.A. Executive Office of the President Website. http://www.whitehouse.gov/omb/circulars_a094/a94_appx-c/ (accessed April 2012).
- Parfit, Derek. 1983. Energy policy and the further future: the social discount rate. In *Energy and the future*, eds. Douglas Maclean, and Peter Brown. Totowa (NJ): Rowman and Littlefield, 31-37.
- Perloff, Jeffrey. 2004. *Microeconomics*. Boston: Pearson Addison Wesley.
- Quiggin, John. 2008. Stern and his critics on discounting and climate change: an editorial essay. *Climatic Change*, 89 (3): 195-205.
- Rabin, Matthew, and Richard Thaler. 2001. Anomalies: risk aversion. *The Journal of Economic Perspectives*, 15 (1): 219-232.
- Ramsey, Frank. 1928. A mathematical theory of saving. *The Economic Journal*, 38 (152): 543-559.
- Rawls, John. 1993. *Political liberalism*. New York: Columbia University Press.
- Resnik, Michael. 1987. *Choices: an introduction to decision theory*. Minneapolis: University of Minnesota Press.
- Scanlon, Thomas. 1982. Contractualism and utilitarianism. In *Utilitarianism and beyond*, eds. Amartya Sen, and Bernard Williams. Cambridge: Cambridge University Press, 103-128.
- Sibley, W. M. 1953. The rational versus the reasonable. *The Philosophical Review*, 62 (4): 554-560.
- Smart, John J. C. 1973. An outline of a system of utilitarian ethics. In *Utilitarianism: for and against*, John J. C. Smart, and Bernard Williams. Cambridge: Cambridge University Press, 3-74.
- Smith, Adam. 1976 [1776-1789]. *An inquiry into the nature and causes of the wealth of nations* (WN), eds. R. H. Campbell, A. S. Skinner, and W. B. Todd. Indianapolis: Liberty Fund.
- Smith, Kathryn. 2011. Discounting, risk and uncertainty in economic appraisals of climate change policy: comparing Nordhaus, Garnaut and Stern. *Commissioned paper for the Garnaut Review update 2011*. Department of Climate Change and Energy Efficiency, Canberra, AU.
- Steele, Katie. 2006. The precautionary principle: a new approach to public decision-making? *Law, Probability and Risk*, 5 (1): 19-31.

Stern, Nicholas. 2007. *The economics of climate change: the Stern review*. Cambridge: Cambridge University Press.

Torgerson, David, and James Raftery. 1999. Economics notes: discounting. *The British Medical Journal*, 319 (7214): 914.

Fabien Medvecky submitted his PhD at University of Sydney on the justifications for the social discount rate in inter-generational decisions. He is currently working as a research fellow at the University of Queensland and teaches in philosophy and economics. His interests are philosophy of economics, environmental philosophy, and ethics.

Contact e-mail: <fabienmedvecky@gmail.com>