

Discounting for public policy: A survey

Hilary Greaves

Abstract

This article surveys the debate over the social discount rate. The focus is on the economics rather than the philosophy literature, but the survey emphasises foundations in ethical theory rather than highly technical details. I begin by locating the standard approach to discounting within the overall landscape of ethical theory. The article then covers the Ramsey equation and its relationship to observed interest rates, arguments for and against a positive rate of pure time preference, the consumption elasticity of utility, and the effect of various sorts of uncertainty on the discount rate. Climate change is discussed as an application.

Keywords: discount rate; climate change; Ramsey equation; pure time preference.

Hilary Greaves

Somerville College, Oxford OX2 6HD
Email: hilary.greaves@philosophy.ox.ac.uk
URL: <http://users.ox.ac.uk/~mert2255/>

1 Introduction

It is now widely accepted that human activity is partly or largely responsible for climate change in the recent past and in the future, that under a “business as usual” approach such climate change is likely overall to be extremely damaging to human life and well-being, and that we can mitigate its impact by taking steps to reduce our emissions of greenhouse gases in the short- and medium-term. This raises the question of *to what extent we ought* to reduce our emissions. Being an ought-question, the question is explicitly normative. Further, it is an *ethical* question, since the people who stand to be damaged the most by anthropogenic climate change, and stand to benefit the most from any mitigative action, are not the same as those on whom most of the responsibility for mitigation would fall. Those who would ‘pay’ for the mitigation in question are largely those living (i) now and (ii) in relatively affluent countries; the beneficiaries of mitigation are primarily those in poorer countries (where climate impacts are expected to be the most severe) and those who are not yet born. The question of discounting relates to the temporal aspect of this issue; to a first approximation, it is the question of the extent to which the fact that some anticipated benefit of mitigation would occur a given length of time into the future reduces the value of that benefit for ethical purposes, compared with an otherwise-similar hypothetical benefit occurring now.

The issues here are enormously important, and have rightly attracted an increasing amount of attention. The inevitable consequence of this is that the debate has become increasingly complex, and it can become difficult to see the wood for the trees. This article is a survey of the literature (mostly: the economics literature) on discounting. The emphasis is on understanding discounting from first principles, organising the issues, and relating the controversies over ‘the discount rate’ to their foundations in matters of ethical theory. My survey thus emphasises the conceptual rather than the technical aspects of discounting, but I have not shied away from the use of mathematical notation where this is the appropriate means of expression; I have (however) tried to make the essential mathematics accessible to those with a minimum of mathematical background. Where possible, the emphasis is on surveying the arguments on all sides, and (for readers who wish to follow up any particular issue in more detail) indicating which positions are taken by which existing authors/articles, rather than on taking sides. Since some of the controversies in the ‘discounting community’ concern the very organisation of the conceptual landscape, however, it has not been possible to remain entirely neutral. The intended audience includes both economists who are interested in (in particular) the normative presuppositions and significance of claims found in the literature on discounting, and philosophers who wish to engage with the economists’ approach to intertemporal ethics, in general and in the context of climate change in particular.

The structure of the article is as follows. Section 2 situates the discussion of discounting within the landscape of ethical theories: in particular, I will follow a sizeable fragment of the literature in taking the fundamental issue to be one of how to maximise a given ‘social welfare function’ or ‘value function’,

so that the direct focus is on the theory of the good, rather than either the theory of the right or e.g. any specifically virtue-ethical considerations. Given that we are going to theorise about some value function, the next question is which; section 3 sets out the ‘discounted-utilitarian’ value function that is standard in the literature on discounting, and explains how this choice of value function embodies both some fundamental (and controversial) ethical principles, and some important simplifications that we will have occasion to revisit later. Section 4 introduces the terminology of ‘discount factors’ and ‘discount rates’ (as applied to consumption). Section 5 covers the ‘Ramsey equation’, an important equation that expresses the conditions for optimising the discounted-utilitarian value function in terms of a discount rate on consumption, and that serves to organise much of the subsequent controversy concerning the choice of discount rate.

Section 6 reviews the relationship of appeal to the Ramsey equation to alternative ways of determining a discount rate. Sections 7 and 8 survey the arguments concerning the value of two of the key inputs to the Ramsey equation: respectively, the discount rate on future *utility* or ‘rate of pure time preference’, and the consumption elasticity of utility. Section 9 considers the extension of the analysis to take into account empirical and evaluative uncertainty; this section includes discussion of the expected-utility approach to uncertainty and its alternatives, and Weitzman’s celebrated uncertainty-based argument for a declining effective discount rate. Section 10 reviews the recent controversy over discounting in the context of climate change, focussing in particular on the Stern review and its aftermath. Section 11 discusses the extent to which the simplifications that are embraced by the standard discounting framework limit that framework’s usefulness for climate-change purposes (specifically, the sense in which the discounting framework is focussed on marginal rather than large-scale changes, and the issues of intratemporal inequality and changing relative prices). Section 12 summarises.

2 Situating the discounting discussion within the landscape of ethical theories

The dominant approach to the issue of discounting takes it that there is some function — a ‘social welfare function’ or ‘value function’ — that is an increasing function of ‘consumption’ (at all times, by all people), and that in some sense we seek to maximise. Within this framework, the key questions are what the social welfare function is, and how one can identify in practice the actions that maximise it. The main task of the present review article is to conduct an overview of the literature regarding the ‘discounting’ aspect of these key questions.

Before doing that, however, we pause briefly to note some aspects of the debate that we thereby set aside. For one thing, the assumption that any such function exists is not trivial. In the first instance, it includes various assumptions of comparability. For example, to hold that all relevant aspects

even of a *single* person's life can be measured by single 'consumption' parameter is to assume that there exists a privileged way to trade off benefits and costs of different types accruing to a given person (for example, changes in the person's level of luxury goods vs subsistence goods, and of standard consumer goods vs environmental goods). Similarly, to hold that there is a single function capturing all evaluatively relevant aspects of *all* people's lives is to assume that one can trade off benefits and costs that accrue to different people (e.g. members of the global rich vs the global poor), and at different times (e.g. present vs. future generations). (This is not, of course, to assume that such rates of tradeoff must all be finite.) Further, to assume that we need only take into account consumption (even in this broad sense of 'consumption') *by people* is to ignore any moral relevance that the interests of *individuals of other species* (especially: other sentient species) might have, and similarly any intrinsic value that 'nature' might have over and above the interests of individual organisms.

Another aspect of the ethical discussion that we tend to set aside by focussing on the maximisation of a social welfare function can be seen if we first take a step back, to place the social-welfare-function framework in the context of an overall picture of the landscape of ethical frameworks. A standard taxonomy of ethics separates ethical questions into those pertaining to the theory of the *good* on the one hand, and those concerning the theory of the *right* on the other. In the first instance, one can enquire as to which *outcomes* (in the present case, involving varying levels of material wealth by various people at various times, and varying levels of climate change) are *better* than which others: this is the question of which theory of the good is correct. Secondly, one can enquire as to what a given agent, faced with a given decision, *ought to do*. The connection between the good and the right is one of the standard and controversial questions of normative ethics. *Maximising consequentialists* hold that one is morally obliged to bring about the best state of affairs one is able to bring about; advocates of *agent-centred prerogatives* argue that while one is always morally *permitted* to bring about the best state of affairs, one is also permitted to give some priority to oneself and one's nearest and dearest at the expense of strangers (for example, spending money on an expensive holiday or gift for one's child rather than donating it to a charity that would do more impartially-measured good with it); advocates of *deontological side-constraints* hold that in some cases, one is morally forbidden from bringing about the best state of affairs (for example, because one has *promised* the money to one's friend or *owes* it to a company that has delivered agreed services, notwithstanding the fact that a carefully selected charity could do more good with the same money). (For a survey of these and related issues, see, e.g., (Sinnott-Armstrong, 2014).)

Returning to our above question: a natural interpretation of the 'social welfare function' is as a function intended to represent overall *good*. On that interpretation, a focus on the questions of what the social welfare function is and how to maximise it amounts to a focus on the 'theory of the good' aspect of our 'how much mitigation ought we to undertake?' question, and a setting-aside of issues pertaining to the theory of the right. This move is in itself innocuous: the diversity of views on the 'theory of the right' notwithstanding, the majority

of ethical frameworks agree that the theory of the good is at least *one important part of* the overall ethical story. It is worth noting, however, that a minority of authors also urge the importance, to the discounting debate, of questions that we thereby set aside, such as questions of whether future generations have *rights* that are not to be violated even if some carefully chosen rights-violations would, on balance, lead to greater overall good. We also set aside questions of how, granted that some specified degree of emissions reduction is the one that *the world as a whole* ought to undertake, the responsibility for that overall degree of reduction is to be divided between countries, either as a matter of moral principle, or as a matter of international political negotiation. (See (Gardiner, 2004) for an overview of these and other broader ethical issues surrounding climate change.)

3 The standard framework for discounting

As advertised in section 2, then: We adopt the perspective of a benevolent social planner, seeking to maximise some value function that assigns numbers to states of affairs, in such a way that better states of affairs are assigned higher numbers than worse states of affairs. More precisely, since all decisions are made in the face of significant uncertainty (regarding, for instance, what the outcome of any given policy intervention would turn out to be), we will judge one course of action (*ex ante*) better than another whenever the first corresponds to a higher *expectation value* of this value function than does the second (cf. section 9).

The standard assumption in the literature on discounting is that the appropriate value function takes the ‘discounted-utilitarian’ form

$$V_1 = \int dt \cdot \Delta(t) \cdot \sum_{i \in N(t)} w(i, t). \quad (1)$$

That is: the overall value of a state of affairs is computed by (1) calculating the amount of well-being present at each time t , by summing momentary well-being levels $w(i, t)$ across all individuals i who are alive at time t ; (2) ‘discounting’ the well-being at each time t by a factor $\Delta(t)$ that represents how important well-being at t is, relative to well-being at other times; (3) summing the resulting discounted quantities across time.¹

This choice of value function is of course not uncontroversial. In particular, the following four remarks are in order.

¹We work in continuous time; thus (1) contains an integral $\int dt$ rather than strictly a sum \sum_t . If we wish to contract a discretised model by splitting time into discrete periods and using average or other representative discount factors and well-being indices for each time period, we would work instead with the expression $\sum_t \Delta(t) \sum_{i \in N(t)} w(i, t)$, in which $\Delta(t)$ and $w(i, t)$ are now (respectively) the discount factor for and the i th person’s well-being level in *period* t . This ‘discretisation’ is common practice, but is both slightly less elegant and slightly less true to reality. We work here with the continuous formula (1), but nothing of substance turns on this choice.

Firstly, note that even by utilitarian standards, the normal way to proceed would be first to compute an index of *lifetime* well-being for each person, corresponding to how well her life goes *as a whole* (that is, taking into account her momentary well-being level at each time in her life and combining them appropriately), and then to sum the resulting quantities across persons (with or without time-discounting). This is consistent with the functional form (1) only if the expression that relates an individual's lifetime well-being to her momentary well-being levels at various times is additively separable, and the latter is a controversial assumption (Broome, 1992, pp. 53–4). The assumption of additive separability, however, has at least a pragmatic justification in the present context, as we typically lack the cross-temporal information that would be needed in order to evaluate a value function in which this separability condition does not hold.

Secondly, the utilitarian claim that overall goodness is represented by *summing* well-being across persons is controversial: many people object to the resulting thesis that an additional unit of well-being contributes just as much to overall value if it accrues to an already-well-off person as it does if it accrues to a badly-off person. *Prioritarians* hold that while there is a quantity that should be summed across persons, that quantity is not well-being, but rather a concave transform of well-being; *egalitarians* deny that the value function exhibits additive separability of persons at all, in which case the correct value function must irreducibly take into account *comparisons* between the well-being levels of different individuals, and there is *no* function of individual well-being that can be simply summed across persons to generate an accurate index of overall value. There is of course an enormous literature on these matters; here we will simply set them aside, and work with the (discounted-)utilitarian value function for the sake of simplicity. (At the general conceptual level, the issues of discounting on which this survey focusses are largely orthogonal to the disputes between utilitarians and prioritarians/egalitarians. The details of the analysis, however, are potentially significantly more complicated in the prioritarian and especially in the egalitarian case; see appendix 212.)

Third, the choice of function for the factor $\Delta(t)$ in the discounted-utilitarian's framework is, as we will see, a matter of significant controversy. This is not a controversy over the appropriateness of including $\Delta(t)$ in the value function at all, however, since, for all we have said so far, $\Delta(t)$ could everywhere take the value 1.

Fourth, the standard discussion assumes that population is exogenous. That is, in using the value function (1), no commitment is here incurred to using the same value function for comparisons of states of affairs involving different sized populations. To use (1) also for variable-population cases would be to commit to *total* utilitarianism (as opposed to, for instance, average utilitarianism or some 'variable value' approach) on questions of population ethics. Questions of population ethics, however, are beyond the scope of this article (for a useful survey and critique, see (Arrhenius, n.d.).)

In practice, the standard framework for discussions of discounting makes three important further amendments and simplifications (some of which we will

need to revisit later, as they become particularly important and potentially seriously misleading in the context of climate change), as follows.

First, for the purposes of practical application we need to be able to relate the abstract quantity $w(i, t)$ — person i 's momentary well-being level at time t — to the various concrete factors that partially determine it. In reality, of course, a person's momentary well-being at any given time is determined by myriad factors, including various aspects of material consumption but also including number and quality of interpersonal relationships, a sense of purpose, education level, physical and mental health, availability of physical exercise and mental stimulation, access to amenities, housing quality, amount and quality of leisure time, and so on. We cannot in practice include all these factors in our analysis; the standard model makes the simplifying assumption that momentary well-being is determined by the individual's 'consumption' of various well-defined resources that are traded on the market (e.g. rice, beans, books, electronic goods), together with certain non-market goods for which 'shadow prices' can reasonably be estimated (air quality, safety, access to national parks). Suppose that there are k such resources. Then we replace $w(i, t)$ in (1) with a function $u_i(c_1(i, t), c_2(i, t), \dots, c_k(i, t))$, where $c_j(i, t)$ is the amount of resource j that person i consumes at time t , and u_i is the function that determines person i 's well-being level on the basis of this vector of consumption levels. In practice, we simplify further by assuming that the utility function is the same for all individuals, so we write simply u rather than u_i . This modifies (1) to

$$V_2 = \int dt \cdot \Delta(t) \cdot \sum_{i \in N(t)} u(c_1(i, t), c_2(i, t), \dots, c_k(i, t)). \quad (2)$$

Second, instead of working directly with the (in principle very long) vector of k different goods, we work in terms of a single variable c that we refer to simply as 'consumption'. This is innocuous *in principle*: for any high-dimensional space representing all possible consumption bundles in the context of k different goods, we can always consider the indifference surfaces that are determined in that space by the individuals' utility function, and construct a real-valued variable c that indexes those indifference surfaces. If we do so in such a way that higher-utility indifference surfaces correspond to higher values of c , we can then work with a utility function $u(c)$ that is an increasing function of its *single* variable c , and no ultimately-relevant information has in principle been lost. We then have

$$V_3 = \int dt \cdot \Delta(t) \cdot \sum_{i \in N(t)} u(c(i, t)). \quad (3)$$

While some aspects of the discounting discussion can be carried out at this level of abstraction, however, any proposal of a *particular functional form* for the utility function $u(c)$ (such as the CRRA utility functions mentioned in section 8), or (relatedly) a particular number of percentage points per annum for the discount rate (such as those surveyed in Table 1 below), is necessarily sensitive to the choice of a particular way of indexing indifference curves by 'consumption

numbers' c . There are several more or less principled labelling techniques. One might, for example, choose a reference vector of relative prices for the commodities in one's model (most naturally, but parochially: the relative prices that are given by marginal rates of substitution between the commodities in question at one's *actual, current* consumption bundle), and index all indifference curves by the minimum expenditure needed to reach the curve in question given the reference prices. (This and other methods of indexing indifference curves are described in e.g. (Deaton & Muellbauer, 1980, section 7.2).) We must then remember that the numerical (as opposed to qualitative/structural) aspects of the discounting discussion are sensitive, at least in principle, to the choice of labelling technique. Relatedly, the use of a single index of 'consumption' can in practice encourage users of the resulting model to neglect the important phenomenon of 'changing relative prices'; we return to this point in section 11.3.

The third simplification is by any standards far from innocuous: the standard framework abstracts away from issues of *intra*temporal inequality. That is: since we wish to examine issues of intertemporal ethics, we of course do not assume that persons' consumption levels are the same *at different times*. But the standard framework *does* proceed as though there were no differences in consumption levels between any two persons at a given time. In that case, the index i in ' $u(c(i, t))$ ' is no longer required, so that the value function V_3 above is further modified to

$$V = \int dt \cdot \Delta(t) \cdot N(t) \cdot u(c(t)), \quad (4)$$

where $N(t)$ is the number of people alive at time t . This raises the question of whether conclusions drawn from analysis of the standard framework are still valid once the reality of *intra*temporal inequality is taken into account.

The quantity V given in (4) is the value function that is used in the vast majority of the literature on discounting. Since its purpose is to engage with that literature, the present article will largely follow suit. Section 11, however, will consider the ways in which the results of the discussion would be substantively altered if these simplifications were not made.

4 Discount factors and discount rates

Suppose we have an opportunity to undertake an investment project, sacrificing $k < 1$ units of consumption today in order to secure an increase of 1 unit in consumption a time interval t later. Our basic question is: what is the threshold value of k at which the status quo becomes socially preferable to such a sacrifice? The answer to this question is the social *discount factor* for consumption at time t , $R(t)$. (One also has a *private* discount factor, corresponding to private as opposed to social preferences; in the remainder of this article, the focus is on the social version.)

One generally expects R to decline with time — we are willing to sacrifice more today to gain an increase in one unit of consumption tomorrow than to gain

an increase of one unit of consumption next year. The decline, however, may be faster or slower, and for current purposes the *rate* of the decline is crucial. We write $\frac{dR}{dt}$ for the *rate of change* of the discount factor (if R is declining, $\frac{dR}{dt}$ is negative). The *discount rate*, r , is proportional to this rate of change, in such a way that the faster the discount factor decreases, the greater the discount rate:

$$r = -\frac{1}{R} \frac{dR}{dt}. \quad (5)$$

We can rearrange the expression (5) to obtain a formula for the discount factor at time t in terms of the discount rates between times 0 and t : $R = \exp\left(-\int_0^t r(t')dt'\right)$. In the special case of a constant discount rate, this simplifies to the standard relationship $R = \exp(-rt)$, illustrated in Figure 1.

INSERT FIGURE 1 HERE
Caption for Figure 1: The discount factor as a function of time, for various constant-discount-rate scenarios $R(t) = \exp(-rt)$. Note, in particular, that with a discount rate of 3% per annum or higher, the discount factor is well below 0.1 for times more than 100 years in the future, and is essentially zero for times more than 200 years in the future. With a lower discount rate (1% per annum), the discount factor declines more slowly, but still reaches very low values for far-future times.

The choice of discount rate is crucial in the evaluation of projects some of whose important effects are long-term. Analyses that use a higher discount rate will tend to favour the short term: projects requiring sacrifices in the short term for the sake of benefits in the further future will be more likely to fail cost-benefit tests. Thanks to the exponential relationship between the discount rate and discount factor, small changes (at all times) to the discount rate can lead to very large changes to the amount by which distant future goods are discounted. Theoretical disagreements over how the discount rate is to be determined therefore have immediate policy importance, in particular in the context of climate change. Notoriously, the disagreement between Stern (2007) and e.g. Nordhaus (2008) over how much action to mitigate climate change is cost-effective is traceable almost entirely to the difference in their discount rates (respectively, 1.4% and 5%). More generally, since the social cost of carbon affects the relative costs of e.g. modes of transport and forms of technology, and is in turn highly sensitive to the discount rate, the latter is highly relevant to project evaluation in all government departments (Rose, 2012).

5 The Ramsey equation

The standard approach to determining the discount rate is via *the Ramsey equation*. This equation arises from the problem of optimising the value function (4) that we arrived at in section 3. To recap, this value function is given by

$$V = \int dt \Delta(t) N(t) u(c(t)),$$

where $\Delta(t)$ is the discount factor for utility at time t , $N(t)$ is population size at time t , $c(t)$ is average per capita consumption at t , and u is the instantaneous utility function for an individual, expressed as a function of consumption.

Relative to any ‘status quo’ consumption path $c(t)$, the ‘investment project’ outlined at the start of section 4 involves two changes: a decrease in consumption by k units now (at time 0), and an increase in consumption by 1 unit at time t . The first change tends to decrease V , while the second tends to increase V ; we seek the conditions under which the *net* effect of these two changes together is to increase (or at least to preserve) V .

The answer² is that in order to increase overall value V , an investment project’s rate of return must exceed (the “welfare-preserving rate of return on savings”)

$$r = \delta + \eta g, \tag{6}$$

where:

- $\delta := \frac{\dot{\Delta}}{\Delta}$ is the (negative of the) proportional time rate of change of the *utility* discount factor, a.k.a. the ‘rate of pure time preference’;
- $\eta := -c \frac{u''}{u'}$, the consumption elasticity of utility, depends on both the utility function and (in principle) the consumption path. Here, $u' \equiv \frac{du}{dc}$ is the rate at which utility increases as consumption increases, while $-u'' \equiv -\frac{d^2u}{dc^2}$ is the rate at which $\frac{du}{dc}$ itself *decreases* as consumption increases. A higher value of η corresponds to a more concave utility function, i.e. one in which the phenomenon of diminishing marginal returns of consumption to utility is more marked;
- $g := \frac{\dot{c}}{c}$ is the consumption growth rate, i.e. the proportional time rate of change of (average per capita) consumption.

(6) is the *Ramsey equation*. The premise that we ought to maximise (4) grounds an argument that we ought to use the quantity (6) as our discount rate when evaluating marginal projects.

Authors agreeing that the Ramsey equation provides correct guidance as to the value of the discount rate r nevertheless obtain different numbers for that rate, as they disagree on the inputs to the Ramsey equation: see Table 1.

INSERT TABLE 1 HERE Caption for Table 1: The values for δ , g and r (in % p.a.), and for η , adopted by a selection authors and public bodies.

Stern’s ‘ $\delta = 0.1\%$ p.a.’ actually corresponds to discounting for existential risk, not pure time preference. See Stern (*ibid.*).

The values of δ and η are discussed in sections 7 and 8 below (respectively). g is a straightforwardly empirical parameter, albeit one about whose value there is great uncertainty; the extension of the model to deal with uncertainty is the topic of section 9.

²The proof is in the Appendix.

6 Two empirical quantities related to the discount rate

Two further arguments urge that the discount rate should be set equal to, respectively, the social rate of return r_{MC} on marginal capital, and the interest rate r_{CM} on credit markets. Since neither of these quantities is guaranteed to equal a discount rate determined directly by the Ramsey equation, potential disagreements arise. It is worth making our peace with these, as their relationship to the Ramsey equation is not immediately clear, and the second in particular has played a significant role in the debate over climate change. (See also the discussion in (Gollier, 2013, chapter 1).)

The social rate of return r_{MC} on marginal capital is defined as the largest value of r for which we have unfunded possible projects (i.e., opportunities not currently being taken) that would deliver an increase of 1 unit of consumption at time t per investment of $\exp(-rt)$ units made now. The argument for using r_{MC} as the discount rate is then a basic arbitrage argument: if the project under consideration requires a sacrifice of more than $\exp(-r_{MC}t)$ units for the same gain, it should not be undertaken, because the same funds could alternatively be invested elsewhere in the market for a greater gain at the same future date. (Thus, for example, in the context of climate change: Lomborg (2001, ch. 24) argues against climate change mitigation measures on the basis that many currently unfunded third world development projects offer higher rates of return.) If, on the other hand, the project under consideration requires a sacrifice of less than $\exp(-r_{MC}t)$ units, then it should be undertaken in preference to some of the projects that are already being funded.

Under the assumption that society's current consumption plan is socially optimal, we do not need to choose between this argument and that of section 5: the social rate of return on marginal capital is then equal to the welfare-preserving rate of return on savings. If, on the other hand, the social rate of return on marginal capital diverges from the welfare-preserving rate of return on savings, then *prima facie* (taking this and the argument of section 5 together) we seem to have a contradiction, since we have arguments for each of two distinct numerical candidates for the socially efficient discount rate.

The appearance of contradiction evaporates when we clarify whether we are asking (i) whether it is better to undertake the proposed project or consume the funds now, or (ii) whether it is better to undertake the proposed project or some alternative investment project. If, say, the social rate of return on marginal capital is higher than the welfare-preserving rate of return on savings, that can only be because society is currently consuming too much (investing too little). In that case it can very well happen that the internal rate of return of some proposed project passes the Ramsey-equation test, but is lower than the social rate of return on marginal capital. Undertaking the project is then better than consuming the resources in question now, but undertaking some other, more productive project is better still ((Goulder & Williams, 2012, sec. 3.4), (Gollier, 2012, section 4)). Provided we understand our question as the first

one, the Ramsey-equation test remains valid. (In a more pluralist vein, Goulder and Williams (2012) also argue that while the Ramsey-equation test determines whether a project improves overall social welfare, it is r_{MC} (in their notation, r_F) that determines whether the project is a ‘potential Pareto improvement’ in the Kaldor-Hicks sense, and that depending on the decision context, either or both criteria may be important.)

Turning now to the relevance (or otherwise) of the credit markets: the argument for setting the discount rate equal to the interest rate r_{CM} prevailing on the credit markets is based on an appeal to democracy. First, an arbitrage argument establishes that the interest rate prevailing on credit markets reflects society’s actual willingness to sacrifice current for future consumption: if (say) investors were willing to sacrifice one unit of current consumption for less than $e^{-r_{CM}t}$ units of consumption a time t hence, then there would be an excess of supply of credit, and the interest rate would be driven down. Second, it is asserted that in a democracy, governments should discount future consumption at the same rate as this empirically observable public willingness to invest.

Again, under certain assumptions of ideal conditions, the interest rate on credit markets is equal to the social rate of return on marginal capital ($r_{CM} = r_{MC}$), and hence we do not need to choose between this and the preceding argument. The ‘idealising’ assumption required in this case is that the *social* rate of return on marginal capital (i.e. the rate of return calculated by taking into account *all* changes in future consumption generated by the project, regardless of to whom they accrue) is equal to the *private* rate of return (i.e. the rate of return calculated by taking into account only future consumption changes accruing to the investor herself). Given that assumption, we can reason as follows. First, an arbitrage argument establishes that the credit-market interest rate r_{CM} must be *at least* the private rate of return on capital: if it were not, then marginal investors in credit, instead of investing in the credit market, would switch to investing directly in productive capital, thus driving up the credit-market interest rate. Second, a similar argument establishes that the credit interest rate must be *at most* the private rate of return on marginal capital: if the former exceeded the latter, investors would switch from direct investment in capital to lending on the credit markets, until the two rates were brought into alignment. We have, then, the intermediate conclusion that $r_{CM} = r_{MC}^{\text{private}}$. Given our further assumption of equality between private and social rates of return, we can therefore conclude that $r_{CM} = r_{MC}$ as defined above (i.e., $r_{CM} = r_{MC}^{\text{social}}$). That further assumption, however, is wildly implausible, since externalities — market imperfections leading to divergences of private from social rate of return — are ubiquitous.

Divergences between r_{CM} and r_{MC} aside, this ‘democracy’ line of thought is often pitted directly against the use of the Ramsey-equation approach, when the latter yields a social discount rate distinct from observed interest rates. This discussion has recently been particularly prominent in the context of climate change; we return to it in section 10.

7 The rate of pure time preference

The present section sets aside the overall discount rate r on *goods*, and focusses on δ , the discount rate on *utility* or ‘rate of pure time preference’. Theorists are divided over whether this rate of pure time preference should be positive on the one hand, or zero on the other.

7.1 Arguments for a zero rate of pure time preference

The basic ‘argument’ for a zero rate of pure time preference is from impartiality. Accepting a zero rate of pure time preference merely amounts to treating utility as equally valuable, regardless of when, where or to whom it occurs. But of *course* (runs the thought) the value of *utility* is independent of such locational factors: there is no possible justification for holding that the value of (say) curing someone’s headache, holding fixed her psychology, circumstances and deservingness, depends upon which year it is. The axiomatic nature of impartiality is endorsed both by many of the seminal articles on the subject (Sidgwick, 1890; Ramsey, 1928; Pigou, 1932; Harrod, 1948; Solow, 1974) and by many current authors (Cline, 1992; Broome, 2008; Dasgupta, 2008; Dietz, Hepburn, & Stern, 2008; Buchholz & Schumacher, 2010; Gollier, 2013).

To this we may add two further arguments. The second argument proceeds from the Pareto principle, and points out that this principle is inconsistent with a nonzero rate of pure time preference (Cowen, 1992). To see the inconsistency, suppose that a particular person — Sarah, say — could live either in this century or the next. Consider two states of affairs that differ over when Sarah lives. Suppose that Sarah’s well-being is slightly better in the state of affairs in which she lives later, while everyone else’s well-being is unchanged. Then according to the Pareto principle, the ‘Sarah lives later’ state of affairs is better. But according to a value function whose rate of pure time preference is positive, this state of affairs may be worse. Thus $\delta \neq 0$ is inconsistent with the Pareto principle.

Thirdly, if $\delta > 0$ for negative (i.e. past) as well as positive (future) times, we have the absurd implication that deaths in the past were worse than deaths now. If, on the other hand, $\delta > 0$ only for positive times, then we have temporal relativity, and this temporal relativity tends to lead to temporal inconsistency in judgments. We pause to explain the latter point (see also (Broome, 2004, section 4.3); (Broome, 2012, pp.148–52)).

Let us define a *schedule of discount factors* to be an assignment of a discount factor to each *pair* consisting of the time of the benefit to be evaluated and the time at which the evaluation is carried out: thus R_j^i is the discount factor used by an evaluator at time t_i when evaluating costs or benefits occurring at time t_j . By definition, a schedule of discount factors is *time-neutral* iff the ratios $\frac{R_j^i}{R_{j'}^i}$ are independent of i ; otherwise it is time-relative. The point is then that if $\frac{R_j^i}{R_{j'}^i} \neq \frac{R_j^{i+1}}{R_{j'}^{i+1}}$, then evaluations carried out at times i and $i + 1$ will disagree with

one another on (some) decisions that involve trading off well-being at time j with well-being at time j' , even in the absence of any new information. Furthermore, the agents doing the evaluating at time i will be in a position to *foresee* that they will change their minds at $i + 1$, despite not having learnt anything new. This sort of time-inconsistency is generally taken to be an indicator of irrationality.

In the present case: Suppose that the rate of pure time preference, at any evaluation time t , is zero for all times prior to t , and some positive constant $\delta > 0$ for all times later than t . Let t_1, t_2 be any pair of times with $t_2 > t_1$. We consider an operation that sacrifices c_1 units of utility at time t_1 in return for an increase of b_2 units of utility at time t_2 . We evaluate this operation, first from the point of view of t_1 , and then from the point of view of t_2 . An evaluator at t_1 applies a discount factor of $\exp(-\delta(t_2 - t_1))$ to utility gains occurring at the later time t_2 , and therefore deems the operation to be an overall improvement iff $n \exp(-\delta(t_2 - t_1)) > m$. The evaluator at t_2 , however, judges one unit of utility to be equally valuable whether it occurs at t_1 or t_2 , and hence favours our operation iff $n > m$. This second criterion is of course less stringent than the first; there are thus operations such that, by the lights of his own evaluation strategy, an actor at t_1 should turn them down, but the same actor at t_2 should regret having done so.

Suppose, for example, that Ben gets the same amount of utility from each doughnut he eats, regardless of which day he eats it on, and (at least for small numbers of doughnuts) regardless of how many other doughnuts he has eaten on the same day. But suppose that each day, Ben discounts his own *future* utility in the manner suggested, while regarding all *present and past* utility as equally valuable: specifically, suppose that he discounts utility at the (extreme) rate of 60% per day. On Monday, he is offered a choice between (option 1) eating one doughnut on Tuesday or (option 2) eating two on Wednesday; given his Monday values, he judges option 1 to be better. He still agrees with this judgment on Tuesday. But come Wednesday, he has changed his mind about the relative value of Tuesday doughnuts and Wednesday doughnuts: he now thinks that Tuesday doughnuts and Wednesday doughnuts are equally valuable, and thus judges option 2 to be better than option 1. It is too late, of course: he cannot now *do* anything about it, but he regrets the decision he made on Monday. Furthermore, when deliberating on Monday, he was already in a position to *foresee* that he would thus later regret the decision he was making.

It is worth noting well, however, that the *only* sort of inconsistency that can result from the discounting structure in question is the phenomenon of foreseeable regret that we have seen in this example. This is arguably not as bad as foreseeable *backtracking*, in which an agent decides at time i to pursue a certain course of action at time $i + 2$, while recognising that at time $i + 1$ she will no longer think that this course of action is best, and so will attempt to reverse her decision. The latter but not the former raises issues of commitment ('tying oneself to the mast'); cf. the discussion of declining discount rates in section 9.2.

7.2 Arguments for a positive rate of pure time preference

Most theorists see some significant *prima facie* force to the intuition of impartiality reported above, but it is far from agreed that that intuition carries the day. We now consider four arguments for a positive rate of pure time preference.

Our first argument is a direct response to the impartiality intuition: that full impartiality, while perhaps a moral ideal, is not a *requirement* of morality. For example, according to common sense morality, it is at least permissible to give the interests of one's friends and family more weight than those of strangers in one's decisions; according to defenders of a positive rate of pure time preference, the latter is merely a manifestation of this permissible partiality in the intertemporal case (K. J. Arrow, 1999; Flanigan, n.d.). In reply, several authors (Cowen & Parfit, 1992; Caney, 2008; Beckstead, 2013) point out that to say that future utility is *just as valuable as* present utility is not to commit to a moral obligation to maximise overall value, and that Arrow's objection seems to be more to the latter account of obligation ('maximising consequentialism'). To this, the proponent of the 'permissible partiality' argument might reply that his 'social welfare function', contrary to the interpretation we dubbed 'natural' in section 1, is not *intended* to be a representation of 'overall value' in the impartial sense; it is, rather, the function such that in his view the course of action that maximises that function is the most advisable one, subject to the constraints of morality. The question then is whether this representation by a single discounted function is a sensible form for the non-maximising-consequentialist view in question (as opposed, say, to the use of a fully impartial value function *supplemented by* principles expressing agent-centred prerogatives); the critics argue that it is not.

Secondly, and relatedly, many theorists hold that the use of a zero rate of pure time preference requires excessive sacrifice on the part of the present generation for the benefit of future generations (again, see Arrow (1999)). The relevant theory here concerns the optimal rate of saving. For example, Arrow calculates that under certain plausible-looking circumstances, a zero rate of pure time preference would require us to save over two thirds of our income for the benefit of future people. (This 'excessive sacrifice' argument is of course closely related to the 'permissible partiality' argument just stated; what the 'excessive sacrifice' argument adds is that impartiality (is not only not required, but also) is not a sensible option.)

It is instructive to compare this "excessive sacrifice" argument in the context of intertemporal ethics with one that often crops up in the assessment of utilitarianism as a moral theory for the *intratemporal* case. According to utilitarianism, rich people (like ourselves) should give away the vast majority of our wealth to the desperately poor. Many (e.g. (Scheffler, 1982)) object that this implication is too demanding, and conclude that utilitarianism is false. The difference in the intertemporal case is that the fully impartial value function asks us to make these large sacrifices for people *who are already richer than us*. Thus, in practice the 'excessive sacrifice' argument is more compelling in the intertemporal case than in the intratemporal case. (The difference arises

because of the possibility of generating a larger benefit in consumption terms for others than one's own sacrifice, thanks to the phenomenon of investment. This effect has no analog in the intratemporal case.)

Whether or not a value function embodying a zero rate of pure time preference does mandate an intuitively excessive level of saving, even under the assumption of a moral requirement to maximise, of course depends on the instantaneous utility function it is based on. Instead of concluding that $\delta > 0$, some authors (for example, Asheim and Buchholz (2003), Dasgupta (2008)) conclude that the utility function must be more concave than those considered by Arrow. Others (e.g. (Dietz et al., 2008, p.382) note that increasing the amount of assumed technical progress will also tend to decrease the optimal level of savings, for fixed δ . Another possibility is that ($\delta = 0$ but) technical progress will push up growth to the point at which (for ordinarily assumed valued of η) optimal savings rates are not excessive.

Thirdly, there is a cluster of worries about adverse implications of time-impartiality in cases where the consumption or utility streams under consideration extend into the infinite future. The basic worry can be seen by considering the natural value function for infinite utility streams in the absence of discounting, $\sum_{i=0}^{\infty} u_i$. Even if there is a common maximum value for each instantaneous utility u_i , the sum $\sum_{i=0}^{\infty} u_i$ has no upper bound, and (in addition) for infinitely many possible utility streams the sum will fail to converge at all. It follows that no utility stream is maximal relative to the preference order over utility streams that is represented by this value function, and (in addition) that the preference 'order' thereby represented is incomplete (strictly speaking, it is a *preorder* rather than an *order* on utility streams).

Developing this basic observation, an extensive literature investigates the consistency of time-impartiality with various other apparently plausible constraints on preference orderings in axiomatic frameworks (as opposed to: in the context of considering particular value functions). Inconsistency results are established by, e.g., Koopmans (1960), Diamond (1965) and Epstein (1986).

Many authors take these concerns relating to infinite contexts to supply sufficient motivation for adopting a positive rate of pure time preference (note that with bounded instantaneous utilities, the *discounted* sum $\sum_{i=0}^{\infty} \beta^i u_i$ ($\beta < 1$) *does* converge for every possible infinite utility stream (u_i)). Alternative responses include the following. Gollier (2013, p. 32) questions the motivation for insisting on the existence of an *optimal* consumption stream, since one can determine whether or not an alternative is an *improvement over the status quo* even if neither status quo nor alternative are optimal; this does not, however, address the second worry that the undiscounted sum may simply fail to converge for either or both alternatives. Svensson (1980) questions the motivation for the 'continuity' assumptions involved in many of the impossibility results, and establishes the consistency of a complete preference ordering and time-impartiality with the remainder of the Koopmans-Diamond axioms (see also Broome (1992, pp.104-105)). A third response is that the utility streams under consideration in real scenarios do not in fact involve an infinite time horizon, since the human race will eventually become extinct due to e.g. the heat death of the Universe;

if so, then (the response continues) it is perverse to base ethics for the actual world on paradoxes of infinity that one would have to face up to in distant alternative possible worlds. From the point of view of practical applications (if not of fundamental principle), perhaps the most powerful response is that of Dasgupta (2008, p. 157): even if the technical considerations in question do induce one to adopt a positive rate of pure time preference, an *arbitrarily small* positive value will suffice. The theorems, for instance, would supply no grounds for any objection to $\delta = 10^{-100}\%$ per annum.

A fourth argument for a positive rate of pure time preference proceeds from the premise that the actions of a government should be selected on the basis of aggregating the preferences of present members of the body politic. The argument starts from the empirical premise is that people do in fact discount future utility (the empirical literature on this issue is surveyed in (Frederick, Loewenstein, & O'Donoghue, 2002)). Therefore, the argument runs, government should respect this preference, and itself discount future utility when evaluating projects. Since a close analogue of this last argument for the discount rate on *goods* has cropped up in the discussion of climate change, we delay discussion of its soundness until section 10.

8 Consumption elasticity of utility, η

Recall (from section 5) that the second input to the Ramsey equation, the consumption elasticity of utility η , is defined by the equation $\eta = -c \frac{u''(c)}{u'(c)}$. The instantaneous value of η thus depends both on the utility function, and on the instantaneous level of consumption. In the literature on the discount rate, it is standardly assumed that the utility function has the ('constant relative risk aversion') form

$$u(c) = \frac{c^{1-\gamma}}{1-\gamma}, \quad (7)$$

where γ is a constant, i.e. is independent of c . (This assumption is merely for mathematical tractability.) In that case, a simple calculation shows that $\eta = \gamma$, so that η itself is also independent of c (and hence also time-independent). A special case of (7) is the logarithmic utility function, $u(c) = \ln c$, obtained in the limit $\gamma \rightarrow 1$.

8.1 Three roles for η

A higher value of η corresponds to a more concave utility function. In the standard approach — that is, the approach that seeks to maximise the expectation value of the value function (4) — this leads η to play three conceptually distinct roles, related to the degree of aversion to (respectively) risk, intratemporal inequality and intertemporal inequality. We pause to explicate these three roles in turn.

On risk aversion: The *expectation value* $E[X]$ of a random variable X is the probability-weighted average of the possible values of X . We then say that a

value function exhibits *risk aversion* iff it ranks a given lottery (in the sense of: an assignment of probabilities to consumption streams) as less good than a guarantee of the consumption stream that is the expectation value of those involved in the lottery: if, for example, it ranks a sure prospect of \$100 as better than a 50:50 lottery yielding either \$50 or \$150. Assuming (as is standard; but see section 9.3) the expected-value approach to uncertainty, this is the condition that $E[V(c)] \leq V(E[c])$. We rewrite (4) as

$$V = \int dt \cdot V_t, \tag{8}$$

where $V_t = \Delta(t) \cdot N(t) \cdot u(c(t))$. Since $E[V] = \int dt \cdot E[V_t]$, we can abstract from intertemporal issues, and consider a representative time t ; we will have risk-aversion in the above sense if the utility function is such that for any lottery over consumption level c , $E[u(c)] \leq u(E[c])$. A standard argument (Eeckhoudt, Gollier, & Schlesinger, 1995, p.8) establishes that this condition obtains if and only if the utility function is concave (that is, while $u' > 0, u'' < 0$).

This tells us when a utility function exhibits risk aversion, but not yet when one utility function exhibits *more* risk aversion than another. That is, we have not yet defined a notion of *degree* of risk aversion. To do so, we define the *risk premium* associated with a given lottery to be the amount of consumption π such that enjoying a sure consumption level that is equal to the expected consumption level in the lottery *less* π is *ex ante* equivalent to facing the lottery: that is, the quantity π such that $E[u(c)] = u(E[c] - \pi)$. We then say that utility function u_1 exhibits *more risk aversion* than utility function u_2 if and only if the risk premium for every lottery is higher according to u_1 than according to u_2 . It is straightforward to show (Pratt, 1964) that one utility function exhibits higher risk aversion than another if and only if the first yields a higher value of η at all consumption levels than the second. (Strictly: If the first has η everywhere at least as high as the second, and strictly higher for at least one consumption level in every interval. See (Pratt, 1964, Thm. 1).)

On aversion to inequality: As a first pass (but see below), we say that a value function exhibits *aversion to inequality* if, given two individuals with unequal consumption levels, a transfer of a certain amount of consumption from the higher-consuming to the lower-consuming individual (which transfer, however, does not reverse the ordering of the two individuals' levels) always increases value. (This is the 'Pigou-Dalton principle' (Sen, 1973a; Fleurbaey, 2012; Adler, 2013).) There are several possible rationales for aversion to inequality. First, consumption exhibits *diminishing marginal utility*, so that a given increment in consumption makes less difference to one's welfare the better off one (already) is. Second, the degree of welfare concomitant on a given level of consumption depends how that consumption level compares to the consumption levels of others with whom one is in contact, both for purely psychological reasons and due to social exclusion effects (Payne, 2000; Alesina, Tella, & MacCulloch, 2004; Di Tella & MacCulloch, 2006; Dynan & Ravina, 2007; Morgan, Burns, Fitzpatrick, Pinfeld, & Priebe, 2007; Clark, Frijters, & Shields, 2008; Oishi, Kesebir, & Diener, 2011). Thirdly, some take inequality of resources simply to

be intrinsically bad, perhaps for reasons of fairness. Finally, some theorists believe that inequality *even of welfare levels* is intrinsically a bad thing (Temkin, 1993), so that inequality of consumption levels is bad to a degree beyond that that would be mediated via by any of the above routes. To probe the relation between inequality aversion and η , we consider the cases of intra- and intertemporal inequality separately.

On aversion to intratemporal inequality: As we noted in section 3, since the value function (4) considers only matters of *average* per capita consumption at a given time, that value function altogether ignores issues of intratemporal inequality. More fundamentally, we would want to consider instead the value function (3):

$$V_3 = \int dt \cdot \Delta(t) \cdot \sum_{i=1}^{N(t)} u(c_i(t)),$$

where i indexes individual people, and $c_i(t)$ is individual i 's consumption at time t . We say that the value function is *averse to intratemporal inequality* if, given any two individuals with differing consumption levels at a given time t , overall value would be increased by giving each of the two individuals a consumption level that is half of their combined consumption, and leaving all others' consumption levels unaffected. An analysis structurally identical to that above then shows that V_3 is averse to intratemporal inequality if and only if u is concave, and that one value function of the form V_3 is 'more averse to intratemporal inequality' than another if and only if the first arises from a utility function u that everywhere has a higher value of η than the second.

Similar considerations apply to the analysis of *intertemporal* inequality, that is, inequality in consumption levels occurring at different times. To see this, return to the simplified value function (4), and consider a case in which two individuals, living at different times, enjoy different levels of consumption $c(t_1), c(t_2)$. The analysis in the intertemporal case is slightly complicated by the presence of the weighting factors $\Delta(t)$, but we can deal with the complications by defining the weighted average $\bar{E}(c(t_1), c(t_2)) := \Delta(t_1)c(t_1) + \Delta(t_2)c(t_2)$. If we say that (4) is *averse to intertemporal inequality* just in case its value is always increased by replacing both consumption levels in the unequal intertemporal pair with their weighted average $\bar{E}(c(t_1), c(t_2))$, then the same analysis as above establishes that (4) is averse to intertemporal inequality iff u is concave, and a similar extension of the reasoning establishes that one value function of the form (4) is 'more averse to intertemporal inequality' than another if and only if the first arises from a utility function u that everywhere has a higher value of η than the second.

It is this third role of η corresponds to its importance in the debate over the discount rate. If (as is standardly assumed) future generations will, climate change or no climate change, enjoy higher levels of average per capita consumption than present generations — if, that is, consumption growth g is positive — then a higher value of η corresponds to a higher discount rate. (If, instead, future generations are *worse* off than the present generation — due to

climate change or otherwise — then of course the relationship between η and the discount rate is reversed.)

8.2 Models that separate the three roles

Several authors (e.g. (Dietz et al., 2008); see also (Selden, 1978; Epstein & Zin, 1989)) have urged that since the three roles that are played by η in the standard approach are conceptually distinct, to represent all three by a single parameter is both to rule out without justification the possibility that the parameters relevant to the three roles vary independently of one another, and to court unnecessary confusion.

On the former point: it does not seem *incoherent* (at least) to exhibit relatively little aversion to risk, but great aversion to any form of interpersonal inequality (intra- or inter-temporal) (Beckerman & Hepburn, 2007). Still less does it seem incoherent to exhibit great aversion to *intratemporal* inequality, but relatively little aversion to *intertemporal* inequality, *if* the reason for inequality aversion is based on the psychological or social exclusion effects mentioned in section 8.1, as opposed to any view that inequality either of resources or of welfare is simply intrinsically bad. One might therefore seek a model that allows parameters for the three types of aversion to be varied independently, so that the resulting positions can be represented within the terms of the model.

On the latter point: As we noted at the outset, discussions of climate change often note that the most significant harms due to climate change will occur in the relatively far future, while the costs of any mitigative action will be incurred much sooner. Assuming that consumption growth is positive, this fact corresponds to a *prima facie tendency* for increasing η to favour less aggressive mitigative action; this is the tendency that is captured in the contribution of η to the discount rate in the basic Ramsey equation (6). But that basic Ramsey model ignores both intratemporal inequality and risk; a full analysis would carry out a line of the reasoning parallel to that of section 5, but starting from the *expectation* value of the more detailed value function (3) rather than from the simplification (4). Notwithstanding the point that *expected average* per capita consumption is higher in the future than it is for those who would stand to bear the costs of mitigation, in such a fuller analysis, increasing η would place greater evaluative emphasis on incremental changes to climate impacts occurring in those future *regions of the world* and/or *possible scenarios* in which the potential beneficiaries of mitigation are actually poorer than the mitigators. Thus the relationships of η to intratemporal inequality and to risk potentially work against its relationship to *intertemporal* inequality, when we ask whether, *overall*, increasing η would favour more or less mitigative action. (See also the discussions in (Dietz et al., 2008, p.378), (Dietz, Hope, & Patmore, 2007).)

To say that the three notions are *conceptually* distinct, however, is not to say that there cannot be close links between them — perhaps close enough that no model of *rational* decision-making can vary them independently. Arguments to this effect concerning the relationship between risk aversion on the one hand and (either type of) inequality aversion on the other are given by Harsanyi (1953;

1955; 1977).

8.3 Determining the value of η

Setting aside the discussion of section 8.2 and assuming the standard model, what then should the value of η be?

Many empirical studies aim to establish the extent to which real people *do* exhibit aversion to risk and interpersonal inequality in their choice behaviour. For the purposes of answering our *normative* question, it is important not to overstate the relevance of such empirical data: our question is what value for η our analysis *should* use, and it does not of course immediately follow, from the fact that people *do* employ a certain value for η , that they, we or anyone else *should* employ that value (Buchholz & Schumacher, 2010, p.378). (It does not *immediately* follow; but that is not to say that it does not follow at all. Some authors would defend the inference using a ‘democracy argument’ that is a close cousin of the one discussed in section 10.) That said, the empirical data may be informative nonetheless: a normative theorist advocating an η value that lies outside the range that empirical studies suggest that people normally use should be aware that he is so deviating from the statistical norm, and that he may be committed to passing a negative judgment on the values of the subjects in the empirical studies in question. We will first survey the more directly normative arguments, and then provide a brief survey of the empirical literature.

On the more directly normative side: Okun (1975, pp.91–5) outlines a ‘leaky bucket’ thought experiment, in which one envisages a transfer of resources from a richer to a poorer person (at specified initial wealth levels), and one asks how ‘lossy’ the transfer can be (that is, by how much the amount gained by the poorer person can fall short of the amount taken from the richer person) before the transfer ceases to constitute an overall improvement. If, for example, one judges that one person having £100k and another having £20k is neither better nor worse overall than one person having £60k and another having £30k (that is, the result of transferring £40k away from the better-off person but with a loss of 25% of the transferred resources), then in the context of the utilitarian model one is implicitly committed to an η value of about 1.2.³ Taking another approach, Buchholz and Schumacher (2010) show that axiomatic conditions intended to capture aspects of ‘circumstance solidarity’ or ‘absence of envy’ variously require $\eta = 1$, $\eta > 1$, $\eta = 2$ and $\eta = \infty$; as witness the variety of these conclusions, however, the motivation for the axioms is open to serious question.

³For *marginal* transfers, we can solve the problem in closed form: a simple calculation (assuming constant η) shows that if the threshold ‘leakiness’ for a marginal transfer is L ($0 \leq L \leq 1$), the corresponding value of η is given by

$$\eta = \frac{\log(1 - L)}{\log \frac{c_1}{c_2}}, \quad (9)$$

where c_1, c_2 are (respectively) the consumption levels of the poorer and the richer person. However, this relationship between the values of L and η can be very far from the mark in the context of non-marginal transfers, so must be used with caution (if at all).

Another approach considers the optimal rate of savings: as discussed in section 7, if one has an independent grip on the optimal balance between consumption and saving, then ‘working backwards’ can provide guidance on the correct value of η .

Turning now to the empirical studies: corresponding to the connections of η to both risk aversion and inequality aversion discussed above, data on either of the latter can be used to infer the values of η that *experimental subjects* are committed to, in their verbal or other behaviour. The ‘experimental subjects’ in these studies are sometimes individuals, sometimes governments. Focussing on risk aversion, (Gollier, 2006) infers $\eta \in [2, 4]$ from studies of individuals’ insurance behaviour. Carlsson, Daruvala and Johansson-Stenman (2005) use a questionnaire involving hypothetical decision-making on behalf of grandchildren to probe individuals’ attitudes to risk aversion and inequality aversion separately; their results suggest a median η of 2–3 from the risk aversion study, and approximately 2 from the inequality aversion study.

Approaches based on *government* behaviour often focus on tax policy. If one assumes that the government accepts a ‘principle of equal sacrifice’ — that taxation is to be organised in such a way that each taxpayer forfeits (say) the same amount of *utility* by his tax contribution — then one can infer an implied value of η from the variation in tax level with income. Using this approach, Stern (1977) argues that UK tax policy in 1973-4 implied an η -value of 1.97; Cowell and Gardiner (1999) note that the analogous data from the late-1990s would imply η in the range 1.3–1.4.

Aside from the above-mentioned worries about the relevance of empirical data to normative questions, the use of these studies to inform a choice of η for policy purposes has been criticised on the grounds that the results of many of the studies are highly sensitive to an ‘arbitrary’ choice of sacrifice principle (Creedy, 2006, p.15), and on the ground that they illicitly assume that the value of η is independent of consumption level (Atkinson & Brandolini, 2010, p.16). It is clear, however, that the more directly normative arguments surveyed above are to a large extent open to the same criticisms.

9 Uncertainty

Section 5 showed that the discount rate is a function of three parameters: the rate of pure time preference δ , the consumption elasticity of utility η and the growth rate of consumption g . But the value of each of these parameters is the subject of significant uncertainty. Predicting the future growth rate is a matter of empirical uncertainty: especially when it comes to the long-run future, it is extremely difficult to predict the combined effects of technological, political and climate change on future consumption levels with any accuracy. The values of δ and η , meanwhile, are matters of *evaluative* uncertainty, as witness the failure of the arguments of sections 7 and 8 to secure anything like unanimous agreement on their values. The question then arises of what discount rate should be used *in practice*, given that we have to carry out our policy assessments while living

with this uncertainty.

As mentioned briefly in section 8, the standard approach to the normative theory of decision-making under uncertainty is *expected utility theory*. (While standard, this theory is not *universally* accepted, especially in the context of ‘deep uncertainty’: we briefly expand on this point, and its significance for discounting, in section 9.3.) Under this approach, the ‘ex ante’ objective is to maximise the *expected* value of the ‘ex post’ value function (4), i.e. the quantity

$$E[V] = E \left[\int dt \cdot \Delta(t) \cdot N(t) \cdot u(c(t)) \right]. \quad (10)$$

Our fundamental question is then under what conditions a given marginal change to the consumption path increases the *expected* value (10).

Let us define the *effective discount factor* for time t , $R_{\text{eff}}(t)$, to be the amount by which marginal changes to consumption at time t should be discounted relative to marginal changes to consumption at time 0 under uncertainty, according to the expected-utility approach to choice under uncertainty. That is, given a marginal project involving a cost B_0 incurred at time 0 and a benefit $+B_t$ enjoyed at time t , $R_{\text{eff}}(t)$ is defined by the condition that the project increases expected value iff $R_{\text{eff}}(t)B_t > B_0$. We can then define the effective discount *rate*, r_{eff} , from R_{eff} in analogy with (5). Our question is then how these ‘effective’ quantities are related to the range of their possible ‘true’ values.

9.1 Weitzman’s argument

In a seminal article, Weitzman (1998) claimed that the correct results are given by using an effective discount *factor* for any given time t that is the probability-weighted average of the various possible values for the true discount factor $R(t)$: $R_{\text{eff}}(t) = E[R(t)]$. From this premise, it is easy to deduce, given the exponential relationship between discount rates and discount factors, that if the various possible true discount rates are constant, the effective discount rate declines over time, tending to its lowest possible value in the limit $t \rightarrow \infty$. Weitzman (2001) uses this analysis to calculate a value for the effective discount rate as a function as time, using as input data the opinions of 2,160 economists on the ‘true’ value of the (constant) discount rate.

Weitzman did not, however, supply any fundamental justification for his assumption that the effective discount factor is the expectation value of the true discount factor. That assumption is equivalent to the claim that a project should be adopted iff its *expected net present value* (ENPV) is positive. Gollier (2004) pointed out that if one adopts instead a criterion of positive expected net *future* value (ENFV), one obtains precisely the opposite of Weitzman’s result: the effective discount rate *increases* over time, tending as $t \rightarrow \infty$ to its *maximum* possible value. Since neither criterion is obviously inferior to the other, the difference in results generated the so-called ‘Gollier-Weitzman puzzle’.

To make progress on this puzzle, we note that more fundamentally, one wants to start from expected utility theory: under conditions of uncertainty, a given

project should be adopted if and only if its adoption increases *expected value*, where the ‘value’ in question aggregates over time as in equation (4). In general, this condition entails *neither* the ENPV nor the ENFV criterion. (To see this, consider a project involving a reduction in consumption $-B_0$ at time 0, in return for a corresponding increase B_t at time t . Such a project increases expected value iff $B_t E \left[\frac{\delta V}{\delta C_t} \right] > B_0 E \left[\frac{\delta V}{\delta C_0} \right]$, where $\frac{\delta V}{\delta C_t}$ (respectively $(\frac{\delta V}{\delta C_0})$), the functional derivative of value with respect to consumption at time t (resp. at time 0), measures the marginal rate at which overall value V increases per unit increase in consumption at time t (resp. 0), holding fixed consumption levels at other times.

The project has positive ENPV, on the other hand, iff $B_t E \left[\frac{\frac{\delta V}{\delta C_t}}{\frac{\delta V}{\delta C_0}} \right] > B_0$; since the random variables $\frac{\delta V}{\delta C_t}$ and $\frac{\delta V}{\delta C_0}$ are not in general independent, these criteria do not in general coincide. Similar remarks apply to the ENFV criterion.) The question, then, is whether the criterion of maximisation of expected utility entails anything like Weitzman’s result *given* certain auxiliary assumptions, or in certain models. In the lively post-2004 debate over the Gollier-Weitzman puzzle (Hepburn & Groom, 2007; Buchholz & Schumacher, 2008; Gollier & Weitzman, 2010; Freeman, 2010; Gollier, 2010; Traeger, 2012a; K. J. Arrow et al., 2014), there is a widespread consensus that something like Weitzman’s original conclusion is correct, although participants to the debate continue to differ significantly over the reasons for this conclusion and the precise conditions under which it holds.

A recent literature in moral philosophy (Lockhart, 2000; Ross, 2006; Sepielli, 2014, 2013b, 2013a; Harman, 2011; Weatherson, 2014; Nissan-Rozen, 2015; Gustafsson & Torpman, 2014; Mason, 2015) explores the possibility that evaluative, as well as empirical, uncertainty should be treated by means of expected utility (or ‘expected value’) theory. While this view is far from uncontroversial, the main sources of dissent stem either from the view that evaluative uncertainty is simply normatively irrelevant, or from the view that under evaluative uncertainty one should use the most-likely evaluative hypothesis, rather than proposing any alternative and non-trivial way to deal with such uncertainty. In the present context, the expected value approach suggests treating the values of δ and η , as well as g , as parameters that vary from one ‘state of nature’ to another. To the author’s knowledge, this approach has not been explored in any detail in the literature on discounting. Weitzman’s original (1998) argument was conducted at a level of generality sufficient to ensure neutrality between empirical and evaluative uncertainty, and his (2001) analysis fairly explicitly averages over evaluative and empirical uncertainty in the same way. Much of the later literature, investigating more detailed models in an attempt to justify one or another approach, tends (however) to restrict the scope of the uncertainty to the empirical. Arrow et al (2014) explicitly claim that normative and empirical uncertainty ‘require different approaches’ (2014, p.378); however, they neither offer any argument for this claim, nor propose any particular alternative treatment of normative uncertainty.

9.2 Declining discount rates

Even before Weitzman’s result, many theorists had suggested that the discount rate appropriate for application to the very far future was lower than that appropriate to the shorter term. The motivations for the suggestion are various (and dependent in part on the authors’ views on the methodological controversies discussed in this article). Those who take interest rates available on the credit markets as a reliable guide to discount rates observe that the interest rates available on government bonds, for instance (Newell & Pizer, 2003; Gollier, Koundouri, & Pantelidis, 2008) decline as the period of the bond increases. Those who are willing to apply data on individual behaviour cite studies suggesting that individuals employ declining discount rates in their private decision-making (Loewenstein & Prelec, 1992; Henderson & Bateman, 1995; Frederick et al., 2002). Quite another type of motivation, however, is unease at the implications of any constant discount rate for evaluation of the sufficiently far future: both the implication that the very far future matters very little indeed, and the implication that the *much further* future matters far less even than the ‘very’ far future. As Weitzman himself put it: ‘To think about the distant future in terms of standard discounting is to have an uneasy intuitive feeling that something is wrong, somewhere’ (Weitzman, 1998, p.201). The result discussed in section 9.1 therefore provides independent motivation for a claim that many had suspected for other reasons. This and other theoretical rationales for declining discount rates are surveyed in (Gollier, 2013, chapters 7 and 8).

As is widely recognised (Strotz, 1955; Thaler, 1981; Loewenstein & Prelec, 1992; Rubinstein, 2003), the use of a non-constant discount rate can lead to time-inconsistency in decision-making. The basic reason for this has already been mentioned (section 7.1 above): if the ratio between the discount factors applicable to times t_1 and t_2 changes between two times $t_A, t_B > t_A$ that are earlier than either t_1 or t_2 , an agent making investment decisions at time t_B may want to reverse decisions that were optimal by the lights of the discount factors applicable at t_A . If the details of the reversal are foreseeable at t_A , game-theoretic issues arise: what should the agent at t_A do *in the light of the knowledge that* if he were to attempt to implement a given policy, the later agent at t_B would attempt to reverse his actions (Phelps & Pollak, 1968; Barro, 1999)? In particular, these issues do arise in the context of a declining discount rate that is rationalised à la Weitzman.

9.3 Alternatives to the expected utility approach

The expected-utility framework requires the decision maker to settle on a (unique) probability distribution to represent his state of uncertainty. Following Ellsberg (1961) and Slovic and Tversky (1974), many authors argue that this is inappropriate when, as in the case of climate change, we face ‘deep uncertainty’: that is, when the available information does not even come close to singling out a single probability distribution as the uniquely rational one. The claim is that the phenomenon of ‘ambiguity aversion’ displayed in Ellsberg’s results

can be (not merely actual, but) perfectly rational: that is, it can be rational to prefer known to unknown odds of significant outcomes, in a manner that is inconsistent with the conjunction of expected utility theory and *any* hypothesis regarding degrees of belief about the unknown-odds cases.

Non-expected utility approaches endorsing the rationality of ambiguity aversion divide into two types: those that eschew the use of probabilities altogether, and those that deal with the non-uniqueness of probabilities by some method of aggregating over the various possible (objective) probability distributions. (For useful surveys, see (Al-Najjar & Weinstein, 2009; Etner, Jeleva, & Tallon, 2012; Heal & Millner, 2013).) The probability-free approaches include a straightforward maximin criterion (Wald, 1945, 1949), a variant on maximin according to which one maximises a weighted sum of the utilities of the worst possible and best possible outcomes (K. Arrow & Hurwicz, 1977), and a principle of ‘minimax regret’. Probability-based approaches include: a ‘maximin expected utility’ approach that evaluates a policy via the expected utility it has according to the prior that is least favourable to it (Schmeidler, 1989; Gilboa & Schmeidler, 1989); a ‘smooth ambiguity’ model that takes an expectation value (with respect to subjective ‘meta-probabilities’) of some increasing transform of expectation values (with respect to objective probabilities) of utility (Klibanoff, Marinacci, & Mukerji, 2005; Epstein, 2010; Klibanoff, Marinacci, & Mukerji, 2012); and a ‘multiplier preferences’ model (Hansen & Sargent, 2001, 2008).

The effect of ambiguity aversion on the discount rate is discussed in (Gierlinger & Gollier, 2008; Traeger, 2012b).

10 Discount rates in the context of climate change: the controversy over the Stern Review

We turn now to the application of the above discussion (of discounting in general) to the particular case of climate change. As we noted at the outset, the choice of discount rate is particularly crucial when, as in the case of climate change mitigation, the project under consideration involves costs and/or benefits that materialise in the distant future. The value placed on costs and benefits 100 years’ hence if a discount rate of 2% per annum is used, for instance, is more than 54 times the value that a 6% p.a. discount rate would place on those same costs and benefits. This fact was thrown into sharp relief by the publication of the Stern Review of the Economics of Climate Change (Stern, 2007), which breathed new urgency into the discounting debate.⁴

Notoriously, Stern applies a discount rate far lower than those advocated by the majority of economists. Stern accepts the ethical arguments for a zero rate of pure time preference, but postulates a 0.1% p.a. risk of human extinction from exogenous (non-climate-change-related) causes; thus he sets $\delta = 0.1\%$ p.a. He assumes a CRRA utility function with $\eta = 1$, and postulates a consumption

⁴For another recent survey of views on discounting in the context of climate change, see the IPCC’s Fifth Assessment Report (Panel on Climate Change, n.d., section 3.6.2).

growth rate $g = 1.3\%$ p.a. The resulting overall discount rate, r , is 1.4% p.a. In stark contrast, other discussions of climate change (Lomborg, 2004; Weitzman, 2007; W. Nordhaus, 2008) have tended to use discount rates in the region $4 - 6\%$ p.a. Stern concludes that *widespread and immediate* action to mitigate climate change would be cost-effective, whereas analyses such as Nordhaus' favour instead a 'climate-policy ramp' approach, with mitigation efforts being stepped up far more gradually. In the wake of the publication of the Stern Review, many commentators (e.g. Weitzman (2007), Nordhaus (2007)) were quick to point out that this difference in conclusions regarding optimal policy depends almost entirely on the difference in discount rates.

The nature of the disagreement between Stern and his critics involves several of the issues touched upon earlier in this article. The objections can be grouped into two main categories. Firstly, many worry about the implications of such a low discount rate for the optimal rate of saving; the arguments here parallel those surveyed in section 7.2. The second cluster of objections centres around a claim that the overall discount rate is simply observable, and that Stern's low figure of 1.4% p.a. is inconsistent with observed data.

We now examine the latter in more detail. Witness, for example, Nordhaus:

The... discount rate on goods... is a *positive* concept... [It] is also called the real return on capital, the real interest rate, the opportunity cost of capital, and the real return. The real return measures the yield on investments corrected by the change in the overall price level. In principle, this is *observable* in the marketplace. (2007, p. 689; emphasis added)

Nordhaus then proceeds to canvass various empirical data on 'the' real rate of return, citing a number of figures clustering around 6% . He then accuses Stern of inconsistency, noting that if r were observed even to be (say) 4% , then Stern's $\eta = 1$ and $g = 1.3$ would require — for consistency — $\delta = 2.7$; similarly, Stern's $g = 1.3$ and $\delta = 0.1$ would require $\eta = 3$. A similar charge is suggested by a comment by the authors of the IPCC Second Assessment Report's chapter on 'Intertemporal equity, discounting, and economic efficiency', defending a 'descriptive' as opposed to a 'prescriptive' approach to the discount rate ('overriding market prices on ethical grounds... opens the door to irreconcilable inconsistencies' (IPCC, 1996, p.133)).

The discussion is simplified somewhat by noting that for present purposes, the details of how any given figure for r breaks down into δ , η and g components is in the end irrelevant, since it is r itself that plugs directly into the cost-benefit analysis. Nordhaus' key claim, and that of other defenders of the 'descriptive' approach, is that this discount rate is itself entirely fixed by empirical data. In contrast to this, however, several discussions of discounting in the context of climate change point out that since the question of discount rate is a question of how society *ought* to discount future goods in policy analysis, the concept is clearly most fundamentally a normative, not a positive, one (e.g. (Broome, 2008), (Goulder & Williams, 2012, p.6)). Puzzlingly, the same authors of the above-quoted chapter of the IPCC Second Assessment Report also comment in

their ‘summary for policymakers’ that ‘[s]election of a social discount rate *is also a question of values* since it inherently relates the costs of present measures, to possible damages suffered by future generations if no action is taken’ (IPCC, 1996, p.8), emphasis added).

As we saw in section 6, this objection from observed interest rates involves two independent and essentially orthogonal issues. The first concerns the coincidence, or lack of it, between the welfare-preserving savings rate and the observed credit-market interest rate r_{CM} . In defence of simply taking observed market interest rates as the discount rate, proponents of this approach correctly point out (recall) that the *opportunity cost* of capital is key: it is not in the interest *even of future people* to accept projects with a low rate of return if this would crowd out alternative, higher-yielding projects. We outlined the general reply to this argument in section 6: the appeal to opportunity costs alone provides no motivation for using r_{CM} as the social discount rate when the status quo consumption path is non-optimal, and/or in the presence of externalities. (In the climate change literature, the relevance of consumption path optimality is pressed by Dasgupta (2008, secs. 7.2 and 7.3); the relevance of ‘market imperfections’ (including externalities) is alluded to in the Stern Review itself (Stern, 2007, p.51), and by Buchholz and Schumacher (2010, p.389).) This reply suggests redirecting Nordhaus’ complaint of inconsistency against the ‘descriptive’ approach: if the ethically and descriptively correct values for the parameters entering the Ramsey equation fail to match observed market interest rates, it is the latter that, for the purposes of selecting a social discount rate, are at fault.

The second issue involved in the discussion of the relevance (or otherwise) of observed interest rates is a charge of anti-democracy that is often levelled against authors who allow ethical issues to affect their determination of a discount rate. Here the basic idea is that the market interest rates encapsulate the degree of de facto willingness of today’s individuals to save for the future, and can therefore be accorded the status of a democratic vote on the value of the discount rate. From this point of view, to adopt any other interest rate is illegitimately to disregard this ‘democratic’ outcome in favour of one’s own individual preferences. Thus, Nordhaus again:

The [Stern] Review takes the lofty vantage point of the world social planner, perhaps stoking the dying embers of the British Empire, in determining the way the world should combat the dangers of global warming. The world, according to Government House utilitarianism, should use the combination of time discounting and consumption elasticity that the Review’s authors find persuasive from their ethical vantage point. (2007, p.691)

In a similar vein, Weitzman (2007, p. 707) accuses Stern of ‘paternalism’, complaining that Stern’s insistence on a zero rate of pure time preference is ‘irrespective of preferences for present over future utility that people seem to exhibit in their everyday savings and investment behavior.’

In reply to this appeal to democracy, four comments are in order. Firstly, we need to distinguish between discounting later utilities within a single life on

the one hand, and discounting utility that occurs in later lives on the other. Whatever one thinks of the ‘democracy’ argument in general, any empirical evidence that people discount *their own* future utility (at any rate) is presumably irrelevant to whether or not governments should discount the utility of *future generations* (e.g. (Cowen & Parfit, 1992, pp.146,155), (Gollier, 2013, p.31), (W. D. Nordhaus, 2007, p.691)). Secondly, even when the project in question is short-term enough that the main concern is the future utility of presently existing people, it is not clear that even a government concerned only with the well-being of its own citizens should respect utility-discounting preferences if, as is generally held to be the case, those preferences are irrational (see, e.g., (Groom, Hepburn, Koundouri, & Pearce, 2005) and references therein)). Thirdly, if (as with climate change) the project in question is longer term, so that the main concern is a trade-off between the utility of current people and the utility of future people, whether or not a democratic government ought to respect utility-discount preferences depends on whether the role of government is only to act as an agent of present voters, or also to act in part as a trustee for future generations (that is, the argument’s basic starting premise can be questioned). The former perspective is well-represented by Stephen Marglin, the latter by Arthur Pigou:

I want the government’s social welfare function to represent 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 members of the body politic. (Marglin, 1963, p.97)

There is wide agreement that the State should protect the interests of the future *in some degree* against the effects of our irrational discounting and of our preference for ourselves over our descendants. . . . It is the clear duty of Government, which is the trustee for unborn generations as well as for its present citizens, to watch over, and, if need, be, by legislative enactment, to defend, the exhaustible natural resources of the country from reckless spoilation. (Pigou, 1932, emphasis in original)

For more general discussion of the appropriate form of democracy when the interests of disenfranchised groups (such as future people) are at stake, see (Dobson, 1996). Fourthly, it has been argued that the idea that market interest rates reflect a democratic determination of the discount rate is anyway based on a misunderstanding of the nature of democracy, and evaporates once we understand the role of such publications as the Stern review as being an input into, rather than an attempt to summarise the output of, the democratic process ((Cowen & Parfit, 1992, p.145–6), (Broome, 2008)).

There is, therefore, a considerable body of opinion in favour of determining the social discount rate directly from the Ramsey equation, whether or not the result coincides with observed market interest rates. None of this, however, is to say that the particular values Stern chooses for the input parameters to the

Ramsey equation are the correct ones. And regardless of whether or not the above objections to Stern’s (Ramseyan) methodology are sound, the fact does remain that Stern’s 1.4% p.a. discount rate is far lower than those employed by most economists *at least for the evaluation of short-term projects*. Many are therefore left with the impression that the Stern Review ‘cooked the books’, imposing an unrealistically low discount rate simply in order to obtain support for the desired policy outcomes.

In response to this, it is worth noting that while it is indeed fairly common to employ a social discount rate in the range 4-6% p.a. for projects on a timescale of 30 years or less, for longer-term environmental projects — partly thanks to the concern that a such high constant discount rates lead to intuitively excessive discounting of environmental impacts in the *far* future — it is actually fairly common to employ a declining schedule of discount rates, so that rates as low as 1 or 2% p.a. apply in the further future. (For example, the UK Green Book (2003) recommends a discount rate r of 2% for consumption occurring after 126-200 years, and of 1% for consumption occurring more than 200 years in the future; the French Lebegue Report (2005) recommends a uniform 2% for all consumption more than 30 years in the future.) An emphasis on comparing opinions on *short*-term discount rates therefore risks overstating the degree of disagreement between Stern and others in the context of climate change.

11 The limits of applicability of the standard discounting framework

This section notes and assesses three respects in which it has been argued that the standard discounting framework misses important insights in the discussion of the appropriate amount of mitigation for climate change, related respectively to the facts that that framework is perturbative, that it ignores intratemporal (as proposed to *intertemporal*) inequality, and that it theorises in terms of a single summary index of consumption (and hence a single discount rate).

11.1 Appropriateness for matters of global climate policy

The standard discounting framework, as set out in sections 4–5, proceeds by means of first order approximation: that is, it analyses the conditions under which *in the limit as the size of the ‘investment project’ under consideration goes to zero* the project generates an improvement over the status quo. (This is clearest in the derivation of the Ramsey equation in the Appendix; it appears in the discussion of sections 4–5 via the stipulation that the investment project under consideration is ‘marginal’.) As a result, the framework is entirely appropriate for analysing e.g. the small-scale policy decisions of individual government departments (who seek to take into account the climate-change implications of their policy decisions when deciding on e.g. transport policies or construction projects), but its usefulness is limited if the question under discussion is that of *global* climate policy — that is, how much mitigation *the world as a whole*

ought to be undertaking. The analytical problem in the latter case is that if we are comparing two scenarios, A and B, that involve *very different* consumption streams, the value difference between A and B is not accurately calculated by evaluating the (marginal) rates of change of value with consumption at any given times, and multiplying by the consumption differences between A and B at the corresponding times — yet that is the closest that a naive application of the discounting framework could come to calculating that value difference. A similar point has been stressed in particular by Stern (2007, 2014), who argues against overly naive application of the standard discounting framework in global-scale analyses of climate change, emphasising that for some purposes, global-scale analyses have to return to simply evaluating the chosen value function at each of the alternative scenarios under consideration, and seeking to identify the highest-value alternative directly (as the Stern Review itself attempts to do).

It is important, however, not to overstate this point. It does not follow, from the acknowledged fact that one cannot calculate the value difference between two very different consumption paths via a ‘marginalist’ framework, that that framework is altogether irrelevant to issues of intertemporal ethics even for non-marginal, global climate policy. There are two reasons for this.

The first reason is already mentioned in Stern’s discussion: it is that, notwithstanding the sense in which choices of global climate policy ultimately require direct (non-perturbative) modelling, many of the same issues as those discussed above in the perturbative framework will also crop up, and will be amenable to much the same arguments, in any such direct modelling exercise. Any such direct analysis still has to decide on (i) the value (positive or otherwise) to assign to the rate of pure time preference in its fundamental value function, (ii) the assumptions it makes regarding future growth rates (or probability distributions over such future growth rates), and (iii) the degree of concavity in the utility function linking consumption level to individual well-being or utility; and both the ethical determinants and the implications of those decisions will still be much as discussed above.

The second reason is that the sense in which the marginalist framework is ‘inapplicable’ to the analysis of non-marginal changes is itself limited: there are, that is, some highly relevant things that one can learn even about *large* changes from examination of a first-order perturbative analysis (provided that the problem we face has certain structural features).

To see this, it is helpful to proceed by means of a simple, abstract analogy. Suppose that ultimately we seek the maximum of some function f ; for simplicity, let f be a function of a single real-valued variable x only. Towards answering our question, let us in the first instance select one possible value (\bar{x}) of this variable to serve as ‘status quo’. We could then start by asking a more limited question: not precisely where the maximum of f lies, but merely whether the optimum is located at a variable-value higher than, lower than or equal to \bar{x} . *If* we know (the structural assumption in this analogy) that f is continuously differentiable, has only one stationary point and that stationary point is a maximum (rather than a minimum or a point of inflection), then a first-order perturbative analysis carried out at the ‘status quo’ point \bar{x} suffices to answer this more limited

question: if the derivative (i.e., the gradient) of f at \bar{x} is positive (respectively, negative, zero), then the maximum must be located at a variable-value higher than (respectively, lower than, equal to) \bar{x} . (The point is clearest in a graphical representation, for which see figure 2.)

INSERT FIGURE 2 HERE
 Caption for Figure 2: The red line shows the graph of f as a function of x ; f has a single maximum, at x_2 . We also consider randomly selected points x_1 and x_3 such that $x_1 < x_2 < x_3$. The gradient of f at x_1 (respectively, at x_3) is positive (respectively, negative); subject to auxiliary structural assumptions, even without e.g. calculating the *value* of f at x_1, x_3 or anywhere else, or calculating the precise location of the maximum, this information on local gradients would suffice for the correct prediction that x_1 (respectively, x_3) is below (respectively, above) the point at which f takes its maximum value.

What this ‘first-order’ perturbative analysis cannot tell us is *by how much* the location of the maximum exceeds, or falls short, of x . It cannot tell us that because the maximum is the (different) variable-value at which the gradient of f is zero (i.e., the graph of f is horizontal), while the perturbative approach looks only at the gradient of f *at the particular point* x . This quantitative matter is certainly an important question. But an answer to the more limited question above is still somewhat informative.

Whether or not the case of climate change is analogous depends on whether or not the climate-change problem satisfies the required structural assumptions, over a sufficiently wide range of mitigation scenarios. This is a substantive question, and not one I will attempt to tackle here.⁵ *If* it does, however, then relative to any ‘status quo’ consumption path — which could be the consumption path that is projected to result from any particular proposal for a course of mitigative action — a first-order perturbative analysis can tell us, given plausible background assumptions, whether the amount of mitigation in that ‘status quo’ proposal is higher than, lower than or equal to the optimum. The perturbative analysis of any *given* path does not tell us by how much *that* path falls short of, or exceeds, the optimum; but by subjecting a variety of possible consumption paths to that more qualitative question in turn, we get a good handle on the quantitative question also. In particular, via this method, the marginal approach would (given the required structural assumptions) easily be informative enough to capture most of the controversy that has raged between Stern and his critics, each of whom, as we have seen, proposes a rough specification of an optimum, and criticises the opponent’s suggestion for involving far too much or far too little mitigation.

⁵A referee has pressed the concern that nonlinearities in the climate system, in particular, could invalidate the required structural assumptions.

11.2 Distributional issues

In section 3, we followed the standard discounting literature in abstracting away from issues of intratemporal inequality: in place of the more fundamental value function 1, we turned to analysing instead the simpler formula 4, which latter deals only with *average* per capita consumption at each time. We must therefore ask whether our eventual conclusions are adversely affected by this simplification.

In the context of global climate change discussions, the answer is surely positive. As many commentators have pointed out, to the extent that we do decide to mitigate climate change, the costs of mitigation will be borne primarily by the richest of those alive today, while the benefits will accrue preferentially to the poorest of present and future people. From the point of view of ‘growth discounting’, the salient question is therefore the comparison in consumption levels between these parties. Therefore, even if the received view of consumption growth (viz., that average per capita consumption will continue to rise over the relevant time period, climate change notwithstanding) is correct, therefore, it may yet be the case that, in the relevant sense, the beneficiaries of mitigation are at best only marginally richer than those who would bear the costs. If so, clearly this would drive down the discount rate that is appropriate to evaluation of climate change mitigation.

This line of argument is strengthened further by considerations of uncertainty: combining the above line of thought with that of the Weitzman argument surveyed in section 9.1 suggests that the relevant comparison is between the worst off possible beneficiaries of mitigation and the best off possible cost-bearers, in which case the applicable effective discount rate is sure to be negative (Fleurbaey & Zuber, 2012).

11.3 Changing relative prices

Again as we noted in section 3, the standard framework proceeds in terms of a single real-valued index of ‘consumption’. Since ‘consumption’ in reality is of course consumption of many different goods (rice, beans, housing, clean air, access to leisure . . .), a significant amount of aggregation has already been carried out, in representing these various widely differing aspects of consumption by a single real value, rather than by a high-dimensional vector of such values.

There is nothing wrong with such aggregation in principle. However, as we flagged in section 3, there is some danger that use of the aggregated framework might lure us into ignoring the phenomenon of changing relative prices. It is now time to explain in more concrete terms what the mistake might be, in the particular case of climate change.

The salient possibility here is that future vs present people will, in general, enjoy consumption bundles of different compositions, in such a way that while future people are (perhaps) better off overall (that is, their consumption bundle is on a higher indifference curve than the consumption bundle enjoyed by present people), the future consumption bundle involves significantly lower standards

of those goods that we expect to be particularly impacted by adverse climate change (perhaps: a pleasant ambient temperature, and relative freedom from natural disasters and from certain diseases). In such a case, it might be that a given improvement to future environmental conditions generates a greater increase in overall value than would be generated by the same improvement if it occurred today — the ‘environmental discount rate’ is negative — despite the fact that future people are better off overall. In such a case, one is misled if one discounts *environmental* goods using the (positive) discount rate that applies to *consumption*.

Slightly more generally, the point is as follows. Suppose (simplifying somewhat) that we have a privileged way of representing all aspects of consumption in terms of just *two* indices m and e , representing, respectively, amount of ‘material’ goods consumed and amount of ‘environmental’ goods enjoyed. In terms of the indices m and e , we have a two-dimensional space of possible goods-bundles, and a utility function defined on pairs (m, e) . As above, we might also define on this space a single ‘consumption’ index c . A *consumption path* is then a map from times to points of this space: it tells us how much consumption of each of e and m there is (and hence what level of c we are at) at each time t . Relative to any such consumption path, we can then define the *shadow prices* of each of e, m and c at any given time t , as the amount by which overall value (as defined by (4)) increases when an extra unit of (respectively) e, m or c becomes available at time t , while other aspects of the consumption path are held fixed. We then have a discount rate for each of the three quantities e, m, c , defined as in section 4 from the corresponding time-profile of shadow prices. In general, all three of these discount rates will be different.

Suppose now that starting from some specified status quo consumption path, we consider a project (for example, a climate-change mitigation project) that offers an increase in environmental goods e at some future time t , in return for some decrease in material goods m at the present time 0. We wish to calculate whether this project represents an increase in overall value. One (perfectly accurate) way of analysing this matter is as follows. Define *relative* shadow prices between any two of e, m and c at any time t , $\lambda_{ij,t}$, as the ratio of the two shadow prices in question at that time. Then, first, convert the increase in future e to an increase in future ‘consumption’ c using the relative price of e relative to c at the future time t ($\lambda_{ec,t}$). Second, discount back to the present time using the discount rate for consumption. Third, to compare the result to the present m -cost, we also convert that m -cost to the welfare-equivalent amount of present ‘consumption’ using the *present* relative prices of m and consumption ($\lambda_{mc,0}$).

The point then is that it is easy, if one is using the aggregated ‘consumption’ framework but not thinking carefully enough, to get the conversions between e, m and c wrong, and hence to arrive at incorrect evaluations for projects of this nature. In particular, one gets them wrong if one tries to use the *present* relative price of environmental goods and consumption ($\lambda_{ec,0}$) as a proxy for their *future* relative price $\lambda_{ec,t}$: in general, and especially in contexts of deteriorating environmental quality but rising material consumption, these two ratios will be very different (‘changing relative prices’). But this is effectively what

one does if one follows the common practice of using willingness-to-pay surveys, conducted on (who else?) *present* people, for the purpose of placing a monetary (and thence ‘consumption’) value on *future* climate damages.

The case for stronger mitigation that might result from this line of thought is developed in e.g. (Guesnerie, 2004), (Sterner & Persson, 2008), (Gollier, 2013, chapter 10), all of whom analyse the issue in a ‘two-good’ framework that works directly in terms of separate indices for environmental and from non-environmental (materialistic) goods throughout and corresponding separate discount rates, and eschews the use of any single, overarching index of ‘consumption’.

12 Summary

The choice of discount rate that is used in a given cost-benefit analysis has an immense impact on the degree of mitigation for climate change mitigation that will be judged cost-effective by that analysis. There is, however, significant disagreement among theorists as to the correct value of the discount rate, both because of methodological disagreements concerning which inputs are relevant to determining it (the ‘prescriptive’ Ramsey-equation-first approach versus direct appeals to observed interest rates), and because of disagreements over the values of the key parameters in the Ramsey equation (notably, the rate of pure time preference and the consumption elasticity of utility). Without attempting to endorse a particular number by way of final conclusion, this article has conducted a survey of the various controversies involved and the arguments on each side, both in general and in the context of climate change and the Stern Review in particular.

Section 2 clarified that the issue of discounting is most naturally understood as a topic within the ‘theory of the good’: that is, the topic of discussion is which states of affairs are better than which others (and by how much). This leaves open the question of whether there might also be other factors relevant to what ‘we’ (whether as private individuals, organisations, countries, or a global community) *ought to do* (for instance, deontological side-constraints or agent-centred permissions); however, it is highly plausible that in large-scale decision contexts in particular, the moral ideal would be for the theory of the good to be the dominant consideration. Section 3 outlined the presuppositions and simplifications that are involved in settling on the particular ‘discounted utilitarian’ value function that is the standard model in the majority of the literature on discounting. Section 4 clarified the two key concepts of ‘discount factors’ and ‘discount rates’ that are central to discussions of discounting, and their mutual relationship. Section 5 outlined how the Ramsey equation for the discount rate arises from the exercise of maximising discounted-utilitarian value.

Section 6 discussed the relationship of the Ramsey equation to (i) the social rate of return on marginal capital and (ii) the interest rate on credit markets. This, in particular, is one place in which I have had to ‘take sides’ regarding the organisation of the debate. Some authors would regard the choice of whether to

locate the Ramsey equation, or instead one of these other two quantities, at the centre of the discounting debate as a key substantive choice point around which any survey of discounting should be organised ('descriptive vs. prescriptive approaches'); I have instead taken the Ramsey equation (more generally, the value-function-first approach) to be clearly the overarching issue, and explored the conditions for the quantities (i) and (ii) to yield the same verdicts as the Ramsey equation (together with reasons for thinking that, when they do not, it is the quantities (i) and (ii) rather than the ideal of maximising value that tend to lose their normative significance).

For applied purposes, under the Ramsey-equation approach the key question is the numerical value of each of the three parameters that appear in that equation: the rate of pure time preference, the consumption elasticity of utility, and the growth rate of consumption. Here in particular, I have not attempted to recommend particular numerical values, but rather to survey the types of arguments that are offered in the course of this exercise by others. Sections 7 and 8 surveyed the arguments concerning the rate of pure time preference and the consumption elasticity of utility respectively (I leave the growth rate of consumption aside as beyond the scope of this survey, since it is an empirical rather than an evaluative matter). Section 9 surveyed the ways in which the 'effective' discount rate is affected by the presence of both empirical and evaluative uncertainty, including Weitzman's influential uncertainty-based argument for a declining effective discount rate, and the possibility of deviating from the standard expected-utility approach to uncertainty.

All of the discussion to this point has been about discounting in purely general terms: i.e., not specific to the use of discounting in the analysis of climate change. Section 10 turned to the discussion of climate change in particular, exploring how the discussion of discounting has played out in response to the Stern Review, and noting where the key moves that have been made in the course of this discussion are reflections of the underlying, more general issues that I have surveyed in sections 2-9. The discussion of climate change also serves as a useful case study to illustrate the application of those more abstract considerations. Section 11 examines the limits of this: that is, the extent to which the general discussion of discounting employs simplifications that, while often reasonable, notably fail in the climate-change scenario. The general message there was that the standard framework remains in principle applicable, but that caution must be exercised, in particular over the issues of intratemporal inequality and changing relative prices, if important insights are not to be lost. In both of the two latter cases, the effect of ignoring the insights in question is a marked tendency to underestimate the case for stronger mitigation.

ACKNOWLEDGEMENTS

I am grateful to an anonymous referee for *Economics and Philosophy* for detailed and very helpful comments.

References

- Adler, M. (2013). *The Pigou-Dalton principle and the structure of distributive justice*. (Working paper)
- Alesina, A., Tella, R. D., & MacCulloch, R. (2004). Inequality and happiness: are europeans and americans different? *Journal of Public Economics*, 88(910), 2009 - 2042. doi: <http://dx.doi.org/10.1016/j.jpubeco.2003.07.006>
- Al-Najjar, N. I., & Weinstein, J. (2009). The ambiguity aversion literature: a critical assessment. *Economics and Philosophy*, 25(03), 249–284.
- Arrhenius, G. (n.d.). *Population ethics: The challenge of future generations*. (Unpublished manuscript)
- Arrow, K., & Hurwicz, L. (1977). An optimality criterion for decision-making under ignorance. *Studies in resource allocation processes*, 482.
- Arrow, K. J. (1999). Discounting, morality, and gaming. In P. R. Portney & J. P. Weyant (Eds.), *Discounting and intergenerational equity*. Washington: RFF Press.
- Arrow, K. J., Cropper, M. L., Gollier, C., Groom, B., Heal, G. M., Newell, R. G., ... others (2014). Should governments use a declining discount rate in project analysis? *Review of Environmental Economics and Policy*, reu008.
- Asheim, G. B., & Buchholz, W. (2003). The malleability of undiscounted utilitarianism as a criterion of intergenerational justice. *Econometrica*, 70, 405-422.
- Atkinson, A. B. (1970). On the measurement of inequality. *Journal of economic theory*, 2(3), 244–263.
- Atkinson, A. B., & Brandolini, A. (2010). On analyzing the world distribution of income. *The World Bank Economic Review*, 24(1), 1–37.
- Barro, R. J. (1999). Ramsey meets Laibson in the neoclassical growth model. *The Quarterly Journal of Economics*, 114(4), 1125-1152. doi: 10.1162/003355399556232
- Beckerman, W., & Hepburn, C. (2007, January). Ethics of the discount rate in the Stern review on the economics of climate change. *World Economics*, 8(1), 187-210.
- Beckstead, N. (2013). *On the overwhelming importance of shaping the far future* (Unpublished doctoral dissertation). Department of Philosophy, Rutgers University.
- Blackorby, C., & Donaldson, D. (1980). A theoretical treatment of indices of absolute inequality. *International Economic Review*, 107–136.
- Broome, J. (1992). *Counting the cost of global warming*. Cambridge: White Horse Press.
- Broome, J. (2004). *Weighing lives*. Oxford University Press.
- Broome, J. (2008, June). The ethics of climate change. *Scientific American*, 298, 96-102.
- Broome, J. (2012). *Climate matters*. New York: W. W. Norton and Company.

- Buchholz, W., & Schumacher, J. (2008). Discounting the long-distant future: A simple explanation for the Weitzman-Gollier-puzzle. *CESifo Working Paper 2357*.
- Buchholz, W., & Schumacher, J. (2010). Discounting and welfare analysis over time: Choosing the η . *European Journal of Political Economy*, 26(3), 372 - 385. Retrieved from <http://www.sciencedirect.com/science/article/pii/S0176268009001001> (Special Issue on Ethics and Economics) doi: <http://dx.doi.org/10.1016/j.ejpoleco.2009.11.011>
- Caney, S. (2008). Human rights, climate change, and discounting. *Environmental Politics*, 17(4), 536–555.
- Carlsson, F., Daruvala, D., & Johansson-Stenman, O. (2005). Are people inequality-averse, or just risk-averse? *Economica*, 72(287), 375–396.
- Clark, A. E., Frijters, P., & Shields, M. A. (2008). Relative income, happiness, and utility: An explanation for the easterlin paradox and other puzzles. *Journal of Economic Literature*, 46(1), 95-144. doi: 10.1257/jel.46.1.95
- Cline, W. R. (1992). *The economics of global warming*. Washington, D.C.: Institute for International Economics.
- Cowell, F., & Gardiner, K. (1999). Welfare weights. *STICERD Research Paper 20*.
- Cowen, T. (1992). Consequentialism implies a zero rate of intergenerational discount. In P. Laslett & J. S. Fishkin (Eds.), *Justice between age groups and generations*. Yale University Press.
- Cowen, T., & Parfit, D. (1992). Against the social discount rate. In P. Laslett & J. S. Fishkin (Eds.), *Justice between age groups and generations*. New Haven: Yale University Press.
- Creedy, J. (2006). Evaluating policy: Welfare weights and value judgments. *University of Melbourne, Department of Economics: Research Paper(971)*.
- Dasgupta, P. (2008, December). Discounting climate change. *Journal of risk and uncertainty*, 37, 141-169.
- Deaton, A., & Muellbauer, J. (1980). *Economics and consumer behavior*. Cambridge University Press.
- Diamond, P. A. (1965, January). The evaluation of infinite utility streams. *Econometrica*, 33(1), 170–177.
- Dietz, S., Hepburn, C., & Stern, N. (2008). Economics, ethics and climate change. In K. Basu & R. Kanbur (Eds.), *Arguments for a better world: Essays in honour of amartya sen. volume ii: Society, institutions and development*. Oxford: Oxford University Press.
- Dietz, S., Hope, C., & Patmore, N. (2007). Some economics of dangerous climate change: Reflections on the *ij* stern review/*ij*. *Global Environmental Change*, 17(3), 311–325.
- Di Tella, R., & MacCulloch, R. (2006). Some uses of happiness data in economics. *Journal of Economic Perspectives*, 20(1), 25-46. doi: 10.1257/089533006776526111

- Dobson, A. (1996). Representative democracy and the environment. In W. LaFerty & J. Meadowcraft (Eds.), *Democracy and the environment*. Cheltenham: Elgar.
- Dynan, K. E., & Ravina, E. (2007). Increasing income inequality, external habits, and self-reported happiness. *American Economic Review*, *97*(2), 226–231. doi: 10.1257/aer.97.2.226
- Eeckhoudt, L., Gollier, C., & Schlesinger, H. (1995). The risk-averse (and prudent) newsboy. *Management Science*, *41*(5), 786–794.
- Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, *75*(4), 643–669.
- Epstein, L. G. (1986). Intergenerational preference orderings. *Social choice and welfare*, *3*, 151–160.
- Epstein, L. G. (2010). A paradox for the smooth ambiguity model of preference. *Econometrica*, *78*(6), 2085–2099.
- Epstein, L. G., & Zin, S. E. (1989). Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, *57*(4), pp. 937–969.
- Etner, J., Jeleva, M., & Tallon, J.-M. (2012). Decision theory under ambiguity. *Journal of Economic Surveys*, *26*(2), 234–270. doi: 10.1111/j.1467-6419.2010.00641.x
- Flanigan, T. (n.d.). *On discount rates in the cost-benefit analysis of climate change*.
- Fleurbaey, M. (2012). Economics and economic justice. *Stanford Encyclopaedia of Philosophy (Summer 2012 edition)*. (Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2012/entries/economic-justice/>.)
- Fleurbaey, M., & Zuber, S. (2012). Climate policies deserve a negative discount rate. *Chicago Journal of International Law*, *13*, 565.
- Fleurbaey, M., & Zuber, S. (2015). Discounting, risk and inequality: A general approach. *Journal of Public Economics*, *128*, 34–49.
- Frederick, S., Loewenstein, G., & O'Donoghue, T. (2002, June). Time discounting and time preference. *Journal of Economic Literature*, *40*(2), 351–401.
- Freeman, M. C. (2010). Yes, we should discount the far-distant future at its lowest possible rate: A resolution of the weitzman-gollier puzzle. *Economics: The Open-Access, Open-Assessment E-Journal*, *4*(2010-13).
- Gardiner, S. M. (2004). Ethics and global climate change*. *Ethics*, *114*(3), 555–600.
- Gierlinger, J., & Gollier, C. (2008). Socially efficient discounting under ambiguity aversion. *IDEI Working Paper*, 561.
- Gilboa, I., & Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of mathematical economics*, *18*(2), 141–153.
- Gollier, C. (2004). Maximizing the expected net future value as an alternative strategy to gamma discounting. *Finance Research Letters*, *1*(2), 85–89.
- Gollier, C. (2006). An evaluation of Stern's report on the economics of climate change. *IDEI Working Paper 464*.
- Gollier, C. (2010). Expected net present value, expected net future value, and the ramsey rule. *Journal of Environmental Economics and Management*,

- 59(2), 142–148.
- Gollier, C. (2012). The debate on discounting: Reconciling positivists and ethicists. *Chi. J. Int'l L.*, 13, 549.
- Gollier, C. (2013). *Pricing the planet's future: The economics of discounting in an uncertain world*. Princeton University Press.
- Gollier, C., Koundouri, P., & Pantelidis, T. (2008, October). Declining discount rates: Economic justifications and implications for long-run policy. *Economic Policy*, 23, 757–795.
- Gollier, C., & Weitzman, M. L. (2010). How should the distant future be discounted when discount rates are uncertain? *Economics Letters*, 107(3), 350–353.
- Goulder, L. H., & Williams, R. C. (2012). The choice of discount rate for climate policy evaluation. *Climate Change Economics*, 3(4).
- Groom, B., Hepburn, C., Koundouri, P., & Pearce, D. (2005). Declining discount rates: the long and short of it. *Environmental and resource economics*, 32, 445–493.
- Guesnerie, R. (2004). *Calcul économique et développement durable* (Vol. 55) (No. 3). Presses de Sciences Po (PFNSP).
- Gustafsson, J. E., & Torpman, O. (2014). In defence of my favourite theory. *Pacific Philosophical Quarterly*, 95(2), 159–174.
- Hansen, L. P., & Sargent, T. J. (2001). Robust control and model uncertainty. *American Economic Review*, 60–66.
- Hansen, L. P., & Sargent, T. J. (2008). *Robustness*. Princeton University Press.
- Harman, E. (2011). Does moral ignorance exculpate? *Ratio*, 24(4), 443–468.
- Harrod, R. (1948). *Towards a dynamic economics*. London: Macmillan.
- Harsanyi, J. (1955, August). Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, 63(4), 309–321.
- Harsanyi, J. C. (1953). Cardinal utility in welfare economics and in the theory of risk-taking. *Journal of Political Economy*, 61(5), 434.
- Harsanyi, J. C. (1977). *Rational behavior and bargaining equilibrium in games and social situations*. Cambridge University Press.
- Heal, G., & Millner, A. (2013). Uncertainty and decision in climate change economics. *NBER Working Paper*(w18929).
- Henderson, N., & Bateman, I. (1995). Empirical and public choice evidence for hyperbolic social discount rates and the implications for intergenerational discounting. *Environmental and Resource Economics*, 5(4), 413–423.
- Hepburn, C., & Groom, B. (2007). Gamma discounting and expected net future value. *Journal of Environmental Economics and Management*, 53(1), 99–109.
- IPCC. (1996). *Climate change 1995: Economic and social dimensions of climate change: Contribution of working group III to the second assessment report of the Intergovernmental Panel on Climate Change*. Cambridge university press. (Edited by James P. Bruce, Hoesung Lee and Erik F. Haites)
- Klibanoff, P., Marinacci, M., & Mukerji, S. (2005). A smooth model of decision making under ambiguity. *Econometrica*, 73(6), 1849–1892.

- Klibanoff, P., Marinacci, M., & Mukerji, S. (2012). On the smooth ambiguity model: a reply. *Econometrica*, 80(3), 1303–1321.
- Koopmans, T. C. (1960, April). Stationary ordinal utility and impatience. *Econometrica*, 28(2), 287–309.
- Lebegue, D. (2005). *Revision du taux d'actualisation des investissements publics*. Commissariat General au Plan.
- Lockhart, T. (2000). *Moral Uncertainty and Its Consequences*. New York: Oxford University Press.
- Loewenstein, G., & Prelec, D. (1992, May). Anomalies in Intertemporal Choice: Evidence and an Interpretation. *The Quarterly Journal of Economics*, 107(2), 573–97.
- Lomborg, B. (2001). *The skeptical environmentalist: measuring the real state of the world*. Cambridge University Press.
- Lomborg, B. (2004). *Global crises, global solutions*. Cambridge university press.
- Marglin, S. A. (1963). The social rate of discount and the optimal rate of investment. *Quarterly Journal of Economics*, 77, 95–111.
- Mason, E. (2015). Moral ignorance and blameworthiness. *Philosophical Studies*, 172(11), 3037–3057.
- Morgan, C., Burns, T., Fitzpatrick, R., Pinfold, V., & Priebe, S. (2007). Social exclusion and mental health conceptual and methodological review. *The British Journal of Psychiatry*, 191(6), 477–483.
- Newell, R. G., & Pizer, W. A. (2003, July). Discounting the distant future: how much do uncertain rates increase valuations? *Journal of Environmental Economics and Management*, 46(1), 52–71.
- Nissan-Rozen, I. (2015). Against moral hedging. *Economics and Philosophy*, 31(03), 349–369.
- Nordhaus, W. (2008). *A question of balance*. New Haven, CT: Yale University Press.
- Nordhaus, W. D. (2007, September). A review of the ‘Stern review on the economics of climate change’. *Journal of Economic Literature*, 45(4), 686–702.
- Oishi, S., Kesebir, S., & Diener, E. (2011). Income inequality and happiness. , 22(9), 1095–1100.
- Okun, A. M. (1975). *Equality and efficiency: The big trade-off*. Washington, D.C.: Brookings Institution.
- Panel on Climate Change, I. (n.d.). *Fifth assessment report*. (Available online from <https://www.ipcc.ch/report/ar5/>)
- Payne, S. (2000). Poverty, social exclusion and mental health: findings from the 1999 pse survey. *Townsend Centre for International Poverty Research, Working Paper*(15).
- Phelps, E. S., & Pollak, R. A. (1968, April). On second-best national saving and game-equilibrium growth. *The Review of Economic Studies*, 35(2), 185–199.
- Pigou, A. C. (1932). *The economics of welfare: Volume 1*. London: Macmillan. (4th edition)

- Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica: Journal of the Econometric Society*, 122–136.
- Ramsey, F. P. (1928). A mathematical theory of saving. *Economic Journal*, 38, 543–559. (Reprinted in F. P. Ramsey, *Foundations: essays in philosophy, logic, mathematics, and economics*, ed. D. H. Mellor.)
- Rose, S. K. (2012). The role of the social cost of carbon in policy. *Wiley Interdisciplinary Reviews: Climate Change*, 3(2).
- Ross, J. (2006). Rejecting Ethical Deflationism. *Ethics*, 116(4), 742–768. doi: 10.1086/505234
- Rubinstein, A. (2003, November). Economics and psychology? the case of hyperbolic discounting. *International Economic Review*, 44(4), 1207–1216.
- Scheffler, S. (1982). *The rejection of consequentialism*. Oxford: Clarendon Press. (Revised edition 1994)
- Schmeidler, D. (1989). Subjective probability and expected utility without additivity. *Econometrica*, 57(3), 571–587.
- Selden, L. (1978). A new representation of preferences over "certain x uncertain" consumption pairs: The "ordinal certainty equivalent" hypothesis. *Econometrica: Journal of the Econometric Society*, 1045–1060.
- Sen, A. (1973a). *On economic inequality*. Oxford University Press.
- Sen, A. (1973b). *On economic inequality*. Oxford University Press.
- Sepielli, A. (2013a). Moral uncertainty and the principle of equity among moral theories1. *Philosophy and Phenomenological Research*, 86(3), 580–589.
- Sepielli, A. (2013b). What to do when you don't know what to do when you don't know what to do. *Nous*. (Forthcoming; published online 2013.)
- Sepielli, A. (2014). What to Do When You Don't Know What to Do. *Nos*, 48(3), 521–544. doi: 10.1111/nous.12010
- Sidgwick, H. (1890). *The methods of ethics*. London: Macmillan.
- Sinnott-Armstrong, W. (2014). Consequentialism. *The Stanford Encyclopaedia of Philosophy (Spring 2014 edition)*. (Edward N. Zalta (ed.). Article available at <http://plato.stanford.edu/archives/spr2014/entries/consequentialism/>)
- Slovic, P., & Tversky, A. (1974). Who accepts Savage's axiom? *Behavioral science*, 19(6), 368–373.
- Solow, R. (1974). The economics of resources or the resources of economics. *American economic review papers and proceedings*, 64, 1–14.
- Stern, N. (1977). Welfare weights and the elasticity of the marginal valuation of income. In M. Artis & R. Nobay (Eds.), *AUTE Edinburgh meeting of April 1976*. Basil Blackwell.
- Stern, N. (2007). *The economics of climate change: The Stern review*. Cambridge University Press.
- Stern, N. (2014, November). Ethics, equity and the economics of climate change paper 2: Economics and politics. *Economics and Philosophy*, 30.
- Sterner, T., & Persson, U. M. (2008). An even sterner review: Introducing relative prices into the discounting debate. *Review of Environmental Economics and Policy*, 2(1), 61–76.

- Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies*, 165–180.
- Svensson, L.-G. (1980, July). Equity among generations. *Econometrica*, 48(5), 1251–1256.
- Temkin, L. (1993). *Inequality*. Oxford University Press.
- Thaler, R. (1981). Some empirical evidence on dynamic inconsistency. *Economics Letters*, 8(3), 201 – 207.
- Traeger, C. P. (2012a, February). *What’s the rate? Disentangling the Weitzman and the Gollier effect* (Department of Agricultural & Resource Economics, UC Berkeley, Working Paper Series). Department of Agricultural & Resource Economics, UC Berkeley.
- Traeger, C. P. (2012b, February). *Why uncertainty matters — discounting under intertemporal risk aversion and ambiguity* (CESifo Working Paper Series). CESifo Group Munich.
- Treasury, H. M. (2003). *The green book — appraisal and evaluation in central government*. London.
- Wald, A. (1945). Statistical decision functions which minimize the maximum risk. *Annals of Mathematics*, 265–280.
- Wald, A. (1949). Statistical decision functions. *The Annals of Mathematical Statistics*, 165–205.
- Weatherston, B. (2014). Running risks morally. *Philosophical Studies*, 167(1), 141–163.
- Weitzman, M. L. (1998). Why the far-distant future should be discounted at its lowest possible rate. *Journal of Environmental Economics and Management*, 36, 201–208.
- Weitzman, M. L. (2001). Gamma discounting. *American Economic Review*, 91(1), 260–271.
- Weitzman, M. L. (2007). The Stern review on the economics of climate change. *Journal of Economic Literature*, 45, 703–724.
- Weymark, J. A. (1981). Generalized gini inequality indices. *Mathematical Social Sciences*, 1(4), 409–430.
- Zuber, S., & Asheim, G. B. (2012). Justifying social discounting: the rank-discounted utilitarian approach. *Journal of Economic Theory*, 147(4), 1572–1601.

Appendix 1. Derivation of the Ramsey equation

The discount *factor* for consumption at time t (relative to consumption at time 0) is, by definition,

$$R(t) := \frac{\frac{\delta V}{\delta C(t)}}{\frac{\delta V}{\delta C(0)}}, \quad (11)$$

where $C(t) = N(t)c(t)$ is aggregate consumption at t , and $\frac{\delta V}{\delta C(t)}$ (respectively $\frac{\delta V}{\delta C(0)}$) is the functional derivative of overall value V with respect to consumption, evaluated at time t (respectively, at time 0).

To see the connection between (11) and the verbal definition given in section 4 of the discount factor, suppose that we have a status quo consumption path c , and consider a change that involves a small reduction $-\Delta c_0$ in consumption at time 0 in return for a small increase Δc_t in consumption at time t (both perturbations lasting for a small time $\Delta\tau$). According to (11), this change will be an improvement to the status quo if and only if $R(t)\Delta c_t > \Delta c_0$: changes in consumption occurring at time t are ‘discounted’ by the factor $R(t)$ for the purposes of comparison with changes in consumption occurring at time 0.

The discount *rate* is therefore (from (5))

$$r = -\frac{\frac{d}{dt} \frac{\delta V}{\delta C(t)}}{\frac{\delta V}{\delta C(t)}}. \quad (12)$$

But, for arbitrary t , we have, from (4),

$$\frac{\delta V}{\delta C(t)} = \Delta(t)u'(c(t)); \quad (13)$$

thus, combining (12) and (13),

$$r = -\frac{\dot{\Delta}}{\Delta} - \frac{u''}{u'} \frac{dc}{dt}. \quad (14)$$

Defining $\delta := -\frac{\dot{\Delta}}{\Delta}$, $\eta := -c \frac{u''}{u'}$, $g := \frac{\dot{c}}{c}$, this can be rewritten as

$$r = \delta + \eta g, \quad (15)$$

as in equation (6).

Appendix 2. Discounting under a prioritarian or egalitarian approach

Intertemporal issues aside, the basic idea of prioritarianism is captured in the value-function form

$$V = \sum_i f(u_i), \quad (16)$$

where, as in the main text, the index i ranges over individuals, and u_i is the utility of the i th person; f is a monotone increasing but strictly concave function, encoding the prioritarian idea of the ‘diminishing marginal moral value of utility [or well-being]’. To frame a prioritarian discussion specifically of discounting, however, we would need to settle some further questions, for it is not immediately obvious how to incorporate the basic ideas expressed by (16) into an overall value function for the intertemporal case.

One reasonably natural intertemporal version of (16), with the possibility of pure time preference included, is

$$V = \int dt \cdot \Delta(t) \cdot \left(\sum_i f(u(c_i(t))) \right). \quad (17)$$

Under the simplifying assumption (made, following the standard literature, in the bulk of the main text) that there is no intratemporal inequality, this reduces to the simpler expression

$$V = \int dt \cdot \Delta(t) \cdot N(t) \cdot f \circ u(c(t)). \quad (18)$$

The prioritarian intertemporal value function (18) leads to an expression for the discount rate that is very similar to the standard Ramsey equation: we simply need to replace u with $f \circ u$ throughout.

The value function (17), however, applies the prioritarian principle of ‘diminishing marginal moral value of utility’ to *instantaneous* well-being. On the moral view represented by that value function, for example, it is of equal moral value to increase the consumption of *any* person at t from a given baseline c_{low} to a given higher amount c_{high} , even if the two persons being compared have very different consumption levels at other times in their lives (and hence very different *lifetime* well-being levels). Note that this can occur even in the absence of intratemporal inequality, since the two persons’ lifespans may overlap without coinciding. This way of applying the basic prioritarian idea is not implausible, but it is also not the only possible way in the intertemporal context. An alternative would be to apply the basic prioritarian principle of ‘diminishing marginal moral value’ to *lifetime* well-being, rather than to instantaneous well-being. On this alternative view, the relevant consideration (for the purpose of computing a prioritarian ‘moral weighting’) would not be how well off a proposed beneficiary is *at the time the benefit is to be delivered*, but rather how well off the proposed beneficiary is *over the course of her life as a whole*. A corresponding *undiscounted* intertemporal value function might be

$$V = \sum_i f \left(\int_{b_i}^{d_i} dt \cdot u(c_i(t)) \right), \quad (19)$$

where b_i, d_i are respectively the dates of birth and death of individual i . There are then two reasonably natural ways in which a notion of pure time preference might be used to generate a modified (‘discounted-prioritarian’) version of (19): we might apply pure discounting directly to utility at the given time t for the purpose of computing a modified index of each individual’s ‘lifetime well-being’, or (alternatively) we might discount the lifetime well-being of each individual (itself calculated in the standard way) according to that individual’s date of birth. The corresponding value functions would be, respectively,

$$V = \sum_i f \left(\int_{b_i}^{d_i} dt \cdot \Delta(t) u(c_i(t)) \right); \quad (20)$$

$$V = \sum_i \Delta(b_i) f \left(\int_{b_i}^{d_i} dt \cdot u(c_i(t)) \right). \quad (21)$$

These expressions lead to different modifications of the Ramsey equation, with the details in each case depending on whether or not, in performing our comparison between ‘consumption increment at time 0’ and ‘consumption increment at time t ’, we assume that we are comparing consumption increments accruing to the *same* person at different times.

The egalitarian case is more complex again. Again temporarily setting aside intertemporal issues, the basic ideas of egalitarianism motivate the value-function structure

$$V = \left(\sum_i u_i \right) + E, \quad (22)$$

where E is an index of betterness with respect to equality.⁶

As in the prioritarian case, there are further questions regarding how the egalitarian idea is to be implemented in the intertemporal case. The most straightforward possibility is to take (22) to represent value at a given time, and to take overall value to be a discounted time-integral of this instantaneous quantity:

$$V = \int dt \cdot \Delta(t) \cdot \left(\sum_i u(c_i(t)) + E(t) \right), \quad (23)$$

where $E(t)$ is the index of equality calculated with respect to instantaneous utilities at t . This value function, of course, differs from the standard discounted-utilitarian one only when we do *not* assume away intratemporal inequality. The usual assumption of the mainstream literature on discounting therefore also assumes away the differences between the discounted-egalitarian value function (23) and the standard discounted-utilitarian value function (4).

Again as in the prioritarian case, however, one might also apply egalitarian principles to lifetime well-being, rather than to instantaneous well-being. To simplify matters, let us deal for the egalitarian case only with a discrete-time model, and suppose that each individual lives for only one period. The corresponding discounted-egalitarian value function would be

$$V = \left(\sum_i \Delta(t_i) u(c(t_i)) \right) + E, \quad (24)$$

where t_i is the time-period during which individual i lives out her life, and E is calculated with respect to the lifetime well-being levels of *all* individuals (i.e.

⁶The expression (22) of course leaves significant freedom over the form of the ‘equality index’ E , and is thus very general. For discussion of the form E might take, see e.g. (Atkinson, 1970; Sen, 1973b; Weymark, 1981; Blackorby & Donaldson, 1980). By way of interpretation of E , note that (it follows from (22) that) the amount of ‘badness with respect to equality’ in a well-being distribution, compared to that in an equally good but perfectly equal distribution, is given by the amount of total well-being that could be given up without making the distribution worse overall, via moving to a distribution of perfect equality (with each person at the ‘equally distributed equivalent’ well-being level). In unpublished work, David McCarthy and Teru Thomas show that the general form (22) is implied by very plausible constraints on an egalitarian value function.

not only those alive at some given time). To derive a Ramsey-like equation for this case, we again need (as in appendix 112) to consider the ratio of the quantities $\frac{\partial V}{\partial c(t)}$, $\frac{\partial V}{\partial c(0)}$; relative to the simpler discounted-utilitarian case treated in appendix 112, this ratio will be complicated in the present case by the appearance of additional summands (on both numerator and denominator) that are proportional to $\frac{\partial E}{\partial c}$.

For a discussion of discounting in a framework that generalises discounted-utilitarianism in related directions, see e.g. (Fleurbaey & Zuber, 2015). The ‘rank-discounted utilitarian’ model investigated by Zuber and Asheim (2012) has some similarities to prioritarianism.

BIOGRAPHICAL INFORMATION

Hilary Greaves is a Professor of Philosophy at the University of Oxford. Her research focusses on various issues in ethics, including: foundational issues in consequentialism (‘global’ and ‘two-level’ forms of consequentialism), issues of aggregation (utilitarianism, prioritarianism, egalitarianism and anti-aggregationist approaches), population ethics, effective altruism, the interface between ethics and economics, the analogies between ethics and epistemology, and formal epistemology. She currently directs the project *Population Ethics: Theory and Practice*, based at the Future of Humanity Institute.