Social welfare functions in global climate-economy models: Methodological Inconsistencies and their Policy Implications

Hans-Martin Füssel

fuessel@stanford.edu

Center for Environmental Science and Policy (CESP),
Stanford University, Stanford, CA, U.S.A

and

Potsdam Institute for Climate Impact Research (PIK), Potsdam, Germany

Abstract

This paper reviews the application of social welfare functions (SWFs) in welfare-maximizing climate policy analysis. We identify several methodological inconsistencies, analyze their policy implications, discuss the theoretical questions raised by them, and provide recommendations for future studies. Our review finds that several SWFs applied in climate policy analysis are internally inconsistent. In particular, different methods for calculating the present values of alternative policy options lead to vastly different cost estimates. This topic has not been discussed in the literature despite the large attention that the discounting problem has generally received in the climate change context. The close link with the index number problem implies that there is no single 'correct' method for comparing the present values of alternative climate policies or other long-term policies involving significantly different economic trajectories. We also find that the uncritical combination of different SWFs can lead to erroneous results since they aggregate differently across time, population groups, states of the world, and components of economic output. We conclude that the translation of non-monetary welfare differences into monetary units is not generally possible. In particular, we show that no discounting scheme for certainty equivalents is consistent with expected discounted utility, which limits the use of certainty equivalents to account for risk aversion in intertemporal problems. Finally, we provide recommendations for avoiding the problems identified here in future climate policy analysis and show how the disregard of these recommendations in some recent studies has lead to artefactual results.

Keywords: climate change; DICE model; FUND model; growth discounting; index number problem; integrated assessment; objective function; social welfare function

Abbreviations: CRRA – constant relative risk aversion; DU – discounted utility; GHG – greenhouse gas; GMT – global mean temperature; GWP – gross world product; PDF – probability density function; PV – present value; PVC – present value of consumption; PVO – present value of output; SWF – social welfare function

Contents

1 Introduction

Simple climate-economy models based on the welfare maximization framework continue to be widely used for a variety of climate policy analyses. Two prominent examples for this type of models are DICE (Nordhaus, 1994) and FUND (Tol, 1999). The DICE model, a Ramsey-type optimal-growth model of the world economy in which a central planner maximizes intertemporal welfare subject to certain constraints, has frequently been adapted by researchers to investigate a wide variety of scientific and policy questions that are beyond the scope of the original model.

This paper reviews the application of social welfare functions (SWFs) in welfare-optimizing climate-economy models where they are used to identify welfare-maximizing climate policies, to quantitatively compare alternative sub-optimal policies, and to characterize the costs of enforcing or lifting certain constraints. In the context of this paper, we use the term 'social welfare function' to refer to any mathematical formulation that assigns a numerical utility to a stream of future economic output or consumption.

Ideally a SWF would be derived from the revealed preferences of the individuals concerned. However, Arrow's Impossibility Theorem (Arrow, 1951) shows that there is no unique method for aggregating individual preferences into social preferences. Even if such an aggregation was theoretically possible, it would not be practical in the context of anthropogenic climate change, which significantly affects future generations who cannot reveal their preferences today. For that reason, the SWFs applied for climate policy analysis are constructed synthetically with the aim of reflecting the implicit or explicit preference structure of current decision-makers. The global climate-economy models considered here assume that individual utility is determined by a single economic good, and that all individuals can be characterized by the same utility function. The SWFs applied by these models for assessing the costs and benefits associated with alternative climate policies are discounted utility of consumption (DU), present value of consumption (PVC), and present value of economic output (PVO), each of which has been defined for different time discounting schemes.

The maximization of intertemporal welfare in the context of anthropogenic climate change has been criticized, among others, for its perfect-market assumption, its assumption of full substitutability between market commodities and environmental goods and services, its neglect of the allocation of rights, its assumption that intergenerational compensation is actually feasible, its assumption of a convex optimization function, its inability to account reliably for deep uncertainty or catastrophic outcomes, and the weak empirical basis of widespread practices such as applying logarithmic utility, exponential time discounting, and assuming representative agents that maximize global intertemporal welfare (Lind *et al.* , 1982; Taylor, 1982; Lind, 1995; Lind & Schuler, 1998; Howarth, 2001; Spash, 2002; Azar & Lindgren, 2003; DeCanio, 2003; Yohe, 2003; Gowdy, 2005; Hall & Behl, 2005). The focus of this paper, in contrast, is on *logical* and *methodological* inconsistencies associated with the application of different SWFs, a topic that has received little attention in the literature so far.

This paper is structured as follows. Sect. 2 presents the SWFs that have been used in applications of the DICE and FUND models, two of the most widely used welfare-maximizing climate-economy models. Sect. 3 discusses inconsistencies of individual SWFs with a focus on the implications of different time discounting schemes, and Sect. 4 examines inconsistencies between several internally consistent SWFs. Based on the results of these theoretical sections, Sect. 5 makes recommendations for the application of SWFs in climate policy analysis, and reviews the application of SWFs in selected policy analyses with the DICE model. Sect. 6 concludes the paper.

2 Notation and definitions

The following notation is used throughout this paper:

Y	≥ 0	economic output
C	≥ 0	consumption
I	≥ 0	investment
L	> 0	population
ρ	≥ 0	pure rate of time preference (or 'utility discount rate')
θ	≥ 0	intertemporal elasticity of substitution
g		*actual* growth rate of per capita consumption
$\tilde{g}$		*assumed* growth rate of per capita consumption
r		social discount rate
s	≥ 0	investment (or savings) rate

These variables and parameters may be supplemented with a time index t, whereby $t = 0$ refers to the present year and $t = T$ to the final year of a time series. For instance, the change in per capita-consumption at time t is defined by

$$g_t = \frac{C_t/L_t}{C_{t-1}/L_{t-1}} - 1. \tag{1}$$

If a time index is missing, the respective variable is assumed to be constant over time. $X_{u\cdots v}$ denotes the stream of variable X from time u to time v (assuming $u \leq v$). For notational convenience, we allow the 'empty product' the value of which is assumed to be unity, *i.e.*, $\prod_{t=1}^{0} X_t = 1$.

The objective function most commonly applied in optimizing climate-economy models is utility from consumption, assuming a constant relative rate of risk aversion (CRRA; see DeCanio, 2003, Table 2.4). CRRA utility functions are also known as isoelastic or constant intertemporal elasticity of substitution utility functions. A CRRA utility function for a single agent and a single point in time is defined by the 'felicity function'

$$\mathrm{U}_\theta(C) \quad = \quad \begin{cases} \frac{C^{1-\theta}-1}{1-\theta}, & \text{if } \theta \neq 1 \\ \ln C, & \text{if } \theta = 1 \end{cases} \tag{2}$$

The values determined by the felicity function do not correspond to any 'real world' units; their unit is sometimes denoted as 'util(s)'. θ is a measure of the curvature of the utility function; it is variably denoted as intertemporal elasticity of substitution, rate of relative risk aversion, or absolute value of the elasticity of the marginal utility of per capita consumption. The lack of a lower bound in the logarithmic utility function (*i.e.*, for $\theta = 1$) can lead to counter-intuitive results when utility is aggregated over multiple agents or across several possible states of the world (see Sect. 3.2).

When multiple agents are concerned, the utilitarian approach defines social welfare as the (weighted) sum of individual utilities. DICE and other globally aggregated models that consider only one 'representative' agent equate total utility with the utility of per capita consumption multiplied by the population size (Nordhaus & Boyer, 2000, App. E):

$$\mathrm{U}_\theta(C, L) = L \cdot U_\theta(C). \tag{3}$$

Obviously, any comparison of social welfare across different population scenarios involves essential value judgements. In addition, we show in Sect. 3.3 that Eq. 3 produces arbitrary results if it is applied to different population scenarios.

The definition of SWFs below includes two parameters that reflect social value judgments about the distribution of wealth within and across generations: ρ and θ. There is a wide range of literature on the most appropriate values of these parameters in economic models of climate change and on the question of discounting in general (*e.g.*, Lind *et al.* ; Arrow *et al.* ; Nordhaus; Heal; Portney & Weyant; Toth; Howarth; Newell & Pizer, 1982; 1996; 1997; 1997; 1999; 2000; 2003; 2004). The standard value for θ in economic models of climate change is unity (Arrow *et al.* , 1996; DeCanio, 2003). The corresponding logarithmic (or Bernoullian) utility function is also applied in the DICE models. However, higher as well as lower values for θ have been suggested by some scholars (*e.g.*, Cline, 1992) and applied in sensitivity analyses of climate-economy models (Gjerde *et al.* , 1998; Azar & Lindgren, 2003). There is more disagreement on appropriate values for ρ, and on the question whether this parameter should be constant over time. The default value in DICE-94 is $\rho = 3\%/yr$ (Nordhaus, 1994, p. 104), DICE-99 model assumes that ρ declines over time from $\rho = 3\%/yr$ in 1995 to $\rho = 1.25\%/yr$ in 2335 (Nordhaus & Boyer, 2000, pp. 15–16), and the adaptation of DICE-99 by Yohe *et al.* (2004) assumes $\rho = 0\%/yr$. Note that discounted utility does not converge over time for $\rho = 0\%/yr$ even though the *finite horizon* approximations discussed here are defined. While declining rates of pure time preference cause time inconsistency in deterministic analyses, the consideration of uncertainties about future interest rates can lead to decreasing expected 'certainty-equivalent' social discount rates without causing time-inconsistent behaviour (Newell & Pizer, 2004; Weitzman, 1998; Newell & Pizer, 2003). The findings in this paper are largely independent of the choice of a constant or declining pure rate of time preference.

Eq. 4 to Eq. 9 define the six SWFs that have been used for comparing alternative policy strategies in connection with the DICE model. All of them take a finite output or consumption stream (expressed in currency, such as dollars) as input and calculate a scalar welfare value (expressed either in currency or in arbitrary 'utils') as output, which is defined as the discounted intertemporal sum of the welfare in each time step.[1]

$$\mathrm{DU}_{\mathrm{DICE}}(C_{0\cdots T}, L_{0\cdots T}; \rho_{1\cdots T}) = \sum_{t=0}^{T} \frac{L_t \cdot \ln(C_t/L_t)}{\prod_{t'=1}^{t} (1+\rho_{t'})} \tag{4}$$

$$\begin{aligned} \mathrm{PVC}_{\mathrm{DICE}}(C_{0\cdots T}, L_{0\cdots T}; \rho_{1\cdots T}) &= \sum_{t=0}^{T} \frac{C_t}{\prod_{t'=1}^{t} (1+\rho_{t'}) \cdot (1+g_{t'})} \\ &= \frac{C_0}{L_0} \cdot \sum_{t=0}^{T} \frac{L_t}{\prod_{t'=1}^{t} (1+\rho_{t'})} \end{aligned} \tag{5}$$

[1]We neglect the difference between the annual specification of the welfare functions defined here and the decadal time step of the DICE model.

$$\mathrm{PVC}_{\mathrm{end}}(C_{0\cdots T}, L_{0\cdots T}; \rho_{1\cdots T}, \theta) = \sum_{t=0}^{T} \frac{C_t}{\prod_{t'=1}^{t}(1+\rho_{t'}+\theta\cdot g_{t'})} \tag{6}$$

$$\mathrm{PVC}_{\mathrm{ex}}(C_{0\cdots T}; \rho_{1\cdots T}, \theta, \tilde{g}) = \sum_{t=0}^{T} \frac{C_t}{\prod_{t'=1}^{t}(1+\rho_{t'}+\theta\cdot \tilde{g})} \tag{7}$$

$$\mathrm{PVO}_{\mathrm{ex}}(Y_{0\cdots T}; \rho_{1\cdots T}, \theta, \tilde{g}) = \sum_{t=0}^{T} \frac{Y_t}{\prod_{t'=1}^{t}(1+\rho_{t'}+\theta\cdot \tilde{g})} \tag{8}$$

$$\mathrm{PVO}_{\mathrm{Yohe}}(Y_{0\cdots T}, L_{0\cdots T}) = \sum_{t=0}^{T} \frac{Y_t}{\prod_{t'=1}^{t}(1+\ln(1+g_{t'}))} \tag{9}$$

$\mathrm{DU}_{\mathrm{DICE}}$ describes the logarithmic utility of consumption based on 'classic' utility discounting at the rate of pure time preference. This utility function is used as objective function in the original DICE-99 model (Nordhaus & Boyer, 2000, p. 181).

The other SWFs express welfare in monetary units. They apply some variant of growth discounting, which focusses on the marginal social utility of consumption today compared with consumption in the future and represents the 'classical' approach to time discounting (Arrow *et al.* , 1996; Nordhaus, 1997; Heal, 1997; Tol, 1999; Toth, 2000). Growth discounting is based on findings about optimal savings in an idealized economy by Ramsey (1928). The conventional formula for social time preference, also known as the 'Ramsey growth discounting rule', is $r = \rho + \theta g$. However, this formula is only an approximate solution of the Ramsey model (see App. A).

$\mathrm{PVC}_{\mathrm{DICE}}$ describes the present value of consumption as calculated in the original DICE-99 model, which applies a variant of growth discounting (note the definition of g_t in Eq. 1).

$\mathrm{PVC}_{\mathrm{end}}$ describes the present value of consumption according to the conventional formulation of the Ramsey rule. In $\mathrm{PVC}_{\mathrm{end}}$, the discount rate is determined based on the *endogenously* determined growth rate of per capita consumption in each year. This SWF has been widely applied in global economic models of climate change (see, *e.g.*, Tol, 1999).

$\mathrm{PVC}_{\mathrm{ex}}$ also describes the present value of consumption according to the Ramsey rule. In contrast to $\mathrm{PVC}_{\mathrm{end}}$, the discount rate is determined based on an *exogenously* specified assumed growth rate of per capita consumption. This welfare measure has been used to determine the total welfare effects of climate policies in DICE-99 (Nordhaus & Boyer, 2000, p. 127).

$\mathrm{PVO}_{\mathrm{ex}}$ describes the present value of economic output according to the Ramsey rule, whereby the discount rate is determined based on an exogenously specified assumed growth rate. $\mathrm{PVO}_{\mathrm{ex}}$ is identical to $\mathrm{PVC}_{\mathrm{ex}}$, except that economic output is substituted for consumption. A special case of this welfare function (assuming $\theta = \rho = 0$) is applied in Fankhauser & Tol (2005), which apparently uses undiscounted gross world product (GWP) calculated by different versions of DICE-94.

This text suggests that monetary values were discounted using the discounting scheme from $\mathrm{PVC}_{\mathrm{DICE}}$ or $\mathrm{PVC}_{\mathrm{end}}$ (the two are identical for $\theta = 1$ and $\rho = 0$). If this had indeed been the case, PVC would have been identical across all consumption scenarios (see Sect. 3.1). However, the model code kindly provided by G. Yohe revealed that the SWF that has actually been used in determining discounted GWP and selecting the costs of alternative policies is $\mathrm{PVO}_{\mathrm{Yohe}}$.

PVO_{Yohe} describes the present value of economic output applying yet another variant of growth discounting.[2] This SWF has been applied in a modified version of DICE-99 (Yohe *et al.* , 2004, and Yohe, pers. comm.), which assumes $\theta = 1$ and $\rho = 0$.[3]

3 Inconsistencies of individual welfare metrics

In this section, we review the various welfare metrics defined in Sect. 2. Sect. 3.1 and Sect. 3.2 examine counterintuitive results and inconsistencies related to the application of growth discounting in ordinal and cardinal welfare metrics, respectively. Ordinal utilities are sufficient for identifying the optimal policy alternative in a deterministic context involving a single 'representative' agent. Cardinal utilities are required for the quantitative comparison of alternative policies and for the aggregation of utilities across different possible states of the world (in a probabilistic analysis) or across different individuals.[4]Sect. 3.3 briefly discusses inconsistencies caused by the application of utility-based welfare metrics across different population scenarios.

3.1 Growth discounting and ordinal welfare metrics

The level of the social discount rate is often critical in determining efficient policy alternatives, in particular for long-term problems such as climate change. Therefore, debates about discounting have always occupied an important place in environmental policy and welfare economics. In this section, we analyze how the discounting schemes applied in the monetary SWFs defined above rank alternative consumption paths.

For the sake of simplicity, this analysis assumes consumption paths with a constant growth rate of the form

$$\mathrm{C}(t; C_0, g) \quad = \quad C_0 \cdot (1+g)^t, \tag{10}$$

whereby C_0 denotes initial consumption at $t = 0$ and g the rate of per capita consumption growth.[5]In addition, we assume the following parameters to be constant: population ($L_t \equiv 1$), investment rate ($C_t/Y_t \equiv s \geq 0$), pure rate of time preference ($\rho_t \equiv \rho \geq 0$), and elasticity of the marginal utility of consumption ($\theta_t \equiv \theta > 0$). Assuming constant population allows us to equate the growth rates for total consumption and for per capita consumption. Assuming a constant investment rate allows to equate the growth rates for economic output and for consumption. Assuming constant time preference, per capita consumption growth, and elasticity of

[2]The second-order Taylor series approximation of the natural logarithm results in a discount rate of $r_t \approx g_t - \frac{g_t^2}{2}$, which does not correspond to any known theoretical model for discounting.

[3]Yohe *et al.* (2004, SOM pp. 1–2) states that *"In this Policy Forum, the pure rate of time preference is set equal to zero. With an elasticity of marginal utility equal to unity, the social discount rate is simply the endogenously determined rate of annual growth of per capita consumption."* This statement suggests that monetary values were discounted according to the discounting scheme of PVC_{DICE} or PVC_{end} (these two SWFs are identical for $\theta = 1$ and $\rho = 0$). However, the model code kindly provided by G. Yohe revealed that PVO_{Yohe} was actually used in calculating discounted GWP and determining the costs of alternative policies.

[4]We note that some economists reject the concept of cardinal utilities because such utilities cannot be solely derived from observations of actual behaviour (see, *e.g.*, Arrow, 1951).

[5]Note that the choice of constant-growth consumption scenarios is solely for the clarity of presentation. All inconsistencies identified for these simple scenarios apply equally to the more complex scenarios typically investigated in economic climate policy analysis.

the marginal utility of consumption leads to a constant social discount rate in the discounting schemes discussed here, which are all based on growth discounting. Furthermore, these simplifying assumptions allow the present analysis to be restricted to a single future time step. Given the time-separability of the SWFs considered here, results for individual points in time can be easily generalized to the whole time series.

Applying the welfare functions defined in Sect. 2 to the constant-growth consumption stream $\mathrm{C}(t; C_0, g)$ defined in Eq. 10, and equating Y_t with C_t, results in the following welfare contributions from consumption in an individual time step t:

$$\mathrm{DU}_{\mathrm{DICE}}(\mathrm{C}(t; C_0, g), t; \rho) = \frac{\ln\left(C_0 \cdot (1+g)^t\right)}{(1+\rho)^t} \tag{11}$$

$$\begin{aligned} \mathrm{PVC}_{\mathrm{DICE}}(\mathrm{C}(t; C_0, g), t; \rho) &= \frac{C_0 \cdot (1+g)^t}{((1+\rho) \cdot (1+g))^t} \\ &= \frac{C_0}{(1+\rho)^t} \end{aligned} \tag{12}$$

$$\mathrm{PVC}_{\mathrm{end}}(\mathrm{C}(t; C_0, g), t; \rho, \theta) = \frac{C_0 \cdot (1+g)^t}{(1+\rho+\theta \cdot g)^t} \tag{13}$$

$$\begin{aligned} \mathrm{PVC}_{\mathrm{ex}}(\mathrm{C}(t; C_0, g), t; \rho, \theta, \tilde{g}) &= \\ \mathrm{PVO}_{\mathrm{ex}}(\mathrm{C}(t; C_0, g), t; \rho, \theta, \tilde{g}) &= \frac{C_0 \cdot (1+g)^t}{(1+\rho+\theta \cdot \tilde{g})^t} \end{aligned} \tag{14}$$

$$\mathrm{PVO}_{\mathrm{Yohe}}(\mathrm{C}(t; C_0, g), t) = \frac{C_0 \cdot (1+g)^t}{(1+\ln(1+g))^t} \tag{15}$$

The analysis in this subsection focuses on the ranking of consumption paths.

To this end, we employ the following monotonicity criterion: *If a consumption scenario A has higher consumption levels at all time steps than scenario B, the SWF should assign higher welfare to scenario A than to B.* The motivation for this criterion is our firm conviction that the vast majority of climate policy-makers seeking advice from optimal-growth models would clearly prefer a policy scenario with consistently higher consumption over a scenario with lower consumption, everything else being equal. This assumption is also made implicitly in most climate policy analyses with optimal-growth models. For instance, Fankhauser & Tol (2005) apply DICE to compare indirect climate impacts under different assumptions, using future loss in *undiscounted* GDP as the main decision criterion. Their conclusions are, therefore, dependent on the assumption that a high-consumption pathway is always preferred. Therefore, we consider SWFs that violate the monotonicity criterion to be inconsistent with the preference structure of the target users and thus unsuitable for comparing alternative climate policies.

We find the following behaviour for the SWFs defined in Sect. 2:

$\mathrm{DU}_{\mathrm{DICE}}$, $\mathrm{PVC}_{\mathrm{ex}}$, $\mathrm{PVO}_{\mathrm{ex}}$ and $\mathrm{PVO}_{\mathrm{Yohe}}$ always prefer the high-growth scenario over the low-growth scenario. In the case of $\mathrm{PVC}_{\mathrm{ex}}$ and $\mathrm{PVO}_{\mathrm{ex}}$, this is true independent of the choice of $\tilde{g}$. Hence, these SWFs do fulfill the monotonicity criterion.

$\mathrm{PVC}_{\mathrm{DICE}}$ is insensitive to the consumption levels after the initial period. This property is obvious in Eq. 5 where PVC is independent of $C_{1 \cdots T}$ and in Eq. 12 where PVC is independent of g. Hence, this SWF does *not* fulfill the monotonicity criterion.

We note that $\mathrm{PVC_{DICE}}$ also violates the 'stationarity axiom' proposed by Koopmans (1960). The stationarity axiom demands that if two sequences have the same start, then eliminating that common start and bringing the rest forward does not change their ranking. For an illustration, consider the two-period sequences [1,1] and [1,2], and assume $\rho = 0\%/yr$. We find that $\mathrm{PVC_{DICE}}([1,1];\rho) = 2 = \mathrm{PVC_{DICE}}([1,2];\rho)$ whereas $\mathrm{PVC_{DICE}}([1];\rho) = 1 \neq 2 = \mathrm{PVC_{DICE}}([2];\rho)$. Hence, the two-period sequences [1,1] and [1,2] have identical PVC but elimination of the common start results the one-period 'sequences' [1] and [2], which have different PVC.

$\mathrm{PVC_{end}}$ requires the distinction of three cases (see the proof below):

- For $\theta < 1+\rho$, higher welfare is assigned to the high-growth scenario. Therefore, the monotonicity criterion is fulfilled.
- For $\theta = 1+\rho$, the discounted welfare is independent of the growth rate of consumption. Therefore, the monotonicity criterion is *not* fulfilled. (For $\theta = 1$ and $\rho = 0$, $\mathrm{PVC_{end}}$ is identical to $\mathrm{PVC_{DICE}}$.)
- For $\theta > 1+\rho$, higher welfare is assigned to the low-growth scenario. Therefore, the monotonicity criterion is *not* fulfilled.

For the proof, we assume C_0, Δg, and t to be positive.

$$\begin{aligned}
\mathrm{PVC_{end}}(\mathrm{C}(t;C_0,g),t;\rho,\theta) &< \mathrm{PVC_{end}}(\mathrm{C}(t;C_0,g+\Delta g),t;\rho,\theta) \\
&\Longleftrightarrow \\
\frac{C_0\cdot(1+g)^t}{(1+\rho+\theta\cdot g)^t} &< \frac{C_0\cdot(1+g+\Delta g)^t}{(1+\rho+\theta\cdot(g+\Delta g))^t} \\
&\Longleftrightarrow \\
\frac{(1+\rho+\theta\cdot(g+\Delta g))^t}{(1+\rho+\theta\cdot g)^t} &< \frac{(1+g+\Delta g)^t}{(1+g)^t} \\
&\Longleftrightarrow \\
\frac{\theta\cdot\Delta g}{1+\rho+\theta\cdot g} &< \frac{\Delta g}{1+g} \\
&\Longleftrightarrow \\
\theta\cdot(1+g) &< 1+\rho+\theta\cdot g \\
&\Longleftrightarrow \\
\theta &< 1+\rho
\end{aligned}$$

We find that $\mathrm{PVC_{end}}$ fulfills the monotonicity criterion for some combinations of θ and ρ but not for others. Note that the *dimensional* rate parameter ρ is an additive factor of the 'threshold value' for the *dimensionless* parameter θ. As a result, whether $\mathrm{PVC_{end}}$ fulfills the monotonicity criterion may depend on the (arbitrary) choice of the time step for its specification. (For instance, the numerical value of ρ is more than ten times larger when it is expressed per decade rather then per year.) The reason for this inconsistency is that growth discounting in $\mathrm{PVC_{end}}$ is based on an approximate solution rather than the exact solution of the Ramsey model (see App. A).

In summary, $\mathrm{PVC_{DICE}}$ and $\mathrm{PVC_{end}}$ violate the monotonicity criterion for many plausible combinations of the normative parameters, including $\theta = 1$ and $\rho = 0$. $\mathrm{PVC_{end}}$ is associated with further inconsistencies due to its lack of a solid theoretical foundation. Consequently, we consider these two SWFs unsuitable for evaluating and comparing alternative climate policies.

3.2 Growth discounting and cardinal welfare metrics

The monotonicity criterion discussed in the previous subsection only considers the ranking of alternative policies, *i.e.*, it regards the welfare metrics as *ordinal*. In this subsection, we analyze how the discounting schemes applied in the various monetary SWFs value the difference in present value (PV) between alternative policies quantitatively. The crucial question addressed here is as follows: *"Is the practice of using different discount factors in the present value calculations for alternative policies consistent, or not?"* This question reflects a yet unnoticed controversy in the integrated assessment modelling community.

One opinion is expressed by the main developer of the DICE models: *"The present values are computed using the base case discount factors."* (Nordhaus & Boyer, 2000, p. 127) and *"In making welfare comparisons between two different policies, the same relative prices should be used to discount the future consumption streams that result from both policies. Thus, in constructing the comparison measures Total abatement cost of policy [. . .], we use the base case relative prices to discount both base case consumption and consumption under current policy."* (Nordhaus, 2001, p. 19). Consequently, the original DICE-99 model determines the monetary welfare associated with different policy alternatives by calculating PVC with the same discount factors (as in $\mathrm{PVC_{ex}}$).

We are not aware of any *explicit* argument in favour of the opposite opinion (*i.e.*, to use different discount factors for determining the PVs of alternative policies). However, several policy analyses with optimizing climate-economy models have applied SWFs that determine the discount factors for each policy option endogenously: $\mathrm{PVC_{end}}$ is applied in FUND (Tol, 1999, 2003), and $\mathrm{PVO_{Yohe}}$ is applied in a variant of DICE-99 (Yohe *et al.* , 2004, and Yohe, pers. comm.). Interestingly, the first authors of these studies are fully aware of the potential problems associated with growth discounting when the discount rates are determined endogenously. They find that the discounting approach applied in $\mathrm{PVC_{end}}$ may lead to infinite expected damages from climate change (and thus infinite expected marginal benefits of mitigation) if there is the possibility for a catastrophic outcome (Tol, 2003), or if the analyst assumes an infinite time horizon (Yohe, 2003). As a consequence, Yohe (2003, p. 243) concludes that *"we have added one more element to our list of reasons why it is inappropriate to use the expected value of discounted net benefits to judge mitigation policy"*.

Let us briefly analyze the practical consequences of applying one or the other discounting approach before we consider the theoretical arguments. Table 1 shows the difference in present value between two finite consumption streams starting at $C_0 = 100$ and growing at $g = 0\%$ or $g = 3\%$ per year over a 10-year period for various discounting schemes. The *undiscounted* PV of these two consumption streams is 1000.0 and 1146.4, respectively (as in the right-most column). For the discounted SWFs, we assume $\theta = 1$ and $\rho = 0\%/yr$ since $\mathrm{PVO_{Yohe}}$ is only defined for these parameter choices. As noted above, $\mathrm{PVC_{DICE}}$ and $\mathrm{PVC_{end}}$ are identical for these parameter values, and they determine the same PV independent of the growth rate of consumption. $\mathrm{PVO_{Yohe}}$ shows a small difference in PV between the two consumption streams but this difference is about 70 times smaller than the difference in undiscounted consumption.

	$\mathrm{PVC}_{\mathrm{DICE}}$, $\mathrm{PVC}_{\mathrm{end}}$	$\mathrm{PVO}_{\mathrm{Yohe}}$	$\mathrm{PVC}_{\mathrm{ex}}$ $[\tilde{g} = 3\%/yr]$	$\mathrm{PVC}_{\mathrm{ex}}$ $[\tilde{g} = 0\%/yr]$
$g = 0\%/yr$	1000.0	1000.0	878.6	1000.0
$g = 3\%/yr$	1000.0	1001.9	1000.0	1146.4
Difference	0%	0.19%	13.8%	14.6%

Table 1: Difference in present value for various discounting schemes between two consumption streams growing at 0% and 3% per year over a 10-year period.

Hence, $\mathrm{PVC}_{\mathrm{ex}}$ is the only SWF considered here that adequately reflects the difference in (undiscounted) consumption between the two scenarios, largely independent of the exogenous choice of the growth rate, $\tilde{g}$.

We will now discuss the discounting question more formally. Present value calculations at the micro-level (*e.g.*, for individual projects) determine how much money has to be put aside today to replicate a given stream of future returns under given (deterministic or probabilistic) assumptions about future economic development.[6] Put another way, the present value states the maximum amount that a rational person would be willing to pay now for a contract that guaranteed to deliver a given utility stream in the future.

Let us do a simple thought experiment. Imagine two rational individuals (A and B) that have the same discounted utility (DU) function (as defined by their pure rate of time preference, the marginal utility of consumption, and the baseline wealth). Each of them stated the maximum price that they are willing to pay now for a contract that will provide a certain payment (C_A and C_B, respectively) at a given time in the future. Let us assume that A offered a lower price for his contract than B (*i.e.*, $PV_A(C_A) < PV_B(C_B)$). If you were offered the choice between the benefits from either of these contracts (without having to pay for the costs), which future payment would you choose? If you knew that A and B had negotiated their contracts based on the same expectations about future economic growth (and thus had used the same interest and discount rates in their PV calculations), it would be safe to choose the higher valued contract (in this case: B's contract), as this contract would provide a higher future payment.

Let us now assume that A and B had negotiated their contracts independently. If B was assuming lower growth (and thus interest and discount) rates than A, B might have offered a higher price than A (*i.e.*, $PV_A(C_A) < PV_B(C_B)$) even if A's contract makes a higher future payment than B's (*i.e.*, $C_A > C_B$). In this case, any rational individual X (including A and B) who knows the future payments of both contracts would assign a higher present value to A's benefits than to B's (*i.e.*, $PV_X(C_A) > PV_X(C_B)$). We thus have $PV_A(C_B) < PV_A(C_A) < PV_B(C_B) < PV_B(C_A)$. Hence, PV_A and PV_B (with fixed discount rates) prefer A's benefits over B's (as would be expected) whereas PV_B assigns a higher value to C_B than PV_A assigns to PV_A. Furthermore, since A and B apply the same DU function, $C_A > C_B$ implies $DU_A(C_A) > DU_B(C_B)$. These findings imply that the use of different discount factors in the PV calculations of alternative policy options may lead to policy rankings based on PVC that are inconsistent with those based on DU.

The next question is *how much* more to value the benefits of A's contract compared to those of B's contract? Since A and B discount future benefits differently, they are likely to assign different

[6] Obviously, it is much more difficult to intuitively interpret PVC or PVO at the global level since putting aside an amount of money equivalent to the future stream of gross world product is not possible.

present values to the difference between A's and B's benefits (*i.e.*, $PV_A(B_A) - PV_A(C_B) \neq PV_B(C_A) - PV_B(C_B)$). The lack an obvious method for determining the monetary difference between two future consumption values is not surprising, given the close relationship between the problem of finding the 'correct' metric for comparing the monetary welfare between alternative policies and the index number problem. The classical index number problem is concerned with the question how to determine a single measure of the deviation in the overall quantity level between different periods (or regions) in a situation in which there are $N \geq 2$ goods and $I \geq 2$ time periods (or regions). The problem discussed here, in contrast, is how to determine a single measure of the deviation in the overall quantity level between different consumption scenarios in a situation in which there are $N \geq 2$ time periods and $I \geq 2$ scenarios. In this context, the discount factors applied to future consumption in the discounting problem correspond to the relative prices of different goods in the index number problem. The index number problem has concerned economists and statisticians since the 19$^{\text{th}}$ century at least (Jevons, 1865), and it has long been known that the index number problem has no unique solution (see, *e.g.*, Edgeworth, 1888). However, there is unanimous agreement that a single set of prices has to be used for a meaningful comparison of quantities between different periods.

The next question is how to determine the 'exogenous' discount factors in the absence of a unique 'correct' method? The recommended method is to select a 'baseline' policy scenario and to calculate the discount factors based on the growth rates of this baseline scenario as in PVC_{ex}. This approach, which is largely equivalent to the choice of a base period or a numeraire country for the multilateral Laspeyres quantity index, has been followed in the original DICE-99 model (Nordhaus, 2001, p. 19).[7] Comparison of the two right-most columns of Table 1 suggests that the choice of the baseline scenario is relatively unimportant for reasonably small welfare differences between alternative scenarios. (For that reason, we neglect in the present discussion that PVC_{ex} applies discount factors that are based on the approximate solution of the Ramsey model rather than the exact solution; see App. A). The choice of baseline scenario may become more important if alternative scenarios involve very large welfare differences with opposite sign at different points in time. In such a situation, we recommend to determine PVs based on selected 'extreme' policy scenarios, and to consider the range of results. We suppose that the uncertainty related to the choice of discount factors is still not very important given the other uncertainties that are necessarily involved in an analysis that comprises such 'extreme' scenarios.

We conclude that all available evidence indicates that PV calculations need to apply the same discount factors for all policy options under consideration, if the resulting PVs shall be used to compare alternative policy options quantitatively. Consequently, we consider PVC_{DICE}, PVC_{end}, and PVO_{Yohe} as unsuitable for comparing the welfare implications of alternative climate policies. Note that PVO_{Yohe} is subject to two other problems as well. First, there is no theoretical basis for the logarithmic relationship between consumption growth rate and discount rate assumed in Eq. 9. Second, PVO_{Yohe} is only defined for $\rho = 0$, which is widely regarded as unrealistically low (Arrow *et al.* , 1996). Consequently, the only SWFs considered here that are consistent with the monotonicity criterion *and* with the other fundamental requirements for SWFs are DU_{DICE}, PVC_{ex}, and PVO_{ex}.

[7]We note that DICE-99 also computes PVC_{DICE}, which is insensitive to the consumption levels after the initial period. However, the intention for computing PVC_{DICE} in DICE is *not* to distinguish between different policy alternatives but only to provide a monetary equivalent to which discounted utility in the no-policy case can be calibrated. (The actual objective function in DICE is $\text{DU}_{\text{DICE,cal}}(\cdot) = \frac{\text{DU}_{\text{DICE}}(\cdot)}{coefopt_1} + coefopt_2$, whereby $coefopt_1$ and $coefopt_2$ are chosen so that $\text{DU}_{\text{DICE,cal}}(\cdot) = \text{PVC}_{\text{DICE}}(\cdot)$ for an optimal strategy in the absence of climate damages.) The original DICE model does not report PVC_{DICE} for scenarios other than the baseline scenario.

3.3 Utility calculation for different population scenarios

We stated in Sect. 2 that utility cannot be objectively compared across different population scenarios since such a comparison involves essential value judgements. In this subsection, we complement these 'philosophical' arguments by showing that the general practice of defining the total utility of several identical agents as the product of population size and per-capita utility provides arbitrary results if it is applied across different population scenarios.

Imagine two scenarios in which a given amount of *total* consumption, $C \cdot L$, is consumed by either L or $k \cdot L$ identical individuals (assuming $k > 1$). Comparison of the total logarithmic utility of these two alternative scenarios yields the following result:

$$\begin{aligned} \mathrm{DU_{DICE}}(C, L) &< \mathrm{DU_{DICE}}(\frac{C}{k}, k \cdot L) \\ &\Longleftrightarrow \\ L \cdot \ln C &< k \cdot L \cdot \ln \frac{C}{k} \\ &\Longleftrightarrow \\ k \cdot \ln k &< (k-1) \cdot \ln C \\ &\Longleftrightarrow \\ k^{\frac{k}{k-1}} &< C \end{aligned} \tag{16}$$

Eq. 16 compares a dimensionless number on the left-hand side with a dimensional value on the right-hand side. Hence, the ranking of the two options is not invariant to the unit of measurement of C.

We are not aware of any welfare-optimizing climate policy analysis that *actually* determines population growth rates endogenously. However, this approach was suggested by Fankhauser & Tol (2005), who motivated their analysis by listing three types of indirect climate impacts on human welfare, including *"health and mortality impacts associated with more widespread diseases [that would] affect population growth"* (p. 4). In contrast to this suggestion, their actual analysis with a modified version of DICE-94 apparently keeps population growth exogenous and implements the health effects of climate change instead by changing the accumulation of human capital (p. 15).

4 Inconsistencies between different welfare metrics

In this section, we investigate differences between the *internally consistent* welfare metrics $\mathrm{DU_{DICE}}$, $\mathrm{PVC_{ex}}$, and $\mathrm{PVO_{ex}}$. Sect. 4.1 investigates inconsistencies between consumption-based and output-based welfare metrics; Sect. 4.2 and 4.3 investigate differences between monetary and non-monetary welfare metrics related to their aggregation of welfare across time and across possible states of the world, respectively.

4.1 Consumption-based *vs.* output-based welfare metrics

The optimal growth models considered here divide net economic output into consumption and investment in productive capital: $Y = C + I$. Gross output also includes the costs of emissions abatement and climate change damages (*i.e.*, net output is defined as gross output minus emissions abatement costs and climate change damages). The fraction of net output devoted to investment is denoted as the investment rate: $s = \frac{I}{Y} = 1 - \frac{C}{Y}$. General equilibrium models such as DICE determine the investment rate endogenously. In the present discussion, we neglect the problem of converting investment into consumption equivalents (see, *e.g.*, Lind & Schuler, 1998).

We analyze the differences between output-based or a consumption-based SWFs in two parts, depending on which of them is used as objective function. When $\mathrm{PVO_{ex}}$ instead of $\mathrm{DU_{DICE}}$ is used as objective function in DICE-99, the optimal (*i.e.*, PVO-maximizing) policy is characterized by an investment rate of $s = 100\%$ over the full time horizon. All economic output is used for investment, none remains for consumption ($\mathrm{PVC_{ex}} = 0$), and logarithmic utility becomes minus infinity ($\mathrm{DU_{DICE}} = -\infty$). Thus the very policy that maximizes $\mathrm{PVO_{ex}}$ minimizes $\mathrm{PVC_{ex}}$ and $\mathrm{DU_{DICE}}$. Since the decision strategy that maximizes $\mathrm{PVO_{ex}}$ is obviously unrealistic, we further conclude that $\mathrm{PVO_{ex}}$ generally is not an appropriate SWF for comparing alternative climate policy options.

This simple example shows that consumption-based and output-based welfare metrics may rank alternative policies completely differently. From a practical point of view, the more relevant question is whether inconsistent rankings also occur for climate policies determined by global climate-economy models in their 'normal' utility-maximizing mode (see Sect. 5.2 for an example of a recent climate policy analysis where this question becomes relevant). Fig. 1 shows presumably utility-maximizing decision strategies determined with a probabilistic version of DICE-99 (implemented in Analytica) for two different probabilistic climate constraints. The more and less stringent constraint limits the probability that global mean temperature (GMT) exceeds 2.5 °C and 3.0 °C above preindustrial levels, respectively, to 1% for a given probability density function of climate sensitivity. The top panel shows that the strategy for the less stringent 3.0 °C constraint has higher consumption levels in the first 120 years but lower output levels in all but one period than the strategy for the 2.5 °C constraint. Consequently, the 3.0 °C strategy has a higher $\mathrm{DU_{DICE}}$ and $\mathrm{PVC_{ex}}$ but a lower $\mathrm{PVO_{ex}}$ than the 2.5 °C strategy. The explanation for the inconsistent ranking of the two policy strategies is shown in the bottom panel. The 3.0 °C strategy is characterized by lower abatement rates (as expected) but also by significantly lower investment rates than the 2.5 °C strategy.

A more detailed analysis reveals that the 3.0 °C strategy shown in Fig. 1 does not actually maximize $\mathrm{DU_{DICE}}$ for that constraint. Even though the initial values of the decision variables were identical for both constraints, the gradient-based solver of Analytica found the *global* optimum for the 2.5 °C constraint but only a *local* optimum for the 3.0 °C constraint. After variation of the initial values, the solver was able to find a decision strategy that satisfies the 3.0 °C constraint with higher utility from consumption *and* higher output than the 2.5 °C strategy. This finding highlights the importance of analyzing the 'optimal' results identified by model solvers carefully in order to avoid that local optima are misinterpreted as global optima. However, it does not affect the main lesson to be learnt from this example: that two policy strategies determined by DICE-99 in utility-maximizing mode may be ranked inconsistently by $\mathrm{DU_{DICE}}$ and $\mathrm{PVC_{ex}}$ compared to $\mathrm{PVO_{ex}}$.

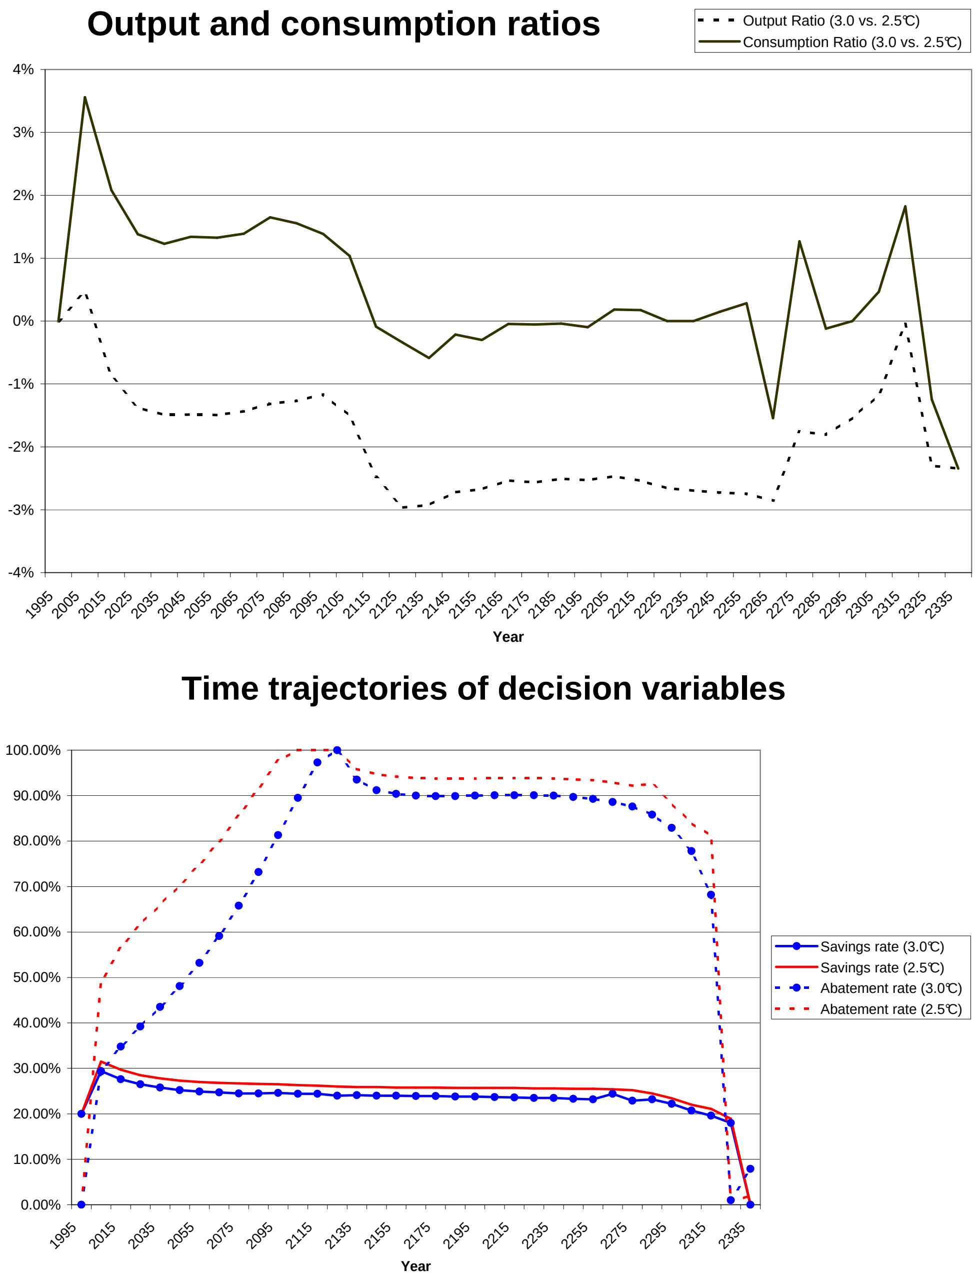

Figure 1: Decision strategies that maximize DU_{DICE} for two probabilistic climate constraints. *Top:* Relative difference of economic output and consumption (in %). *Bottom:* Investment (or savings) rate and abatement rate.

4.2 Discounting dollars *vs.* discounting utils

In this subsection, we investigate the difference between discounting monetary values (*e.g.*, consumption expressed in dollars or other currency) and discounting utility expressed in 'utils'. This distinction is often blurred, as in the following quote from DeCanio (2003, p. 168): *"According to the discounted utility formulation, an individual's subjective time rate of discount δ is determined by the relationship $x = (1+\delta)^t$, where the individual would be indifferent between \$1 in consumption today and \$x in consumption at time t in the future."* In this sentence, the text refers to discounted *utility* whereas the mathematical definition refers to discounting *consumption.*

We first analyze the relationship between the two discount rates for an agent that lives for two periods only. Since logarithmic utility is undefined if consumption in *any* time period is zero, we actually need to compare consuming *additional* \$1 now with consuming *additional* $\$x = 1+\delta$ in the next period.

The logarithmic utility of consuming C_0 at $t=0$ and C_1 at $t=1$, discounted at rate r, is

$$\mathrm{DU}(C_0, C_1; r) \quad = \quad \ln C_0 + \frac{\ln C_1}{1+r}. \tag{17}$$

Assuming a baseline consumption of C in both periods and denoting additional consumption at $t=0$ as aC (*e.g.*, $a = \frac{\$1}{C}$ for an additional consumption of \$1), we find the following relationships between the equivalent discount rates for utility, r, and for consumption, δ:

$$\begin{aligned}
\mathrm{DU}(C+aC, C; r) \quad &= \quad \mathrm{DU}(C, C+aC\cdot(1+\delta); r) \\
&\Longleftrightarrow \\
\ln(C+aC) + \frac{\ln C}{1+r} \quad &= \quad \ln C + \frac{\ln(C+aC\cdot(1+\delta))}{1+r} \\
&\Longleftrightarrow \\
(1+r)\cdot\ln(1+a) \quad &= \quad \ln(1+a\cdot(1+\delta)) \\
&\Longleftrightarrow \\
r \quad &= \quad \frac{\ln(1+a\cdot(1+\delta))}{\ln(1+a)} - 1
\end{aligned} \tag{18}$$

Solving Eq. 18 for δ yields

$$\delta \quad = \quad \frac{(1+a)^{1+r} - 1}{a} - 1. \tag{19}$$

Fig. 2 shows that the two discount rates are very similar for marginal changes in baseline consumption (*i.e.*, $a \ll 1$). For non-marginal consumption differences, however, discounting utility is inconsistent with the use of a single discount rate for consumption (and vice versa).

$\mathrm{DU}_{\mathrm{DICE}}$ and $\mathrm{PVC}_{\mathrm{ex}}$ produce identical rankings for the constant-growth consumption paths considered in Sect. 3.1. The next example shows, however, that the difference between discounting consumption and logarithmic utility may lead to inconsistent rankings between those two welfare metrics if more complex consumption paths are involved. In this example, we are considering consumption paths that involve a one-time reduction in consumption relative to a reference path with a constant growth rate

$$\mathrm{C}(t; g, a, t_a) \quad = \quad \begin{cases} (1+g)^t, & \text{if } t \neq t_a \\ (1+g)^t\cdot(1-a), & \text{if } t = t_a \end{cases} \tag{20}$$

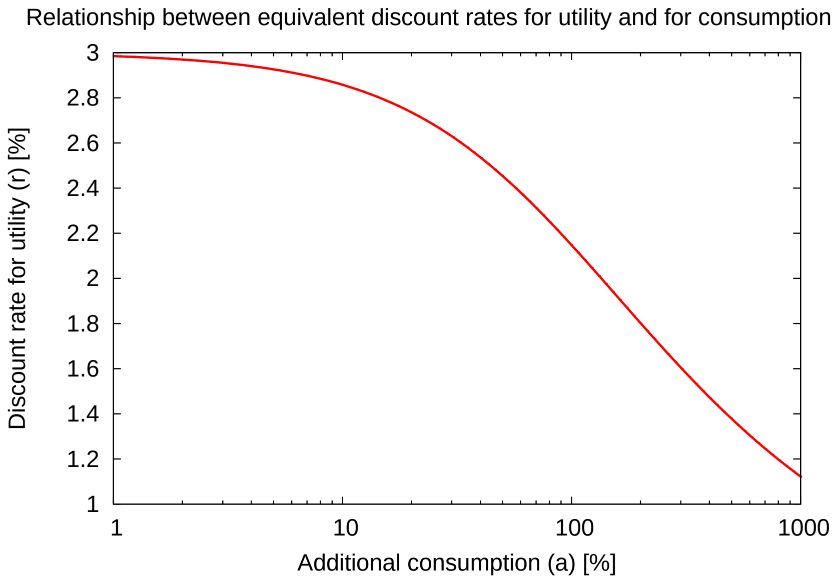

Figure 2: Discount rate for logarithmic utility that is equivalent to a discount rate for consumption of $\delta = 3\%$.

whereby t is time, g is the 'default' growth rate in consumption, and a is the relative loss in consumption compared to the default path at time t_a. The two paths of specific interest here are identical except for $t = 0$ and $t = T$. $\mathrm{C}(\cdot; g, a_0, 0)$ is characterized by a relative reduction in consumption of a_0 at $t = 0$ whereas $\mathrm{C}(\cdot; g, a_T, T)$ involves a reduction of a_T at $t = T$. We assume logarithmic discounting ($\theta = 1$), constant population ($L_t \equiv 1$), and constant time preference ($\rho_t \equiv \rho$). These assumptions allow the following simplified description of the two SWFs considered here:

$$\mathrm{DU}_{\mathrm{DICE}}(C_{0\cdots T}; \rho) = \sum_{t=0}^{T} \frac{\ln C_t}{(1+\rho)^t} \tag{21}$$

$$\mathrm{PVC}_{\mathrm{ex}}(C_{0\cdots T}; \rho, \tilde{g}) = \sum_{t=0}^{T} \frac{C_t}{(1+\rho+\tilde{g})^t} \tag{22}$$

We are now interested which reduction in consumption at $t = 0$ is 'equivalent' (in terms of these two SWFs) to a given reduction at $t = T$? For DU equivalence, we determine a_0 (depending on a_T, T, and ρ) by:

$$\begin{aligned}
\mathrm{DU}_{\mathrm{DICE}}(\mathrm{C}(\cdot; g, a_0, 0); \rho) &= \mathrm{DU}_{\mathrm{DICE}}(\mathrm{C}(\cdot; g, a_T, T); \rho) \\
&\Longleftrightarrow \\
\sum_{t=0}^{T} \frac{\ln \mathrm{C}(t; g, a_0, 0)}{(1+\rho)^t} &= \sum_{t=0}^{T} \frac{\ln \mathrm{C}(t; g, a_T, T)}{(1+\rho)^t} \\
&\Longleftrightarrow \\
\ln(1-a_0) + \frac{\ln\left((1+g)^T\right)}{(1+\rho)^T} &= \frac{\ln\left((1+g)^T \cdot (1-a_T)\right)}{(1+\rho)^T} \\
&\Longleftrightarrow \\
\ln(1-a_0) &= \frac{\ln(1-a_T)}{(1+\rho)^T}
\end{aligned}$$

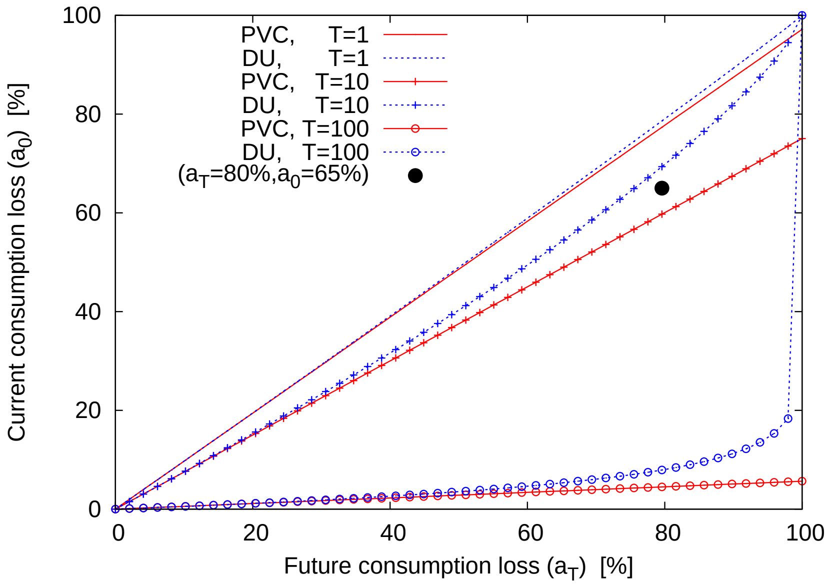

Figure 3: Current and future consumption losses (relative to a constant-growth reference path) that are equivalent according to $\mathrm{DU_{DICE}}$ and $\mathrm{PVC_{ex}}$.

$$\Longleftrightarrow$$

$$a_0 = 1-(1-a_T)^{\left((1+\rho)^{-T}\right)}$$

For PVC equivalence, we determine a_0 (depending on a_T, T, g, ρ, and $\tilde{g}$) as follows:

$$\begin{aligned}
\mathrm{PVC_{ex}}(\mathrm{C}(\cdot;g,a_0,0);\rho,\tilde{g}) &= \mathrm{PVC_{ex}}(\mathrm{C}(\cdot;g,a_T,T);\rho,\tilde{g}) \\
&\Longleftrightarrow \\
\sum_{t=0}^{T}\frac{\mathrm{C}(t;g,a_0,0)}{(1+\rho+\tilde{g})^t} &= \sum_{t=0}^{T}\frac{\mathrm{C}(t;g,a_T,T)}{(1+\rho+\tilde{g})^t} \\
&\Longleftrightarrow \\
(1-a_0)+\frac{(1+g)^T}{(1+\rho+\tilde{g})^T} &= 1+\frac{(1+g)^T}{(1+\rho+\tilde{g})^T}\cdot(1-a_T) \\
&\Longleftrightarrow \\
a_0 &= a_T\cdot\left(\frac{1+g}{1+\rho+\tilde{g}}\right)^T
\end{aligned}$$

The relationship between a reduction in future consumption, a_T, and the equivalent reduction in current consumption, a_0, for $\mathrm{DU_{DICE}}$ and $\mathrm{PVC_{ex}}$ is thus characterized as follows:

$$\mathrm{a}_0^{\mathrm{DU}}(a_T,T,\rho) = 1-(1-a_T)^{\left((1+\rho)^{-T}\right)} \tag{23}$$

$$\mathrm{a}_0^{\mathrm{PVC}}(a_T,T,g,\rho,\tilde{g}) = a_T\cdot\left(\frac{1+g}{1+\rho+\tilde{g}}\right)^T \tag{24}$$

Fig. 3 depicts $\mathrm{a}_0^{\mathrm{DU}}(\cdot)$ and $\mathrm{a}_0^{\mathrm{PVC}}(\cdot)$ across the whole range of a_T for three different points in time: $T=1$, $T=10$, and $T=100$ (assuming $\rho=\tilde{g}=g=3\%/yr$). We observe that $\mathrm{a}_0^{\mathrm{DU}}(\cdot)$ and $\mathrm{a}_0^{\mathrm{PVC}}(\cdot)$ are very similar for small values of a_T and/or T. However, $\mathrm{a}_0^{\mathrm{PVC}}$ is linear in a_T whereas $\mathrm{a}_0^{\mathrm{DU}}$ is non-linear, with $\lim_{a_T\to 1}\mathrm{a}_0^{\mathrm{DU}}(\cdot)=1$. Fig. 3 also shows for which combinations of a_0 and

State of the world	SOW 1		SOW 2		Aggregated measures		
Welfare measure	C	U(C)	C	U(C)	E(C)	E(U(C))	C^*
Policy A	**1.5**	**0.405**	0.5	-0.693	**1.0**	-0.144	0.866
Policy B	1.0	0.000	**0.8**	**-0.223**	0.9	**-0.112**	**0.894**

Table 2: Consumption and logarithmic utility for two policies and two equally likely states of the world (see text). The preferred policy is indicated by bold face.

a_T each of the two welfare metrics prefers the 'late-loss path' $\mathrm{C}(\cdot; g, a_T, T)$ over the 'early-loss path' $\mathrm{C}(\cdot; g, a_0, 0)$, depending on whether an (a_T, a_0) pair is above or below the corresponding curve. For instance, the black point $(a_T = 80\%, a_0 = 65\%)$ lying between the lines marked 'DU, T=10' and 'PVC, T=10' indicates that a 80% reduction of consumption at $t = 10$ is preferred over a 65% reduction at $t = 0$ according to $\mathrm{PVC}_{\mathrm{ex}}$ but not for $\mathrm{DU}_{\mathrm{DICE}}$. We conclude that $\mathrm{DU}_{\mathrm{DICE}}$ and $\mathrm{PVC}_{\mathrm{ex}}$ may produce inconsistent rankings of consumption paths if these paths involve significant welfare deviations at different points in time. We suppose, however, that this inconsistency has minor implications for the integrated assessment of climate change, compared to the other inconsistencies discussed here.

4.3 Expected consumption *vs.* expected utility

In this subsection, we investigate the difference between expected consumption and expected utility in probabilistic analyses. Table 2 provides a simple example showing that expected consumption and expected logarithmic utility may produce inconsistent rankings of policies since they aggregate differently across possible states of the world. The four left columns show consumption (C) and logarithmic utility ($\mathrm{U}(C) = \ln C$) for two equally likely states of the world (SOW 1 and 2) and for two different policies (A and B). Policy A is associated with higher consumption (and utility) under SOW 1 whereas policy B leads to higher consumption (and utility) under SOW 2. The three right columns show three aggregated welfare measures (E denotes the expected value). We find that policy A has higher expected consumption but policy B has higher expected utility. The right-most column shows the certainty equivalent C^* of the two policies, which was calculated such that $\mathrm{U}(C^*) = \mathrm{E}(\mathrm{U}(C))$. A certainty equivalent is the certain result that would make an individual indifferent between it and the uncertain outcome.

The intuitive reason for the inconsistent policy rankings produced by expected logarithmic utility and expected consumption is that the former accounts for risk aversion whereas the latter does not. In this example, Policy A is more risky than policy B in the sense that the consumption difference between the two states of the world is much larger and, as a result, the certainty equivalent is considerably lower than the expected consumption. The importance of this inconsistency for a particular analysis requires careful examination by the analyst.

According to Arrow *et al.* (1996, p. 130), *"Most economists believe that considerations of risk can be treated by converting outcomes into certainty equivalents, [...] and discounting these certainty equivalents"*. However, Lind & Schuler (1998, p. 66) note that *"we have never seen [a cost-benefit analysis] that has systematically converted costs or benefits to certainty equivalents"*, arguing further that *"In the case of most public projects of policies it is virtually impossible to do so"*. Furthermore, sequential decision-making produces results that are superior to those of 'classical' cost-benefit analysis using certainty equivalents when there is the possibility of obtaining new

	Policy	A		B		C	
	Period	0	1	0	1	0	1
SOW 1	C_t	2	1	1	4	4	1/4
	$\mathrm{U}(C_t)$	$\ln 2$	0	0	$2 \cdot \ln 2$	$2 \cdot \ln 2$	$-2 \cdot \ln 2$
SOW 2	C_t	2	1	1	1	4	1
	$\mathrm{U}(C_t)$	$\ln 2$	0	0	0	$2 \cdot \ln 2$	0
Expected	$\mathrm{E}(\mathrm{U}(C_t))$	$\ln 2$	0	0	$\ln 2$	$2 \cdot \ln 2$	$-\ln 2$
utility	$\mathrm{E}(\mathrm{U}(C_0, C_1))$	$\ln 2$		$\ln 2$		$\ln 2$	
Certainty	C_t^*	2	1	1	2	4	1/2
equivalent	$\mathrm{PV}(C_0^*, C_1^*)$	$2 + R$		$1 + 2 \cdot R$		$4 + R/2$	

Table 3: Consumption and logarithmic utility for three policies and two equally likely states of the world. $C_t^* = \mathrm{U}^{-1}(\mathrm{E}(\mathrm{U}(C_t)))$ denotes the certainty equivalent of consumption, and R denotes the discount factor for certainty equivalents in the second period

information over time (Dixit & Pindyck, 1994). Finally, while it is straightforward to compute certainty equivalents for individual time steps (as in Table 2), there exists no discounting scheme for certainty equivalents such that the ranking of alternative policies according to the present value of certainty equivalents is consistent with the ranking according to expected discounted utility.

We prove the last statement by demonstrating that such a discounting scheme does not even exist for a simple example involving logarithmic utility, no pure time preference ($\rho = 0$), two time steps ($t = 0$ and $t = 1$), two equally likely states of the world (SOW 1 and 2), and three policy options (Policy A, B, and C). Table 3 shows the consumption and utility values in this example for each policy, time step, and SOW as well as the resulting expected (discounted) utility and (present value of) certainty equivalents. The third row from the bottom shows that expected discounted utility is equal to $\ln 2$ in all three policies. For the two welfare metrics (present value of certainty equivalents and expected discounted utility) to be equivalent, the discount factor for certainty equivalents in the second period, R, needs to be chosen such that the present value of certainty equivalents is identical for all three policies, *i.e.*, $2 + R = 1 + 2 \cdot R = 4 + R/2$. The first equivalence requires $R = 1$ whereas the second equivalence requires $R = 2$. Since these two requirements are incompatible with each other, we conclude that no discounting scheme for certainty equivalents values all three policies equally.

This finding implies that the above citation from Arrow *et al.* (1996) appears too optimistic. At least when "considerations of risk" are reflected by the very widely applied logarithmic utility function (implying a rate of relative risk aversion of unity), "converting outcomes into certainty equivalents, and discounting these certainty equivalents" is not able to reflect the policy ranking defined by the discounted utility function.

5 Implications for climate policy analysis

In this section, we apply the findings from Sect. 3 and 4 to welfare-maximizing climate policy analyses. Sect. 5.1 makes several recommendations for the application of SWFs, and Sect. 5.2 reviews the use of SWFs in two recent climate policy analyses. We note again that the uncritical application of welfare economics in the climate change context has been criticized for various reasons not addressed in this paper (see Sect. 1 for selected references). Our aim in this section is to help preventing the introduction of *additional* inconsistencies that may be caused by choosing flawed SWFs or by combining SWFs inappropriately.

5.1 Recommendations

The inconsistencies identified above can be avoided if the application of social welfare functions in welfare-optimizing climate policy analysis follows regards the following recommendations:

1. The choice of SWFs should be based on the (supposed) preferences of target decision-makers. This choice should be made explicit.
2. If the results of monetary SWFs are to be compared across different policy options, the same discounting factors need to be applied to all policy options,
3. The same SWF should preferably be used in all optimizations as well as to report the relative 'desirability' of alternative policies. If different SWFs are combined in an analysis, the analyst needs to demonstrate that the inconsistencies between them do not affect the conclusions of the analysis.
4. In a probabilistic analysis, the monetary value of a policy option with non-marginal welfare differences between different states of the world (caused, for instance, by potential abrupt large-scale changes to the climate system) should be determined as its certainty equivalent.
5. Social welfare cannot be objectively compared across different population scenarios.

What do these recommendations mean for the choice of SWFs for climate policy analysis involving global welfare maximization? PVC_{ex} is an appropriate SWF if (and only if!) an analyst assumes risk-neutral decision-makers. However, economists find very little empirical support for risk-neutral behaviour in individuals (Arrow *et al.* , 1996). While the Arrow-Lind theorem (Arrow & Lind, 1970) holds that if risk can be pooled or spread in such a way that aggregate risk is negligible, governments can be considered risk-neutral these conditions are generally not met in climate policy analysis. DU_{DICE} is an appropriate SWF if an analyst believes that the degree of risk aversion and other preferences of target decision-makers are adequately reflected by the logarithmic utility function. This SWF is indeed very commonly used in economic models of climate change (DeCanio, 2003, Table 2.4), and it can in principle be modified to accommodate different degrees of risk aversion.

A potential disadvantage associated with non-monetary SWFs, such as DU_{DICE}, is that they are expressed in arbitrary utility units. As a result, *"Any economist doing this work will obviously feel a strong urge to discount the difference in the consumption streams to a present value"* (Lind & Schuler, 1998, p. 80). Following this "strong urge", some analysts have attempted to convert utility differences into monetary costs, defined as the difference in present value between alternative policies (see Sect. 5.2). However, the combination of different welfare metrics in an analysis is likely to introduce inconsistencies because different welfare metrics aggregate differently across components of economic output (Sect. 4.1), across time (Sect. 4.2), across possible states of the world (Sect. 4.3), and across regions and population groups (not discussed in this paper). Analysts who nevertheless combine different welfare metrics (*e.g.*, by maximizing expected utility and reporting the present value of the certainty equivalents of consumption over time) need to demonstrate convincingly that the inconsistencies between these metrics do not affect the policy conclusions of the analysis. In the absence of such a demonstration, the analysis must be regarded as potentially inconsistent.

5.2 Review of recent climate policy analyses

In this section, we review the use of welfare metrics in two recent applications of DICE that have not followed the recommendations made in Sect. 5.1.

Yohe *et al.* (2004) presents a hedging analysis that aims to identify the optimum short-term policy under uncertainty about climate change and the long-term stabilization target. This uncertainty is described by discrete 'policy cases', which are characterized by a specific value for the climate sensitivity and an upper bound for the greenhouse gas (GHG) concentration level. Each policy case is assigned a probability based on an empirical probability density function (PDF) for climate sensitivity, assuming that all considered GHG stabilization levels are equally likely. It is further assumed that the 'true' policy case will be revealed in 2035. A modified version of DICE-99 is used to determine the 'optimal' decision strategy for each policy case by maximizing DU_{DICE} for different initial levels of the carbon tax (until 2035) and without such a constraint. For each of those utility-maximizing strategies, the discounted gross world product (GWP) is calculated according to PVO_{Yohe}. The "discounted adjustment costs" for each policy case and initial carbon tax level are then defined as the difference in discounted GWP between the utility-optimal strategies with and without prescribing the initial carbon tax level. Finally, the "optimal" initial carbon tax level is determined by minimizing the *expected* discounted adjustment costs for each tax level, considering the probability of the various policy cases.

We argue that the use of welfare metrics in Yohe *et al.* (2004) involves several inconsistencies, with important implications for the results presented. Our first argument recalls the findings from Sect. 3.2 that PVO_{Yohe} is internally inconsistent and underestimates the undiscounted welfare differences between alternative policies by about two orders of magnitude. Hence, Yohe *et al.* (2004) violate the second recommendation from Sect. 5.1. Fig. 4, which reproduces two diagrams from Yohe *et al.* (2004), shows the relevance of this problem for practical climate policy analysis. The discounted GWP difference between a 400 ppm and a 900 ppm CO_2 concentration target in this figure is only 0.025% (left panel), which is about two orders of magnitude smaller than the cost estimates from most other studies (Metz *et al.* , 2001). Furthermore, variation in expected discounted GWP across all considered tax levels is a mere 0.0004% (right panel). We argue that the internal inconsistencies associated with PVO_{Yohe} are the main reason for the surprisingly low variation in discounted GWP depicted in Fig. 4. Our second argument recalls the findings from Sect. 4.1 that DU_{DICE} and PVO_{ex} may produce inconsistent policy rankings. In Yohe *et al.* (2004), DU_{DICE} is initially maximized but PVO_{Yohe} is later used as the basis for selecting the 'optimal' policy for each policy case. In violation of the third recommendation, there

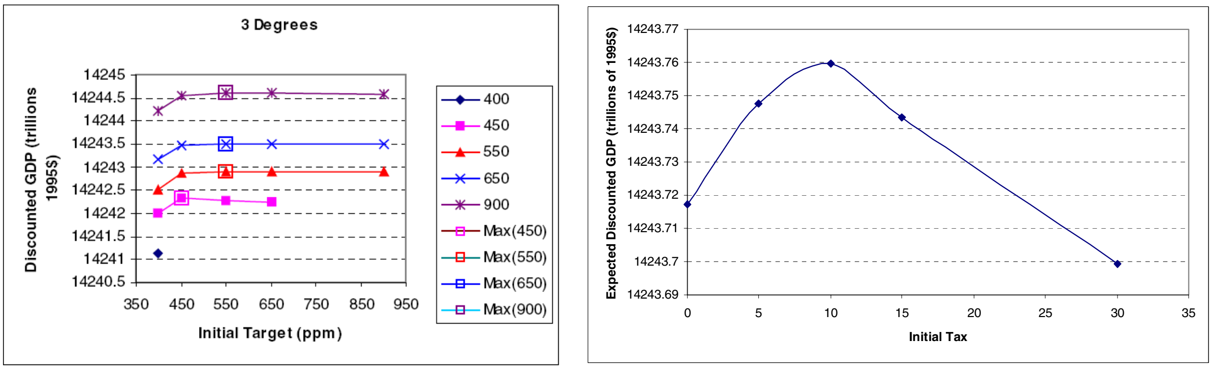

Figure 4: Discounted GWP for a range of stabilization targets and tax levels (reprinted from Yohe *et al.* , 2004). *Left:* Discounted GWP for a range of greenhouse gas stabilization targets. *Right:* Expected value of discounted GWP for a range of initial carbon tax levels.

is no discussion whether these two SWFs produce similar rankings for the policies considered in that analysis, or what the potential implications of the inconsistent rankings could be. Our third argument recalls the findings from Sect. 4.3 that expected consumption and expected logarithmic utility may produce inconsistent policy rankings due to different degrees of risk aversion. The same arguments hold in relation to Yohe *et al.* (2004), where expected output is maximized within a limited set of policy strategies that were initially determined by utility maximization in a deterministic context.

What are the implications of these flaws for the policy conclusions reported in Yohe *et al.* (2004)? The study concludes that *"An initial \$10 tax policy is remarkably robust across the remaining possibilities"*, noting further that it is *"surprising that climate insurance over the near term can be so inexpensive and that an economically efficient near-term hedging policy can be so robust across a wide range of futures in comparison with doing nothing"*. We have argued above that the reported costs of policies depicted in Fig. 4 are incorrect, most likely by about two orders of magnitude. Furthermore, even if the values depicted in Fig. 4 were correct, the tiny GWP variation across different policies would hardly support such a strong conclusion. Determining the correct optimal carbon tax level (subject to the assumptions of this particular analysis) would require a rerun of the whole modelling exercise in accordance with the recommendations from Sect. 5.1. Since such a reanalysis is beyond the scope of this paper, we cannot say for sure whether its results would still support the conclusions cited above.

Fankhauser & Tol (2005) apply DICE-94 to compare indirect climate impacts under different assumptions. This analysis defines several modifications to the production function of DICE-94, determines the optimal decision strategy for each model variant by maximizing the standard discounted utility function of DICE-94, $\mathrm{DU}_{\mathrm{DICE}}$, and presents the time paths and growth rates of undiscounted GDP (corresponding to $\mathrm{PVO}_{\mathrm{ex}}$ for $\theta = \rho = 0$) for these decision strategies. Given the findings from Sect. 4.1 that $\mathrm{DU}_{\mathrm{DICE}}$ and $\mathrm{PVO}_{\mathrm{ex}}$ may produce inconsistent policy rankings, we have to consider the results of Fankhauser & Tol (2005) as potentially flawed. We also note that the endogenous determination of population growth suggested but not performed in Fankhauser & Tol (2005) would violate the last recommendation from Sect. 5.1 and introduce another inconsistency to the analysis (see Sect. 3.3).

6 Summary and Conclusions

This paper reviews social welfare functions (SWFs) that have been applied in welfare-maximizing climate-economy models and discusses the methodological questions and policy implications that arise from the findings of this review.

The review of individual SWFs in Sect. 3 found several SWFs to be internally inconsistent. In particular, different implementations of growth discounting in monetary SWFs can lead to cost estimates for achieving given climate policy targets that differ by several orders of magnitude. The question of appropriate discounting schemes for determining monetary welfare differences between different climate policies has not received attention in the pertinent literature despite the intense discussion about appropriate discount rate(s) in the climate change context. The most important conclusion from the discussion here is that the same discount factors should be used across all climate policies under consideration. However, the close link with the index number problem implies that there is no single 'correct' method for comparing the present values of alternative climate policies (or other long-term policies involving significantly different economic trajectories).

The comparison of SWFs in Sect. 4 shows that the various internally consistent SWFs are generally not interchangeable since they aggregate differently across time, population groups, states of the world, and components of economic output. Several climate policy analyses have uncritically translated non-monetary welfare differences into monetary units. However, the various inconsistencies between different SWFs imply that there is no generally applicable method for doing so, and that the combination of different SWFs in an analysis can lead to erroneous results. In this context, it has been suggested that considerations of risk can be treated by converting outcomes into certainty equivalents and discounting these certainty equivalents. We show that there exists no discounting scheme for certainty equivalents such that the ranking of alternative policies according to the present value of certainty equivalents is consistent with the ranking according to expected discounted logarithmic utility.

Sect. 5 discusses the implications of our findings for welfare-maximizing climate policy analysis. We present several recommendations for the application of SWFs in climate policy analysis that would avoid the inconsistencies identified in the previous sections. A review of recent analyses with the DICE model that have not followed these recommendations finds key results to be artefacts of the choice of flawed SWFs or of their inconsistent combination.

The various inconsistencies identified in this paper are not only of theoretical interest. Each of them can strongly affect the policy recommendations drawn from a particular analysis. Hence, our findings indicate that the application of SWFs in climate policy analyses requires considerably more caution than has been exercised in the past. While this paper focuses on logical inconsistencies associated with the application of different SWFs in climate policy analysis, it is important to note that there are many other conceptual, empirical, and philosophical problems associated with the identification of 'optimal' climate policies by global welfare maximization.

Acknowledgements

I am grateful to Gary Yohe for his willingness to make the code of the DICE-99 model applied in Yohe *et al.* (2004) available, and to discuss several aspects of the work presented in this paper. Stephen DeCanio, Stephane Hallegatte, Charles F. Mason, Mike Mastrandrea, and Jürgen Scheffran have provided helpful comments on an earlier version of this paper.

The work presented here has been funded by a Marie Curie Outgoing International Fellowship of the European Commission within the Sixth Framework Programme for Research.

References

Arrow, K. J. 1951. *Social Choice and Individual Values.* New Haven, CT: Yale University Press.

Arrow, K. J., & Lind, R. C. 1970. Uncertainty and the Evaluation of Public Investment Decisions. *American Economic Review*, **60**, 364–378.

Arrow, K. J., Cline, W. R., Mäler, K.-G., Munasinghe, M., Squitieri, R., & Stiglitz, J. E. 1996. Intertemporal equity, discounting, and economic efficiency. *Chap. 4 of:* Bruce, J. P., Lee, H., & Haites, E. F. (eds), *Climate Change 1995: Economic and Social Dimensions of Climate Change.* Cambridge, UK: Cambridge University Press.

Azar, C., & Lindgren, K. 2003. Catastrophic Events and Stochastic Cost–benefit Analysis of Climate Change. *Climatic Change*, **56**, 245–255.

Cline, W. R. 1992. *The Economics of Global Warming.* Washington, CD: Institute for International Economics.

DeCanio, S. J. 2003. *Economic Models of Climate Change.* New York, NY: Palgrave Macmillan.

Dixit, A. K., & Pindyck, R. S. 1994. *Investment under Uncertainty.* Princeton, NJ: Princeton University Press.

Edgeworth, F. Y. 1888. Some new methods of measuring variations in general prices. *Journal of the Royal Statistical Society*, **51**, 346–368.

Fankhauser, S., & Tol, R. S. J. 2005. On climate change and economic growth. *Resource and Energy Economics*, **27**, 1–17.

Gjerde, J., Grepperud, S., & Kverndokk, S. 1998. *Optimal Climate Policy under the Possibility of a Catastrophe.* Discussion Paper 209. Research Department of Statistics Norway. available at `http://ideas.repec.org/p/ssb/dispap/209.html`.

Gowdy, J. 2005. Toward a new welfare economics for sustainability. *Ecological Economics*, **53**, 211–222.

Hall, D. C., & Behl, R. J. 2005. Integrating economic analysis and the science of climate instability. *Ecological Economics*, [**in press**].

Heal, G. 1997. Discounting and Climate Change. *Climatic Change*, **37**, 335–343.

Howarth, R. B. 2001. Intertemporal social choice and climate stabilization. *International Journal of of Environment and Pollution*, **15**, 386–405.

Howarth, R. B. 2003. Discounting and Uncertainty in Climate Change Policy Analysis. *Land Economics*, **79**, 369–381.

Jevons, W. S. 1865. Variations of Prices and the Value of Currency since 1762. *Journal of the Royal Statistical Society*, **28**, 294–325.

Koopmans, T. 1960. Stational Ordinal Utility and Impatience. *Econometrica*, **28**, 287–309.

Lind, R. C. 1995. Intergenerational equity, discounting, and the role of cost-benefit analysis in evaluating global climate policy. *Energy Policy*, **23**, 379–389.

Lind, R. C., & Schuler, R. E. 1998. Equity and Discounting in Climate-Change Decisions. *Chap. 3, pages 59–96 of:* Nordhaus, W. D. (ed), *Economics and Policy Issues in Climate Change.* Washington, DC: Resources for the Future.

Lind, R. C., Arrow, K. C., Corey, G. R., Dasgupta, P., Sen, A. K., Stauffer, T., Stiglitz, J. E., Stockfisch, J. A., & Wilson, R. 1982. *Discounting for Time and Risk in Energy Policy.* Washington, DC: Resources for the Future.

Metz, B., Davidson, O., Swart, R., & Pan, J. (eds). 2001. *Climate Change 2001: Mitigation.* Cambridge: Cambridge University Press.

Newell, R. G., & Pizer, W. A. 2003. Discounting the distant future: how much do uncertain rates increase valuations? *Journal of Environmental Economics and Management*, **46**, 52–71.

Newell, R. G., & Pizer, W. A. 2004. Uncertain discount rates in climate policy analysis. *Energy Policy*, **32**, 519–529.

Nordhaus, W. D. 1994. *Managing the Global Commons.* Cambridge, MA: MIT Press.

Nordhaus, W. D. 1997. Discounting in Economics and Climate Change. *Climatic Change*, **37**, 315–328.

Nordhaus, W. D. 2001. *Documentation of the DICE-99 Excel spreadsheet.* Available at `http://www.econ.yale.edu/~nordhaus/homepage/dice99doc.doc`.

Nordhaus, W. D., & Boyer, J. 2000. *Warming the World: Economic Models of Global Warming.* Cambridge, MA: MIT Press.

Portney, P. R., & Weyant, J. P. (eds). 1999. *Discounting and Intergenerational Equity.* Washington, DC: Resources for the Future.

Ramsey, F. 1928. A Mathematical Theory of Saving. *Economic Journal*, **38**, 543–559.

Spash, C. L. (ed). 2002. *Greenhouse Economics: Values and Ethics.* New York: Routledge.

Taylor, C. 1982. The diversity of goods. *Pages 129–144 of:* Sen, A., & Williams, B. (eds), *Utilitarianism and Beyond.* Cambridge, UK: Cambridge University Press.

Tol, R. S. J. 1999. Time discounting and optimal emission reduction: an application of FUND. *Climatic Change*, **41**, 351–362.

Tol, R. S. J. 2003. Is the Uncertainty about Climate Change Too Large for Expected Cost-Benefit Analysis? *Climatic Change*, **56**, 265–289.

Toth, F. L. 2000. Intergenerational equity and discounting. *Integrated Assessment*, **1**, 127–136.

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.

Yohe, G., Andronova, N., & Schlesinger, M. 2004. To Hedge or Not Against an Uncertain Climate Future. *Science*, **306**, 416–417. (Includes Supporting Online Material.

Yohe, G. W. 2003. More Trouble for Cost-benefit Analysis. An Editorial Comment. *Climatic Change*, **56**, 235–244.

A Different versions of the Ramsey growth discounting rule

Most economic models of climate change calculate the costs and benefits of alternative climate policies based on the present value of economic output or consumption. In the 'classical' approach to time discounting (see, *e.g.*, Tol, 1999), future discount factors are based on the 'Ramsey rule' for optimal saving (Ramsey, 1928). The 'Ramsey growth discounting rule' is typically stated as

$$r = \rho + \theta \cdot g \tag{25}$$

(see Sect. 2 for an explanation of the parameters). This formulation of the Ramsey rule is, however, only an approximation of the exact solution of the Ramsey model (see, *e.g.*, DeCanio, 2003, Section 3.3.1):

$$\begin{aligned} r &= (1+\rho)\cdot(1+g)^{\theta} - 1 && (26)\\ &\approx (1+\rho)\cdot(1+\theta\cdot g) - 1 \\ &= \rho + \theta\cdot g + \rho\cdot\theta\cdot g && (27)\\ &\approx \rho + \theta\cdot g && (28) \end{aligned}$$

Eq. 26 gives the exact solution under the assumption that all parameters are constant over time. Eq. 26 and Eq. 27 are identical for $\theta = 1$, the standard assumption in integrated assessment models of climate change (see DeCanio, 2003, Table 2.4). Eq. 27 and Eq. 28 are identical for $\rho = 0$, a value that was originally recommended by Ramsey (1928) but that is only occasionally assumed in economic assessments of climate change (for an example, see Yohe *et al.* , 2004). While the differences between Eq. 27 and Eq. 28 are generally negligible, they are relevant in the context of this paper where they explain the difference between the two welfare metrics $\mathrm{PVC}_{\mathrm{DICE}}$ (based on Eq. 27) and $\mathrm{PVC}_{\mathrm{end}}$ (based on Eq. 28).